LEADERSHIP FOR IT SECURITY & PRIVACY ACROSS HHS

HHS CYBERSECURITY PROGRAM

OFFICE OF INFORMATION SECURITY

A.I. Application and Security Implications in the Healthcare Industry

Report #: 202002061000

02/06/2020

Agenda

- Overview
- A.I. Types
- A.I. Application: Patients and Families
- A.I. Application: Clinician Care Teams
- A.I. Expanding Existing Threats
- A.I. Introducing New Threats
- A.I. Risks
- A.I. Mitigations
- Questions

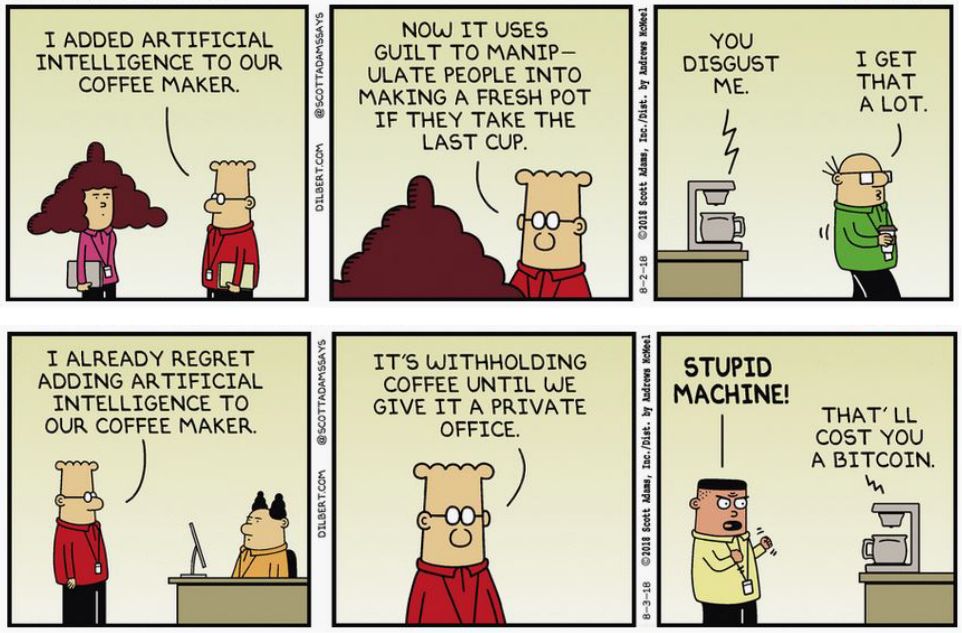

Snoopnet

Slides Key:

Non-Technical: managerial, strategic and high-level (general audience)

Technical: Tactical / IOCs; requiring in-depth knowledge (sysadmins, IRT)

Overview

- The emergence of artificial intelligence (AI) as a tool for better health care offers unprecedented opportunities to improve patient and clinical team outcomes, reduce costs, and impact population health. IoT sensor networks are characteristically different to conventional networks. Sensor devices are low powered and often use batteries as their primary source of energy.
  - Examples include but are not limited to automation; providing patient, “fRamily” (friends and family unpaid caregivers), and health care professionals’ information synthesis; and recommendations and visualization of information for shared decision making.
  - Data from Accenture estimates AI in healthcare will be a $6.6 billion market by 2021 and healthcare providers are projected to save almost $150 billion by 2026 with the help of AIs that can prevent medication dosing errors. (Forbes)

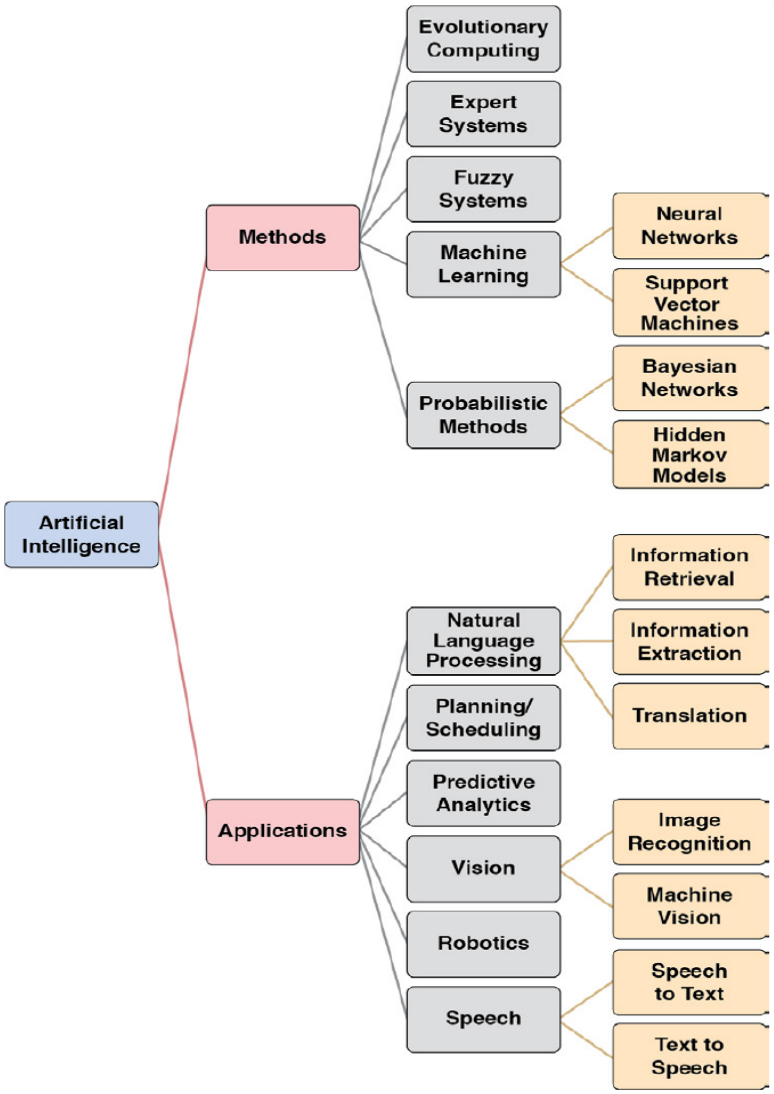

National Academy of Medicine

National Academy of Medicine

Overview Continued

- AI algorithms must be trained on population-representative data to achieve performance levels necessary for scalable “success.” Trends such as the cost for storing and managing data, data collection via electronic health records, and exponential consumer health data generation have created a data-rich health care ecosystem.
    - However, this growth in health care data struggles with the lack of efficient mechanisms for integrating and merging these data beyond their current silos.

- The challenges in operationalizing AI technologies into the health care systems are countless in spite of the fact that this is one of the strongest growth areas in biomedical research and impact.
    - The AI community must develop an integrated best practice framework for implementation and maintenance by incorporating existing best practices of ethical inclusivity, software development, implementation science, and human–computer interaction.

A.I Types: Machine Learning

- **Machine Learning** — It is an application of artificial intelligence that provides the AI System with the ability to automatically learn from the environment and applies that learning to make better decisions. There are a variety of algorithms that Machine Learning uses to iteratively learn, describe and improve data in order to predict better outcomes. These algorithms use statistical techniques to spot patterns and then perform actions on these patterns.
- **Supervised Machine Learning** — In supervised learning, the objective is to come up with a mapping function (f) that will best describe the input data (x) to conclude the output data (Y). We know x and we know Y. But, we have to find the mapping function (f) that will achieve a certain level of performance. Then, we can apply the mapping function (f) to new data to gain similar results. Training data is used to find the function f.
    - There are two types of Supervised Machine Learning problems: Classification and Regression depending on the type of output variable. If the output variable is categorical, then it is a **classification problem**. (Example: Color can be red, blue, purple, etc…) If the output variable is a real value, then it is a **regression problem.** (Example: Height can be on a scale of 0ft to 10ft)

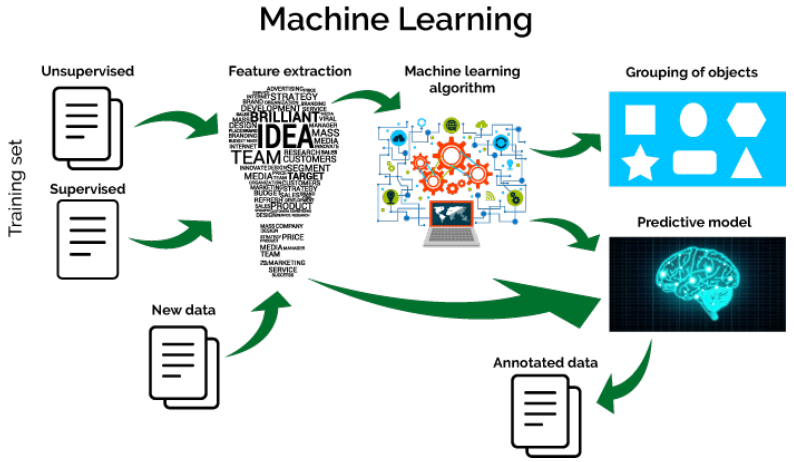

Pantech Solutions

Towards Data Science

A.I Types: Machine Learning continued

Unsupervised Machine Learning — Unlike Supervised Machine Learning, unsupervised machine learning does not assume a correct set of output “Y”. There are no outputs. The objective here is to present the most interesting structure that best describes the input data.

There are two types of Unsupervised Machine Learning problems: Clustering and Association**. Clustering** problems are when you discover groupings inside the input data. (Example: grouping voting behaviors by gender**) Association** is when you discover rules inside the input data. (Example: female voters tend to vote for female candidates)

Machine Learning algorithms can be used by **smart Electronic Health Records, Drug Discover and Development, and wearable devices that can assist in disease prediction and treatment.**

A.I Types: Deep Learning

- Deep learning is a subset of machine learning. Usually, when people use the term deep learning, they are referring to deep artificial neural networks, and somewhat less frequently to deep reinforcement learning.
    - Neural networks are a set of algorithms, modeled loosely after the human brain, that are designed to recognize patterns. They interpret sensory data through a kind of machine perception, labeling or clustering raw input. The patterns they recognize are numerical, contained in vectors, into which all real-world data, be it images, sound, text or time series, must be translated. (Pathmind)
- Deep Learning is the next generation of machine learning algorithms that use multiple layers to progressively extract higher level features (or understanding) from raw input. For instance, in image recognition applications, instead of just recognizing matrix pixels, deep learning algorithms will recognize edges at a certain level, nose at another level, and face at yet another level.
    - With the ability to understand data from the lower level all the way up the chain, a deep learning algorithm can improve its performance over time and arrive at decisions at any given moment in time.
    - The power of deep learning algorithm lies in its ability to take on both supervised learning tasks as well as unsupervised learning tasks. It also approximates many brain development theories of the human brain.
    - Deep learning algorithms are now used by **computer vision systems, speech recognition systems, natural language processing systems, audio recognition systems, bioinformatics systems and medical image analysis systems.**

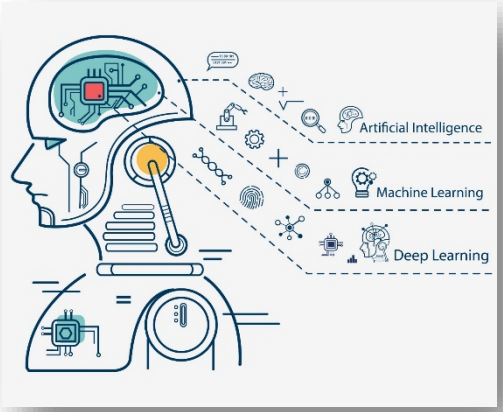

Semi

A.I. Application: PATIENTS AND FAMILIES

Conversational Agents

- Conversational agents can engage in two-way dialogue with the user via speech recognition, natural language processing (NLP), natural language understanding, and natural language generation.
  - These interfaces may include text-based dialogue, spoken language, or both. They are called, variously, virtual agents, chatbots, or chatterbots. Some conversational agents present a human image (e.g., the image of a nurse or a coach) or nonhuman image (e.g., a robot or an animal) to provide a richer interactive experience.
- The future potential for conversational agents in self-management seems high. While simulating a real-world interaction, the agent may assess symptoms, report back on outputs from health monitoring, and recommend a course of action based on these varied inputs.
  - Most adults say they would use an intelligent virtual coach or an intelligent virtual nurse to monitor health and symptoms at home.
- In other applications, conversational agents can be used to increase the engagement and effectiveness of interventions for health behavior change. Most studies of digital interventions for health behavior change have included support from either professionals or peers.
  - Professionally supported interventions cost two to three times what technology interventions cost. A conversational agent could provide some social support and increased engagement while remaining scalable and cost-effective.

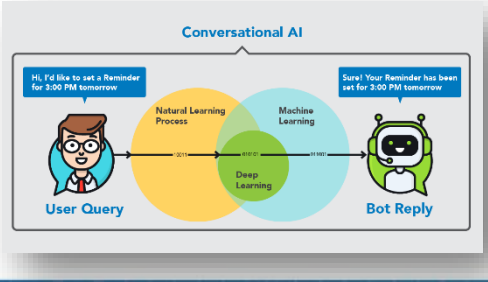

National Academy of Medicine

Chat Bots

A.I. Application: PATIENTS AND FAMILIES

Timely Personalized Interventions

- AI-driven adaptive interventions are called JITAIs, “just-in-time adaptive interventions.” These are learning systems that deliver dynamic, personalized treatment to users over .
- The JITAI makes decisions about when and how to intervene based on response to prior intervention, as well as on awareness of current context, whether internal (mood, anxiety, blood pressure), or external (e.g., location, activity).
- JITAI assistance is provided when users are most in need of it or will be most receptive to it. These systems can also tell a clinician when a problematic pattern is detected.
    - A JITAI intervention might detect when a user is in a risky situation for substance abuse relapse—and deliver an intervention against it.
    - As sensors become more ubiquitous in homes, in smartphones, and on bodies, the data sources for JITAIs are likely to continue expanding. AI can be used to allow connected devices to communicate with one another. (Perhaps a glucometer might receive feedback from refrigerators regarding the frequency and types of food consumed.)
- Leveraging data from multiple inputs can uniquely enhance AI’s ability to provide real-time behavioral management.

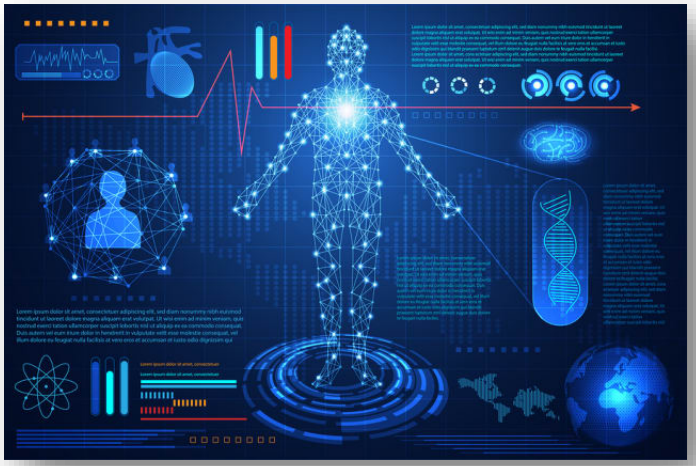

Engineering.com

National Academy of Medicine

A.I. Application: Clinician Care Teams

Diagnosis

- There are a number of demonstrations of AI in diagnostic imaging. Diagnostic image recognition can differentiate between benign and malignant melanomas, diagnose retinopathy, identify cartilage lesions within the knee joint, detect lesion-specific ischemia, and predict node status after positive biopsy for breast cancer.
- Although some believe AI will replace physicians in diagnostic imaging, it is more likely that these techniques will mainly be assistive, sorting and prioritizing images for more immediate review, highlighting important findings that might have been missed, and classifying simple findings so that the humans can spend more time on complex cases.

Surgery

- AI is becoming more important for surgical decision making. It brings to bear diverse sources of information, including patient risk factors, anatomic information, disease natural history, patient values and cost, to help physicians and patients make better predictions regarding the consequences of surgical decisions.
- In addition to planning and decision making, AI may be applied to change surgical techniques. Remote-controlled robotic surgery has been shown to improve the safety of interventions where clinicians are exposed to high doses of ionizing radiation and makes surgery possible in anatomic locations not otherwise reachable by human hands

GE.com

National Academy of Medicine

A.I. Application: Clinician Care Teams

Personalized Management and Treatment

- Precision medicine allows clinicians to tailor medical treatment to the individual characteristics of each patient. Clinicians are testing whether AI will permit them to personalize chemotherapy dosing and map patient response to a treatment so as to plan future dosing.
- AI-driven natural language processing (NLP) has been used to identify polyp descriptions in pathology reports that then trigger guideline-based clinical decision support to help clinicians determine the best surveillance intervals for colonoscopy exams. Other AI tools have helped clinicians select the best treatment options for complex diseases such as cancer.

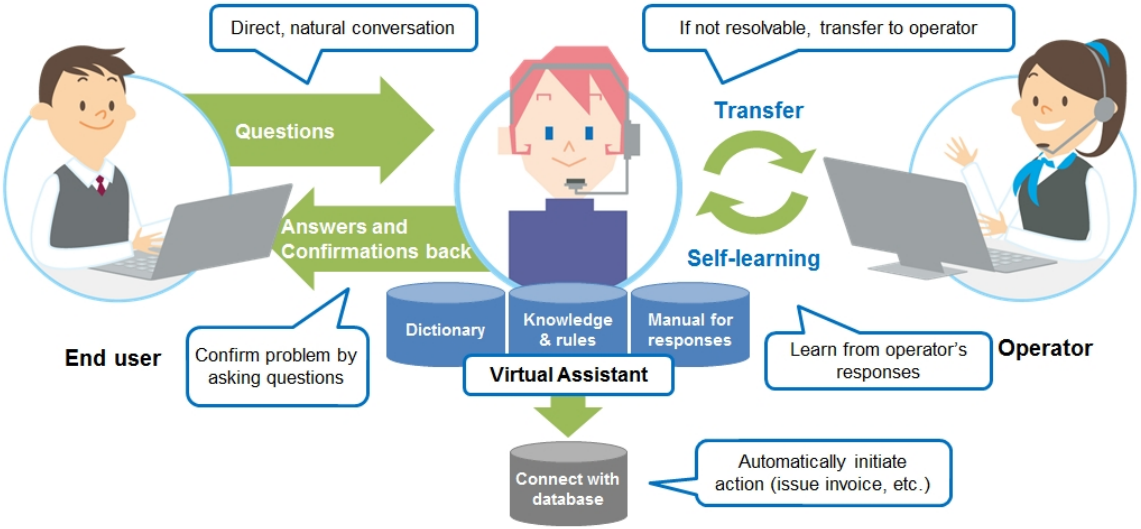

Business Wire

National Academy of Medicine

A.I. Application: Clinician Care Teams Continued

Information Management

- EHR systems and regulatory requirements have introduced significant clinical documentation responsibilities to providers, without necessarily supporting patient care decisions. AI has the potential to improve the way in which clinicians store and retrieve clinical documentation.
    - The role of voice recognition systems in clinical documentation is well known. However, such systems have been used mostly to support clinicians' dictation of narrative reports, such as clinical notes and diagnostic imaging reports.
- As mentioned previously, AI-enabled conversational and interactive systems could be used in EHR systems to support various information management tasks.
    - Clinicians could ask a conversational agent to find specific information in the patient's record, enter orders, and launch EHR functions. Instead of clicking through multiple screens to find relevant patient information, clinicians could verbally request specific information and post orders while still looking at and talking to the patient or caregivers.
- In the near future, this technology has the potential to improve the patient–provider relationship by reducing the amount of time clinicians spend focused on a computer screen.

A.I. Application: Clinician Care Teams

Cognitive Support

- AI has the potential to not only improve existing clinical decision support (CDS) modalities but also enable a wide range of innovations with the potential to disrupt patient care.
- Improved cognitive support functions include smarter CDS alerts and reminders as well as better access to peer-reviewed literature.

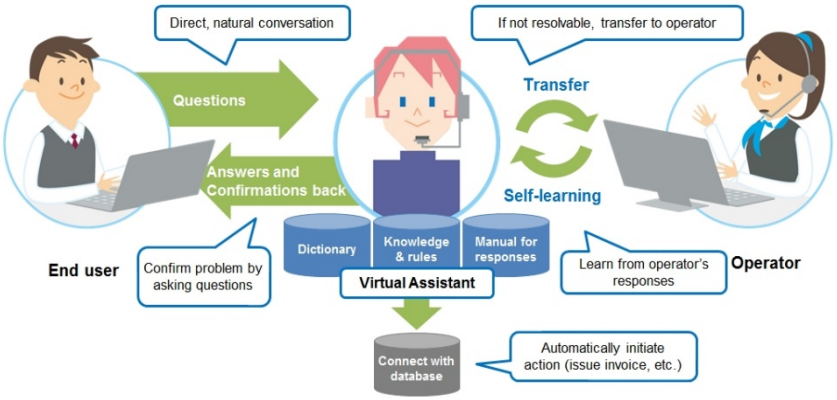

Business Wire

Improved Access to Biomedical Literature to Support Clinical Decision Making

- Recent advances in AI show promising applications in clinical knowledge retrieval.
  - Mainstream medical knowledge resources are already using machine learning algorithms to rank search results, including algorithms that continuously learn from users' search behavior.
  - In addition, "living systematic reviews" can continuously update clinical evidence as soon as the results of new clinical trials become available, with EHRs presenting evidence updates that may warrant changes to the treatment of specific patients

National Academy of Medicine

A.I. Threat Analysis: Expanding Existing Threats

- For many familiar attacks, we expect progress in AI to expand the set of actors who are capable of carrying out the attack, the rate at which these actors can carry it out, and the set of plausible targets. This claim follows from the efficiency, scalability, and ease of diffusion of AI systems.
  - In particular, the diffusion of efficient AI systems can increase the number of actors who can afford to carry out particular attacks.
- If the relevant AI systems are also scalable, then even actors who already possess the resources to carry out these attacks may gain the ability to carry them out at a much higher rate.
- Finally, as a result of these two developments, it may become worthwhile to attack targets that it otherwise would not make sense to attack from the standpoint of prioritization or cost-benefit analysis.

CSO

hindustantimes

A.I. Threat Analysis: Introducing New Threats

- Progress in AI will enable new varieties of attacks. These attacks may use AI systems to complete certain tasks more successfully than any human could, or take advantage of vulnerabilities that AI systems have but humans do not.
    - Most people are not capable of mimicking others' voices realistically or manually creating audio files that resemble recordings of human speech. However, there has recently been significant progress in developing speech synthesis systems that learn to imitate individuals' voices (a technology that's already being commercialized).
- AI systems could also be used to control aspects of the behavior of robots and malware that it would be infeasible for humans to control manually.
    - No team of humans could realistically choose the flight path of each drone in a swarm being used to carry out a physical attack.
    - Human control might also be infeasible in other cases because there is no reliable communication channel that can be used to direct the relevant systems; a virus that is designed to alter the behavior of air-gapped computers, as in the case of the 'Stuxnet' software used to disrupt the Iranian nuclear program, cannot receive commands once it infects these computers.
- Second, the property of possessing unresolved vulnerabilities implies that, if an actor begins to deploy novel AI systems, then they may open themselves up to attacks that specifically exploit these vulnerabilities.
    - If multiple robots are controlled by a single AI system run on a centralized server, or if multiple robots are controlled by identical AI systems and presented with the same stimuli, then a single attack could also produce simultaneous failures on an otherwise implausible scale.

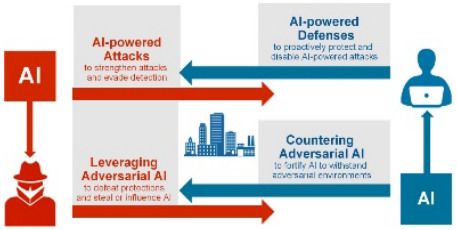

National Academies Press

Cornell University

A.I. Risks

- **Injuries and error.** The most obvious risk is that AI systems will sometimes be wrong, and that patient injury or other health-care problems may result. If an AI system recommends the wrong drug for a patient, fails to notice a tumor on a radiological scan, or allocates a hospital bed to one patient over another because it predicted wrongly which patient would benefit more, the patient could be injured.
  - Patients and providers may react differently to injuries resulting from software than from human error.
  - If AI systems become widespread, an underlying problem in one AI system might result in injuries to thousands of patients—rather than the limited number of patients injured by any single provider's error.
- **Data availability.** Training AI systems requires large amounts of data from sources such as electronic health records, pharmacy records, insurance claims records, or consumer-generated information like fitness trackers or purchasing history. Data are typically fragmented across many different systems. Even aside from the variety just mentioned, patients typically see different providers and switch insurance companies, leading to data split in multiple systems and multiple formats.
  - This fragmentation increases the risk of error, decreases the comprehensiveness of datasets, and increases the expense of gathering data—which also limits the types of entities that can develop effective health-care AI.

Brookings

RE-Work

A.I. Risks continued

- **Privacy concerns.** The requirement of large datasets creates incentives for developers to collect such data from many patients. Some patients may be concerned that this collection may violate their privacy, and lawsuits have been filed based on data-sharing between large health systems and AI developers. AI could implicate privacy in another way: AI can predict private information about patients even though the algorithm never received that information.
  - An AI system might be able to identify that a person has Parkinson's disease based on the trembling of a computer mouse, even if the person had never revealed that information to anyone else (or did not know). Patients might consider this a violation of their privacy, especially if the AI system's inference were available to third parties, such as banks or life insurance companies.

A.I. Mitigations

- **Data generation and availability.** Several risks arise from the difficulty of assembling high-quality data in a manner consistent with protecting patient privacy.
  - One set of potential solutions turns on government provision of infrastructural resources for data, ranging from setting standards for electronic health records to directly providing technical support for high-quality data-gathering efforts in health systems that otherwise lack those resources.
  - A parallel option is direct investment in the creation of high-quality datasets. Reflecting this direction, both the United States' All of Us initiative and the U.K.'s BioBank aim to collect comprehensive health-care data on huge numbers of individuals. Ensuring effective privacy safeguards for these large-scale datasets will likely be essential to ensuring patient trust and participation.

C4ISRNET

Brookings

A.I. Mitigations continued

- **Quality oversight.** Oversight of AI-system quality will help address the risk of patient injury. The Food and Drug Administration (FDA) oversees some health-care AI products that are commercially marketed. The agency has already cleared several products for market entry, and it is thinking creatively about how best to oversee AI systems in health.
- Many AI systems in health care will not fall under FDA's purview, either because they do not perform medical functions or because they are developed and deployed in-house at health systems themselves—a category of products FDA typically does not oversee.
    - Increased oversight efforts by health systems and hospitals, professional organizations like the American College of Radiology and the American Medical Association, or insurers may be necessary to ensure quality of systems that fall outside the FDA's exercise of regulatory authority.
- **Provider engagement and education.** The integration of AI into the health system will undoubtedly change the role of health-care providers. A hopeful vision is that providers will be enabled to provide more-personalized and better care, freed to spend more time interacting with patients as humans.
- A less hopeful vision would see providers struggling to weather a monsoon of uninterpretable predictions and recommendations from competing algorithms. In either case—or in any option in-between—medical education will need to prepare providers to evaluate and interpret the AI systems they will encounter in the evolving health-care environment.

Brookings

Reference Materials

References

- The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation
  - https://arxiv.org/ftp/arxiv/papers/1802/1802.07228.pdf
- Risks and Remedies for Artificial Intelligence in Healthcare
  - https://www.brookings.edu/research/risks-and-remedies-for-artificial-intelligence-in-health-care/
- Governing Artificial Intelligence: Upholding Human Rights & Dignity
  - https://datasociety.net/wp-content/uploads/2018/10/DataSociety_Governing_Artificial_Intelligence_Upholding_Human_Rights.pdf
- ARTIFICIAL INTELLIGENCE IN HEALTH CARE: THE HOPE, THE HYPE, THE PROMISE, THE PERIL
  - https://nam.edu/artificial-intelligence-special-publication/
- Guide to Neural Networks and Deep Learning
  - https://pathmind.com/wiki/neural-network
- Artificial Intelligence Market worth $190.61 billion by 2025 with a Growing CAGR of 36.6%
  - https://www.marketsandmarkets.com/PressReleases/artificial-intelligence.asp

Questions

Questions

Upcoming Briefs

- Electronic Healthcare Records
- PyXie RAT

Product Evaluations

Recipients of this and other Healthcare Sector Cybersecurity Coordination Center (HC3) Threat Intelligence products are highly encouraged to provide feedback to **HC3@HHS.GOV**.

Requests for Information

Need information on a specific cybersecurity topic? Send your request for information (RFI) to **HC3@HHS.GOV** or call us Monday-Friday, between 9am-5pm (EST), at **(202) 691-2110.**

About Us

HC3 works with private and public sector partners to improve cybersecurity throughout the Healthcare and Public Health (HPH) Sector

Products

Sector & Victim Notifications

Directed communications to victims or potential victims of compromises, vulnerable equipment or PII/PHI theft and general notifications to the HPH about currently impacting threats via the HHS OIG

White Papers

Document that provides in-depth information on a cybersecurity topic to increase comprehensive situational awareness and provide risk recommendations to a wide audience.

Threat Briefings & Webinar

Briefing document and presentation that provides actionable information on health sector cybersecurity threats and mitigations. Analysts present current cybersecurity topics, engage in discussions with participants on current threats, and highlight best practices and mitigation tactics.

Need information on a specific cybersecurity topic or want to join our listserv? Send your request for information (RFI) to HC3@HHS.GOV or call us Monday-Friday, between 9am-5pm (EST), at (202) 691-2110.

Contact

Health Sector Cybersecurity Coordination Center (HC3)

(202) 691-2110

HC3@HHS.GOV