



E. What if My Doctor Is a Robot?

Gary Marchant
Regents' Professor; Lincoln Professor of Emerging
Technologies, Law & Ethics
Arizona State University
Phoenix, AZ

What If My Doctor Is a Robot? Artificial Intelligence and the Practice of Medicine

**AHLA Annual Conference
June 25 and 26, 2018**


Gary E. Marchant, Ph.D., J.D.
gary.marchant@asu.edu

Outline


- What is AI?
- Applications in Healthcare:
 - Administration
 - Diagnosis
 - Treatment
- Concerns and Issues:
 - Transparency/Bias
 - Hype
 - Liability
 - Privacy
 - Strategy

What is AI?


- A machine that displays intelligent behavior, such as reasoning, learning and sensory processing.
- AI involves tasks that have historically been limited to humans and intelligent animals, such as decision-making and problem-solving.
- Two types:
 - Narrow AI
 - Artificial General Intelligence (AGI)



A Short History of A.I.



Components of AI



AI Component: Natural Language Processing




Original Article

Correlating Mammographic and Pathologic Findings in Clinical Decision Support Using Natural Language Processing and Data Mining Methods

Todd A. Patel, MD^{1,2,3}, Sherry Papadopoulos, MD^{1,2,3}, Richard D. Ogden, MD^{1,2,3}, Joe E. Enos, MD^{1,2,3}, Tiancheng He, MD^{1,2,3}, Jason B. Shewchuk, MD^{1,2,3}, Joseph A. Stewart, MD^{1,2,3}, Virginia G. Stadelman, MD^{1,2,3}, Angel A. Rodriguez, MD^{1,2,3}, Stephen T. C. Wong, MD^{1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17,18,19,20,21,22,23,24,25,26,27,28,29,30,31,32,33,34,35,36,37,38,39,40,41,42,43,44,45,46,47,48,49,50,51,52,53,54,55,56,57,58,59,60,61,62,63,64,65,66,67,68,69,70,71,72,73,74,75,76,77,78,79,80,81,82,83,84,85,86,87,88,89,90,91,92,93,94,95,96,97,98,99,100}

BACKGROUND: A key challenge in using electronic health records for mammography research is the interpretation of natural language text, which often contains complex medical terminology. This study aims to develop a natural language processing (NLP) system that can automatically extract and correlate mammographic and pathologic findings in clinical decision support. **METHODS:** The authors searched the literature for data on NLP systems for mammography. They identified 10 studies that met the inclusion criteria. The authors then developed a NLP system that can automatically extract and correlate mammographic and pathologic findings in clinical decision support. **RESULTS:** The authors found that the NLP system was able to extract and correlate mammographic and pathologic findings in clinical decision support. **CONCLUSIONS:** The authors conclude that the NLP system was able to extract and correlate mammographic and pathologic findings in clinical decision support. **KEYWORDS:** NLP, mammography, pathologic findings, clinical decision support.



AI Component: Vision




cit:1711.05225v1 [cs.CV] 14 Nov 2017

CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning

Pranav Rajpurkar^{1,2} Jeremy Irvin^{1,2} Kayle Zhu¹ Brandon Yang¹ Hershel Mehta¹
Tony Duan³ Daley Ding⁴ Aarti Bagul¹ Curtis Langlotz⁵ Katie Spinkovska²
Matthew P. Lungren² Andrew Y. Ng¹

Abstract

We develop an algorithm that can detect pneumonia from chest X-rays at a level exceeding practicing radiologists. Our algorithm, CheXNet, is a 121-layer convolutional neural network trained on ChestX-ray14, currently the largest publicly available chest X-ray dataset, containing over 100,000 frontal-view X-ray images with 14 diseases. Four practicing academic radiologists annotate a test set, on which we compare the performance of CheXNet to that of radiologists. We find that CheXNet exceeds average radiologist performance on pneumonia detection on both sensitivity and specificity. We retuned CheXNet to detect all 14 diseases in ChestX-ray14 and achieve state of the art results on all 14 diseases.



1. Introduction

More than 1 million adults are hospitalized with pneumonia and around 50,000 die from the disease every year in the United States.

AI Component: Machine Learning



How Google's AlphaGo Beat a Go World Champion

Inside a man-versus-machine showdown
CHRISTOPHER WICKER — MARCH 20, 2016

Source: The Atlantic

AlphaGo Zero: Google DeepMind supercomputer learns 3,000 years of human knowledge in 40 days



Chinese Go player Ke Jie competes against Google's artificial intelligence (AI) program, AlphaGo
CREDIT: BANGORCHINA/REUTERS/GETTY IMAGES

Source: The Telegraph

AI: Machine Learning

Computer Programming


- Machine mechanically implements human made code
- Bad outcomes are attributable to faults or oversights in code
- Human programmer can be questioned to explain why machine did what it did

Machine Learning

- Humans provide data sources and specify overall goal for machine
- Machine self-learns and adapts its approach to maximize specified goal
- Only machine knows why it did what it did

A doctor reads about a half dozen medical research papers in a monthwhereas Watson can read a half million in about 15 seconds.”





Plus Many Big Companies (Plus IBM) Investing Heavily in AI Health

Amazon has a secret health tech team called 1492 working on medical records, virtual doc visits

- Amazon has a secret skunkworks lab called 1492, dedicated to health care tech.
- Areas of exploration include a platform for electronic medical record data, telemedicine and health apps for existing devices like the Amazon Echo.

Eugene Kim | Christina Farr
Published 5:30 PM ET Wed, 26 July 2017 | Updated 2:30 PM ET Thu, 27 July 2017

CNBC

Microsoft announces new AI-powered health care initiatives targeting cancer

By Lauren Hsu | Updated 1:00 PM ET Wed, 26 July 2017

Facebook Twitter LinkedIn



Connected Health

Timeline: How Apple is piecing together its secret healthcare plan

From the launch of its HealthKit API in September 2014 to revelations this month of EHR-like work with a tiny startup, it's clear that healthcare is in Apple's eye.



By Bill Shwartz | June 23, 2017 | 11:16 AM

Facebook Twitter LinkedIn

Apple Acquisition Timeline



Rumors are at a fever pitch that Apple has big plans for healthcare, including putting a medical record on the iPhone, possibly acquiring its way into the EHR market.



- “Health care is a promising market for AI. There is enormous potential in its ability to draw inferences and recognize patterns in large volumes of patient histories, medical images, epidemiological statistics, and other data. AI has the potential to help doctors improve their diagnoses, forecast the spread of diseases, and customize treatments. Artificial intelligence combined with health care digitization can allow providers to monitor or diagnose patients remotely as well as transform the way we treat the chronic diseases that account for a large share of health-care budgets.”



FDA moves to encourage A.I. in medicine, drug development

[cnbc.com/2018/04/26/fda-moves-to-encourage-a-i-in-medicine-drug-development.html](https://www.cnbc.com/2018/04/26/fda-moves-to-encourage-a-i-in-medicine-drug-development.html)

Meg Tirrell

March 3, 5877521




AI in Health Care Administration

Health Care Billing/Records

- Machine learning systems starting to be applied to payments and claims management
 - Automated entry and integration of medical codes, patient contact, patient insurer info and coverage info
- Voice recognition systems can help doctors and their assistants automatically enter EHR data





Artificial Intelligence

AI chatbots might be the money-savers hospitals are looking for

But the technologies, which tap into natural language processing, knowledge management and sentiment analysis capabilities, require commitment and discipline to use well.

e.g., scheduling appointments
e.g., hospital admission
e.g., referral to specialists

By **Bill Siwicki** | September 18, 2017 | 11:16 AM




Like other innovations, AI chatbots in healthcare will be a crawl-walk-run endeavor, where the easier tasks will move to chatbots while awaiting the tech to evolve for more complex work.

Source:
HealthCareIT News

Scanadu cofounders' new project, doc.ai, is a conversational robot that explains lab results

 healthcareitnews.com/news/scanadu-launches-new-project-using-conversational-robot-explain-lab-results

8/25/2017



Thanks to AI, you'll never be ignored at a Hospital Again



Source Cbsistatic.com

AI helping in our Hospitals

IBM's Watson takes on the scut work at a Philadelphia hospital, so nurses can focus on what matters.

Disease Diagnosis

FDA News Release

FDA permits marketing of artificial intelligence algorithm for aiding providers in detecting wrist fractures

For Immediate Release

May 24, 2018

Summary

FDA permits marketing of artificial intelligence algorithm for aiding providers in detecting wrist fractures

Release

Today, the U.S. Food and Drug Administration permitted marketing of Imagen OsteoDetect, a type of computer-aided detection and diagnosis software designed to detect wrist fractures in adult patients.

"Artificial intelligence algorithms have tremendous potential to help health care providers diagnose and treat medical conditions," said Robert Ochs, Ph.D., acting deputy director for radiological health, Office of In Vitro Diagnostics and Radiological Health in the FDA's Center for Devices and Radiological Health. "This software can help providers detect wrist fractures more quickly and aid in the diagnosis of fractures."

Stanford algorithm can diagnose pneumonia better than radiologists

news.stanford.edu/2017/11/15/algorithm-outperforms-radiologists-diagnosing-pneumonia/

By Taylor Kubota


Stanford researchers have developed an algorithm that offers diagnoses based off chest X-ray images. It can diagnose up to 14 types of medical conditions and is able to diagnose pneumonia better than expert radiologists working alone. A paper about the algorithm, called CheXNet, was published Nov. 14 on the open-access, scientific preprint website arXiv.




AI detects stroke, dementia from brain scans

[UPI upi.com/Health_News/2018/05/16/AI-detects-stroke-dementia-from-brain-scans/1271526481113](https://www.upi.com/Health_News/2018/05/16/AI-detects-stroke-dementia-from-brain-scans/1271526481113)

May 16, 2018



May 16 (UPI) -- Artificial intelligence has been used to detect the most common causes of dementia and stroke -- small vessel damage, according to a study.



Artificial intelligence may help prevent heart failure.

Devrim/Stockphoto

Self-taught artificial intelligence beats doctors at predicting heart attacks

By Matthew Hutson | Apr. 14, 2017, 3:30 PM

Doctors have lots of tools for predicting a patient's health. But—as even they will tell you—they're no match for the complexity of the human body. Heart attacks in particular are hard to anticipate. Now, scientists have shown that computers capable of teaching themselves can perform even better than standard medical guidelines, significantly increasing prediction rates. If implemented, the new method could save thousands or even millions of lives a year.

"I can't stress enough how important it is," says Elsie Ross, a vascular surgeon at Stanford University in Palo Alto, California, who was not involved with the work, "and how much I really hope that doctors start to embrace the use of artificial intelligence to assist us in care of patients."

26 Mar 2018 | 17:00 GMT

AI Cardiologist Aces Its First Medical Exam

A neural network outperforms human cardiologists in a task involving heart scans

By Eliza Strickland



Photo-illustration: Rima Arnaout

Source:
IEEE Spectrum

JAMA | Original Investigation

Diagnostic Assessment of Deep Learning Algorithms for Detection of Lymph Node Metastases in Women With Breast Cancer

Babak Ehteshami Bejnordi, MS; Mitko Veta, PhD; Paul Johannes van Diest, MD, PhD; Bram van Ginneken, PhD; Nico Karssemeijer, PhD; Geert Litjens, PhD; Jeroen A. W. M. van der Laak, PhD; and the CAMELYON16 Consortium

IMPORTANCE Application of deep learning algorithms to whole-slide pathology images can potentially improve diagnostic accuracy and efficiency.

OBJECTIVE Assess the performance of automated deep learning algorithms at detecting metastases in hematoxylin and eosin-stained tissue sections of lymph nodes of women with breast cancer and compare it with pathologists' diagnoses in a diagnostic setting.

DESIGN, SETTING, AND PARTICIPANTS Researcher challenge competition (CAMELYON16) to develop automated solutions for detecting lymph node metastases (November 2015–November 2016). A training data set of whole-slide images from 2 centers in the Netherlands with (n = 110) and without (n = 160) nodal metastases verified by immunohistochemical staining were provided to challenge participants to build algorithms. Algorithm performance was evaluated in an independent test set of 129 whole-slide images (49 with and 80 without metastases). The same test set of corresponding glass slides was also evaluated by a panel of 11 pathologists with time constraint (WTC) from the Netherlands to ascertain likelihood of nodal metastases for each slide in a flexible 2-hour session, simulating routine pathology workflow, and by 1 pathologist without time constraint (WOTC).

← Editorial page 2184

← Related articles page 2211 and page 2250

+ Supplemental content

+ CME Quiz at jamanetwork.com/learning and CME Questions page 2252

AI beats doctors at visual diagnosis, observes many times more lung cancer signals

By Graham Templeton on August 10, 2016 at 1:00 pm | 26 Comments

725 shares [f](#) [t](#) [G+](#) [e](#) [Y](#)



A new study from Stanford University could ruffle feathers in the medical community, as the researchers report that their newly developed machine learning algorithm can identify tissue slides exhibiting a specific type of cancer with far greater accuracy than human epidemiologists. It's one of the first indications that computers aren't just capable of addressing the "subjective" portions of medicine, but that they're actually better suited to such problems than human doctors, in some cases.

FRONT PAGE

ALL NEWS

TOPICS

MULTIMEDIA

Email

Share 1.2K

Tweet

Computers trounce pathologists in predicting lung cancer type, severity

Automating the analysis of slides of lung cancer tissue samples increases the accuracy of tumor classification and patient prognoses, according to a new study.

AUG 16

2016

Computers can be trained to be more accurate than pathologists in assessing slides of lung cancer tissues, according to a new study by researchers at the [Stanford University School of Medicine](#).

The researchers found that a machine-learning approach to identifying critical disease-related features accurately differentiated between two types of lung cancers and predicted patient survival times better than the standard approach of pathologists classifying tumors by grade and stage.

"Pathology as it is practiced now is very subjective," said [Michael Snyder](#), PhD, professor and chair of genetics. "Two highly skilled pathologists assessing the same slide will agree only about 60 percent of the time. This approach replaces this subjectivity with sophisticated, quantitative measurements that we feel are likely to improve patient outcomes."

Assessing a biopsied slice of tissue to determine the grade and severity of a tumor can be highly subjective, but Stanford researchers found that computers could be trained to make accurate assessments of lung cancer tissue.

Science Photo/Shutterstock

SCIENCE \ TECH \ ARTIFICIAL INTELLIGENCE \

Artificial intelligence can spot skin cancer as well as a trained doctor

By [James Vincent](#) | [@jvincent](#) | Jan 26, 2017, 7:45am EST

f

t

SHARE

Photo by Joe Raedle/Getty Images

Researchers at Stanford University have created an [AI algorithm](#) that can identify skin cancer as well as a professional doctor. The program was trained on nearly 130,000 images of moles, rashes, and lesions using a technique known as deep learning. It was then tested head-to-head against 21 human dermatologists, where its creators say it performed with an accuracy on par with humans ("at least" 91 percent as good). In the future, they suggest it could be used to create a mobile app for spotting skin cancer at home.

News » UK » UK Politics

Theresa May says AI revolution will help NHS prevent thousands of cancer-related deaths by 2033

Medical records, along with information about patients' habits and genetics, will be cross-referenced with national data to spot those at an early stage of cancer

Ashley Cowburn Political Correspondent | @ashcowburn | 2 days ago | 73 comments



Like Click to follow
Indy Politics



The prime minister will say that late diagnosis of otherwise treatable illnesses is one of the biggest causes of avoidable deaths. PA

Zebra Medical offers affordable access to deep-learning, image-analysis tools

by Aminah Al Idrus | Oct 27, 2017 11:33am



This is the latest in Zebra Medical's quest to meet the growing demand for radiology services.

+
f
t
in

Zebra Medical Vision is on a mission to bridge the gap between supply and demand for radiology. To that end, the company is rolling out a new offering that includes all of its image analytics algorithms to healthcare providers for a flat rate per scan.

"With this new model, we hope to facilitate adoption globally, especially in countries where access to radiology is difficult," said Zebra CEO Elad Benjamin, in a statement. "We are making a commitment to provide our current tools, and all future ones, for a flat \$1 USD per scan. By doing so we believe that a true difference can be made in the provision of radiology services worldwide."

Using AI to Invent New Medical Tests

by John J. Dillon and Paul A. Friedman

MARCH 26, 2018



FDA approves AI-powered software to detect diabetic retinopathy

The software can be used by health care providers who don't typically deal with eye care.



Rob LeFebvre, @roblef
04.11.18 in [Medicine](#)

0
Comments

356
Shares



New Machine Learning Program Shows Promise For Early Alzheimer's Diagnosis

Tue, 08/15/2017 - 1:37pm by Case Western Reserve University



A new machine learning program developed by researchers at Case Western Reserve University appears to outperform other methods for diagnosing Alzheimer's disease before symptoms begin to interfere with every day living, initial testing shows.

AI can identify Alzheimer's disease a decade before symptoms appear

[T thetimes.co.uk/edition/news/ai-can-identify-alzheimer-s-a-decade-before-symptoms-appear-9b3qdrf7](https://www.thetimes.co.uk/edition/news/ai-can-identify-alzheimer-s-a-decade-before-symptoms-appear-9b3qdrf7)

Mark Bridge, Technology Correspondent



Artificial intelligence can identify the signs of Alzheimer's in scans a decade before symptoms show, scientists say.

Google can predict your risk of a heart attack by scanning your retina



A system has been developed by Google that can see the risks of cardiovascular disease through the eyes.

Image: REUTERS/Staffan Wierum

This article is published in collaboration with **Business Insider**

22 Feb 2018

Kevin Loria
Writer, Business Insider



Latest Articles

This is the biggest threat to Latin America's digital transformation
Jordi Boffill 13 Mar 2018

Meet IBM's Simon, the world's first smartphone
Rob Smith 13 Mar 2018

Your eyes might be the perfect windows into your heart.

At least, they're windows that Google-created artificial intelligence software can use to calculate your risk factors for heart disease.

According to a study recently published in the *Nature Biomedical Engineering* journal, an AI algorithm created by Google AI and Verily Life Sciences (an Alphabet subsidiary that spun off from Google) can predict whether a patient is likely to suffer a major cardiovascular event like a heart attack or stroke within five years, based on a photo of their retina.

So far, the predictions work about as well as presently accepted methods that are more invasive, according to the study.

Learning to predict heart disease

The fact that disease can be spotted in the retina isn't a surprise. Doctors often spot medical conditions including diabetes, extreme high blood pressure, high cholesterol, and some cancers during eye exams.

To mimic that ability, the Verily and Google researchers trained AI software to identify cardiovascular risks by having the system analyze retina photos and

Methodology (Results)

Predicting the Risk of Suicide by Analyzing the Text of Clinical Notes

Chris Poulin^{1,2,3,4}, Brian Shiner^{1,2}, Paul Thompson^{1,2}, Linus Vehtari², Yinyoung Young^{2,3}, Benjamin Goertzel², Bradley Wirth², Laura Flashman², Thomas McAllister²

¹ The Great School of Medicine at Dartmouth College & The Thayer School of Engineering at Dartmouth College, Hanover, New Hampshire, United States of America, ² The Dartmouth-Hitchcock Medical Center, Hanover, New Hampshire, United States of America, ³ Center for Data Science, Dartmouth College, Hanover, New Hampshire, United States of America, ⁴ Center for Data Science, Dartmouth College, Hanover, New Hampshire, United States of America

Abstract

We developed linguistics-driven prediction models to estimate the risk of suicide. These models were generated from unstructured clinical notes taken from a national sample of U.S. Veterans Administration (VA) medical records. We created three matched cohorts: veterans who committed suicide, veterans who used mental health services and did not commit suicide, and veterans who did not use mental health services and did not commit suicide during the observation period (n=72 in each group). From the clinical notes, we generated datasets of single keywords and multi-word phrases, and constructed prediction models using a machine-learning algorithm based on a genetic programming framework. The resulting inference accuracy was consistently 85% or more. Our data therefore suggests that computational text analysis can be applied to unstructured medical records to estimate the risk of suicide. The resulting system could allow clinicians to potentially screen seemingly healthy patients at the primary care level, and to continuously evaluate the suicide risk among psychiatric patients.

DOI: [10.1371/journal.pone.0180000](https://doi.org/10.1371/journal.pone.0180000)

Editor: Nathan Bress, Drexel University, UNITED STATES OF AMERICA

Received: July 10, 2017; **Accepted:** November 20, 2017; **Published:** January 26, 2018

Copyright: © 2018 Poulin et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Funding: This research was funded by the National Institutes of Health (NIH) through the Department of Veterans Affairs (VA) Center for Data Science (CDCS). The authors thank the VA for providing the data and for supporting this research. The authors also thank the VA for providing the data and for supporting this research. The authors also thank the VA for providing the data and for supporting this research.

Competing Interests: Dr. Poulin is President of Veterans and Physicians, which competes in competition with P1 and P2. He has a patent pending in relation to the research described in this paper. There are no other competing interests in development or related products to disclose. This does not alter our adherence to PLOS ONE policies on sharing data and materials.

^{*} cpoulin@dartmouth.edu

Predicting the risk of suicide by analyzing the text of clinical notes. PLOS ONE, March 20, 2014


Artificial Intelligence in Behavioral and Mental Health Care

Edited by David D. Luxton

Copyright © 2015 Elsevier Inc. All rights reserved.

Artificial Intelligence in Behavioral and Mental Healthcare, Elsevier/Academic Press, October 2, 2015

Figure. The Process of Digital Phenotyping




Thomas Insel
JAMA, 2017

What is the Human Diagnosis Project?

The Human Diagnosis Project (also referred to as "Human Dx" or "the Project") is a worldwide effort created with and led by the global medical community to build an online system that maps the best steps to help any patient. By combining collective intelligence with machine learning, Human Dx intends to enable more accurate, affordable, and accessible care for all.

[Learn More](#)



Partners include:

American Medical Association
American Board of Internal Medicine
American Board of Medical Specialties
Association of American Medical Colleges

Treatment

IBM's Watson artificial intelligence system could improve cancer treatment


IANS | Updated: Aug 08, 2017, 01:21 PM IST



Watson has the ability to rapidly identify multiple variations in BRAF along with variations in nearly a dozen other genes known to contribute to the skin cancer.

NEW YORK: IBM's Watson artificial intelligence system could help persons with cancer by providing a rapid, comprehensive report of the genetic mutations at the root of their specific disease and the therapies that target them, says an Indian-origin pathologist.

"We have to change our whole behaviour in looking at tumours. We are missing too much and too often treatment does not work for patients," said Ravindra Kolhe, breast and molecular pathologist at the Medical College of Georgia at Augusta University in the US.



NEUROREHABILITATION

A multidirectional gravity-assist algorithm that enhances locomotor control in patients with stroke or spinal cord injury

Jean-Baptiste Mignardot,^{1,2*} Camille G. Le Goff,^{1,2*} Rubia van den Brand,^{1,2*} Marco Capogrosso,^{1,2} Nicolas Fumeaux,¹ Heike Vallery,³ Selin Anil,¹ Jessica Lanini,⁴ Isabelle Fodor,⁵ Grégoire Eberle,⁵ Auke Ijspeert,⁴ Brigitte Schurch,⁶ Armin Curt,⁷ Stefano Carda,^{2,5} Jocelyne Bloch,^{2,8†} Joachim von Zitzewitz,^{1†} Grégoire Courtine^{1,8†*}

Gait recovery after neurological disorders requires remastering the interplay between body mechanics and gravitational forces. Despite the importance of gravity-dependent gait interactions and active participation for promoting this learning, these essential components of gait rehabilitation have received comparatively little attention. To address these issues, we developed an adaptive algorithm that personalizes multidirectional forces applied to the trunk based on patient-specific motor deficits. Implementation of this algorithm in a robotic interface reestablished gait dynamics during highly participative locomotion within a large and safe environment. This multidirectional gravity-assist enabled natural walking in nonambulatory individuals with spinal cord injury or stroke and enhanced skilled locomotor control in the less-impaired subjects. A 1-hour training session with multidirectional gravity-assist improved locomotor performance tested without robotic assistance immediately after training, whereas walking the same distance on a treadmill did not ameliorate gait. These results highlight the importance of precise trunk support to deliver gait rehabilitation protocols and establish a practical framework to apply these concepts in clinical routine.

Improving Palliative Care with Deep Learning

Anand Avati*, Kenneth Jung†, Stephanie Harman‡, Lance Downing†, Andrew Ng* and Nigam H. Shah‡

*Dept of Computer Science, Stanford University

Email: {avati,ang}@cs.stanford.edu

†Center for Biomedical Informatics Research, Stanford University

Email: {kjung,l Downing,nigam}@stanford.edu

‡Dept of Medicine, Stanford University School of Medicine

Email: {smharman}@stanford.edu

Abstract— Improving the quality of end-of-life care for hospitalized patients is a priority for healthcare organizations. Studies have shown that physicians tend to over-estimate prognoses, which in combination with treatment inertia results in a mismatch between patients wishes and actual care at the end of life . We describe a method to address this problem using Deep Learning and Electronic Health Record (EHR) data, which is currently being piloted, with Institutional Review Board approval, at an academic medical center. The EHR data of admitted patients are automatically evaluated by an algorithm, which brings patients who are likely to benefit from palliative care services to the attention of the Palliative Care team. The algorithm is a Deep Neural Network trained on the EHR data from previous years, to predict all-cause 3-12 month mortality of patients as a proxy for patients that could benefit from palliative care. Our predictions enable the Palliative Care team to take a proactive approach in reaching out to such patients, rather than relying on referrals from treating physicians, or conduct time consuming chart reviews of all patients. We also present a novel interpretation technique which we use to provide explanations of the model's predictions.


sive care. Second, a shortage of palliative care professionals makes proactive identification of candidate patients via manual chart review an expensive and time-consuming process.

The criteria for deciding which patients benefit from palliative care can be hard to state explicitly. Our approach uses deep learning to screen patients admitted to the hospital to identify those who are most likely to have palliative care needs. The algorithm addresses a proxy problem - to predict the mortality of a given patient within the next 12 months - and use that prediction for making recommendations for palliative care referral. This frees the palliative care team from manual chart review of every admission and helps counter the potential biases of treating physicians by providing an objective recommendation based on the patient's EHR. Currently existing tools to identify such patients have limitations, and they are discussed in the next section.

II. RELATED WORK


Legal/Policy Issues

Transparency/Bias



Big Data Cautions

- Machine learning approaches powered by strong, but theory free, associations in the data
- Can be very powerful tool for discerning previously unrecognized associations and relationships
- But:
 - may give wrong answer if future different from past
 - Data may contain unrealized biases
 - Chaos theory



A Cautionary Example

Intelligible Models for HealthCare: Predicting Pneumonia Risk and Hospital 30-day Readmission

Rich Caruana
Microsoft Research
rcaruana@microsoft.com

Yin Lou
LinkedIn Corporation
ylou@linkedin.com


Johannes Gehrke
Microsoft
johannes@microsoft.com

Paul Koch
Microsoft Research
paulkoch@microsoft.com

Marc Sturm
NewYork-Presbyterian Hospital
mas9161@nyp.org

Noémie Elhadad
Columbia University
noemie.elhadad@columbia.edu

- Hospital used AI system to predict probability of death for patients with pneumonia
- Overall results better predicted outcomes than existing systems
- But system predicted asthmatic pneumonia patients were low-risk and should be treated as outpatients
- In fact, such patients were immediately put into intensive care, which explains better outcome



Hype



Note: PaaS = platform as a service; UAVs = unmanned aerial vehicles

Source: Gartner (July 2017)

A STAT INVESTIGATION

IBM pitched its Watson supercomputer as a revolution in cancer care. It's nowhere close

By CASEY ROSS @byCaseyRoss and IKE SWETLITZ @ikeswetlitz / SEPTEMBER 5, 2017

Breathlessly promoting its signature brand — Watson — IBM sought to capture the world's imagination, and it quickly zeroed in on a high-profile target: cancer.

But three years after IBM began selling Watson to recommend the best cancer treatments to doctors around the world, a STAT investigation has found that the supercomputer isn't living up to the lofty expectations IBM created for it. It is still struggling with the basic step of learning about different forms of cancer. Only a few dozen hospitals have adopted the system, which is a long way from IBM's goal of establishing dominance in a multibillion-dollar market. And at foreign hospitals, physicians complained its advice is biased toward American patients and methods of care.

Liability

New Medical Technology & Liability

- “Although technology is generally seen as a boon to safety, no other factor historically has surpassed it as a stimulus for litigation. Gains in clinical competence redefine success upward and make delay actionable.”

- William Sage, Medical Liability and Patient Safety


Why Do New Health Technologies Increase Medical Malpractice Liability?

- Glitches and defects in the new technology
- Providers unfamiliar and inadequately trained on new technologies
- Unrealistic patient expectations
- The more you can do, the more you can do wrong (eg kidney dialysis)
- Too slow uptake by doctors
- Gaps between early- and late-adopters

Three Important Liability Issues

- Learning curves
- Tipping points
- Standard of care

Learning Curves



CONTACT US FOR A **FREE** CONSULTATION 855-ROBOT-101

SCHLESINGER
LAW OFFICES, P.A.

HAVE YOU UNDERGONE A
SURGICAL PROCEDURE USING
THE **DAVINCI ROBOT** AND
SUFFERED COMPLICATIONS?



**DID YOU SUFFER AN INJURY
AS A RESULT OF THE
DA VINCI SURGICAL ROBOT?
CONTACT US TODAY!**

Da Vinci
Robot
Lawsuits

Have you been injured by
Robotic Surgery?
[Click here for a free case review](#)

No Win, No Fee

We've Been
Featured On

Bloomberg

WABC
Radio

The
New York
Times

THE WALL STREET
JOURNAL

FOX
NEWS


THE SACRAMENTO BEE

UNIVERSITY
OF CALIFORNIA

San
Antonio
Express-News

THE OREGONIAN

Tipping Point



Standard of Care

- With AI, machines are replacing physicians
- What standard of care applies to a machine?



FEDERAL TRADE COMMISSION
PROTECTING AMERICA'S CONSUMERS

Data Privacy

Data is the key input for
machine learning AI

--

Als “learn” from
massive, messy data sets

Data Privacy Issues


Obtaining the data
Storing the data securely
Using the data



Strategy


How Should Health Care Entities Respond to AI?

- If adopt AI too early, invest time, resources & data into system and/or vendor that is dead end
- If adopt AI too slowly, fail to keep up with competition
 - Especially given learning curve of machine learning



--	--	--

Conclusion



Dr. Robot

Jane R. Bambauer*

TABLE OF CONTENTS

I. THE SEARCH FOR THE APPROPRIATE BASELINE.....	385
II. THE SPECIAL DUTIES OF DR. ROBOT	391
A. <i>Competence</i>	391
B. <i>Confidentiality and Duties to Warn</i>	394
C. <i>Research and Informed Consent</i>	395
D. <i>Conflicts of Interest</i>	397
CONCLUSION.....	397

Predicting the future is a surefire way to embarrass oneself. But it is a relatively safe bet that Artificial Intelligence (“AI”) will transform the practice of healthcare to some appreciable degree. My confidence in this prediction stems from the fact that healthcare is already being transformed by AI. (Predicting the present, it turns out, is somewhat easier than predicting the future.) Mobile apps provide customized instructions to patients suffering from a range of illnesses,¹ and IBM’s Watson co-counsels patients (alongside their doctors) on particularly complex cases.² The prediction is also sound because modern

* Copyright © 2017 Jane R. Bambauer. Professor of Law, University of Arizona James E. Rogers College of Law. I am grateful to the UC Davis Law Review and to Anupam Chander for the opportunity to contribute to the Future-Proofing Law symposium. Special thanks are also owed to Derek Bambauer, Glenn Cohen, Margot Kaminski, Paul Ohm, and Ryan Calo for their feedback and comments on this project.

¹ See C. Lee Ventola, *Mobile Devices and Apps for Health Care Professionals: Uses and Benefits*, 39 *Pharmacy & Therapeutics* 356, 356-64 (2014); see also *Digital Health*, FDA, <https://www.fda.gov/MedicalDevices/DigitalHealth/default.htm> (last updated Sept. 6, 2017) (providing an overview of the various technologies that make up “digital health”); cf. Sy Mukherjee, *VC Funding for Mobile Health Apps Hit an All-Time Record in 2016*, *FORTUNE* (Jan. 17, 2017), <http://fortune.com/2017/01/17/vc-funding-mobile-health-record>.

² See Mallory Locklear, *IBM’s Watson Is Really Good at Creating Cancer Treatment Plans*, *ENGADGET* (June 1, 2017), <https://www.engadget.com/2017/06/01/ibm-watson-cancer-treatment-plans>.

medicine offers a lot of room for improvement, both in price and in quality of care.

Health technology is starting to incorporate data-driven intelligence. It is moving beyond measuring and storing information to spinning the information into actionable advice. The U.S. Food & Drug Administration ("FDA") and the Federal Trade Commission ("FTC") have been monitoring this development closely, and have used the authority available to them to steer the marketing and development of our emerging robot healers. For the FDA's part, its regulatory oversight derives from existing regulations on medical devices, and its willingness to apply these old rules to new technology has generated criticism.³ One problem is analogical: the most important contributions from health and medical AI will substitute not devices, but doctors. As Ryan Calo has observed, "robots begin to blur the line between people and instrument."⁴

This Essay explores whether health and medical AI should be regulated more like doctors than like devices, and what difference it would make. The punch line is that the public safety concerns that are at the heart of medical device regulations are going to be less relevant in the context of medical AI than some of the other, more ancillary duties that doctors usually owe to their patients and to society: duties to provide confidentiality, to warn, to provide informed consent, and to avoid conflicts of interest. In most cases, treating robots like doctors rather than machines reveals a flaw in the assumptions and fundamental goals of longstanding fiduciary rules. This case study can teach us something about future-proofing law: while most legal scholars have focused on adjustments to the law in order to optimize our future robots, it is just as plausible that robots will help us adjust and optimize our aging laws.

The Essay proceeds in two parts. Part I describes the legal machinery currently in use to regulate medical AI and explains why it is a poor match. Part II analyzes the policy future of medical AI by applying the laws that currently regulate its nearest substitutes: physicians. It discusses the benefits and drawbacks of treating AI as professionals.

³ See, e.g., Richard Epstein, Opinion, *Manhattan Moment: FDA Overreach Has Heavy Costs*, WASH. EXAMINER (Nov. 29, 2013, 12:00 AM), <http://www.washingtonexaminer.com/manhattan-moment-fda-overreach-has-heavy-costs/article/2539939>.

⁴ Ryan Calo, *Robots in American Law* 5 (U. Wash. Sch. L., Research Paper No. 2016-04, 2016), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2737598.

I. THE SEARCH FOR THE APPROPRIATE BASELINE

Today, new medical apps and technologies fall into the jurisdiction of the Food & Drug Administration. FDA regulations are complex, so by necessity the summary here is over-simplified.

In a nutshell, the FDA regulates devices by requiring all manufacturers to register a new device.⁵ For devices that pose a greater than nominal risk, the manufacturer also has to perform premarket testing similar to testing for drugs before publicly releasing the device.⁶ After a device is on the market, the FDA may continue to assess its risk and utility.⁷ A key factor during pre-screening and after-market analysis is the estimated impact on a typical user, which is done by comparing it to the other available options. What would happen to the patient if this device was not available?⁸

For example, when a particular brand of home pregnancy test started having a high rate of false positives, the FDA removed it from the market because there were other home pregnancy tests that worked just as early in the female fertility cycle and gave more accurate results.⁹ By contrast, when a steam sterility monitor used in operating rooms started to have higher error rates, the FDA left it on

⁵ See 21 C.F.R. § 807.20(b)(a) (2017); *Device Registration and Listing*, FDA, <https://www.fda.gov/MedicalDevices/DeviceRegulationandGuidance/HowtoMarketYourDevice/RegistrationandListing/default.htm> (last updated Sept. 28, 2017).

⁶ Class III devices, which have the most onerous pre-market clearance procedures, must be shown by their manufacturers that they are “substantially equivalent” to a device that has already been approved for the market based on its demonstrated safety and efficacy. See 21 C.F.R. § 807.87(h)-(k) (2017).

⁷ Sometimes the FDA requires manufacturers to conduct studies of their products after they have been approved for the market. See FDA, GUIDANCE FOR INDUSTRY AND FDA STAFF: PROCEDURES FOR HANDLING POST-APPROVAL STUDIES IMPOSED BY PMA ORDER (2009), <https://www.fda.gov/downloads/MedicalDevices/DeviceRegulationandGuidance/GuidanceDocuments/ucm071013.pdf>. Other times, the FDA reviews the safety and efficacy of a medical device using information from mandatory or voluntary adverse event reports submitted to the agency. *Medical Device Reporting (MDR)*, FDA, <https://www.fda.gov/MedicalDevices/Safety/ReportaProblem/default.htm> (last updated Nov. 7, 2016).

⁸ FDA, GUIDANCE FOR INDUSTRY AND FDA STAFF: FACTORS TO CONSIDER WHEN MAKING BENEFIT-RISK DETERMINATIONS IN MEDICAL DEVICE PREMARKET APPROVAL AND DE NOVO CLASSIFICATIONS 13, 15-16, 44 (2016), <https://www.fda.gov/ucm/groups/fdagov-public/@fdagov-meddev-gen/documents/document/ucm517504.pdf> [hereinafter *FACTORS TO CONSIDER*].

⁹ FDA, GUIDANCE FOR INDUSTRY AND FDA STAFF: FACTORS TO CONSIDER REGARDING BENEFIT-RISK IN MEDICAL DEVICE PRODUCT AVAILABILITY, COMPLIANCE, AND ENFORCEMENT DECISIONS 19-20 (2016), <https://www.fda.gov/downloads/MedicalDevices/DeviceRegulationandGuidance/GuidanceDocuments/UCM506679.pdf>.

the market because there were no alternatives — no other manufacturers — and using the monitor was better than nothing.¹⁰

The basis on which safety is assessed is imminently sensible — it looks at the marginal risks and benefits that each device offers to the healthcare ecosystem. But the FDA procedures add considerable cost and delay, and particularly in the form of preapproval (which takes five months, on average, even when the device is a newer version of an already-approved device).¹¹ There are also FDA requirements beyond proving safety and efficacy, including formal registration and reporting, and branding requirements.¹²

Until recently, medical devices were used to either take measurements or administer treatments. But software-driven medical devices are not so limited. While some of them take measurements and record data, an increasing portion of smart medical apps are knowledge devices. They analyze data (either preexisting or newly measured) and interpret the data to form opinions.

For example, the following post-traumatic stress disorder (“PTSD”) app helps fine-tune a patient’s therapy and daily activities based on real time inputs. Patients can use the app to assess and track their symptoms, and to manage them through breathing exercises, calming imagery, and positive self-talk.¹³

¹⁰ *Id.* at 17-18.

¹¹ Drew Simshaw et al., *Regulating Healthcare Robots: Maximizing Opportunities While Minimizing Risks*, 22 RICH. J.L. & TECH. 3, 16 (2016); Adam Candeub, Digital Medicine, the FDA, and the First Amendment 13 (Nov. 11, 2015) (unpublished manuscript) (on file with Georgia Law Review).

¹² Candeub, *supra* note 11, at 14.

¹³ See *Mobile App: PTSD Coach*, U.S. DEP’T VETERANS AFF., <https://www.ptsd.va.gov/public/materials/apps/ptsdcoach.asp> (last updated May 31, 2017); *PTSD Coach*, iTUNES PREVIEW, <https://itunes.apple.com/us/app/ptsd-coach/id430646302?mt=8> (last updated Aug. 1, 2017).



23andMe's genetic health reports¹⁴ and the IBM Watson for Oncology project¹⁵ likewise suggest that health AI is headed down a path that will provide diagnoses, prognoses, and treatment instructions. When we consider what a patient's alternatives to a new device will be, it is helpful to distinguish between measurement and knowledge applications.¹⁶ Measurement apps will have traditional instruments as their nearest conceptual neighbor, while knowledge apps emulate doctors or, perhaps, the patients' informal networks of health advisers.

¹⁴ See *Find Out What Your DNA Says About Your Health, Traits and Ancestry*, 23ANDME, <https://www.23andme.com/dna-health-ancestry> (last visited Sept. 7, 2017).

¹⁵ See *Watson for Oncology*, IBM, <https://www.ibm.com/watson/health/oncology-and-genomics/oncology> (last visited Sept. 7, 2017).

¹⁶ Adam Candeub divided health apps into three categories: Non-Invasive Measurement Apps, Invasive Measurement Apps, and Prediction Calculators ("Physician Avatars"). See Candeub, *supra* note 11, at 11-12.

Today, the FDA has claimed jurisdiction over both types of medical apps under its device regulations. For example, it intervened in 23andMe's business model when the home genomics kit manufacturer began to report health risks associated with their customers' personal genome.¹⁷ The kit and subsequent sequencing of the customers' personal genome fell under the FDA's purview only when the company began to provide customized health reports. This is a break from typical medical device oversight because the quality of the measuring component of 23andMe kits is undisputed. The company's DNA sequencing works as well as any other commercially available option. It is the health reports — the knowledge component — that concerned the FDA. Citing concerns that the home DNA test results could spur users to self-manage serious medical conditions or to over-react to 23andMe's reports, the FDA ordered 23andMe to discontinue its health reporting services in a 2013 warning letter.¹⁸

Similarly, in its guidance for mobile health apps, the FDA has claimed jurisdiction for any software that “transforms” a phone into a medical device by having the intent to diagnose, cure, mitigate, treat, or prevent a disease.¹⁹

By the way, the FDA uses the manufacturer's marketing and promotion as a primary means of determining whether a service developer has the requisite intent to diagnose or treat a disease.²⁰ This is the only thing preventing Google's search bar from being a medical device. Google certainly has knowledge that people use its search engine to diagnose and treat diseases, and it even makes use of that knowledge by analyzing medical search histories and creating tools like Flu Trends.²¹ But presumably Google avoids FDA jurisdiction

¹⁷ *Warning Letter: 23andMe, Inc. 11/22/13*, FDA (Nov. 22, 2013), <https://www.fda.gov/ICECI/EnforcementActions/WarningLetters/2013/ucm376296.htm>.

¹⁸ See *id.* The letter references risks of “false positives,” suggesting it is worried that the genome sequencing itself is flawed, but 23andMe has always been permitted to offer genome sequencing services to the general public. For a more thorough description of the anticipated problems with connecting a genome sequence to health diagnoses or risk factors, see Rob Arthur, *What's in Your Genes?: Some Companies Analyzing Your DNA Rely on Junk Science*, SLATE (Jan. 20, 2016, 3:10 PM), http://www.slate.com/articles/health_and_science/medical_examiner/2016/01/some_personal_genetic_analysis_is_error_prone_and_dishonest.html.

¹⁹ FDA, GUIDANCE FOR INDUSTRY AND FOOD AND DRUG ADMINISTRATION STAFF: MOBILE MEDICAL APPLICATIONS 7-9 (2015), <https://www.fda.gov/downloads/MedicalDevices/.../UCM263366.pdf> [hereinafter GUIDANCE FOR INDUSTRY].

²⁰ For a discussion of this in the context of the definition of a pharmaceutical drug, see Christopher Robertson, *The Tip of the Iceberg: A First Amendment Right to Promote Drugs Off-Label*, 78 OHIO ST. L.J. 8-14 (forthcoming 2017).

²¹ Alexis Madrigal, *In Defense of Google Flu Trends*, ATLANTIC (Mar. 27, 2014),

because it does not encourage these health uses of its search bar in its promotional materials. And if this is so, advertising is the basis for regulatory enforcement. This is one of many practices that leaves the FDA vulnerable to First Amendment challenges. Since the only difference between a medical device and a generic information service is commercial speech, the FDA's regulations may at some point have to withstand constitutional scrutiny.²² Indeed, medical AI may have multiple free speech problems since some of the services targeted by the FDA (including its enforcement against 23andMe's health reports) consist entirely of information exchanges.²³ These are interesting developments in constitutional law, and to some extent they dovetail with the regulation of doctors' professional speech.²⁴

Perhaps because of the tricky constitutional issues, the FDA has promised to use discretionary forbearance over mobile apps that are strictly informational and provide only "simple tools . . . to organize and track health information without providing recommendations to alter or change a previously prescribed treatment."²⁵ Apps like the PTSD Coach can therefore avoid the regulatory barriers that would ordinarily apply. But the FDA will and has applied the usual rules for medical devices on mobile apps that use a sensor of any sort to collect information about the user in order to provide health advice.²⁶

Is the FDA's decision to apply the rules of medical devices to mobile apps consistent with its past practices of considering substitution effects? For measurement apps, I think it is. A mobile Electrocardiogram ("ECG") app can be compared, in terms of accuracy, cost, availability, etc., to a standard ECG machine in order to determine whether users are better off having it on the market or not. But for knowledge apps, the FDA is using the wrong baseline. The comparison for health AI is *not* other existing devices. Right now, it is

<https://www.theatlantic.com/technology/archive/2014/03/in-defense-of-google-flu-trends/359688> (describing Flu Trends, a flu prevalence predictor based on Google key word searches, and the skepticism about its value).

²² Commercial speech generally receives intermediate free speech scrutiny. See *Cent. Hudson Gas & Elec. Corp. v. Pub. Serv. Comm'n.*, 447 U.S. 557, 573 (1980) (Blackmun, J., concurring).

²³ See Candeub, *supra* note 11, at 39-45 (arguing that informational output platforms like Caracal Diagnosis, Isabel, iLiver, and arguably 23andMe are "pure speech" entitled to the highest level of constitutional protections).

²⁴ See generally Claudia E. Haupt, *Professional Speech*, 125 YALE L.J. 1238 (2016) (arguing for First Amendment protection of professional speech frustrated by limitations on doctors' interactions with patients).

²⁵ GUIDANCE FOR INDUSTRY, *supra* note 19, at 16.

²⁶ *Id.* at 27.

doctors. AI will pose danger to consumers only if the costs, risks, and inaccuracy of its advice are out of line, and a bad deal, compared to the costs, risks, and inaccuracy of these human advice-givers.

This type of marginal comparison is consistent with what the FDA has done outside Artificial Intelligence.²⁷ But the FDA's track record with knowledge products like 23andMe seems to be driven by an abundance of caution rather than by utilitarianism.²⁸ Preapproval requirements are getting in the way of improved health by, for example, forbidding 23andMe from informing the small but identifiable group of Ashkenazi Jews whose mutations across three BRCA genes can be reliably identified by 23andMe for breast cancer risk.²⁹ And 23andMe's health reports are just the beginning of a potentially pivotal knowledge movement. Pharmacogenomics information that can inform what a drug's absorption, metabolism, binding, transport, and excretion rate is for an individual patient is not used in medicine at all, even for the set of drugs for which this sort of customizable information is known.³⁰

The Federal Trade Commission, which more directly regulates claims that are made in advertising, also seems to apply impractically high standards. The FTC fined the manufacturer of a health app called Mole Detective because of concern about the app's accuracy and false negative errors.³¹ That is, the FTC was worried that users who have melanoma may rely on the app when it fails to detect the cancer. But

²⁷ See FACTORS TO CONSIDER, *supra* note 8, at 13 (discussing availability of alternative treatments or diagnostics when reviewing a medical product for premarket approval).

²⁸ As Gary Marchant put it, "this is the last shoe to drop in the FDA's effort to wipe out the right of consumers to discover their own genetic information, some of the most important, private, useful, and interesting information about our own health and well-being." Gary Marchant, *The FDA Could Set Personal Genetics Rights Back Decades*, SLATE (Nov. 26, 2013, 12:39 PM), http://www.slate.com/articles/technology/future_tense/2013/11/_23andme_fda_letter_premarket_approval_requirement_could_kill_at_home_genetic.html.

²⁹ ERIC TOPOL, *THE PATIENT WILL SEE YOU NOW* 66-76 (2015).

³⁰ *Id.* at 85.

³¹ Stipulated Final Judgment and Order for Permanent Injunction and Other Equitable Relief Against Defendant Avrom Boris Lawarow at 6-8, *FTC v. Lasarow*, No. 15-cv-1614 (N.D. Il. Apr. 30, 2015) (documenting the fine and other remedies); see Complaint for Permanent Injunction and Other Equitable Relief at 10, *FTC v. Lasarow*, No. 1:15-cv-01614 (N.D. Il. Feb. 23, 2015) (alleging that the app made false or misleading claims about Mole Detective's capacity to accurately analyze moles for melanoma). But see Dissenting Statement of Commissioner Maureen K. Ohlhausen at 1-2, *FTC v. Lasarow*, No. 132-3210 (N.D. Il. Feb. 23, 2015) (objecting to the majority's implication that the app makers must show the app detects cancer as accurately as dermatologists).

as Adam Candeub has pointed out, while the Mole Detector app was less accurate than trained dermatologists, it actually performed as well or better than primary care physicians.³² Primary care physicians may be the more appropriate baseline comparison if patients typically use their generalist doctors as an initial screen. (And consultations by both primary care and dermatology specialists are, of course, much more expensive than the app.)

So what would happen if medical AI were regulated like their closest substitutes — doctors — instead of like devices?

II. THE SPECIAL DUTIES OF DR. ROBOT

Applying the law of doctors to health AI rather than the law of medical products leads to some surprising differences in regulation. States regulate doctors using a combination of licensing statutes and tort rules that apply to special, fiduciary relationships.³³ With the exception of the training and licensing exam, regulation comes retrospectively — after something has gone wrong — rather than prospectively through pre-market approval. Doctors do, of course, owe their patients a duty of competence, which closely matches the safety concerns that drive the regulation of medical devices. Future health-based knowledge apps will be able to comply with this duty fairly easily. The other duties that doctors owe — duties of confidentiality, duties to warn, duties to provide informed consent about research, and duties to avoid conflicts of interest — will be more difficult. In fact, some will so badly interfere with the prospective advantages of medical AI that policymakers should begin to question the virtues not of the robots but of the law itself.

I will briefly sketch the application of each of these special duties to AI to show where our current law may lead our future medical industry astray.

A. Competence

The duty of competence incorporates both the safety and efficacy goals that form the core goals of FDA medical device approval. There is little reason to doubt that AI will eventually flourish and outperform physicians in many aspects of medical care. In fact, much of what the

³² See Candeub, *supra* note 11, at 45-46 (citing a study where apps performed with a seventy percent success rate compared to a ninety percent success rate of specialists and a sixty percent to seventy-five percent success rate of general practitioners).

³³ See Simshaw, *supra* note 11, at 2 (“[S]tate licensing statutes oversee the conduct of doctors and nurses who, heretofore, have all been human beings.”).

training and licensing process does for doctors is train them to act more like algorithms.³⁴ Decision trees and probabilistic decision-making are critical to diagnosis and treatment recommendations. They do not come naturally to people, but are easily automated. Other skills, like memorizing everything that is tested in the Board exams would also be trivially easy for a computer. Of course AI still is not well developed in certain ways — we have not figured out how to replicate human eye sight and perception without tragic or comic error.³⁵ But computers can do the reasoning part of the practice of medicine very well.³⁶

Consider missed diagnosis, which is one of the principal fears with health apps. Malpractice based on missed diagnosis is usually proved by showing a doctor failed to do a proper differential diagnosis where every ailment that could be explained by the presented symptoms are systematically tested and ruled out in order of probability (weighted by urgency).³⁷ In the T.V. show *House*, the doctors sat around talking through this process, and often struggled to recall rare conditions that would present with the patients' symptoms. Watson would be able to beat that crackerjack team in a matter of seconds.³⁸

Some health law experts doubt that robots will be able to replace human doctors anytime soon,³⁹ but there is wide agreement that AI will supplement doctor care, and replace it in some critical aspects.⁴⁰

³⁴ For example, new doctors are trained to avoid heuristics and biases that are commonly used in human judgment. KEVIN BARRACLOUGH ET AL., AVOIDING ERRORS IN GENERAL PRACTICE 15-17 (2013).

³⁵ See Shimon Ullman et al., *Atoms of Recognition in Human and Computer Vision*, 113 PROC. NAT'L ACAD. SCI. 2744, 2744 (2016) (finding that the human visual system uses features and processes that current computer simulation models do not have).

³⁶ PEDRO DOMINGOS, THE MASTER ALGORITHM: HOW THE QUEST FOR THE ULTIMATE LEARNING MACHINE WILL REMAKE OUR WORLD 13 (2016) ("Machine learning is the scientific method on steroids.").

³⁷ See Dan Minc, *Differential Diagnosis Can Be Used in Proving Medical Malpractice*, RMFW LAW (Sept. 19, 2014), <http://www.medicalmalpractice.net/blog/2014/09/differential-diagnosis-can-be-used-in-proving-medical-malpractice.shtml>.

³⁸ Campbell McLaren, an editor of this journal, pointed out a delightful irony: the character Dr. House is based on Sir Arthur Conan Doyle's Sherlock Holmes, so the fact that Watson (named after Sherlock Holmes' obedient companion and biographer) will best Dr. House is particularly amusing.

³⁹ See Frank Pasquale, *Automating the Professions?*, 8-11 (U. Md. Francis King Carey Sch. L., Research Paper No. 2016-21) (reviewing RICHARD & DANIEL SUSSKIND, THE FUTURE OF THE PROFESSIONS (2015), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2775397) (disagreeing with Susskind on the facility with which computer technology can quickly enhance the field of medicine).

⁴⁰ See Vimla L. Patel et al., *The Coming of Age of Artificial Intelligence in Medicine*, 46 ARTIFICIAL INTELLIGENCE MED. 5, 14 (2009) (concluding that AI in medicine is coming of

Geoffrey Hinton, a computer scientist at the University of Toronto who works on health-related AI, already believes it is time to stop wasting energy attempting to train radiologists to read medical images. “It’s just completely obvious that in five years deep learning [by computers] is going to do better than radiologists . . . It *might* be ten years.”⁴¹ Even the more “human” aspects of care like emotional support and therapy can be greatly improved by including or substituting in AI.⁴²

That is not to say medical AI requires no safety regulations. There is always a role for law to ensure that care is optimized based on safety, efficacy, and price. But whether those regulations come in the form of premarket clearance or instead use a post-market liability rule, the reference for comparison should be doctors at the highest.⁴³ This is, after all, the standard of care that medical malpractice law currently applies to other doctors.⁴⁴

The other fiduciary duties and professional limitations pose far more difficult to the development and adoption of medical AI. What makes them difficult is that AI will be much more concentrated in just a small set of firms, whereas care through doctors is distributed across

age as a discipline); see, e.g., Daniela Hernandez, *Artificial Intelligence Is Now Telling Doctors How to Treat You*, WIRED (June 2, 2014, 6:30 AM), <https://www.wired.com/2014/06/ai-healthcare> (illustrating an instance in which a doctor’s diagnosis is facilitated by an AI technology); Paul Hsieh, *AI in Medicine: Rise of the Machines*, FORBES (Apr. 30, 2017, 12:10 PM), <https://www.forbes.com/sites/paulhsieh/2017/04/30/ai-in-medicine-rise-of-the-machines/#783fde6dabb0> (concluding that AI algorithms are beginning to perform medical work which, until very recently, was thought capable of being performed only by humans); Stephen F. Weng et al., *Can Machine-Learning Improve Cardiovascular Risk Prediction Using Routine Clinical Data?*, PLOS ONE (Aug. 4, 2017), <http://journals.plos.org/plosone/article?id=10.1371/journal.pone.0174944> (showing that machine learning predicted cardiovascular risk more effectively than traditional methods).

⁴¹ Siddhartha Mukherjee, *The Algorithm Will See You Now*, NEW YORKER, Apr. 3, 2017, at 46; see also Todd C. Frankel, *New Machine Could One Day Replace Anesthesiologists*, WASH. POST (May 11, 2015), http://www.washingtonpost.com/business/economy/new-machinecould-one-day-replace-anesthesiologists/2015/05/11/92e8a42c-f424-11e4-b2f3-af5479e6bbdd_story.html [<https://perma.cc/X6TW-Y7U6>] (discussing a technology which may potentially replace anesthesiologists).

⁴² See Barbara Peters Smith, *Robots and More: Technology and the Future of Elder Care*, HERALD TRIB. (May 27, 2013), <http://www.heraldtribune.com/article/LK/20130527/News/605195720/SH>; Tony Rousmaniere, *What Your Therapist Doesn’t Know*, ATLANTIC (Apr. 2017), <https://www.theatlantic.com/magazine/archive/2017/04/what-your-therapist-doesnt-know/517797>.

⁴³ And even then, it only makes sense to compare AI to doctor care if patients would have sought help from doctors in the absence of the AI.

⁴⁴ DAN DOBBS ET AL., *THE LAW OF TORTS: PRACTITIONER TREATISE SERIES* § 292 (2d ed. 2011). In time, comparisons to other existing AI will be more appropriate once medical technology consistently outperforms doctors on the relevant task.

hundreds of thousands of highly variable, small practices. The law was designed for the diffuse and messy organization of human professionals, and the same rules have surprisingly perverse effects when applied to large producers of AI.

B. Confidentiality and Duties to Warn

In addition to performing their jobs as advisors well, doctors also have obligations to keep their patients' confidentiality. This duty is partly codified in the privacy rule of the Health Insurance Portability and Accountability Act ("HIPAA").⁴⁵ Those statutory provisions, which require confidentiality and data security for health records, will apply to health AI to the extent they are being used as part of the care provided by a "covered entity" (basically a traditional healthcare provider or insurer). But even independent of HIPAA, doctors have long had a common law duty to keep their patients' confidentiality.⁴⁶

In most respects, promises of confidentiality will not be hard for current and future forms of AI to adopt. In fact, promises of data privacy may be necessary to gain the trust of new, hesitant users. But medical AI companies are going to rely on a lot more pooling, sharing, and analysis of semi-anonymized patient data than the current practice of medicine typically does. Pooling and mining data is what gives new health technology an opportunity to learn about hidden patterns that can guide better medical advice. For example, a program called LIONsolver is being used to pool data for Parkinson's patients to learn and provide customized health advice.⁴⁷ And the Global Alliance for Genomics and Health ("GA4GH") is trying to pool genome data for a broad swath of patients and make the data broadly available for research purposes.⁴⁸ These efforts run into conflict with HIPAA guidance that considers data too vulnerable to reidentification attacks unless significant, utility-compromising precautions are taken.⁴⁹ And

⁴⁵ 45 C.F.R. § 164 (2016).

⁴⁶ *Alberts v. Devine*, 479 N.E.2d 113, 120 (Mass. 1985); *Hague v. Williams*, 181 A.2d 345, 349 (N.J. 1962); *McCormick v. England*, 494 S.E.2d 431, 436 (S.C. Ct. App. 1997).

⁴⁷ See Candeub, *supra* note 11, at 6 (describing LIONsolver's ability to predict the progression of Parkinson's over a ninety-day timeframe, and its potential, with increased data collection, to monitor and advise patients on drug dosages, food intake, and sleeping habits).

⁴⁸ See TOPOL, *supra* note 29, at 165.

⁴⁹ The fine-grained, longitudinal data that is required for significant advances in medicine will not be able to meet the standards of these statistical confidentiality experts. See U.S. DEP'T HEALTH & HUMAN SERVS., GUIDANCE REGARDING METHODS FOR DE-

the Federal Trade Commission has recommended that all Internet of Things technologies incorporate consent mechanisms anytime personal data is going to be used for an unanticipated purpose.⁵⁰ Thus, some scholars are already worried about the privacy and security protocols for medical AI.⁵¹

Doctors also have a duty to disclose or warn third parties about the risks their patients pose to others in certain circumstances.⁵² This duty obviously conflicts with the obligation to maintain confidentiality, so they must be carefully managed and weighed against each other. But when a patient is likely to spread a communicable disease or harm a known, identifiable person, the duty of confidentiality yields to the duty to warn.

Today these duties to warn are only rarely activated, but they may become more common for AI that keeps an ongoing information-gathering relationship with its users all day and all year long. This is particularly true for services that have data on many other users and therefore can support a broader range of inferences about whether a user poses danger to others.

AI companies will be inclined to avoid duties to warn by simply not looking for statistical patterns that could reveal otherwise latent dangers. But if a company has relevant information, this sort of willful blindness combined with a particularly sad set of facts could motivate a court to interpret the duty to warn as an obligation for companies to look for signs that their users are dangerous. A strongly enforced duty to warn may interfere with patient trust in medical AI more than it has for doctors.

C. Research and Informed Consent

The concentration of medical AI to a small number of companies that oversee the care of a large number of patients creates an

IDENTIFICATION OF PROTECTED HEALTH INFORMATION IN ACCORDANCE WITH THE HEALTH INSURANCE PORTABILITY AND ACCOUNTABILITY ACT (HIPAA) PRIVACY RULE 8-22 (2012), https://www.hhs.gov/sites/default/files/ocr/privacy/hipaa/understanding/coveridentities/De-identification/hhs_deid_guidance.pdf.

⁵⁰ FTC, INTERNET OF THINGS: PRIVACY & SECURITY IN A CONNECTED WORLD 38-39, 43 (2015), <https://www.ftc.gov/system/files/documents/reports/federal-trade-commission-staff-report-november-2013-workshop-entitled-internet-things-privacy/150127iotrpt.pdf>.

⁵¹ Simshaw et al., *supra* note 11, at 14.

⁵² *Tarasoff v. Regents of the Univ. of Cal.*, 551 P.2d 334, 342 (Cal. 1976) (therapist's duty to warn known third party victim of attack); *Reisner v. Regents of the Univ. of Cal.*, 37 Cal. Rptr. 2d 518, 522 (Cal. Ct. App. 1995) (doctor's duty to warn third party about risk of communicable disease from HIV positive partner).

unprecedented opportunity for health research. But it is also likely to strain the limits of the laws regulating research and informed consent.

Here is how I figure: There are many areas in the practice of medicine where there is not one single standard treatment, or even one particular order in which treatments should be tried. Instead, doctors choose the treatment that they think is best within the range of seemingly reasonable options.⁵³ Because of the variation among their choices, the doctors' idiosyncratic decisions can advance the state of medical knowledge by providing some field evidence about what works best, or what works better for whom. The law does not treat any one of the doctors' individual treatment decisions as research. Each one is just a therapeutic treatment that can later create observational evidence about which therapies work better than others.

With a dominant AI making treatment decisions for a large swath of patients, developers will have two options: vary the treatments through "A/B-testing," or treat every similar patient exactly the same.

At a society level, we are of course much better off if medical AI does the same sort of A/B testing that other tech firms use all the time, as long as the testing stays within the range of reasonable treatments.⁵⁴ In fact, when we compare the kind of testing that a medical AI system could do across a larger swath of patients to the haphazard sort of experimentation that we permit in today's healthcare, clearly formal, randomized testing would be much better than the unscientific experiments that are currently performed on patients every day. But this testing triggers legal obligations and liability risks that cluster around formal research.⁵⁵ Generally speaking, doctors who experiment on patients must first provide informed consent and formally enroll them in clinical trials.

I suspect many patients would resist when confronted with informed consent about research nearly every time they interact with their medical AI. Doctors and policymakers are in part responsible for this resistance because they have promoted a fiction that patients are not already experimented on routinely (albeit in haphazard,

⁵³ For a longer discussion of this problem, see Jane R. Bambauer, *All Life Is an Experiment (Sometimes It Is a Controlled Experiment.)*, 47 LOY. U. CHI. L.J. 487, 507 (2015).

⁵⁴ For a description of A/B testing and its benefits, see SETH STEPHENS-DAVIDOWITZ, *EVERYBODY LIES: BIG DATA, NEW DATA, AND WHAT THE INTERNET CAN TELL US ABOUT WHO WE REALLY ARE* 209-21 (2017).

⁵⁵ See Robert J. Morse & Robin Fretwell Wilson, *Realizing Informed Consent in Times of Controversy: Lessons from the SUPPORT Study*, 44 J.L. MED. & ETHICS 402, 403 (2016) (describing a doctor's duty to obtain consent for medical procedure imposed by common law and the Federal Policy for the Protection of Human Subjects).

disorganized ways), and that their treatment was chosen especially for them. The law of informed consent for research could be an area that could be improved with the help of robots rather than the other way around.

D. Conflicts of Interest

Finally, the doctor-patient relationship includes a duty of loyalty that requires a doctor to disclose conflicts of interest. Disclosure is useful if the patient has a realistic option to leave and choose another care-giver, but if there is only one or two dominant forms of AI for a particular disease, each with financial motives independent of the patient's health, choice is not realistic.

It is tempting to conclude that the law should evolve to require not merely disclosure but avoidance of conflicts. But this may not be in the patients' best interests, particularly if the companies that get into this industry (IBM, Google) have other highly successful products and tools that not only could be used, but should be used in conjunction with the patient's medical care. On the other hand, when a medical AI producer exploits a relationship with a patient to promote products and services that are not superior, or when AI stands to benefit from a user's prolonged illness, these will be occasions for legal intervention.

CONCLUSION

This thought experiment exposes where I believe the legal headaches from mechanical healers truly lie: not in the quality of the treatment itself, but in ancillary issues.

Before closing the thought experiment, I would like to comment about the political economy of these sorts of emerging health technologies.

Right now, the FDA is the primary regulator, but Congress may get involved in the future, too. My impression is that even when the FDA is more risk-averse than I would prefer, their decisions come from an independent process with high integrity. But at Congress (and perhaps at administrative agencies as well), there is a risk that overly high safety standards may be erected as a means to protect doctors. Every year, the American Medical Association is one of the leading spenders for lobbying efforts in Washington, D.C.⁵⁶

⁵⁶ See, e.g., *Top Spenders*, OPENSECRETS.ORG, <https://www.opensecrets.org/lobby/top.php?showYear=a&indexType=s> (last updated Aug. 7, 2017).

There is also a risk of anti-competitive maneuvering by the health tech firms themselves, at least the ones that come to the table first. IBM's Watson team already has a very close working relationship with the FDA. The two are collaborating on a Blockchain project.⁵⁷ Even 23andMe, which was originally resistant to FDA's preclearance requirements for its health reports, is now working with the FDA's standards.⁵⁸ Whether this was a goal or not, it will have the effect of keeping competitors in the consumer genome sequencing space at a disadvantage until they go through the same preclearance that 23andMe has passed. Regulation of our future robot doctors is perfectly appropriate. But we should be on the lookout for protectionist rules and safety standards that are unnecessarily high or complex.

⁵⁷ IBM Watson Health Announces Collaboration to Study the Use of Blockchain Technology for Secure Exchange of Healthcare Data, IBM (Jan. 11, 2017), <https://www-03.ibm.com/press/us/en/pressrelease/51394.wss>.

⁵⁸ See FDA Accepts 510(k) Application for 23andMe Health Report on Bloom Syndrome, GENOMEWEB (June 20, 2014), <https://www.genomeweb.com/clinical-genomics/fda-accepts-510k-application-23andme-health-report-bloom-syndrome>; see also FDA Allows Marketing of First Direct-to-Consumer Tests that Provide Genetic Risk Information for Certain Conditions, FDA (Apr. 6, 2017), <https://www.fda.gov/newsevents/newsroom/pressannouncements/ucm551185.htm>.