Machine Learning and Artificial Intelligence (3/3)

Introduction to Biomedical & Health Informatics

William Hersh

Oregon Health & Science University

Machine learning (ML) and artificial intelligence (AI)

- Overview
- Methods
- Results
- Future directions

Future directions

- Trust
- Fairness
- Educating clinicians
- Regulation and liability
- Implementation challenges
- Ethical and equitable AI
- Organizational statements
- Implementing in the real world

How do we ensure trustworthy AI? (Ammanath, 2022)

Fair / Impartial
AI applications should include checks from internal and external stakeholders to help ensure equitable application across all participants

Transparent / Explainable
All relevant individuals should understand how their data is being used and how AI systems make decisions; algorithms, attributes, and correlations should be open to inspection

Responsible / Accountable
Policies should outline governance and who is held responsible for all aspects of the AI solution (e.g., initiation, development, outputs, decommissioning)

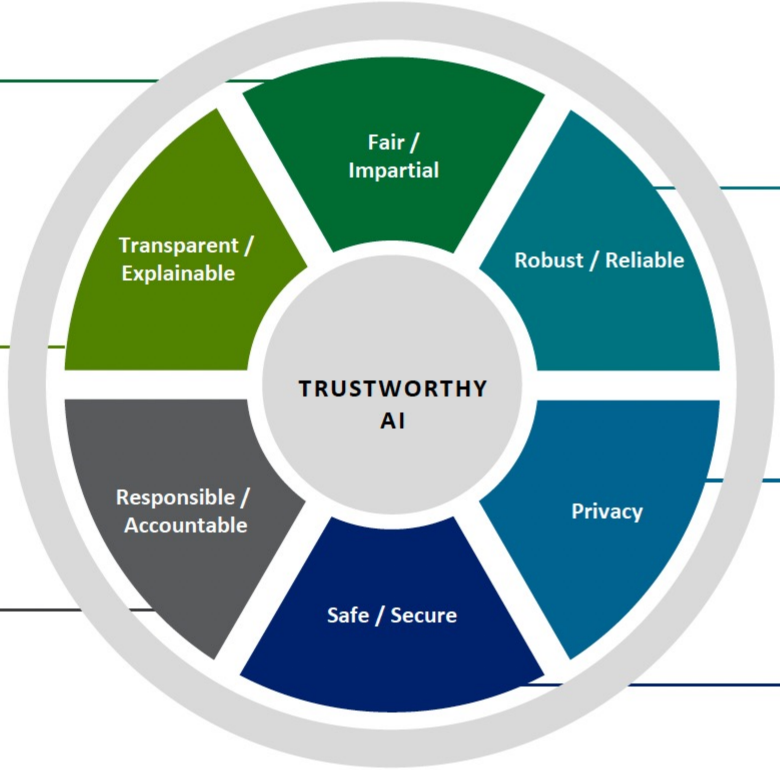

Robust / Reliable
AI systems should have the ability to learn from humans and other systems and produce accurate and reliable outputs consistent with the original design

Privacy
Individual, group, or entity privacy should be respected, and their data should not be used beyond its intended and stated use; data used has been approved by the data owner or steward

Safe / Secure
AI systems should be protected from risks (including Cyber) that may directly or indirectly cause physical and/or digital harm to any individual, group, or entity

HHS Trustworthy AI Playbook (2021)

How do we ensure trustworthy AI (cont.)?

- Recommendations from US government
  - Trustworthy AI playbook (HHS, 2021)
  - AI risk management framework (NIST 2021)
- Recommendations in specific uses
  - Clinical (Grote, 2020)
  - Clinical research (Volovici, 2022)
  - Genomics (Whalen, 2022)
- Other identified needs
  - "Nutrition fact sheet" about model and validation (Sendak, 2020; Cohen, 2022)
  - "Medical algorithm audit" (Liu, 2022)
  - Continuous monitoring and improvement (Feng, 2022)

AI Facts

Machine-learning algorithm to identify disease

Last updated: 03/07/2022

Training data	Dataset size, racial makeup
Model development	Algorithm type
Performance	False positives, false negatives
Assessments	Fairness, bias attestations
Validation studies	Safety, efficacy

How do we ensure fairness?

- Roadmap to fairer algorithms (Kusner, 2020)
- Framework (Xu, 2022) and toolkit (Bellamy, 2019) to insure fairness and reduce bias for individuals and groups
- Data ownership
  - Who owns patient photographs? (Davis, 2020; Davis, 2022)
  - De-identified data can be re-identified with other data; who should own resulting data? (Ross, 2022)
  - Data should be “democratized” (Allen, 2019), especially for disabled (Newman-Griffis, 2022) and historically underserved groups (Gottschalk, 2022; Hossain, 2023)

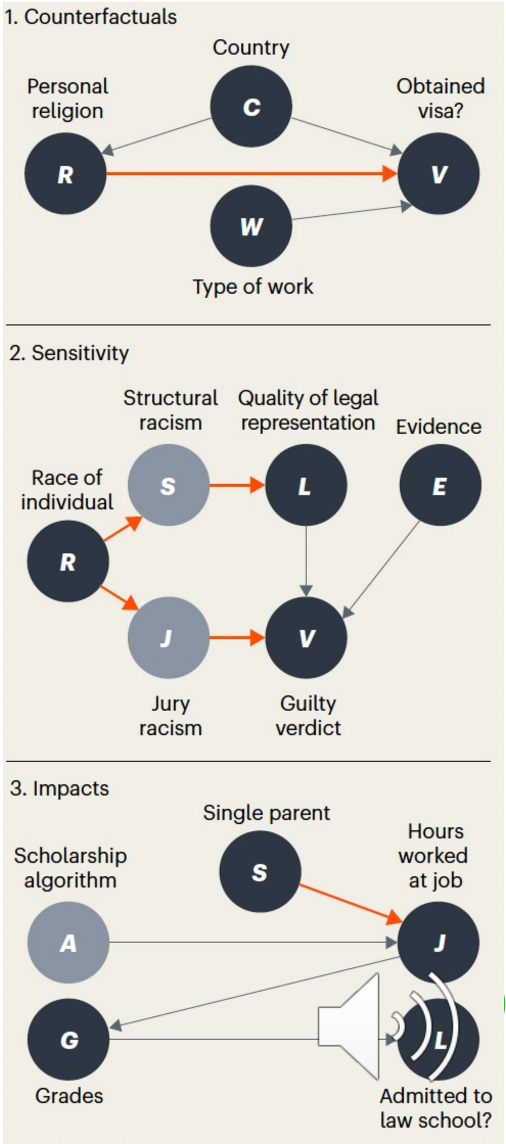

How do we educate clinician users of AI?

- Added as another informatics competency for clinicians (Hersh, 2020)
- Clinicians must be prepared for a clinical world influenced by AI (James, 2022)
- Medical schools may be “missing the mark” on AI (Palmer, 2023)
- AI should be taught as a “fundamental toolset of medicine” (Ötleş, 2022)
- AI in radiology education (Tejani, 2023)

Regulation and liability

- What is liability for use or non-use (Price, 2019)?
- Need lifecycle-based regulation – when to allow incremental changes vs. wholesale renewal (Hwang, 2019)
- AMIA position paper on "adaptive clinical decision support" (Petersen, 2021)
- For first 130 medical AI devices approved by US FDA, only 4 had undergone prospective evaluation, with number of evaluation sites and sample sizes often not reported (Wu, 2021)
- Over 500 and growing medical AI devices approved
  - https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medical-devices

FDA guidelines for when software is or is not medical device

Your Clinical Decision Support Software: Is It a Device?

FDA

The FDA issued a guidance, Clinical Decision Support Software, to describe the FDA's regulatory approach to Clinical Decision Support (CDS) software functions. This graphic gives a general and summary overview of the guidance and is for illustrative purposes only. Consult the guidance for the complete discussion and examples. Other software functions that are not listed may also be device software functions. *

Your software function must meet all four criteria to be Non-Device CDS.

Summary interpretation of CDS criteria

1. Your software function does **NOT** acquire, process, or analyze medical images, signals, or patterns.
2. Your software function displays, analyzes, or prints medical information normally communicated between health care professionals (HCPs).
3. Your software function provides recommendations (information/options) to a HCP rather than provide a specific output or directive.
4. Your software function provides the basis of the recommendations so that the HCP does not rely primarily on any recommendations to make a decision.

Your software function may be non-device CDS.

Non-Device Examples

Non-Device examples display, analyze, or print the following examples of medical information, which must also not be images, signals, or patterns:

- Information whose relevance to a clinical decision is well understood
- A single discrete test result that is clinically meaningful
- Report from imaging study

AND

Non-Device examples provide:

- Lists of preventive, diagnostic, or treatment options
- Clinical guidelines matched to patient-specific medical info
- Relevant reference information about a disease or condition

AND

Non-Device examples provide:

- Plain language descriptions of the software purpose, medical input, underlying algorithm
- Relevant patient-specific information and other knowns/unknowns for consideration

Device Examples

Device examples acquire, process, or analyze:

- Signal acquisition systems
- In vitro diagnostics
- Magnetic resonance imaging (MRI)
- Next Generation Sequencing (NGS)
- Continuous Glucose Monitoring (CGM)
- Computer aided detection/diagnosis (CADe/CADx)

OR

Device examples display, analyze or print:

- Continuous signals/patterns
- Medical images
- Waveforms (ECG)
- More continuous sampling (aka – a signal or pattern)

OR

Device examples provide:

- Risk scores for disease or condition
- Probability of disease or condition
- Time-critical outputs

OR

Device examples:

- Basis of recommendations is not provided

Your software function is a device.

https://www.fda.gov/medical-devices/software-medical-device-samd/your-clinical-decision-support-software-it-medical-device

WhatIs09

Implementation challenges

- Overcome issues of privacy, interoperability, adequate data for research, and more (Ching, 2018; Price, 2019)
- As models change healthcare practice, they must be adapted (Lenert, 2019)
- Healthcare organizations may not be able to implement (Panch, 2019)
  - May not be able to re-engineer care
  - Issues of data – who owns, can use, or maintain responsibility?
- Too much emphasis in research on algorithm success (“winner’s curse”) above all else (Sculley, 2018)?
- ML system more accurate than humans in robustly simulated mammography interpretation (McKinney, 2020), but may be wrong question, since value of screening is uncertain (Aschwanden, 2020)

Additional concerns

- Computationally intensive ML algorithms can generate notable carbon emissions (Lacoste, 2019)
  - About 3% of global energy use from computationally-intensive ML processing (Knowles, 2021)
  - Training single AI model can emit as much carbon as five cars in their lifetimes (Strubell, 2020)
  - Emerging principles for “green AI” (Schwartz, 2020)
  - Google best practices may reduce energy use 100-1000 fold (Patterson, 2022)
- Preserving privacy in learning data sets (Kaissis, 2020)
  - Methods for obfuscating facial images (Yang, 2021)
- Adversarial attacks on ML that might not be detectable by humans (Finlayson, 2019)
- Deepfakes create fake images, videos, and documents (Greengard, 2019)
  - Deepfake bot submissions to federal public comment websites cannot be distinguished from human submissions (Weiss, 2019)

Ethical AI

- AI and data ethics must be guided by justice and transparency (Basl, 2021)
- Call for responsible AI (Shneiderman, 2021; Shneiderman, 2022)
- Framework (Chen, 2020)

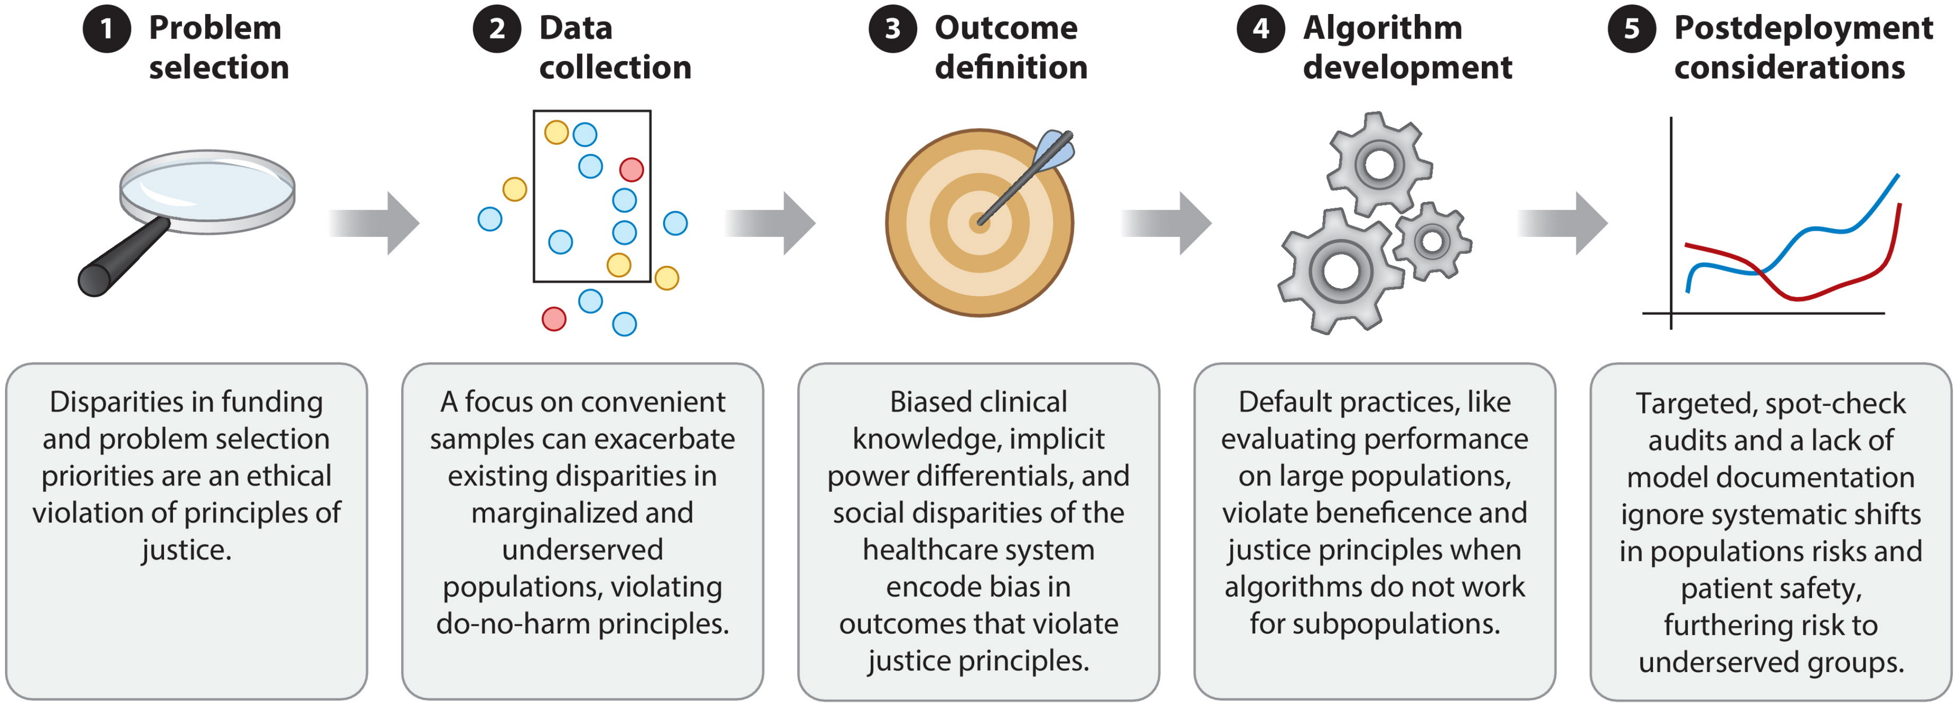

WhatIs09

Equity in AI

- Need framework to incorporate ethical AI principles into development process in ways that intentionally promote racial health equity and social justice (Dankwa-Mullan, 2021)
- NIH Artificial Intelligence/Machine Learning Consortium to Advance Health Equity and Researcher Diversity (AIM-AHEAD)
  - https://datascience.nih.gov/artificial-intelligence/aim-ahead
  - https://aim-ahead.net/
- Democratizing immunology datasets (Bhattacharya, 2021)
- Algorithmic Bias Playbook – how to define, measure, and mitigate racial bias in live algorithms (Ross, 2021; Obermeyer, 2021)

How do we move forward?

- Need "algorithmovigilance" (Embi, 2021)
- ML in healthcare must have trustworthiness, explainability, usability, and transparency (Cutillo, 2020)
- Moving beyond prediction of diagnostic labels to "wayfinding" (Adler-Milstein, 2021)
- Must shift ML for healthcare from development to deployment and from models to data (Zhang, 2022)
- AI can "illuminate the dark spaces of healthcare with ambient intelligence" (Haque, 2020)

Organizational statements about AI

- AMIA principles (Solomonides, 2021)
- WHO Ethics and Governance of AI for Health (2021)
  - https://apps.who.int/iris/rest/bitstreams/1352854/retrieve
- US White House, Blueprint for an AI Bill of Rights
  - https://www.whitehouse.gov/ostp/ai-bill-of-rights/
- ACM, Statement on Principles for Responsible Algorithmic Systems
  - https://www.acm.org/articles/bulletins/2022/november/tpc-statement-responsible-algorithmic-systems
- US GAO 22-104629, AI in Health Care
  - https://www.gao.gov/assets/gao-22-104629.pdf

AMIA principles for AI (Solomonides, 2022)

Rule	Principle	Definitions
I.	Autonomy	AI systems must protect the autonomy of all people and treat them with courtesy and respect including facilitating informed consent.
II.	Beneficence	AI systems must be helpful to people modeled after compassionate, kind, and considerate human behavior.
III.	Nonmaleficence	AI systems shall "do no harm" by avoiding, preventing, and minimizing harm or damage to any stakeholder.
IV.	Justice	AI systems must include equity for people in representation and access to AI, its data, and its benefits. AI must support social justice.
V.	Explainability	AI developers must describe AI systems in context-appropriate language so that their scope, proper application, and limitations are understandable.
VI.	Interpretability	AI developers must endow their systems with the functionality to provide plausible reasoning for decisions or advice in accessible language.
VII.	Fairness	AI systems must be free of bias and must be nondiscriminatory.
VIII.	Dependability	AI systems must be robust, safe, secure, and resilient. Failure must not leave any system in an unsafe or insecure state.
IX.	Auditability	AI systems must provide and preserve a performance "audit trail" including internal changes, model state, input variables, and output for any system decision or recommendation.
X.	Knowledge management	AI systems must be maintained including retraining of algorithms. AI models need listed creation, revalidation, and expiration dates.
Organizations deploying or developing AI		
XI.	Benevolence	Organizations deploying or developing AI must be committed to use AI systems for positive purposes.
XII.	Transparency	AI must be recognizable as such or must announce its nature. AI systems do not incorporate or conceal any special interests and deal even-handedly and fairly with all good faith actors.
XIII.	Accountability	AI systems must be the subject of active oversight by the organization, and any risk attributed to AI must be reported, assessed, monitored, measured, and mitigated as needed. Complaints and redress must be guaranteed.
Special considerations		
XIV.	Vulnerable populations	AI applied to vulnerable populations requires increased scrutiny to avoid worsening the power differential among groups.
XV.	AI research	Academic and industrial research organizations must continue to research AI to address inherent dangers as well as benefits.
XVI.	User education	AI developers have a responsibility to educate healthcare providers and consumers on machine learning and AI systems.

OHSU

Considerations for AI development

- Learning from earlier experience of mammography computer-aided diagnosis (Elmore, 2022)
  - AI-human interface
  - Reimbursement for improved outcomes, not usage
  - Incorporate ongoing improvements in systems
  - Cognizance of legal risk
- Bridging chasm between AI and clinical implementation (Aristidou, 2022)
  - Transparency in systems and data used
  - Deconstruct neural networks so clinicians understand predictions
  - Allow addition of local data for fine-tuning

Implementing in the real world

- Account for data variability across institutions (Balachandar, 2020)
- Consider impact on clinical workflows – especially nurses (Brodwin, 2020; Jung, 2020)
- Provide short overviews of purpose and potential harms (Sendak, 2020)
- Stewardship of algorithms for efficacy and safety (Eaneff, 2020)
- Methods to assess deployment into new settings (Nicora, 2022)
- Need "translational" AI from basic to clinical science (Hersh, 2021)

How will ML and AI impact healthcare?

- Physicians (Jha, 2016; Jha, 2018; Shah, 2019) and ML (Verghese, 2018) must adapt
- “AI won’t replace radiologists, but radiologists who use AI will replace radiologists who don’t,” (Langlotz, 2019)
  - True for all physicians, even Dr. McCoy?