



CODE



BRIEFING PAPER

Sharing and Utilizing Health Data for AI Applications

April 2019

Table of Contents

- EXECUTIVE SUMMARY 1**
- INTRODUCTION 2**
- BACKGROUND AND TERMINOLOGY 3**
 - Machine Learning, Deep Learning & Natural Language Processing 3
 - Algorithms 3
 - Augmented Intelligence 3
 - Data 4
- EMERGING APPLICATIONS OF AI IN HEALTHCARE 5**
 - Deep Learning for Diagnostics 5
 - Large-Scale Predictive Analytics 5
 - Electronic Health Record (EHR) Management 6
 - Risk Stratification 6
- EMERGING APPLICATIONS OF AI WITHIN HHS 7**
 - HHS Supporting AI Applications in Healthcare 7
 - Improving Administrative Processes and Addressing Fraud 8
- DATA REQUIREMENTS, ACCOUNTABILITY, AND ETHICS 10**
 - Data Quality & Access 10
 - Data Privacy & Security 10
 - Accuracy 10
 - Accountability & Ethics 11
- NATIONAL POLICY CONTEXT AND INTERNATIONAL AI STRATEGIES 12**
 - Towards a Department-Wide Strategy at HHS 12
 - The State of Data Sharing at HHS 12
 - International AI Strategies and Principles 13
- NEXT STEPS 15**
- REFERENCES 16**

EXECUTIVE SUMMARY

Artificial Intelligence (AI) is being increasingly adopted across a variety of sectors in the United States, becoming a major priority for the U.S. Government. The healthcare industry has been one of the leading adopters of AI, as healthcare providers have applied AI to improve patient outcomes, increase organizational efficiency, and reduce costs. Much of this growth in AI can be attributed to the expanding availability and sharing of health data. The Office of the Chief Technology Officer (CTO) at the U.S. Department of Health and Human Services (HHS) is hosting a **Roundtable on Sharing and Utilizing Health Data for AI Applications** on April 16, 2019. This Roundtable will address priorities for applying AI at HHS and opportunities for HHS to enable the continued growth of AI in the healthcare sector.

This Briefing Paper serves as a primer for participants attending that Roundtable, and is divided into the following sections:

- Introduction (pg. 3)
- Background and terminology (pg. 3-4)
- Emerging applications of AI in healthcare (pg. 4-6)
- Emerging applications of AI within HHS (pg. 6-8)
- Data requirements, accountability, and ethics (pg. 8-10)
- National policy context and international AI strategies (pg. 10-12)

This briefing paper incorporates journal articles, blog posts, and news articles that feature overarching trends, applications, and challenges. This document is not intended to provide a complete overview of all AI applications and research in the healthcare space, but highlights useful background information for Roundtable participants that may help to frame the day's discussions.

INTRODUCTION

Artificial Intelligence (AI) has the ability to help transform health care by improving diagnosis, treatment, and the provision of patient care. Much of this progress depends on sharing and utilizing large amounts of health data. While the private sector has driven much of the innovation in this field, the federal government can play a major role. As a major federal data steward, the U.S. Department of Health and Human Services (HHS) can support this transformation by enabling the application of AI inside and outside of government.

The HHS Office of the Chief Technology Officer (CTO) is now exploring the potential for a department-wide AI strategy to help realize these opportunities, and to establish policies and practices for facilitating AI development. This strategy comes in tandem with the White House Executive Order on Artificial Intelligence and the “State of Data Sharing at the US Department of Health and Human Services Report” published in September 2018. This Roundtable will bring together HHS leaders, and experts in AI and health data from other federal and state government agencies, industry, academia, and patient-centered research organizations. Together, they will identify high priority health applications of AI and the key issues for an HHS AI strategy to address. The Roundtable will inform an HHS AI strategic plan in two ways:

- 1. Priorities for Applying AI within HHS:** AI approaches to help HHS manage its own data, facilitate HHS research, or help HHS achieve its mission in other ways.
- 2. Opportunities to Support External AI Development:** Activities to help support data-driven AI applications in industry, academia, and research institutions, such as the release of specific datasets, issuing AI challenges, and other actions.

This Briefing Paper is not a complete overview of all AI applications and research in the healthcare space, but highlights useful background information for Roundtable participants.

BACKGROUND AND TERMINOLOGY

Artificial Intelligence denotes the process by which a computer can be trained to successfully perform tasks that traditionally have been conducted by humans. This includes learning how to perform certain tasks either under supervision or autonomously using large quantities of data.ⁱ

Machine Learning, Deep Learning & Natural Language Processing

Machine Learning is the process by which computers are trained to 'learn' by exposing them to data. Machine learning is a subset of AI and Deep Learning is a further subset of machine learning. Deep Learning is the process by which algorithms can learn to identify hierarchies within data that allow for truly complex understandings of data. Natural Language Processing (NLP) refers to the subfield of machine learning designed to allow computers to examine, extract, and interpret data that is structured within a language.

Algorithms

Most AI applications depend on algorithms, which describe a logical process that follows a set of rules. Computers can be taught a series of steps in order to process large amounts of data to produce a desired outcome. There are two forms of algorithm:

- 1. Supervised algorithms** use 'training data sets' in which the input factors and output are known in advance. Supervised algorithms produce highly accurate algorithms because the 'right answers' are already known. For example, scientists may feed a dataset of retina images into the algorithm in which board-certified physicians have already identified and agreed upon diagnoses for each image.
- 2. Unsupervised algorithms** refer to a process whereby data is fed into the algorithm and the computer has to 'learn' what to look for. Unlike the training data sets fed into supervised algorithms, the data fed into unsupervised algorithms do not necessarily include the 'right answers.' Unsupervised algorithms are adept at finding clusters of relationships between observations in the data, but may identify erroneous relationships because they are not instructed what to look for.

ⁱ There are a variety of artificial intelligence definitions but for the purposes of this briefing paper, we define Artificial intelligence as the ability of a computer to be trained to conduct tasks that are normally done by a human. Please consult the References section for more information on additional AI definitions.

Augmented Intelligence

Augmented Intelligence is a form of AI that enhances human capabilities rather than replaces physicians and healthcare providers. Augmented Intelligence has been embraced by physician organizations to underscore that emerging AI systems are designed to aid humans in clinical decision-making, implementation, and administration to scale health care.¹ In a recent White Paper, Intel framed Augmented Intelligence as the AI tools that perform specific tasks and are designed to support users, rather than replace human users.²

Data

At the core of AI is the need for high quality data. Algorithms are designed around data, which underscores the importance of providing clean and accurate data. Researchers emphasize the need for big, multifaceted datasets that allow machine learning processes to incorporate as many factors as possible into analysis.³ Doing so will allow researchers to better understand their data, which in turn allows them to improve their training of the algorithm. Artificial intelligence also demands clear, accountable data governance with defined data elements and processes for ensuring data quality and access.⁴ Researchers now are attempting to funnel large troves of health data - from electronic health records (EHRs) to data collected from wearable devices and sensors - to improve diagnostics and predictive analytics. More connected and interoperable data in the healthcare system will enable more transformative AI applications in the future.

EMERGING APPLICATIONS OF AI IN HEALTHCARE

Deep Learning for Diagnostics

One of the more promising avenues for progress in healthcare is the application of AI for assisting physicians in accurately diagnosing potentially tricky medical conditions.

- **Diagnosing Cancers:** High-capacity computing has led to a reduction of error in the diagnosis of breast cancer by 85 percent.⁵ High powered computing, in combination with intelligent AI software, allows physicians to better identify metastasized cancerous lymph nodes, and will support future researchers to diagnose and treat oncology patients early.
- **Diabetic Retinopathy:** AI has been used to learn how to diagnose diabetic retinopathy with high levels of accuracy. Diabetic retinopathy is the world's leading cause of blindness, and researchers at Google have trained algorithms to analyze images of retinas and diagnose with over 90 percent accuracy.⁶
- **Cardiology:** Some cardiologists have integrated AI into their work, and recent applications of cardiology-oriented AI have passed medical exams when diagnosing echocardiograms.⁷

As diagnostic tools, AI algorithms are emerging as powerful partners for physicians supplementing their own diagnoses with an immediately-available 'second opinion.'

Large-Scale Predictive Analytics

In addition to diagnosing active problems, AI and big data can be used to engage in predictive analytics designed to reduce infections and hospital readmissions.⁸ Below is a short list of some of the most notable use cases that apply predictive analytics:

- **Diagnosing Sepsis:** Diagnosing and treating sepsis is extremely difficult. Recent research by Imperial College London demonstrated an AI application where algorithms learned how to best treat individuals with sepsis by analyzing data from many medical records and treatment plans rather than relying on the data found within the individual's EHRs.⁹ This tool, known as the AI Clinician, was the result of collaboration between five Intensive Care Units (ICUs) and over 17,000 observed admissions for sepsis in the United States.¹⁰ In a recent trial, research showed that patients treated by a combination of medical doctors and AI algorithms had the lowest mortality rates of all sepsis patients within the study.¹¹
- **Avoiding Hospital Infections:** Health Catalyst has focused on helping hospitals avoid infections that patients acquire while in the hospital.¹² Another private firm, H2O, uses predictive analytics to identify when patients should be transferred to and from intensive care units (ICUs), predict hospital acquired infections (HAI), and prevent hospital readmissions by comparing individual-level data against a large dataset.¹³

- **Treatment Paths:** Recent work by AI company Ayasdi produced their Clinical Variation Management tool, which developed a set of clusters of treatment paths to make recommendations about optimal treatment paths for future patients.¹⁴

Electronic Health Record (EHR) Management

Health providers and researchers use NLP and other AI applications to better organize and analyze electronic health records (EHRs). For example, researchers at UC San Francisco, Stanford Medicine, and The University of Chicago Medicine, along with the Google AI team, published “Scalable and Accurate Deep Learning with Electronic Health Records” in *Nature Partner Journals: Digital Medicine*.¹⁵ The researchers propose integrating EHRs into the existing Fast Healthcare Interoperability Resources (FHIR) standard that was developed by HL7, a healthcare IT standards organization.¹⁶ The researchers demonstrate that by automatically integrating EHRs into the FHIR framework using machine learning, the AI application could make predictions about medical events on a per-patient basis with 93 percent accuracy.

Risk Stratification

AI-driven risk stratification has already garnered significant attention in the research community and the private sector. According to Eric Just, the Senior Vice President of Product Development at Health Catalyst, “[Accountable care organizations] have to be able to pinpoint which heart failure patients are at high risk for readmission as well as for sudden cardiac arrest. Armed with this knowledge, clinicians can schedule follow-up appointments and ensure those patients understand their medications and other aspects of the care plan.”¹⁷ Within the past few years, technology companies have released AI-driven risk stratification tools including:

- **Lumiata’s Risk Matrix:** This tool is trained on 175 million patient records to produce a comprehensive evaluation of patient risk of disease and medical condition.¹⁸ Recent research by De Beule et al. shows that AI can classify patients into ‘low risk’ categories with nearly 90 percent accuracy, helping reduce the likelihood of future illness or complications, and to facilitate recovery.¹⁹
- **NLP Applications:** Amazon Web Services has used NLP to extract and interpret hand-written notes and text from medical records.²⁰ NLP is particularly well-suited to deciphering physician input and patient records, since EHRs do not follow a single, unified structure, yet contain important information for understanding diagnostic trends and risk profiles of individuals.²¹

The implementation of AI, when combined with access to large scale datasets, can vastly improve the ability of the medical community to categorize risk at the individual level. Several private sector companies are leading the path to deciphering unstructured health data.

EMERGING APPLICATIONS OF AI WITHIN HHS

Emerging applications refer to late-stage prototypes and early-stage implementations of AI-driven tools. HHS already has a number of tools and is currently exploring additional applications for both internal and external use, including streamlining administrative processes and improving fraud detection. These are detailed below.

HHS Supporting AI Applications in Healthcare

HHS Tech Sprint for Health Innovation: HHS recently completed a “tech sprint” engaging outside teams on the use of AI to match patients to appropriate clinical trials, improving physicians’ ability to critically test health solutions.²² The challenge was modeled on The Opportunity Project (TOP), a government project that uses agile development methods to build digital tools with federal data. In the 14-week sprint led by the HHS Office of the CTO, tech teams developed seven AI-driven solutions that use federal data to “improve clinical trials, experimental therapies, and data-driven solutions for complex challenges from cancer to Lyme and tick-borne diseases.”

The Centers for Disease Control and Prevention (CDC) Workbench Web Services: The CDC in partnership with The Food and Drug Administration (FDA) has created a suite of open source web services that use NLP transform narrative text from clinical trials, laboratory reports, and medical records into coded, structured data.²³ Given the time constraints of workforce members translating narrative text, the web services can free up staff time to better code text. Moreover, the CDC’s Informatics Innovation Unit hopes the initiative will connect engineers and medical end users to jointly develop Machine Learning and NLP tools to decode narrative text.

NIH Artificial Intelligence Workshop: The National Institutes of Health (NIH) hosted “Harnessing Artificial Intelligence and Machine Learning to Advance Biomedical Research” last June. The meeting convened healthcare and AI leaders to discuss practical applications of AI in different healthcare settings. Experts reviewed the use of AI to further in-home non-invasive patient monitoring, its application to radiology, and new biomedical research opportunities. Participants underscored the value of open data and the need for AI training tools.

BARDA DRIVe Program: The Biomedical Advanced Research and Development Authority (BARDA) currently spearheads the DRIVe program, which identifies and enables the use of disruptive technology to further solutions to today's most pressing public health issues.²⁴ As an example, the Early Notification to Act, Control, and Treat (ENACT) uses new, disruptive technologies, including AI, to improve the prediction and treatment of infectious diseases. The ENACT program employs AI-powered solutions like chatbots and telemedicine to “bring clinical care into the home” and provide patients with quick diagnosis and treatment for diseases like influenza.

FDA's Software Precertification Program: With the advent of Software as a Medical Device, or SaMD, the FDA has established a nimbler set of principles to evaluate the integrity of medical organizations and streamline the premarket review for new medical products.²⁵ The more iterative process ensures that products that meet the basic requirements of safety and effectiveness can quickly be adopted by patients who are seeking to manage their data. Many of these software functions will increasingly leverage machine learning and algorithmic decision-making processes to improve their efficacy.

Improving Administrative Processes and Addressing Fraud

BUYSMARTER: The BUYSMARTER initiative is a transformative, data-driven initiative leveraging the collective purchasing power of HHS to secure lower prices, achieve operational efficiencies, and generate cost savings on goods and services.²⁶ BUYSMARTER uses AI technology to analyze department requirements based on current HHS-wide spend data. This helps identify opportunities to consolidate contract vehicles across agencies within HHS to leverage overlapping requirements at significant cost savings for the federal government.

CMS Improper Payment Fraud Review: The Centers for Medicare and Medicaid Services (CMS) aims to use statistical analysis to identify fraudulent and improper payments made to healthcare providers. In 2018, CMS determined that 8.12 percent of all Medicare payments were improper. In order to address this problem, CMS employs a testing methodology called the Comprehensive Error Rate Testing (CERT) and uses AI to predict fraudulent and improper healthcare payments.

The use of AI to predict fraudulent Medicare and Medicaid claims has saved the government approximately \$42 billion, according to CMS. In addition to improper payments, AI can also be used for detecting potentially fraudulent billing practices.²⁷ In the United States, over 3.2 trillion dollars are spent on healthcare, leading to a large market for potential fraud. Companies are working to reduce fraud via the application of AI for evaluating claims and predicting fraud.²⁸

HHS Office of the Inspector General: The HHS Office of the Inspector General (OIG) has focused on harnessing the power of big data to spot fraudulent claims. By developing an Integrated Data Repository, OIG has collected “petabytes of data on patient claims, providers, risk scores and other topics which investigators continuously analyze to spot potential wrongdoers.”²⁹ In June 2018, OIG was able to identify and arrest doctors inappropriately prescribing opioids.³⁰ The advent of big databases on prescriptions and medical claims allows OIG to classify providers and patients according to its Fraud Risk Indicator scale.³¹

CMMI Innovation Center: In its early stages, CMS has developed an Innovation Center designed to improve healthcare and reduce costs associated with Medicare and Medicaid programs.³² Two initiatives are worth mentioning:

1. *Quality Payment Program (QPP):* The QPP seeks to gather data from healthcare providers to evaluate how doctors use technology to support their medical practice and decisions. QPP allows for two-way communications between practitioners and CMS and seeks to improve its ability to pay practitioners by collecting and analyzing countrywide data.
2. *CMS Innovation Center Artificial Intelligence Health Outcomes Challenge:* The challenge invites participants to use AI to “predict health outcomes that are important to patients and clinicians, and to enhance care delivery.” HHS has used similar initiatives, such as the Opioid Symposium and Code-a-thon, to raise awareness of the importance of tackling the opioid problem and develop innovative solutions using existing data on opioids.³³

DATA REQUIREMENTS, ACCOUNTABILITY, AND ETHICS

Data Quality & Access

Machine-learning algorithms require access to high-quality, unbiased data that can be used to teach the algorithm what to look for. Big Data is playing an increasingly central role in healthcare, but health data poses a unique challenge, as much of it is unavailable to the public due to privacy laws. Moreover, datasets from different sources are often structured differently to one another, and are not easily collated into a larger dataset. Moving forward, researchers, practitioners, and healthcare companies will need to develop methods for collecting and integrating data while protecting individual privacy.

Data Privacy & Security

One major risk in using machine-learning algorithms is that the identities of the individuals within datasets may be inadvertently revealed. Berkeley Professor Anil Aswani noted in his recent research that “It is possible to identify individuals by learning daily patterns in step data (like that collected by activity trackers, smartwatches and smartphones) and correlating it to geographic and demographic data. The mining of two years’ worth of data covering more than 15,000 Americans led to the conclusion that the privacy standards associated with 1996’s HIPAA (Health Insurance Portability and Accountability Act) legislation need to be revisited and reworked.”³⁴

Data de-identification has also emerged as a promising area to maintain individual privacy, allowing providers to analyze highly granular individual-level data without revealing the identity of each patient in the dataset.³⁵ This approach would allow researchers to access large datasets of public and private health data to help train algorithms. As the needs of patient privacy grow, AI applications will also need to operate in environments in which data security is paramount. A recent report by the Healthcare Industry Cybersecurity Task Force indicated that data security is insufficient at most hospitals, leaving patient data open to exploitation.³⁶

Accuracy

Beyond privacy and data concerns, AI must also address issues with hazardous outcomes and incorrect suggestions. Because algorithms are developed by individuals and use data collected by individuals, they are susceptible to human error and may make erroneous predictions, diagnoses, and suggestions. In some cases, this has included developing incorrect and unsafe treatment plans for patients.³⁷ These outcomes suggest that AI algorithms should be vetted by board-certified clinicians who can learn from algorithmic suggestions but ignore them when necessary.

Accountability & Ethics

While AI has positive potential, it can also cause damage through miscalibrated algorithms or by researchers overestimating the quality of the data. Policymakers and healthcare providers must develop a system of governance and accountability as AI becomes more integrated into healthcare settings. This framework would assess the trustworthiness of algorithmic decisions, use path-tracing to understand how algorithms produced their outcome, and ensure that healthcare providers are responsible for how they utilize AI in their practice. This growing grey area of accountability is exemplified by decisions made in part by algorithms, including concerns over accountability with machine-assisted and machine-led surgeries.³⁸ This has raised questions about accountability in healthcare as AI plays an increasing role in managing patient records, diagnosing illnesses, and developing treatment plans and a focus on the limits of algorithms.³⁹

Algorithm auditing is a potential solution whereby the outcomes of AI and machine learning are checked for bias and other potential systemic concerns. Recent legislation by the European Union, for example, “requires that organizations be able to explain their algorithmic decisions.”⁴⁰ The Institute of Internal Auditors has released a framework for auditing artificial intelligence to discover how and why algorithms arrive at the solutions that they do.⁴¹

The development of an “Ethics & Algorithms Toolkit” and an “Algorithmic Impact Assessments toolkit” provide government officials and agencies with the necessary information to understand bias and algorithmic risks that could have tangible social impacts.⁴² A 2018 blog from Stanford Medicine highlighted four major ethical challenges to the use of AI in healthcare:⁴³

1. The possibility of racial or gender bias baked in to an algorithm;
2. The possibility that doctors will rely too heavily on AI applications that have not properly been audited, and cannot be easily interpreted by practitioners;
3. The possibility that data collected using machine learning, deep learning, or natural language processing will be integrated into a larger ‘collective’ dataset without express permission of the patients whose information makes up the dataset; and
4. The introduction of an intermediary (the algorithm) that complicates the chain of responsibility for medical decisions.

These concerns, and others, are now being echoed throughout the medical profession. These considerations will guide HHS as it moves through the process of designing its own AI strategy.

NATIONAL POLICY CONTEXT AND INTERNATIONAL AI STRATEGIES

As AI has proliferated both inside and outside of government, the White House Office of Science and Technology Policy (OSTP) has responded by laying the framework for a national AI strategy. In 2016, the National Science and Technology Council published The National Artificial Intelligence Research and Development Strategic Plan.⁴⁴ The Plan recommends that the Federal government develop a coordinated approach to maximize the impact of AI technology as it grows in scope. The Plan outlines the overarching principles for R&D foundations and identifies strategies that will spur long-term investment in key AI areas such as medicine, education, and transportation.

In February 2019, the Trump Administration signed an Executive Order on Maintaining American Leadership in Artificial Intelligence.⁴⁵ In coordination with this announcement, the OSTP published a fact sheet on AI.⁴⁶ The Executive Order outlines five core principles to guide the core AI strategy and calls for an action plan to preserve the U.S. advantage in AI. Moreover, it requires agencies to “enhance access to high-quality and fully traceable Federal data, models, and computing resources to increase the value of such resources for AI R&D, while maintaining safety, security, privacy, and confidentiality protections consistent with applicable laws and policies.” The Executive Order also highlights connections with the President’s Management Agenda and the Cross-Agency Priority Goal: Leveraging Data as a Strategic Asset, emphasizing the need for agencies to “consider methods of improving the quality, usability, and appropriate access to priority data identified by the AI research community.”

Towards a Department-Wide Strategy at HHS

To align itself with national priorities, HHS has also shifted its data gathering priorities and mapped out the initial steps towards a comprehensive AI strategy. Undertaken in April 2017, Relmagine has become the guiding strategic plan for HHS that aligns with the priorities of the OMB and outlines the foundations to build an HHS better equipped to serve the American people.⁴⁷ As part of this new strategic platform, HHS intends to leverage the power of data and disruptive technologies to make HHS more innovative and responsive. As part of leveraging data as a strategic asset, HHS has embarked on a strategic path to establish its data governance regime.

The State of Data Sharing at HHS

In September 2018, the Data Initiative Team of HHS completed its first comprehensive report of data sharing across HHS’s eleven operating divisions.⁴⁸ By understanding the ecosystem of available data across its divisions, HHS aims to establish an enterprise data management system to harmonize data across its divisions and create a decision-making structure for the management of these data assets. The report identified 27 high-value data assets and their functionality for both HHS and the broader public. The report also identified five core challenges for HHS to address:

1. **Process for Data Access:** HHS does not have a consistent data requesting process from one agency to another, causing delays and a lack of accountability.
2. **Technology for Data Access and Analysis:** There is wide variance of data collected across agencies and its technical utility by the public varies widely.
3. **Regulatory Environment:** All data collection efforts have statutes and policies to govern the collection and access to that data, but many limit access and use.
4. **Disclosure Risk Management:** The risk of violating individual privacy increases as more granular data is collected and shared, which leads to limits on microdata access.
5. **Norms and Resource Constraints:** HHS agencies currently do not see the value of sharing restricted or non-public data and rarely address data discrepancies with other agencies.

International AI Strategies and Principles

Among the international community, Canada and the United Kingdom have led the way in adopting both national and healthcare-focused AI principles. These examples could inform future efforts by HHS to craft an AI strategic plan.

The United Kingdom's Department of Health and Social Care has adopted a Code of Conduct to guide the adoption of its data and AI strategies.⁴⁹

United Kingdom's Code of Conduct for Data-Driven Health and Care Technology

1. Understand Users, Their Needs, Their Context
2. Define the outcome and how the technology will contribute to it
3. Use data that is in line with appropriate guidelines for the purpose for which it is being used
4. Be fair, transparent and accountable about what data is being used
5. Make use of open data standards
6. Be transparent about the limitations of the data used
7. Show what type of algorithm is being developed or deployed, the ethical examination of how the data is used, how its performance will be validated and how it will be integrated into health and care provision
8. Generate evidence of effectiveness for the intended use and value for money
9. Make security integral to the design
10. Define the commercial strategy

The Canadian government has also established a set of guiding principles for the responsible use of Artificial Intelligence in government.⁵⁰

Canada's Responsible Use of Artificial Intelligence

1. **Understand and measure** the impact of using AI by developing and sharing tools and approaches
2. **Be transparent** about how and when we are using AI, starting with a clear user need and public benefit
3. **Provide meaningful explanations** about AI decision making, while also offering opportunities to review results and challenge these decisions
4. **Be as open as we can** by sharing source code, training data, and other relevant information, all while protecting personal information, system integration, and national security and defence
5. **Provide sufficient training** so that government employees developing and using AI solutions have the responsible design, function, and implementation skills needed to make AI-based public services better

NEXT STEPS

This Roundtable will bring diverse stakeholders together to inform the development of an HHS AI strategic plan. Drawing on the background research and use cases in this briefing paper, and the Roundtable's interactive breakout sessions, participants will identify high-priority health applications of AI and key issues for an HHS AI strategy to address. The structure of the Roundtable will focus on prioritizing applications of AI in the healthcare sector, utilizing and improving health data for these applications, and outlining objectives for an HHS AI Strategy. The Roundtable will conclude with a presentation of highlights to the full group and HHS leadership. Following the Roundtable, CODE will prepare a public report of findings and recommendations from the work of the day.

REFERENCES

- ¹ “Appendix: Key Terminology for AI in Health” ConnectedHealth, ACT: The App Association n.d. <http://actonline.org/wp-content/uploads/Artificial-Intelligence-in-Health-Appendix.pdf>.
- ² “Intel’s Recommendations for the U.S. National Strategy on Artificial Intelligence.” White Paper. Intel, n.d. <https://newsroom.intel.com/wp-content/uploads/sites/11/2019/03/intel-ai-white-paper.pdf>.
- ³ Ismail, Nick. 2018. “The Success of Artificial Intelligence Depends on Data.” Information Age (blog). April 23, 2018. <https://www.information-age.com/success-artificial-intelligence-data-123471607/>.
- ⁴ Thomas, Rob. “The Road to AI Leads through Information Architecture.” VentureBeat (blog), January 12, 2018. <https://venturebeat.com/2018/01/12/the-road-to-ai-leads-through-information-architecture/>.
- ⁵ Kontzer, Tony. “Deep Learning Cuts Error Rate for Breast Cancer Diagnosis | NVIDIA Blog.” The Official NVIDIA Blog, September 19, 2016. <https://blogs.nvidia.com/blog/2016/09/19/deep-learning-breast-cancer-diagnosis/>.
- ⁶ Varma, Rohit. “Using Artificial Intelligence to Automate Screening for Diabetic Retinopathy.” Ophthalmology Times. Accessed March 19, 2019. <https://www.ophthalmologytimes.com/article/using-artificial-intelligence-automate-screening-diabetic-retinopathy>.
- ⁷ Johnson, Kipp W., Jessica Torres Soto, Benjamin S. Glicksberg, Khader Shameer, Riccardo Miotto, Mohsin Ali, Euan Ashley, and Joel T. Dudley. “Artificial Intelligence in Cardiology.” Journal of the American College of Cardiology 71, no. 23 (June 12, 2018): 2668. <https://doi.org/10.1016/j.jacc.2018.03.521>.
- ⁸ Bartley, Andrew. “Predictive Analytics in Healthcare: A Data-Driven Approach to Transforming Care Delivery.” White Paper. Intel, n.d. <https://www.intel.ai/nervana/wp-content/uploads/sites/53/2018/05/Predictive-Analytics-in-Healthcare.pdf>.
- ⁹ “AI Doctor Could Boost Chance of Survival for Sepsis Patients.” ScienceDaily. Accessed March 20, 2019. <https://www.sciencedaily.com/releases/2018/10/181022123945.htm>.
- ¹⁰ Abbasi, Jennifer. “Artificial Intelligence Tools for Sepsis and Cancer.” JAMA 320, no. 22 (December 11, 2018): 2303–2303. <https://doi.org/10.1001/jama.2018.19383>.
- ¹¹ Komorowski, Matthieu, Leo A. Celi, Omar Badawi, Anthony C. Gordon, and A. Aldo Faisal. “The Artificial Intelligence Clinician Learns Optimal Treatment Strategies for Sepsis in Intensive Care.” Nature Medicine 24, no. 11 (November 2018): 1716. <https://doi.org/10.1038/s41591-018-0213-5>.
- ¹² “Catalyst.AiTM.” Health Catalyst. Accessed March 20, 2019. <https://www.healthcatalyst.com/catalyst-ai/>.
- ¹³ “Predicting Hospital Acquired Infections (HAIs).” Open Source Leader in AI and ML (blog). Accessed March 20, 2019. <https://www.h2o.ai/solutions/usecases/predicting-hospital-acquired-infections-hais/>.
- ¹⁴ “Clinical Variation Management.” Ayasdi (blog). Accessed April 4, 2019. <https://www.ayasdi.com/applications/clinical-variation-management/>.

- ¹⁵ Rajkomar, Alvin, Eyal Oren, Kai Chen, Andrew M. Dai, Nissan Hajaj, Michaela Hardt, Peter J. Liu, et al. “Scalable and Accurate Deep Learning with Electronic Health Records.” *Npj Digital Medicine* 1, no. 1 (December 2018). <https://doi.org/10.1038/s41746-018-0029-1>.
- ¹⁶ “Overview - FHIR v4.0.0.” Accessed March 20, 2019. <https://www.hl7.org/fhir/overview.html>.
- ¹⁷ Just, Eric. “Understanding Risk Stratification and Comorbidities.” *Health Catalyst*, November 25, 2016. <https://www.healthcatalyst.com/understanding-risk-stratification-comorbidities/>.
- ¹⁸ Lumiata. “Lumiata Launches Its AI-Powered Risk Matrix That Improves Care Coordination and Risk Management.” Accessed March 20, 2019. <https://www.prnewswire.com/news-releases/lumiata-launches-its-ai-powered-risk-matrix-that-improves-care-coordination-and-risk-management-300326848.html>.
- ¹⁹ De Beule, M., E. Maes, O. De Winter, W. Vanlaere, and R. Van Impe. “Artificial Neural Networks and Risk Stratification: A Promising Combination.” *Mathematical and Computer Modelling* 46, no. 1–2 (July 2007): 88–94. <https://doi.org/10.1016/j.mcm.2006.12.024>.
- ²⁰ Arndt, Rachel. “Amazon Technology Deciphers Text in Electronic Health Record.” *Modern Healthcare*, November 27, 2018. <https://www.modernhealthcare.com/article/20181127/TRANSFORMATION01/181129951/amazon-technology-deciphers-text-in-electronic-health-record>.
- ²¹ Miliard, Mike. “EHR Natural Language Processing Isn’t Perfect, but It’s Really Useful.” *Healthcare IT News*, May 18, 2017. <https://www.healthcareitnews.com/news/ehr-natural-language-processing-isnt-perfect-its-really-useful>.
- ²² “‘TOP Health’ Tech Sprint Unleashes the Power of Open Data and AI.” Text. HHS.gov, January 17, 2019. <https://www.hhs.gov/cto/blog/2019/1/17/top-health-tech-sprint-unleashes-the-power-of-open-data-and-ai.html>.
- ²³ “Natural Language Processing Workbench Web Services | CDC,” March 12, 2019. <https://www.cdc.gov/cancer/npcr/informatics/nlp-workbench/index.htm>.
- ²⁴ “DRIVE.” Accessed March 21, 2019. <https://drive.hhs.gov/impactareas.html?id=enact#tab1>.
- ²⁵ “Developing Software Precertification Program: A Working Model.” U.S. Food & Drug Administration, June 2018. <https://www.fda.gov/downloads/MedicalDevices/DigitalHealth/DigitalHealthPreCertProgram/UCM611103.pdf>.
- ²⁶ Division of Acquisition, Office of Grants and Acquisition Policy and Accountability (OGAPA). “BUYSMARTER.” Text. HHS.gov, August 27, 2018. <https://www.hhs.gov/grants/contracts/buysmarter/index.html>.
- ²⁷ McCarthy, Jack. “CMS Snags \$42 Billion in Medicare and Medicaid Fraud with Predictive Analytics.” *Healthcare IT News*, July 22, 2016. <https://www.healthcareitnews.com/news/cms-snags-42-billion-medicare-and-medicare-fraud-predictive-analytics>.
- ²⁸ “AI Strategies in Healthcare.” H2O.ai, 2018. https://www.h2o.ai/wp-content/uploads/2018/08/Healthcare-AI-Exec-Brief-8_18.pdf.
- ²⁹ Alboum, Jonathan. “Big Data and Enterprise Data Governance Are the Cures for Cancer.” *GovLoop*. Accessed April 2, 2019. <https://www.govloop.com/community/blog/big-data-revolutionizing-healthcare/>.

- ³⁰ Corrigan, Jack. "HHS Is Using Data to Fight Opioid Abuse, Fraud and Just About Everything Else." Nextgov.com. Accessed April 2, 2019. <https://www.nextgov.com/analytics-data/2018/07/hhs-using-data-fight-opioid-abuse-fraud-and-just-about-everything-else/149226/>.
- ³¹ "Fraud Risk Indicator | Healthcare Compliance | Office of Inspector General | U.S. Department of Health and Human Services." Accessed April 2, 2019. <https://oig.hhs.gov/compliance/corporate-integrity-agreements/risk.asp>.
- ³² Medicare, the U. S. Centers for, Medicaid Services 7500 Security Boulevard, Baltimore, MD 21244 7500 Security Boulevard Baltimore, and Md 21244 Usa. "About the CMS Innovation Center | Center for Medicare & Medicaid Innovation." Accessed April 2, 2019. <https://innovation.cms.gov/About/>.
- ³³ "HHS Challenge: Opioid Symposium & Code-a-Thon." Text. HHS.gov, October 17, 2017. <https://www.hhs.gov/challenges/index.html>.
- ³⁴ December 21, John Hickey|, 2018 December 22, and 2018. "Advancement of Artificial Intelligence Opens Health Data Privacy to Attack." Berkeley News, December 21, 2018. <https://news.berkeley.edu/2018/12/21/advancement-of-artificial-intelligence-opens-health-data-privacy-to-attack/>.
- ³⁵ Cohen, Jordan, and Dianne Bourque. "Strategies to Unlock AI's Potential in Health Care, Part 3: HIPAA and Other Privacy Considerations." The National Law Review. Accessed March 21, 2019. <https://www.natlawreview.com/article/strategies-to-unlock-ai-s-potential-health-care-part-3-hipaa-and-other-privacy>.
- ³⁶ Davis, Jessica. "Where Does AI Automation Fit Into Health Data Security?" HealthITSecurity, January 7, 2019. <https://healthitsecurity.com/news/where-does-ai-automation-fit-into-health-data-security>.
- ³⁷ Ross, Casey. "IBM's Watson Recommended 'unsafe and Incorrect' Cancer Treatments." STAT, July 25, 2018. <https://www.statnews.com/2018/07/25/ibm-watson-recommended-unsafe-incorrect-treatments/>.
- ³⁸ Hart, Robert David. "Who's to Blame When a Machine Botches Your Surgery?" Quartz. Accessed March 22, 2019. <https://qz.com/1367206/whos-to-blame-when-a-machine-botches-your-surgery/>.
- ³⁹ O'Leary, D. E. "Artificial Intelligence and Big Data." IEEE Intelligent Systems 28, no. 2 (March 2013): 96–99. <https://doi.org/10.1109/MIS.2013.39>.
- ⁴⁰ Guszczka, James, Iyad Rahwan, Will Bible, Manuel Cebrian, and Vic Katyal. "Why We Need to Audit Algorithms." Harvard Business Review, November 28, 2018. <https://hbr.org/2018/11/why-we-need-to-audit-algorithms>.
- ⁴¹ "Global Perspectives and Insights: The IIA's Artificial Intelligence Auditing Framework." The Institute of Internal Auditors, 2018. <https://na.theiia.org/periodicals/Public%20Documents/GPI-Artificial-Intelligence-Part-III.pdf>.
- ⁴² "Ethics & Algorithms Toolkit Feedback." Google Docs. Accessed March 22, 2019. https://docs.google.com/forms/d/e/1FAIpQLSf36lmeTDDTPCijZ6lz8bpkFAQTb06wBEJ0fqN2931oU4Jgww/viewform?embedded=true&usp=embed_facebook.
- ⁴³ Hannon, Patricia. "Researchers Say Use of Artificial Intelligence in Medicine Raises Ethical Questions." News Center. Accessed March 22, 2019. <http://med.stanford.edu/news/all-news/2018/03/researchers-say-use-of-ai-in-medicine-raises-ethical-questions.html>.

- ⁴⁴ “National Artificial Intelligence Research and Development Strategic Plan.” National Science and Technology Council, October 2016. https://www.nitrd.gov/PUBS/national_ai_rd_strategic_plan.pdf.
- ⁴⁵ “Executive Order on Maintaining American Leadership in Artificial Intelligence.” The White House. Accessed April 2, 2019. <https://www.whitehouse.gov/presidential-actions/executive-order-maintaining-american-leadership-artificial-intelligence/>.
- ⁴⁶ “Accelerating America’s Leadership in Artificial Intelligence.” The White House, February 11, 2019. <https://www.whitehouse.gov/articles/accelerating-americas-leadership-in-artificial-intelligence/>.
- ⁴⁷ Wayne Dooley. “Reimagine HHS: Moving Towards a 21st Century Workforce.” Text. Office of Human Resources, March 2, 2018. <https://hr.nih.gov/about/events/hhs-hr-clinic/presentations/reimagine-hhs-moving-towards-21st-century-workforce>.
- ⁴⁸ “The State of Data Sharing at the U.S. Department of Health and Human Services.” Washington, D.C: Office of the Chief Technology Officer, U.S. Department of Health and Human Services, September 2018. https://www.hhs.gov/sites/default/files/HHS_StateofDataSharing_0915.pdf.
- ⁴⁹ “Code of Conduct for Data-Driven Health and Care Technology.” GOV.UK. Accessed March 19, 2019. <https://www.gov.uk/government/publications/code-of-conduct-for-data-driven-health-and-care-technology/initial-code-of-conduct-for-data-driven-health-and-care-technology>.
- ⁵⁰ Secretariat, Treasury Board of Canada. “Responsible Use of Artificial Intelligence (AI).” aem, November 22, 2018. <https://www.canada.ca/en/government/system/digital-government/responsible-use-ai.html>.