

## 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.

# Response to the Request for Information on the National Artificial Intelligence Research and Development Strategic Plan

*University of California San Diego and University of Southern California*  
October 26, 2018

## Executive Summary

The current NSTC document is two years old and in a field such as Artificial Intelligence and Machine Learning (AI/ML), significant progress has occurred during this time. In addition, the current NSTC document mainly addresses issues focused on AI/ML labor force considerations and not the key challenges and opportunities directly related to research and delivery of AI/ML solutions. There is a strong need to consider the main drivers to utilization of AI/ML and the associated R&D challenges to deliver such solutions. R&D is required, but to a large degree, it must be driven by societal, economic, and security considerations. How can we secure continued US leadership in AI/ML and its application to ensure continued economic growth for American society and institutions?

The University of California San Diego (UC San Diego) and the University of Southern California (USC) recommend that the NSTC National Artificial Intelligence Research and Development Strategic Plan be revised we offer three main recommendations:

1. **Application Context Is Important:** The AI/ML strategy should be based on a holistic view of the impact of AI/ML across multiple business/societal domains, such as manufacturing, healthcare, defense, finance, etc. These drivers will motivate the required investment.
2. **Human Context Is Important:** Core technologies behind AI/ML advances must rise to meet new requirements related to robustness, security, fairness and explainability introduced largely due to human context in which these technologies are deployed in systems that are making extraordinarily critical decisions related economy and society.
3. **Socio-Economic Context Is Important:** There is a need for a holistic strategy for AI/ML across government, industry and society in a broader sense to remain competitive with respect to other countries such as a China, Canada and Europe. It is essential that the broader societal impact of data security, workforce, ethics and education is considered as part of a national strategy.

It is recommended that the Strategic Plan may be improved to better address three key dimensions:

1. Currency: The rapidity of change and understanding in the field of Artificial Intelligence/Machine Learning (AI/ML) requires a strategy that both reflects current capabilities, but more importantly projects into the coming 5 years.
2. Research and Development of AI/ML: Key societal issues (economic, security, etc.) need to be factored into an R&D investment portfolio that captures social and behavioral dimensions.
3. Delivery and Application of AI/ML: The economical and societal drivers must be addressed in tandem with the associated core technologies that enable the design of state of the art AI systems.

This white paper introduces critical AI/ML application examples and presents the associated core R&D challenges. In addition, a number of key societal issues to be considered are also discussed. This expanded scope reflects the breadth of the AI/ML that should be addressed in this next generation AI Strategic Plan.

## Introduction

UC San Diego and University of Southern California present a set of comments in response to the request for information for the National Science and Technology Council's Strategy for Artificial Intelligence (AI). The current document is nearly two years old, which is considerable in a domain such as AI where we are seeing exponential progress. It is considered essential to have a comprehensive view of AI as it will impact most aspects of our daily lives. A strategy must consider many dimensions, including: what are the application requirements; how can new technology be applied; what are the basic R&D questions to be addressed; how can the US sustain its leadership role in the broader field of AI; and how does progress in AI impact issues such as data security, education, and employment.

There is a proliferation of terminology used with respect to AI. It is important to be clear about the use of these terms to avoid confusion. In this document, AI is considered the overarching theme and machine learning (ML) is considered a key sub-area of AI. Related to both AI and ML, there is also an emergence of Data Science (DS), where data is leveraged as the core to build application systems. An essential part of Data Science are data-sets that form the foundation for the derivation of information/knowledge that is utilized to reach actionable conclusions to be used by users of an AI system.

It is suggested that a revised strategic plan consider a comprehensive view of AI/ML that includes the applications drivers where AI/ML is utilized, and the core technologies that much be addressed to make progress on fundamental AI/ML questions. We will discuss a few application drivers and also some of the core technologies that should be considered.

## Application Drivers

AI is permeating many different areas of application. Some of these major areas include: i) manufacturing; ii) healthcare; iii) environmental monitoring; iv) finance/real estate; v) defense/security; vi) services; vii) policy; and viii) education. Three of these domains are discussed below to highlight how AI/ML impacts these domains.

### Manufacturing

Manufacturing is an essential part of the US economy. Today, US Manufacturing represents about 20% of GDP or about \$2 trillion. To reach this magnitude, the US Manufacturing industry has led several technology revolutions that drove fundamental shifts in manufacturing efficiencies and corresponding economic growth. The 1st revolution was the move from agriculture to factories/use of steam power, the 2nd was the assembly line, and the 3rd was the utilization of automation/computers. The next big push in manufacturing, identified as Industry 4.0 or the 4th industrial revolution, is based on connectivity and AI to streamline production. Today, we are seeing extreme optimization of manufacturing lines through analysis of the flow of components. Some of the most efficient automotive factories have reduced the storage capacity to 12 minutes (i.e., components are mounted in a car within 12 minutes of arriving at the factory). In addition, there is an increasing move from manufacturing-to-storage to manufacturing-on-demand. Amazon manufactures 45% of the books sold after they have been purchased. A local distribution center prints the book and the associated cover and delivers it the next day to the customer without any need for storage or long-range transportation. Industry 4.0 and AI methods are key to ensuring the future of lean manufacturing and supply chain orchestration.

A significant challenge in manufacturing is monitoring of assets to determine predictive maintenance and diagnostics of systems. There is a significant need to develop more mature models as the cost of downtime of a line can be significant. The typical cost of a stop for large-scale manufacturing line is \$50,000/minute.

Close to 90% of US manufacturing is performed in small and medium-sized companies. These companies represent 60% of the US manufacturing industry export. These companies typically have very small production series with a high mix. For such situations, automated programming is essential to success and continued competitiveness. There is a clear need to use AI techniques to automate the programming and to simplify the operation of machines so that it can be performed with a minimum of skills.

Manufacturing is a great example of a domain where there are strong models for products and processes and as such is a great counterexample to the increasing use of data-only models in AI. It is important that both model-based and model-free AI is considered for the future.

## Healthcare

Healthcare is another great example where AI does and can make a major difference. Today, research results are published faster than doctors can read the literature. In addition, the cornerstone of practicing medicine is experience. Through experience, doctors acquire knowledge about symptoms and diagnosis. Through utilization of AI/ML it is possible to improve access to the research literature and assist doctors arrive at a diagnosis. A simple example is a male patient with 36 hours of lower right quadrant pain in combination with fever and elevated white blood cell count is a 90% indication of appendicitis. Once a CT scan has confirmed the diagnosis, the patient is sent to surgery. Through an analysis of electronic medical records and the utilization of literature and other sources, it is slowly becoming possible to provide diagnostic support in radiology, pathology and other diagnostic domains. The use of AI enables utilization of broader data sources to make state of the art care accessible to many more patients. Through the use of this technology improved care can be provided across larger metropolitan areas and remote areas, where there might be limited access to healthcare. The utilization of knowledge-based techniques also has the potential to improve triage and reduce cost through elimination of unnecessary medical tests.

Through the use of AI/ML, it is possible to provide assistance not only in diagnostics, but also in delivery, such as surgery. It is possible to utilize real-time data to provide moment-to-moment decision support to surgeons, which will improve the quality of care and reduce morbidity. The integration of both prior data and real-time data will allow future surgeons to have a level of support that is parallel to the autopilots in the aerospace industry. Certain aspects of a procedure may be automated such as suturing, whereas other parts are directly controlled by the surgeons. In addition,

through the use of imaging, it is possible to assist surgeons in new ways during surgery to improve accuracy and reduce complications.

Finally, the integration of AI/ML with new techniques in augmented/virtual reality will allow for improved training of future doctors to allow them to experience a broader set of procedures and cases before they perform procedures on their own, creating opportunities to do certification using a rich set of tools.

For healthcare, the use of electronic medical records, a rich literature base, and a more diverse set of data sources has tremendous potential to significantly improve the quality of care at a reduced cost.

## Environment

Environmental monitoring has a broad set of applications from weather forecasting to prediction of long-term changes. Increasingly, model-based prediction enables accurate prediction not only for the short-term (minutes/hours), but for longer timeframes (days/months). There is tremendous growth in accessibility to data sources such as satellite images and temperature measurements. The “weather underground system” is an example of a grassroots movement to provide a dense set of weather information, which in turn allows for improved forecasting. Major institutions, such as Scripps Institution of Oceanography at UC San Diego, have more than 3000 sensors deployed to measure sea water temperatures across the world, which are critical for forecasting, but also for strategic applications, such as the impact of temperature variations across depth on submarine sensor systems and navigation. The integration of these diverse data-sources allows prediction of El-Nino, early forecasting of a major weather system that allows for early evacuation, reduction on water usage for crops, etc. Changes in weather and climate have deep economic and societal implications, and it is essential to forecast these changes so that appropriate measures can be taken in a timely and impactful way.

Multi-sensory data has recently become much more prolific, and it is only now that adequate models are available and viable to integrate these data for reliable longer-term forecasting. There is a significant need to consider how more detailed data, stronger models and improved forecasting can be leveraged.

# Core Technologies

It is important to consider the economical and societal drivers, but also the associated core technologies that enable the design of such systems. Figure 1 is an illustration of the relationship between applications and core technologies. As the current NSTC Strategic Plan discusses only a limited set of these technologies, there is a need for a comprehensive view of all the required technologies. As an example, data processing is not only a software challenge, but there is a need to consider the technologies to make them possible in terms of processor technology (from traditional CPUs over GPUs to new types of architectures and quantum computing), storage technology and data communication from local networking to cyber infrastructure.

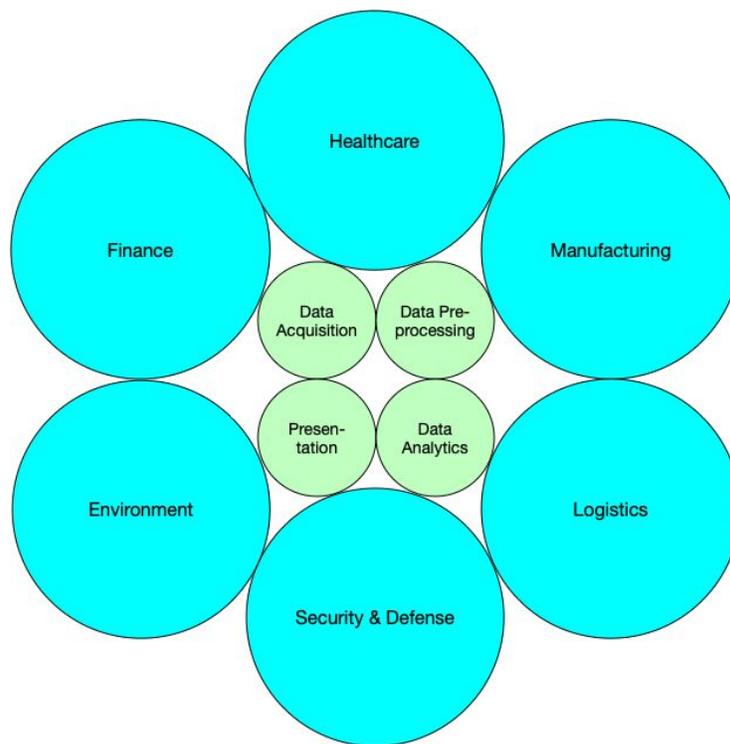


Figure 1. For a strategy on AI it is essential to consider both application drivers (blue circles) and core technologies (green circles) .

## Data Analytics - Natural Language Processing

Natural Language Processing is a great example of data analytics. In many cases, data have associated semantic descriptions that are captured in a textual form. The potential set of sources is very rich such as diagnosis of CT-scans, stock variation, maintenance

records, social media feeds, and news sources. The accurate processing of these systems requires access to major data-sets as it is not only a question of classifying data, but also an issue of analyzing the contextual framework within which it occurs. There is a need to study techniques that allow the transfer of knowledge across domains to enable scale up to a broader set of applications.

Most human knowledge is captured in a textual form through publications, records, notes, etc. It is consequently essential to consider the design of effective strategies for the transformation of these textual data-sets into knowledge bases for future applications.

As an example, it is possible to use learning to classify radiology reports according to diagnosis. Once generated, these diagnosis labels can be used for training an image classifier for the associated radiology scans. The combined system has an improved performance and is more cost effective for technician training. Labeling texts is generally simpler than labeling other types of data (i.e., it is simpler to read a radiology report than an analysis of radiology scans). Natural Language Processing has tremendous potential for utilization across a wide range of applications.

## Presentation - Human-Computer Interaction

Data analysis is essential to AI systems. Equally important is how to present the actionable results of a system to users of the system whether machine or human. Particularly for collaboration with humans, there is a need to have efficient techniques for the presentation of the data, and in many cases, the basis for the results. As an example in medical diagnostics, the indication of appendicitis should be presented with the associated data that led to the conclusion, or similarly, a stockbroker would want to know why the acquisition of a stock is recommended. Contextualizing the recommendation is essential to broader acceptance.

Another important application of AI is in education, where adaptive systems have significant untapped potential. As an example, it is possible to automatically determine the affective state of a student - motivated, frustrated, engaged, etc. Using these systems, it is possible to determine if the brain is in a good state for learning. Using game theoretic models, it is possible to modulate the interaction to engage the student such that they are much more motivated to learn new content.

## Presentation - Cyber-Physical Systems

Cyber-physical (CP) systems utilize a deeply embedded plane of distributed computing, communications and storage fabric to make our societal infrastructure for transportation, work, energy and environment more resilient, efficient and useful. The grand challenge of this domain is accurate and timely prediction of CP system state and its effect on humans. Uncertainty -- of inadequate physics and humans -- must be part of the reasoning framework used to make predictions. Sensing and contextual sociological data holds the promise of filling the gaps left by inadequate physical models that are either too expensive or too time consuming to construct. Instead, sensing or perturbational sensing can be used to arrive at estimates of current state and for improving predictions. There are a number of technical challenges to creating data-enabled autonomously intelligent CP systems. Precise and inexpensive localization in real-time is one of the essential ingredients to contextual awareness of humans and assets. Automated construction of data-models of interacting human-CP systems requires advances in sensory data analytics, in-network summarization and automated hypothesis formulation and testing. Sound methods for experimental design, especially those formulated for use by computational methods and tools, are needed.

## Annotated Data-Sets

The paucity of large datasets along with the necessary metadata is a fundamental roadblock for developing and testing AI/learning methods. Especially, even when there are large data sets, there is little or no metadata associated, which makes validation of any learning impossible. The National Institutes of Health has begun to recognize this fundamental problem, especially in the large scale measurements on human subjects, and have initiated an initiative called Data Commons. Any national initiative such as NSTC should embrace the NIH initiative and make available such large data sets along with metadata publicly using the FAIR (findable, accessible, interoperable, reusable) principle. Several such data sets are available in UCSD (the Signaling Gateway, the LIPIDMAPS gateway, the Metabolomics Workbench, etc.).

## Recommendations

The current NSTC strategy document has a strong focus on the impact of AI on workforce, data security, etc. It is recommended that a comprehensive view of AI be adopted for a revision of the strategy. The strategy should not only consider AI on its own contextualize it in terms of its application domains.

It is recommended that the strategy consider how the utilization of AI/ML will impact across different areas of business, such as manufacturing, healthcare, logistics, finance, and real-estate. It is anticipated that most areas of society and business will be impacted within the next few years. At the same time, these areas all provide key constraints, data and resources to contextualize R&D in artificial intelligence.

It is recommended that a revised strategy provides a unified US strategy for investment in AI/ML. Other countries/regions such as China, Canada, and Europe are making major investments in AI/ML. In almost all cases, there is a coherent strategy to advance AI/ML. In the US, the strategy is fragmented across multiple government agencies and private companies. It is not clear that there is a federated strategy to ensure that US continues to be the leader in AI/ML. This is of major importance both for industrial growth, as well as from a security/defense perspective.

Obviously, the strategy should also consider how we address essential societal issues, such as data security/ethics, but also the impact on the workforce and the need for continuing education.

Without a comprehensive strategy (and investments), there is a considerable risk that the US will lose its leadership position in AI/ML.