ARTIFICIAL INTELLIGENCE IN HEALTH CARE: HOPE NOT HYPE, PROMISE NOT PERIL

This publication synthesizes the major literature to date, in both the academic and general press, to create a reference document for health care AI model developers, clinical teams, patients, "fRamilies", and regulators and policy makers to: 1. identify the current and near-term uses of AI within and outside the traditional health care systems (see Chapters 2 and 3); 2. highlight the challenges and limitations (see Chapter 4) and the best practices for development, adoption, and maintenance of AI tools (see Chapters 5 and 6); 3. understand the legal and regulatory landscape (see Chapter 7); 4. ensure equity, inclusion, and a human rights lens for this work; and 5. outline priorities for the field.

The authors of the eight chapters are experts convened by the National Academy of Medicine's Digital Health Learning Collaborative to explore the field of AI and its applications in health and health care, consider approaches for addressing existing challenges, and identify future directions and opportunities. This final chapter synthesizes the challenges and priorities of the previous chapters, highlights current best practices, and identifies key priorities for the field.

Contents

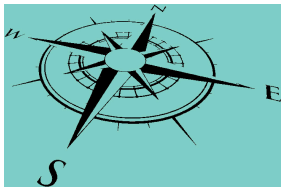

Digital Health Learning Collaborative (DHLC)

Stakeholder(s):

National Academy of Medicine :
The National Academy of Medicine is one of three Academies constituting the National Academies of Sciences, Engineering, and Medicine (the National Academies). The National Academies provide independent, objective analysis and advice to the nation and conduct other activities to solve complex problems and inform public policy decisions. The National Academies also encourage education and research, recognize outstanding contributions to knowledge, and increase public understanding in matters of science, engineering, and medicine.

Michael Matheny :
Co-Author - Vanderbilt University Medical Center and the Department of Veterans Affairs

Sonoo Thadaney Israni :
Co-Author - Stanford University

Danielle Whicher :
Co-Author - National Academy of Medicine

Mahnoor Ahmed :
Co-Author - National Academy of Medicine

Mission

To synthesize the challenges and priorities, highlight current best practices, and identify key priorities for the application of AI in health and health care

Values

Data: One example of a principles declaration that promotes data robustness and quality is the FAIR (findability, accessibility, interoperability, and reusability) Principles (Wilkinson et al., 2016). These principles, put forth by molecular biology and bioinformatics researchers, are not easily formalized or implemented. However, for health care AI to mature, a similar set of principles should be developed and widely adopted.

Robustness

Quality

Findability

Accessibility

Interoperability

Reusability

1. Data

Provide high quality, population-representative, and diverse data

Stakeholder(s)

Glenn Cohen :
In a recent commentary, Glenn Cohen and Michelle Mello propose that "it is timely to reexamine the adequacy of the Health Insurance Portability and Accountability Act (HIPAA), the nation's most important legal safeguard against unauthorized disclosure and use of health information. Is HIPAA up to the task of protecting health information in the 21st century?" (Cohen and Mello, 2018). When entities bring data sources together, they face ethical, business, legislative, and technical hurdles. There is a need for novel solutions that allow for robust data aggregation while promoting transparency and respecting patient privacy and preferences.

Michelle Mello

It is widely accepted that the successful development of an AI system requires high quality, population-representative, and diverse data (Shrott, 2017; Sun et al., 2017).

1.1. Access

Promote Data Access

Stakeholder(s):

Health Care Community :
The health care community should continue to advocate for policy, regulatory, and legislative mechanisms that improve the ease of data aggregation. These would include (but are not limited to) a national patient health care identifier and mechanisms to responsibly bring data from multiple sources together. The debate should focus on the thoughtful and responsible ability of large-scale health care data resources to serve as a public good and the implications of that ability.

The Public :
Discussions around wider and more representative data access should be carefully balanced by stronger outreach, education, and consensus building with the public and patients in order to address where and how their data can be reused for AI research, data monetization, and other secondary uses; which entities can reuse their data; and what safeguards need to be in place.

Patients

1.2. Standardization

Promote Data Standardization

Figure 8-1 outlines a standardized pathway for the collection and integration of multiple data sources into a common data model (CDM), which efficiently feeds the transformation to a feature space for AI algorithm training. However, some of the standardization tools and data quality assessments and methodologies for curating the data do not yet exist... We cannot disregard the fact that there are varying data requirements for the training of AI and for the downstream use of AI.

Stakeholder(s):

Health Care Community :
Some initiatives do exist and are driving the health care community in the direction of interoperability and data standardization, but they have yet to see widespread use (HL7, 2018; Indiana Health Information Exchange, 2019; NITRD Program Workshop, 2019; OHDSI, 2019).

HL7

Indiana Health Information Exchange

NITRD

OHDSI

EHR Vendors :
Interoperability is critical at all layers, including across the multi-vendor electronic health record (EHR) and ancillary components of a health care system, between different health care systems, and with consumer health applications.

1.3. Quality

Report Data Quality

Methods to assess data validity and reproducibility are often ad hoc. Ultimately, for AI models to be trusted, the semantics and provenance of the data used to derive them must be fully transparent, unambiguously communicated, and available, for validation at least, to an independent vetting agent. This is a distinct element of transparency, and the conflation of data transparency with algorithmic transparency complicates the AI ecosystem's discourse. We suggest a clear separation of these topics.

1.4. Bias

Minimize Data Bias

2. Fairness & Inclusiveness

Prioritize Equitable and Inclusive Health Care

Stakeholder(s)

Karl Pearson :
"That which is measured, improves," opined Karl Pearson, famed statistician and founder of mathematical statistics.

Therefore, prioritizing equity and inclusion should be a clearly stated goal when developing and deploying AI in health care.

In addition these solutions need to be equitable to avoid a potential conundrum (see Chapters 1 and 4), in which patients, especially those who are the least AI-savvy are unaware of how their data is monetized.

2.1. Biases

Unpack the underlying biases in the data used to develop AI tools

It is imperative for developers and implementers to consider the data used to develop AI tools and unpack the underlying biases in that data.

2.2. Deployment

Consider how the tool should be deployed

It is also essential to consider how the tool should be deployed, and whether the range of deployment environments could impact equity and inclusivity.

2.3. Environments

Consider whether the range of deployment environments could impact equity and inclusivity

There are widely recognized inequities in health outcomes due to the social determinants of health (BARHII, 2015) and the perverse incentives in existing health care systems (Rosenthal, 2017). Unfortunately, consumer-facing technologies have often exacerbated historical inequities in other fields, and the digital divide continues to be a reality for wearables deployment and the data-hungry plans they require, even if the initial cost of the device is subsidized.

Stakeholder(s):

Health Care Systems

Cathy O'Neil :
As Cathy O'Neil reported in Weapons of Math Destruction, AI and related sciences can exacerbate inequity on a monumental scale. The impact of a single biased human is far less than that of a global or national AI (O'Neil, 2017).

2.4. Transparency

Ensure data transparency to assess biases and whether the data are representative of the population in which the AI tool will be deployed

Data transparency is key to ensuring AI adopters can assess the underlying data for biases and to consider whether the data are representative of the population in which the AI tool will be deployed.

Stakeholder(s):

Indiana Health Information Exchange :
The United States has some population-representative datasets, such as national claims data, and high levels of data capture in certain markets (such as the Indiana Health Information Exchange).

Social Media :
But, in many instances AI is being developed with data that is not population-representative, and while there are efforts to link health care data to the social determinants of health, environmental, and social media data to obtain a comprehensive profile of a person, this is not routine.

2.4.1. Quality & Representation

Develop and standardize approaches for evaluating and reporting on data quality and representativeness

Nor are there ethical or legal frameworks for doing so. It is imperative that we develop and standardize approaches for evaluating and reporting on data quality and representativeness.

2.4.2. Diversity

Ensure and report on the diversity of gender, race, age, and other human characteristics of AI development teams

It is equally vital that we ensure and report on the diversity of gender, race, age, and other human characteristics of AI development teams to benefit from their much-needed diverse knowledge and life experiences (see Chapters 1 and 5).

Stakeholder(s):

AI Development Teams

2.5. Governance

Develop a new governance framework

Executing and delivering on equity and inclusion will require a new governance framework.

Stakeholder(s):

Technology Companies :
Current self-governance efforts by technology companies are plagued with numerous struggles and failures, Google's April 2019 Ethics Board dissolution being one recent example (Piper, 2019).

Google

Mark Latonero :
Mark Latonero suggests, "In order for AI to benefit the common good, at the very least its design and deployment should avoid harms to fundamental human values. International human rights provide a robust and global formulation of those values" (Latonero, 2018).

2.5.1. Agency or Committee

Create a new neutral agency or a committee to manage the review of health care AI products and services

For objective governance, a new neutral agency or a committee within an existing governmental or nongovernmental entity, supported by a range of stakeholders, could own and manage the review of health care AI products and services while protecting developers' intellectual property rights.

Stakeholder(s):

New Model for Industry-Academic Partnerships :
One example of this type of solution is the New Model for Industry-Academic Partnerships, which developed a framework for academic access to industry (Facebook) data sources: The group with full access to the data is separate from the group doing the publishing, but both are academic, independent, and trusted.

Facebook

Social Science Research Council :
The group with full access executes the analytics and verifies the data, understands the underlying policies and issues, and delivers the analysis to a separate group who publishes the results but does not have open access to the data (Social Science Research Council, 2019).

New Model Project Funders :
To ensure partisan-neutrality, the project is funded by ideologically diverse supporters, including the Laura and John Arnold Foundation, the Democracy Fund, the William and Flora Hewlett Foundation, the John S. and James L. Knight Foundation, the Charles Koch Foundation, the Omidyar Network, and the Alfred P. Sloan Foundation.

Laura and John Arnold Foundation

Democracy Fund

William and Flora Hewlett Foundation

John S. and James L. Knight Foundation

Charles Koch Foundation

Omidyar Network

Alfred P. Sloan Foundation

Research Projects :
Research projects use this framework when researchers use Facebook social media data for election impact analysis, and Facebook provides the data required for the research but does not have the right to review or approve the research findings prior to publication.

2.6. Quintuple Aim

Add equity and inclusion as a dimension to the quadruple aim

Perhaps the best way to ensure that equity and inclusion are foundational components of a thriving health care system is to add them as a dimension to the quadruple aim, expanding it to a Quintuple Aim for health and health care: better health, improved care experience, clinician well-being, lower cost, and health equity throughout. (see Figure 8-2).

2.6.1. Health

Aim for better health

2.6.2. Experience

Aim for improved care experience

2.6.3. Clinician Well-Being

Aim for clinician well-being

Stakeholder(s):
Clinicians

2.6.4. Cost

Aim for lower cost

2.6.5. Fairness

Aim for health equity

3. Accuracy, Risk, Liability & Trust

Promote a Spectrum of Transparency-Based Trust, Based on Considerations of Accuracy, Risk, and Liability

Stakeholder(s)

Health Care AI Stakeholders :
All stakeholders should prioritize equity and inclusion, requiring transparency on how AI tools are monitored and updated. Many of these are shared, not siloed, responsibilities. In all cases, the transparency of the underlying data used for AI model generation should be endorsed.

A key challenge to the acceptance and widespread use of AI is the tension between data and algorithmic transparency, accuracy, perceived risk, and tort liability. One of the priorities identified in this publication is the need to present each health care AI tool along with the spectrum of transparency related to the potential harms and context of its use.

3.1. Sub-Domains

Evaluate, address, and report transparency in each sub-domain of data, algorithms, and performance

Evaluating and addressing appropriate transparency, in each sub-domain of data, algorithms, and performance, and systematically reporting it, must be a priority.

3.2. ROI, Risks & Benefits

Understand the return on investment and the risks and benefits of adoption

In addition, health system leaders must understand the return on investment and the risks and benefits of adoption, including the risks of adverse events post-implementation ...

Stakeholder(s):
Health System Leaders

3.3. Culture & Workflows

Understand the culture and workflows where AI tools will be used

and informatics implementers must understand the culture and workflows where AI tools will be used so the algorithms can be adjusted to reflect their needs.

Stakeholder(s):
Informatics Implementers

3.4. Data Sources

Routinely publish information on the sources from which patient-level data were aggregated

While granular, patient-level data should not be publicly shared, publishing information on the data sources from which they were aggregated; how the data were transformed; data quality issues; inclusion and exclusion criteria that were applied to generate the cohort; summary statistics of demographics; and relevant data features in each source should be conventional practice. This information could be a supporting document and would tremendously improve the current understanding of and trust in AI tools.

Stakeholder(s):

Data Aggregators

3.5. Context

Adjust and align algorithmic transparency with the use context

The need for algorithmic transparency is largely dependent on the use context. For applications that have immediate clinical impact on patient quality of life or health outcomes, the baseline requirement for transparency is high. However, the level of transparency could be different depending on the (1) known precision accuracy of the AI; (2) clarity of recommended actions to end users; (3) risk to the patient or target; and (4) legal liability.

Stakeholder(s):

End Users :

For example, if an AI tool has high-precision accuracy and low risk, provides clear recommendations to the end user, and is unlikely to impose legal liability on the institution, manufacturer, or end user, then the need for complete algorithmic transparency is likely to be lower. See Figure 8-3 for additional details on the relationships of transparency and these axes within different conceptual domains.

Institutions

Manufacturers

3.5.1. Accuracy

Take into account the known precision accuracy of the AI

3.5.2. Clarity

Take into account the clarity of recommended actions to end users

Stakeholder(s):

End Users

3.5.3. Risks

Take into account the risk to the patient or target

Stakeholder(s):

Patients

3.5.4. Liability

Take into account legal liability

4. Augmentation

Focus of Near-Term Health Care AI: Augmented Intelligence Versus Full Automation

Stakeholder(s)

Clinicians :
The opportunity for augmenting human cognition is vast, from supporting clinicians with less training in performing tasks currently limited to specialists to filtering out normal or low-acuity clinical cases so specialists can work at the top of their licensure.

Medical Personnel :
Additionally, AI could help humans reduce medical error due to cognitive limits, inattention, micro-aggression, or fatigue.

Surgeons :
In the case of surgery, it might offer capabilities that are not humanly possible.

Although some AI applications for health care business operations are likely to be poised for full automation, most of the near-term dialogue around AI in health care should focus on promoting, developing, and evaluating tools that support human cognition rather than replacing it. Popular culture and marketing have overloaded the term "AI" to the point where it means replacing human labor, and as a result, other terms have emerged to distinguish AI that is used to support human cognition. Augmented intelligence refers to the latter, which is the term the authors of this chapter endorse.

4.1. Business Processes

Avoid the trough of disillusionment in automating business processes

Opportunities exist for automating some business processes, and greater automation is possible as the field matures in accuracy and trust. But it would not be prudent to deploy fully automated AI tools that could result in inaccuracy when the public has an understandably low tolerance for error, and health care AI lacks needed regulation and legislation. This is most likely to create a third AI Winter or a trough of disillusionment as seen in the Gartner Hype Cycle (see Chapter 4).

4.2. Consumer Applications

Track and surveil potential harms that could result from usage of consumer health applications

Differential levels of automation are even more relevant to consumer health applications because they are likely to have more automation components, but are regulated as entertainment applications, and their standards and quality controls are much more variable. The quandaries here are perhaps even more dire given consumer health applications' widespread use and the difficulties of tracking and surveilling potential harms that could result from their use in the absence of expert oversight.

Stakeholder(s):
Consumers

5. Training & Education

Develop Appropriate Professional Health Training and Educational Programs to Support Health Care AI

Stakeholder(s)

Curt Langlotz :
Stanford's Curt Langlotz, offered the following question and answer: "Will AI ever replace radiologists? I say the answer is no—but radiologists who use AI will replace radiologists who don't" (Stanford University, 2017).

Health Care AI Stakeholders :
In order to sustain and nurture health care AI, we need a sweeping, comprehensive expansion of relevant professional health education focused on data science, AI, medicine, humanism, ethics, and health care. This expansion must be multidisciplinary and engage AI developers, implementers, health care system leadership, frontline clinical teams, ethicists, humanists, and patients and "fRamilies", because each brings essential expertise and AI progress is contingent on knowledgeable decision makers balancing the conflicting pressures of the relative ease of implementing newly developed AI solutions while understanding their validity and influence on care.

AI Developers

AI Implementers

Health Care System Leaders

Frontline Clinical Teams

Ethicists

Humanists

Patients

fRamilies :
friends and family unpaid caregivers

Universities :
To begin addressing challenges, universities such as the Massachusetts Institute of Technology, Harvard, Stanford, and The University of Texas have added new courses focused on the embedding ethics into their development process.

Massachusetts Institute of Technology

Harvard

Stanford

The University of Texas

Mehran Sahami :
Mehran Sahami, a Stanford computer science faculty member, who formerly worked at Google as a senior research scientist said, "Technology is not neutral. The choices that get made in building technology then have social ramifications" (Singer, 2018).

Health Care Professionals :
Health care professionals have requirements for continuing education as part of their scope of practice; we suggest that new continuing education AI curricula be developed and delivered.

Professional Health Educators :
Professional health education should incorporate how to critically evaluate the utility and risk of these AI tools in clinical practice.

Curriculum Developers :
Curricula should provide an understanding of how AI tools are developed, the criteria and considerations for the use of AI tools, how best to engage and use such tools while prioritizing patient needs, and when human oversight is needed.

Health Care Leaders :
For health care system leadership and AI implementers, it is important to have training on the importance and lenses of the multiple disciplines that must be brought together to evaluate, deploy, and maintain AI in health care.

AI Implementers

Clinical Training Programs :
Current clinical training programs bear the weight of growing scientific knowledge within a static time window of training.

Clinicians :
We recognize the impracticality of each clinician or team being an expert on all things health care–AI related.

Clinical Teams :
Instead, we propose that each team have a basic and relevant understanding as described and add an AI consult when and where needed.

AI Consultants :
Such consults could be done virtually, supporting the team effort and group decision making, and costing less than if they were done on-site. Regional or content-expert AI consults could be leveraged across many health care systems.

Undiagnosed Diseases Network (UDN) :
One example of such regional consults is the National Institutes of Health–funded Undiagnosed Diseases Network (UDN), which seeks "to improve and accelerate diagnosis of rare and undiagnosed conditions (NIH, 2019). The UDN uses both basic and clinical research to improve the level of diagnosis and uncover the underlying disease mechanisms associated with these conditions." National (or global) efforts like this can support the building and deployment of responsible AI solutions for health care.

— continued next page

Stakeholders (continued)

National Institutes of Health

Health Care Workers :
It is necessary to develop retraining programs to target job categories that are likely to be the most susceptible to a shift in desired skill sets with AI deployment. It is unlikely that many health care jobs will be lost, but skill and knowledge mismatches are to be expected (see Chapter 4).

Some important topics that should be covered are how to [address the following objectives] As the field evolves, the nature and emphasis of these topics will change, necessitating periodic review and updating.

5.1. Need, Validity & Applicability

Assess the need, validity, and applicability of AI algorithms in clinical care

5.2. Performance & Impact

Understand algorithmic performance and the impact on downstream clinical use

5.3. Error & Liability

Navigate medical liability and the ways in which AI tools may impact individual and institutional liability and medical error

5.4. Standardization & Transparency

Advocate for standardization and appropriate transparency for a given use case

5.5. Emerging Technologies

Discuss emerging AI technologies, their use, and their dependence on patient data with patients and "fRamilies" and the patient–clinician relationship

5.6. Equity, Inclusion & Impact

Ensure the Quintuple Aim of Equity & Inclusion when measuring impact

5.7. Consultation

Know when and how to bring in AI experts for consults

Stakeholder(s):

AI Experts

6. Success Factors

Articulate Success Factors for the Development, Adoption, and Maintenance of AI in Health Care

Stakeholder(s)

Health Care System Leaders

AI Developers

AI Implementers

Regulators

Humanists

Patients

fRamilies

The success factors for development, adoption, and maintenance of AI tools will need clarity, acknowledging that practices will differ depending on the physical, psychological, or legal risk to the end user, the adoption setting, the level of augmentation versus automation, and other considerations. Dissonance between levels of success and users' expectations of impact and utility are likely to create harm and disillusionment. Below, we summarize the key components that must be wrangled.

6.1. Understanding & Expectations

Build a shared understanding and expectations

In order to implement AI tools in health care settings with sustained success, it is important that system leadership, AI developers, AI implementers, regulators, humanists, and patients and "fRamilies" collaboratively build a shared understanding and expectations.

6.2. Best Practices

Develop integrated best-practice frameworks for AI implementation and maintenance

The global health care AI community must develop integrated best-practice frameworks for AI implementation and maintenance, balancing ethical inclusivity, software development, implementation science, and human–computer interaction. These frameworks should be developed within the context of the learning health care system and can be tied to various targets and objectives. Earlier chapters provide summaries and considerations for both technical development (see Chapter 5) and health care system implementation (see Chapter 6). However, the AI implementation and deployment domain is still in a nascent stage, and health systems should maintain appropriate skepticism about the advertised benefits of health care AI.

6.3. Deployment

Deploy AI to address problems for which simpler or more basic solutions are inadequate

It is important to approach health care AI as one of many tools for supporting the health and well-being of patients. Thus, AI should be deployed to address real problems that need solving, and only among those problems in which a simpler or more basic solution is inadequate. The complexity of AI has a very real cost to health care delivery environments.

6.3.1. Individuals & Communities

Address individual patient and communal needs

Health care AI could go beyond the current limited, biology-focused research to address individual patient and communal needs. The current medical enterprise is largely focused on the tip of the iceberg (i.e., human biology), lacking meaningful and usable access to relevant patient contexts such as social determinants of health and psychosocial risk factors. AI solutions have the potential (with appropriate consent) to link personal and public data for truly personalized health care.

Stakeholder(s):

Patients

Communities

UnitedHealthcare :
The April 2019 collaborative effort by UnitedHealthcare and the American Medical Association to create nearly 2 dozen International Classification of Diseases, Tenth Revision codes to better incorporate social determinants of health into health care delivery is a laudable and responsible step in the right direction (Commins, 2019).

American Medical Association

6.3.2. Scale & Needs

Consider AI where scale is important and resources are insufficient for current needs

AI should be considered where scale is important and resources are insufficient for current needs.

Stakeholder(s):

Complex Patients :
Some of these environments include complex patients with multiple co-morbid conditions, such as chronic disease sufferers and the elderly, or low-resource settings.

Telehealth Users :
For innovative telehealth—disaster relief and rural areas—when resources are limited and access difficult, triaging or auto-allocating resources can be powered by AI solutions.

Victims of Disasters

Rural Residents

Mobile Technology Service Providers :
Current mobile technology allows for critical imaging at the local site ...

U.S. Department of Veterans Affairs :
and the U.S. Department of Veterans Affairs has operationalized a robust telehealth program that serves their very diverse population (VA, 2016).

Veterans

6.4. IT Governance Strategy

Institute an information technology (IT) governance strategy

We strongly suggest that a robust and mature underlying information technology (IT) governance strategy be in place within health care delivery systems prior to embarking on substantial AI deployment and integration.

6.5. Resources

Provide the required resources

The needs for on- or off-site hardware infrastructure, change management, inclusive stakeholder engagement, and safety monitoring all require substantial established resources. Systems that do not possess these infrastructure components should develop them before significant AI deployment.

7. Regulation & Legislation

Balance Regulation and Legislation for Health Care Innovation

Stakeholder(s)

Regulators :
Regulators should remain flexible, but the potential for lagging legislation remains an issue.

The regulatory and legislative considerations for AI use in consumer and professional health care domains are documented in Chapter 7. AI applications have great potential to improve patient health but could also pose significant risks, such as inappropriate patient risk assessment, treatment recommendations, privacy breaches, and other harms (Evans and Whicher, 2018). Overall, the field is advancing rapidly, with a constant evolution of access to data, aggregation of data, new developments in AI methods, and expansions of how and where AI is added to patient health and health care delivery.

7.1. Approach

Take a graduated approach to the regulation of AI

In alignment with recent congressional and U.S. Food and Drug Administration developments and guidance, we suggest a graduated approach to the regulation of AI based on the level of patient risk, the level of AI autonomy, and how static or dynamic certain AI tools are likely to be.

Stakeholder(s):

Congress **U.S. Food and Drug Administration**

7.2. Surveillance

Adopt postmarket surveillance mechanisms to ensure continuing high-quality performance

To the extent that machine learning–based models continuously learn from new data, regulators should adopt postmarket surveillance mechanisms to ensure continuing (and ideally, improving) high-quality performance.

7.3. Liability, Risks & Benefits

Understand regulation and liability to evaluate risks and benefits.

Liability accrued within the deployment of various contexts of AI will continue to be a developing area as regulators, courts, and the insurance industry weigh in. Understanding regulation and liability is essential to evaluating risks and benefits.

Stakeholder(s):

Regulators **Insurance Industry**

Courts

7.4. Collaboration & Evaluation

Collaborate with stakeholders and experts to continuously evaluate deployed clinical AI for effectiveness and safety based on real-world data

The linkages between innovation, safety, progress, and regulation are complex. Regulators should engage in collaborative efforts with stakeholders and experts to continuously evaluate deployed clinical AI for effectiveness and safety based on real-world data.

Stakeholder(s):

Regulators **AI Experts**

7.5. Infrastructure

Invest in infrastructure that promotes wider data collection and access to data resources for building AI solutions

Throughout that process, transparency can help deliver well-vetted solutions. To enable both AI development and oversight, governmental agencies should invest in infrastructure that promotes wider data collection and access to data resources for building AI solutions, within a framework of equity and data protection (See Figure 8-4).

Stakeholder(s):

Governmental Agencies

Administrative Information

Start Date: 2019-12-31
End Date:

Publication Date: 2020-05-01
Source: https://nam.edu/wp-content/uploads/2019/12/AI-in-Health-Care-PREPUB-FINAL.pdf

Submitter:

Given Name: Owen
Surname: Ambur
Email: Owen.Ambur@verizon.net
Phone: