AI, Machine-learning & Clinician Interactions

Doug McNair MD PhD

President, Cerner Math

November 30, 2017

Introduction

***"The CMIO role** is gradually being transformed, particularly in more advanced health organizations, from its early 'tech-head doc' role,
to a management role in EHR implementation,
to a transformational leadership role,
to a role catalyzing discovery and translational research. Skills are needed to help lead AI discovery, validation, curation, and authorization."***

***** and so also the roles of specialty Boards and Societies***

—Remark of AMDIS CMIO conference organizer, Ojai, California, 22-JUN-2017

Definitions & issues

- Artificial Intelligence (AI)
  - Computational emulation of effective [human] goal-oriented behavior, often with autonomy and emulating deliberation, intention
- Machine Learning (ML)
  - Several dozen methods of empirical pattern detection
- All of these have issues of quality, safety, clinical competence, efficiency, cost-effectiveness, fairness
- Derived from existing corpora of expertise and as-treated de-identified clinical data
  - Tend to exalt ‘conventional’ decisions, reasons, indications (high-prevalence data)
  - Tend to under-value innovations (low-prevalence data)

Possibilities for regulating clinical AI and their current chief limitations

- Traditional peer-review publication process
  - too slow to keep pace with rates of discovery and deployment
  - fraught with reviewer preconceptions/biases/alacrity
  - inadequate *specificity* (rejects things that are in fact effective, safe, beneficial)
- Traditional FDA CDRH de novo/510(k) process for IVDMIA devices
  - too slow, under-resourced
  - unreceptive to dynamic ML and AI
- Traditional CLIA process for LDT services
  - inadequate *sensitivity* to detect safety deficits/defects in AI (high Type II error rate, false-negatives, accepts things that are in fact ineffective, unsafe, low utility)
  - limited relation of certification rigor to different hazard levels and hazard mitigations
- Traditional court system (civil, criminal)
  - unclear allocation of liability, 'agency' issues
  - oriented to dealing with few cases (*qui tam* mass torts, e.g., airbags, autonomous self-driving cars, blockchain)
- Traditional court system (discrimination)
  - inadequate statistical methodologies to ascertain fairness of ML
  - not enough Administrative Law Judges who have statistical expertise
- Traditional standards organizations
  - too slow, under-resourced
  - unreceptive to numerous "long-tail" use-cases and dynamic ML and AI
- Traditional ABMS specialty Board certification (and MOC)
  - patient simulations not machine-executable/automated
  - insufficient simulations case databases' depth, breadth, and temporal detail

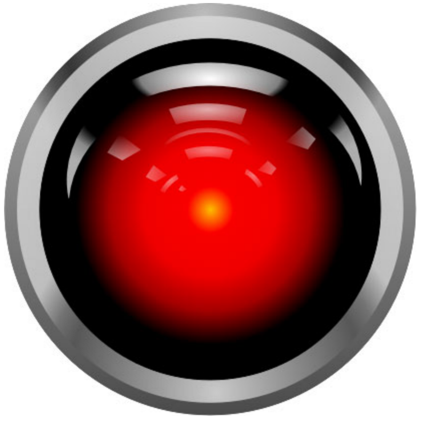

Additional possibilities

- CAP-style QC “CheckSample™” round-robin testing challenges (user fees)
- AAMI-style V&V testing against corpora of historical data (Physionet.org)
- Kaggle-style crowd-sourcing of testing, with CLIA oversight (awards)
- Amazon Mechanical Turk-style crowd-sourcing of testing (contingent labor)
- Other (“soft law” and standards, e.g., ISO, IEC, IEEE, ANSI, ASTM, NIST)

Autonomous self-driving cars regulation (NHTSA)

- **Level 0** – No Automation: Human driver is in complete and sole control. AI only plays advisory role. (examples: forward-collision warning, lane-departure warning, blind-spot monitoring).

- **Level 1** – Function-specific Automation: One or more specific control functions; if multiple functions are automated, they operate independently from each other. Human driver has overall control and is solely responsible for safe operation, but can choose to cede limited authority (adaptive cruise control, dynamic brake support in emergencies).

- **Level 2** – Combined-Function Automation: Automation of at least two primary control functions designed to work in unison to relieve the driver of control (adaptive cruise control in combination with lane-centering).

- **Level 3** - Limited Self-Driving Automation: Automation enabling the driver to (a) cede full control of all safety-critical functions under certain traffic or environmental conditions; in those conditions, to (b) rely heavily on the automated vehicle system to monitor for changes (driver does not constantly monitor the roadway while in motion); and (c) to resume control if change.

- **Level 4** - Full Self-Driving Automation: Automation where vehicle is designed to perform all safety-critical driving functions (includes both human-occupied and unoccupied vehicles; safe operation rests solely on the automated vehicle system).

Impressions

- Current legal and regulatory systems' methods of addressing effectiveness and injury from devices and services achieve appropriate balance of innovation and liability for benefit and injury
- Current E&O insurance for 'strict product liability' and for 'negligence' services liability will be underwritable for AI; possible 'no-fault' policies
- Systems are neither flawed nor need fundamental change
- Conclude that:
  - Legal, regulatory, and insurance systems will continue to foster innovation of reasonably-safe sophisticated AI for health use-cases
  - Additional resources are needed to insure *responsiveness* and *efficiency* of certification

Recommendations

- Multi-modal regulation of clinical ML-based AI
  - FDA-style for ‘devices’ and CLIA-style certification for ‘services’
  - Participative governance by NAM, ABMS
  - Voluntary standards (‘soft law’) + public simulations datasets for V&V
  - Evolve statistical methods for evaluating fairness of ML, AI
  - Align with changes in tort law and policy re: self-driving cars

Thank you!

1. A robot may not injure a human being, or, through inaction, allow a human being to come to harm.

2. A robot must obey the orders given it by human beings except where such orders would conflict with the First Law.

3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.

—Isaac Asimov, The Naked Sun (1957), 'Laws of Robotics,' pp.31-33.