

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.



October 26, 2018

Attn: Faisal D'Souza
NCO
2415 Eisenhower Avenue
Alexandria, VA 22314

Subject: (2018-20914; 83 FR 48655) Request for Information on Update to the 2016 National Artificial Intelligence Research and Development Strategic Plan; Comments of the American College of Radiology

The American College of Radiology (ACR)—a professional organization representing more than 38,000 radiologists, radiation oncologists, interventional radiologists, nuclear medicine physicians, and medical physicists—appreciates the opportunity to respond to the Networking and Information Technology Research and Development (NITRD) National Coordination Office's (NCO's) Request for Information (RFI) on Update to the 2016 National Artificial Intelligence Research and Development Strategic Plan published in the *Federal Register* on September 26, 2018 (document number: 2018-20914; 83 FR 48655). The ACR supports the federal government's efforts to update the National Artificial Intelligence Research and Development Strategic Plan, and we urge continued federal support and collaboration with professional associations and other stakeholders to ensure a safe and efficacious use of this technology.

The following comments are specifically focused on the healthcare domain, and particularly on medical imaging. The feedback was compiled by members of the ACR Data Science Institute, ACR Commission on Research, and ACR Government Relations. Individual contributing members are listed at the end of this submission.

ACR Responses to RFI Topics

Since the release of the 2016 Strategic Plan, the ACR has become involved in educating radiologists about the value of artificial intelligence (AI) applications in healthcare and in

facilitating research, development, and deployment of AI tools that will help radiology professionals improve patient care and add value to the health system. In May 2017 the ACR created the Data Science Institute (DSI) to accomplish these goals. We are leveraging radiologists' expertise in identifying the clinical challenges amenable to AI solutions and promoting development of artificial intelligence algorithms to address those challenges through structured AI use case development and collaborations with researchers, health systems, industry and governmental agencies. DSI's intent is to facilitate the translation of AI research to clinical practice in a manner that protects the safety of patients and the public. ACR's comments reflect the experience and knowledge we have gained in the DSI over the past 18 months with respect to each of the seven strategies in the 2016 Strategic Plan.

Strategy 1: Make Long-Term Investments in AI Research

In terms of long-term investment in AI research, we feel that one of the main deterrents to implementing AI in the healthcare setting is the lack of integrated, scalable systems. As evidence of this, an audience poll at the August 2018 NIH-National Institute for Biomedical Imaging and Bioengineering (NIBIB) workshop on "AI in Medical Imaging" revealed that the majority of attendees were involved in the research or creation of AI algorithms; however, very few were actually using AI in their clinical practices.

While the development of individual AI algorithms related to medical imaging is progressing at a rapid pace, translating this research to routine clinical practice has not yet occurred. We believe that investments in research and infrastructure that enhance the ability to deploy AI in the clinical space are needed. Current long-term federal investment needs include:

- *Standards for clinical integration and care management.* We are concerned that the current pathway of single institutions developing algorithms that are designed to work in their institutions may not be generalizable to routine practice unless they can be implemented across the entirety of existing HIT resources including specialty health IT systems (e.g., Radiology Information Systems and Picture Archiving and Communication Systems) and electronic health record technology.
- *Standards for training data focused on ensuring diversity in training sets.* The goal should be to eliminate unintended bias in the algorithms, thereby increasing generalizability to various populations and settings.
- *Structured use cases that are clinically effective and readily integrated into routine clinical use.* Structured use cases that define: (a) parameters for training data, (b) pathways for validation, and, (c) mechanisms for deployment and monitoring in clinical practice, allow algorithms for a specific clinical purpose to be developed across a variety of institutions and by numerous developers with a standard output than can be consumed by a variety of HIT applications.

- *Standard pathways for validation and monitoring AI use in clinical practice through registry reporting, which is needed to ensure patient safety.* These areas have become a major focus of activity for the ACR DSI. We are taking a leading role in developing standardized AI use cases for the radiological sciences and we are working with other radiological organizations and standards bodies to develop common data elements corresponding with standards for clinical integration and care management that can be used in AI development.

In summary, we believe that the Strategic Plan should promote continued federal investments in infrastructure and tools that promote generalizability of foundational AI research to clinical practice as this will be required to facilitate the use of AI in the healthcare environment.

Strategy 2: Develop Effective Methods for Human-AI Collaboration

While using artificial general intelligence solutions (general AI) for healthcare, (where machines function independently and duplicate the cognitive activity of humans) may be a laudable goal for the distant future, thus far, general AI solutions have not produced usable healthcare tools. Whether it is lack of algorithm explicability or the overall poor quality and unstructured nature of healthcare data for algorithm development, we do not believe general AI solutions are ready for use in healthcare, nor do we believe that strong support of general AI should be a major priority of the Strategic Plan at this time. It can be argued that in instances where general AI has been applied to healthcare, the results have actually hindered the human-AI collaboration rather than promoting it.

Explanations for these failures may include the overall quality of healthcare data for use in AI and lack of explicability of algorithm inference, among others. The poor results of these projects indicate that the Strategic Plan should focus on federal support for the development of narrow AI solutions to help physicians solve specific healthcare challenges. Physicians need AI tools that can be used in tandem with clinical practice and produce outputs within clinical workflows that are readily explicable, so that the final decisions are made by physicians, not the machines. AI algorithms for healthcare need to be designed to transparently solve specific tasks (narrow AI) in explicable ways (explicable AI) so that AI and human intelligence are combined.

As such we believe the Strategic Plan should focus primarily on promoting narrow and explicable AI solutions in the near term (10 to 20 years) until general AI solutions for healthcare become more reliable and/or clinically usable and government agencies are able and positioned to provide appropriate regulatory controls.

Strategy 3: Understand and Address the Ethical, Legal, and Societal Implications of AI

Ensuring the ethical use of AI in healthcare should continue to be a focus within the Strategic Plan. Key areas of data ethics include informed consent for data use, privacy and data protection, ownership of patient data, objectivity, and the gap between those who have and those who lack the resources to manage and analyze large datasets. Other issues include mitigating bias against group-level subsets of individuals such as women, specific ethnic or economic groups. The importance of trust in assessing data ethics and providing meaningful access rights to individual patients should also be addressed.

Numerous reports have shown AI algorithms can generate or amplify bias and that developers can inadvertently introduce ethnic, racial, geographic, or economic bias into algorithm development when single institutional data for algorithm training developers are used for its development. Structured use cases for AI algorithm development that include data elements defining datasets for training and testing, such as those being developed at the ACR DSI, will allow multiple institutions to create these datasets ensuring researchers and developers will have access to technically, geographically and demographically diverse data for AI development. **The Strategic Plan should promote solutions that make diverse data from multiple institutions widely available to developers for training and testing AI algorithms to ensure the algorithms perform as expected for the population as a whole, and not merely subsets of the population, which could lead to algorithm bias in clinical practice.^{1,2}**

Strategy 4: Ensure the Safety and Security of AI Systems

Ensuring the safety and security of AI systems continues to be an important strategy that requires focus on several fronts. In addition to the strategies to diminish algorithm bias and promote explicability through saliency maps and other tools discussed above, we believe the Strategic Plan should also focus on promoting sound mechanisms for data anonymization. Currently there is no single mechanism for anonymization and developers are increasingly concerned that by not using a standardized mechanism for anonymization, they may be at risk if confidential patient data were exposed. Patients should be made aware that their personal health data may be used in AI research and development. **Therefore, the Strategic Plan should promote the development of a**

¹ Mittelstadt, B.D., Floridi, L. The ethics of big data: current and foreseeable issues in biomedical contexts. *Sci Eng Ethics*. 2016;22:303–341

² Ethics, Artificial Intelligence, and Radiology. Kohli, Marc, Geis, Raym. *Journal of the American College of Radiology*, Volume 15 , Issue 9 , 1317 - 1319

standard method for anonymization and data use agreements/standards to protect patient privacy.

Consideration should also be given to how AI algorithms will be used in clinical practice. Whether the majority of solutions will be on-premises or cloud-based (or more likely a hybrid) is yet to be determined, but when AI algorithms interact with other HIT tools there is a risk for data breaches including exposing patient information. While standard APIs are needed so that all of the tools will be able to interact with each other, **the Strategic Plan should promote pathways for standardization with mechanisms to prevent exposure of confidential information as data are shared among various networks.**

Strategy 5: Develop Shared Public Datasets and Environments for AI Training and Testing

When polling developers, the ACR DSI found that one of the top three reasons why AI technology may have difficulty becoming embedded in routine clinical practice is that large annotated datasets are difficult to create. In a clinical setting, developers have difficulty in obtaining annotated datasets for training and testing due to patient privacy issues and reluctance of health systems to allow patient information to leave their premises. Making structured data available for AI algorithm testing and training is a critical need for the developer community. Current models of making loosely annotated datasets publicly available are limited. Using unstructured data for algorithm training and testing can diminish algorithm explicability and potentially introduce bias. Since a significant amount of work goes into creating annotated datasets for algorithm training and testing, many developers have demanded exclusivity agreements for those producing annotated datasets for AI development. We recognize that intellectual property rights needs to be recognized; however, we also believe there should be mechanisms to make these annotated data sets more widely available by developing federally-supported strategies to mitigate the costs of AI training and testing through data sharing. **The Strategic Plan should continue to prioritize the development of publicly available, richly annotated datasets.**

Strategy 6: Measure and Evaluate AI Technologies through Standards and Benchmarks

Before enabling widespread use of AI algorithms, there should be federally mandated and/or supported processes in place to ensure artificial intelligence algorithms are safe, reliable, and effective. This process begins with the creation of structured AI use cases that include data elements which allow measurement of an algorithm's performance in a validation environment. Based on these data elements, institutions can develop validation

datasets (which have not been used in algorithm training) to assess algorithm performance according to the metrics specified in the use case. These validation data sets should be created around a strict set of standards for data quality to ensure ground truth consistency between sites, and standard metrics for measuring performance. Multi-institution datasets should be used for this purpose as they contain geographic, technical, and patient demographic that will help ensure algorithms are generalizable to widespread clinical deployment and free of unintended bias. Independent honest brokers, such as medical specialty societies, can play an active role in the curation of these embargoed datasets and in validation of the algorithms according to the metrics define in the use cases. Validation reports can be provided to the developers ahead of FDA or other regulatory review before clinical deployment.

Continued monitoring of an algorithm's effectiveness in clinical practice will also be important to ensure patient safety and reliability of the algorithm in a variety of practice settings and the adaptability of the algorithm to changes in equipment or protocols. It is important in diagnostic imaging to have an infrastructure that captures the algorithm's performance in clinical practice as well as specified metadata about the exam, such as the equipment manufacturer, the field strength or number of detectors, and other relevant examination parameters. AI use cases can define parameters that should be monitored in clinical practice. Algorithm performance – as judged by the end user radiologists – can be captured real-time through the reporting software or other means, stored in a clinical data registry and aggregated nationally. Once the data are aggregated, performance reports can be developed and provided to the developer, the FDA, or other regulatory bodies for defining areas where the algorithm might fail to perform as expected. These reports can be used for product labelling, clinician notifications, and subsequent algorithm improvement.

The ACR DSI is working with the FDA to test the viability of both the validation and monitoring processes in clinical practice with positive initial results. **The Strategic Plan should continue to support the independent validation of AI technologies prior to deployment and the use of real-world data through clinical data registry data reporting to monitor its performance in routine clinical use.**

Standards for interoperability are also crucial to the development of a viable AI ecosystem. The development of structured use cases for AI requires the use of common data elements (CDEs) to define a multitude of parameters in a standardized machine-readable format. CDEs define the attributes and allowable values of a unit of information, so that information can be collected and stored uniformly across institutions and studies. The CDEs are defined in a data dictionary, which specifies attributes including the item's name, the way the item is collected, valid values and coding, and data type (e.g., number or text). The radiology community, having recognized this need, is collaborating to develop and house a compendium of these data elements for diagnostic radiology (ACR and Radiological Society of North America (RSNA)

RadElements project) to optimize clinical interoperability and implementation of AI tools by ensuring standardization of the input and output elements of the algorithms.

In addition, standard methods of transferring the data will also be important to allow algorithms from many different developers to be deployed quickly and information to be exchanged in a consistent, interoperable manner. If such existing standards can be modified and used for this purpose, they will undoubtedly be adopted more readily and easily.

The ACR believes **the Strategic Plan should continue to vigorously promote and support the development of standards including common data elements, many of which may be novel specifically for AI algorithm development, to ensure interoperability and seamless workflow integration of AI tools into routine clinical practice.**

Strategy 7: Better Understand the National AI R&D Workforce Needs

From the ACR's perspective, it is too early to adequately understand and predict the implications of AI on the healthcare workforce. While AI will certainly improve quality and efficiency, the implications for human resources are less clear. Likewise, AI is poised to bring more data to bear on individual patients' problems; how that information will be consumed by the clinical team is unclear. While some have predicted a drastic reduction in the need for some physician specialists—including radiologists—actual use in clinical diagnostic imaging is showing that AI requires *more* radiologist time to review and incorporate AI inferences into clinical reporting of imaging examinations.

It is also likely that, as AI advances, the role of some physician specialists may change as their tools automate some of the necessary, but potentially repetitive and more mundane physician tasks. Physicians could then potentially use this time to allocate their specialized expertise to other areas of need.

The ACR is monitoring and educating our members regarding how AI will be used in clinical practice and will work to ensure radiology professionals have the knowledge and skill necessary to adapt to and implement data science solutions in their practices.

We believe **the Strategic Plan should continue to monitor workforce issues in the healthcare domain and the changing practice roles of physicians and other healthcare professionals.**

Sincerely,

Cynthia Moran
Executive Vice President
Government Relations, Economics and Health Policy
American College of Radiology

Contributors:

ACR Data Science Institute
Keith Dreyer, DO, PhD, FACR
Bibb Allen, MD, FACR
Mike Tilkin, ACR CIO & EVP
Laura Coombs, PhD
Chris Treml, ACR EVP

ACR Commission on Research
Etta Pisano, MD, FACR
Charlie Apgar

ACR Government Relations
Gloria Romanelli, JD
Mike Peters