UNIVERSITY
OF
CALIFORNIA
HEALTH

Precisely Practicing Medicine from 700 Trillion Points of University of California Health Data

Atul Butte, MD, PhD

Chief Data Scientist, University of California Health (UC Health)

Priscilla Chan and Mark Zuckerberg Distinguished Professor

Director, Bakar Computational Health Sciences Institute, UCSF

atul.butte@ucsf.edu ▪ @atulbutte

Conflicts of Interest

- Scientific founder
  - Personalis
  - NuMedii
  - Carmenta (Progenity)
  - Genstruct
- Honoraria for talks
  - Lilly
  - Pfizer
  - Siemens
  - Bristol Myers Squibb
  - AstraZeneca
  - Roche
  - Genentech
  - Warburg Pincus
  - CRG
  - AbbVie
  - Westat
- Past or present consultancy
  - Personalis
  - NuMedii
  - Lilly
  - Johnson and Johnson
  - Roche
  - Genstruct
  - Tercica
  - Ecoeos
  - Helix
  - Ansh Labs
  - uBiome
  - Prevendia
  - Samsung
  - Assay Depot
  - Regeneron
  - Verinata (Illumina)
  - Pathway Diagnostics
  - Geisinger Health
  - Covance
  - Wilson Sonsini Goodrich & Rosati
  - Orrick
  - 10X Genomics
  - GNS Healthcare
  - Gerson Lehman Group
  - Coatue Management
- Other corporate relationships
  - Northrop Grumman
  - Genentech
  - Johnson and Johnson
  - Optum
- Shares or Ownership
  - NuMedii (major)
  - Personalis (major)
  - Apple
  - Facebook
  - Alphabet (Google)
  - Microsoft
  - Amazon
  - Snap
  - 10x Genomics
  - Illumina
  - Nuna Health
  - Assay Depot (Scientist.com)
  - Vet24seven
  - Regeneron
  - Sanofi
  - Royalty Pharma
  - AstraZeneca
  - Moderna
  - Biogen
  - Paraxel
  - Sutro
- Speakers' bureau
  - None
- Companies started by students
  - Carmenta
  - Serendipity
  - Stimulomics
  - NunaHealth
  - Praedicat
  - MyTime
  - Flipora
  - Tumbl.in
  - Polyglot
  - Iota Health
  - Ongevity Health

nature
THE BITER BIT
Viral infections for viruses
TROPICAL CYCLONES
The strong get stronger
BLACK HOLE PHYSICS
A new window on the Galactic Centre
BIG DATA
NATUREJOBS
Minnesota musings
Harvard Business Review
OCTOBER 2012
The True Measures Of Success
10 Rules for Managing Global Innovation
What Ever Happened To Accountability?
GETTING CONTROL OF BIG DATA
How vast new streams of information are changing the art of management
PAGE 59
WORLD'S BEST SUBCUSTODIANS 2015
GLOBAL FINANCE
BIG DATA
The Future Is Now
Page 16
FORTUNE
BIG DATA
MEETS BIOLOGY
Inside Tech's Next
A tale of two Libyas
Cannibals ... They're Everywhere!
SCIENCE FOR THE CURIOUS
APRIL 2017
The Economist
Obama the warrior
Misgoverning Argentina
The economic shift from West to East
Genetically modified crops blossom
The right to eat cats and dogs
The data deluge
AND HOW TO HANDLE IT: A 14-PAGE SPECIAL REPORT
What it takes to end an AIDS epidemic
Polar bears suffer through lean summers
Sperm produced in ovary of mutant fish
Science
$10
17 JULY 2015
sciencemag.org
AAAS
SPECIAL ISSUE
The Economist
Theresa May v Brussels
Ten years on: banking after the crisis
South Korea's unfinished revolution
Biology, but without the cells
The world's most valuable resource

Artificial Intelligence and Machine Learning

- **Artificial Intelligence**: aspects of human intelligence modeled by computers
- **Machine Learning**: implementing aspects of AI through processing data
  - Supervised or unsupervised learning
- **Deep Learning**: one type of ML, modeling brain architecture with layers of individual classifiers, adding non-linearity

JAN 20, 2017 @ 02:11 AM 63,945

First FDA Approval For Clinical Cloud-Based Deep Learning In Healthcare

Be
FU
Opi

The first FD
clinical setti
and industr

Arterys's me
help doctors
network wh
improve its
new case it

Viz.ai Granted De Novo FDA Clearance for First Artificial Intelligence Triage Software

A new era of intelligent stroke care begins as regulatory approval is granted for the Viz.ai LVO Stroke Platform

The platfor
Occlusion (
access to lif

NEWS PROVID
Viz.ai, Inc. ⟶
Feb 15, 2018, 0

FDA permits marketing of AI software autonomously detects diabetic retinop

By Dave Muoio | April 12, 2018

SH

The FDA has granted diagnostics company IDx's De Novo request

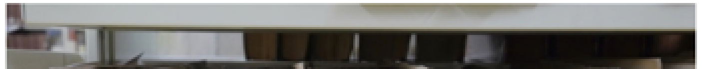

As of
mid-2019

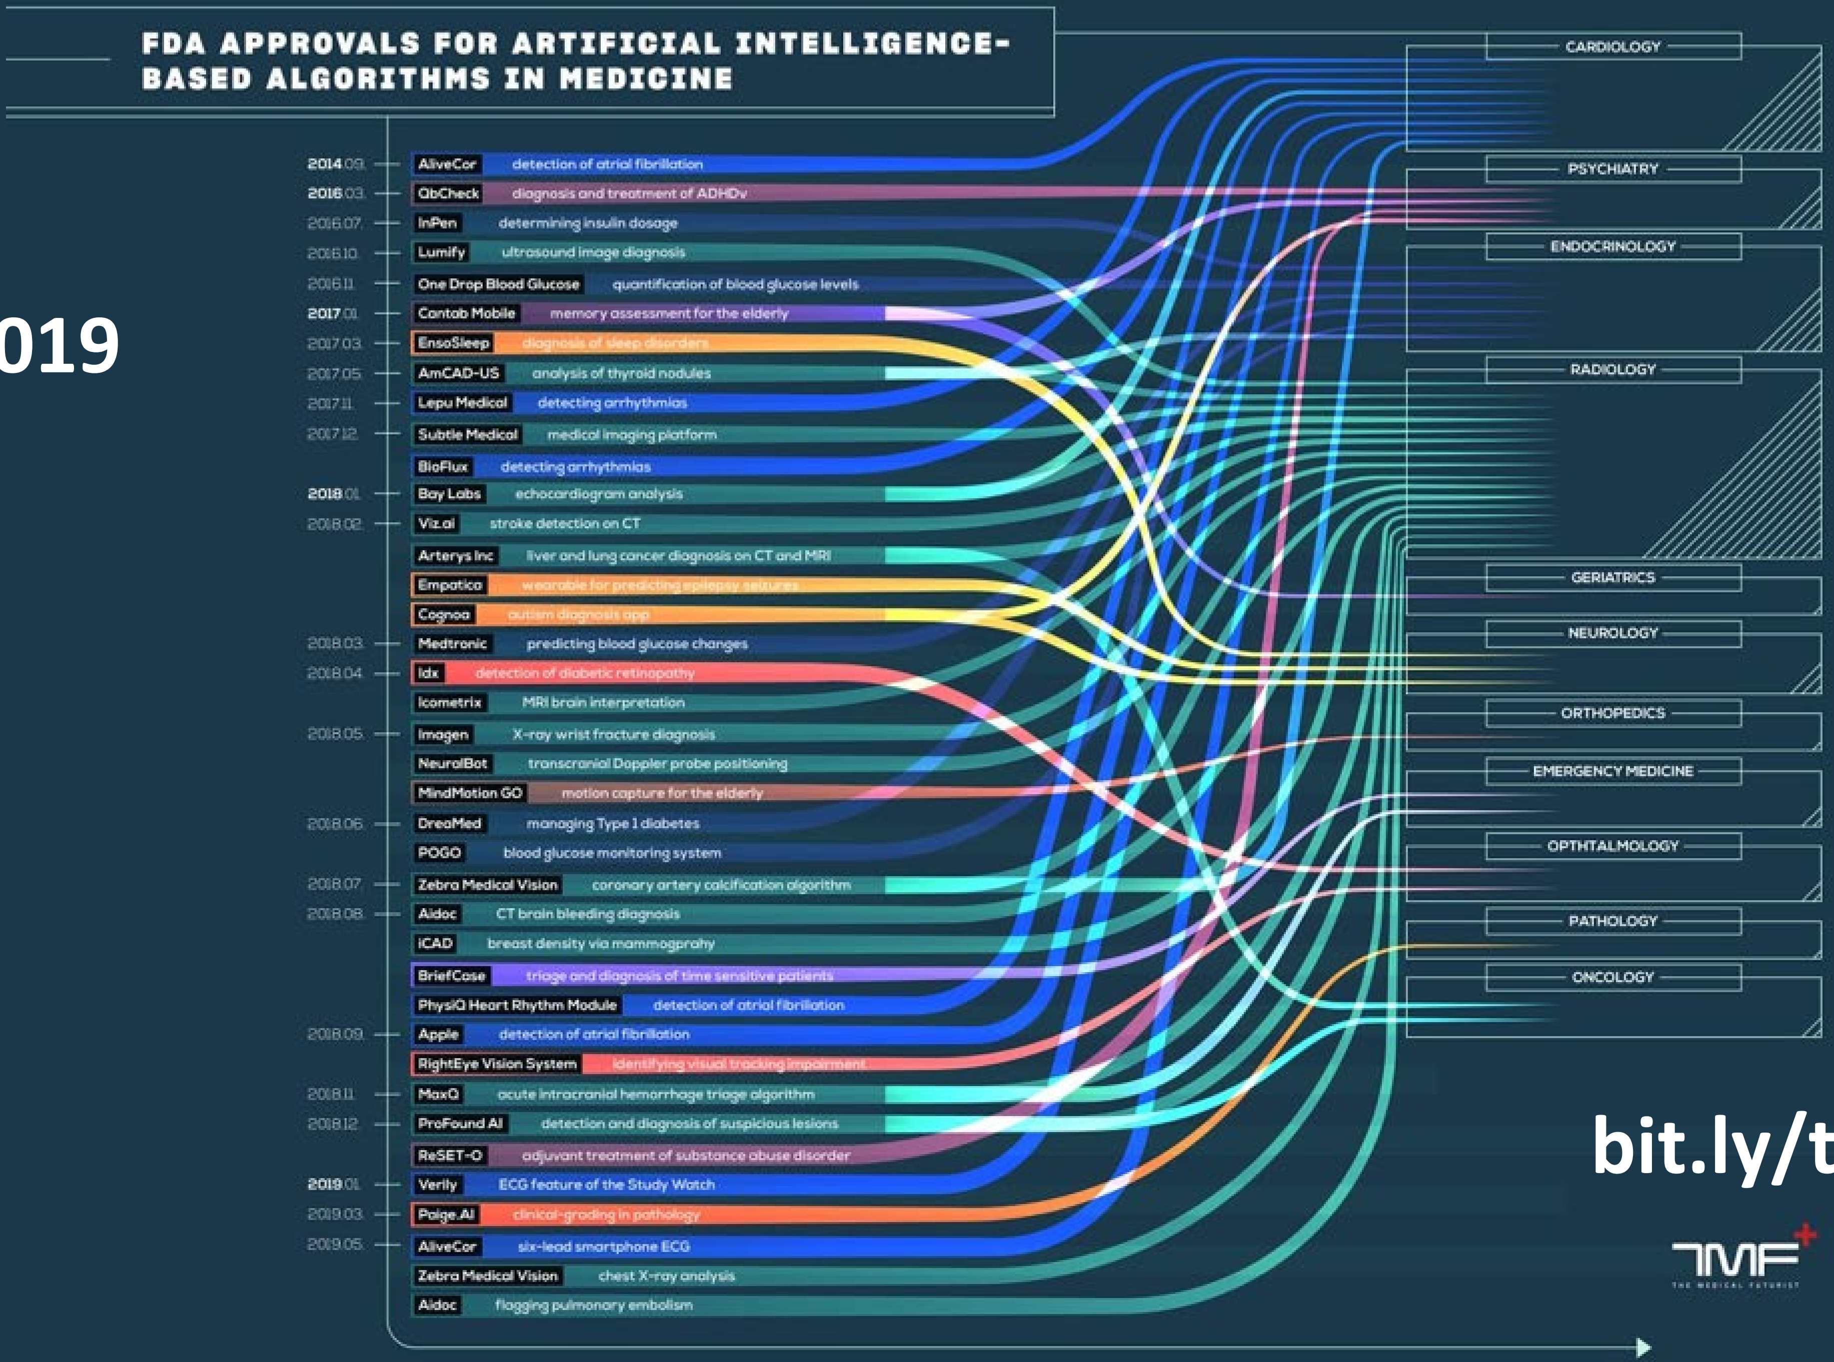

bit.ly/tmfFDA19

Over 100 now approved just counting radiology!

FDA Cleared AI Algorithms

Our list of FDA cleared AI algorithms provides valuable details on each model, bringing all of the relevant information together for easy access. Convenient summaries for each algorithm include model manufacturer, FDA product code, body area, modality, predicate devices, evaluation related to product performance, and clinical validation. Our Define-AI use cases match many of the models and those are listed under Related Use Cases. For other details, clicking on the model will take you directly to the FDA summary.

Check back regularly to see which new algorithms are available and have been added to the list. Send information on AI algorithms that are not listed and report missing information to DSI@acr.org.

Search | Company | Subspeciality | Body Area | Modality | Date Cleared

Product	Company	Subspeciality	Body Area	Modality	Date Cleared
qp-Prostate	Quibim	Abdominal Imaging	Prostate	MR	02/04/2021
Visage Breast Density	Visage Imaging GmbH	Women's Imaging	Breast	MAM	01/29/2021
uAI EasyTriage-Rib	Shanghai United Imaging Intelligence Co., Ltd.	Chest Imaging	Chest	CT	01/15/2021
BrainInsight	Hyperfine Research, Inc.	Neuroradiology	Brain	MR	01/07/2021
SQuEEZ Software	Cardiowise, Inc.	Cardiac Imaging	Heart	CT	12/18/2020
EchoGo Pro	Ultromics Ltd.	Cardiac Imaging	Heart	US	12/18/2020
HepaFat-AI	Resonance Health Analysis Service Pty Ltd.	Abdominal Imaging	Liver	MR	12/07/2020
HealthJOINT	Zebra Medical Vision Ltd.	Musculoskeletal Imaging	Knee	XRAY	12/04/2020
HALO	NiCo-Lab B.V.	Neuroradiology	Brain	CT	11/20/2020
Genius AI Detection	Hologic, Inc.	Women's Imaging	Breast	XRAY	11/18/2020
PROView	GE Medical Systems	Abdominal Imaging	Prostate	MR	11/17/2020
FastStroke, CT Perfusion 4D	GE Medical Systems	Neuroradiology	Brain	CT	11/12/2020
Neuro.AI Algorithm	TeraRecon, Inc.	Neuroradiology	Brain	CT,MR	11/06/2020
WRDensity	Whiterabbit.ai Inc.	Women's Imaging	Breast	MAM	10/30/2020
LSN	Imaging Biometrics, LLC	Abdominal Imaging	Liver	CT	10/29/2020
AVIEW LCS	Coreline Soft Co., Ltd	Chest Imaging	Lung	CT	10/16/2020
Syngo.CT Neuro Perfusion	Siemens Healthineers	Neuroradiology	Brain	CT	10/11/2020
Quantib Prostate	Quantib BV	Abdominal Imaging	Prostate	MR	[illegible]
Cleerly Labs V2.0	Cleerly, Inc.	Cardiac Imaging	Coronary Arteries	CT,CTA	10/02/2020

http://bit.ly/acr_ai

U.S. FOOD & DRUG ADMINISTRATION

The United States is spending $billions on electronic health records, and too few are using any of this data

Sutter's $1 Billion Boondoggle-New Electronic Records System Goes Dark

California Nurses Association Press Release, 8/27/13

Contact Information | Media Center

Yet Another Ris

A controversial e
billion went comp
additional risk be

For several mont
care delivery tha
over 100 reports
Oakland, docum

Sutter managem

Partners' $1.2b patient data system seen as key to future

Aims for one file per person, fewer errors

f

How Kaiser bet $4 billion on electronic h
records -- and won

Kaiser Permanente CIO Philip Fasano explains how electronic records have paid off and the health care giant's embrace of mobile technology

University of California

UNIVERSITY OF CALIFORNIA HEALTH

- 10 campuses and 3 national labs
- ~200,000 employees, ~250,000 students/yr

UC Health

- 20 health professional schools (6 med schools)
- Train half the medical students and residents in California
- ~$2 billion NIH funding
- $13+ billion clinical operating revenue
- 5000 faculty physicians, 12000 nurses
- UCSF and UCLA are in US News top 10
- 5 NCI Comprehensive Cancer Centers, 5 NIH CTSA
- IRB reliance, centralized contracting

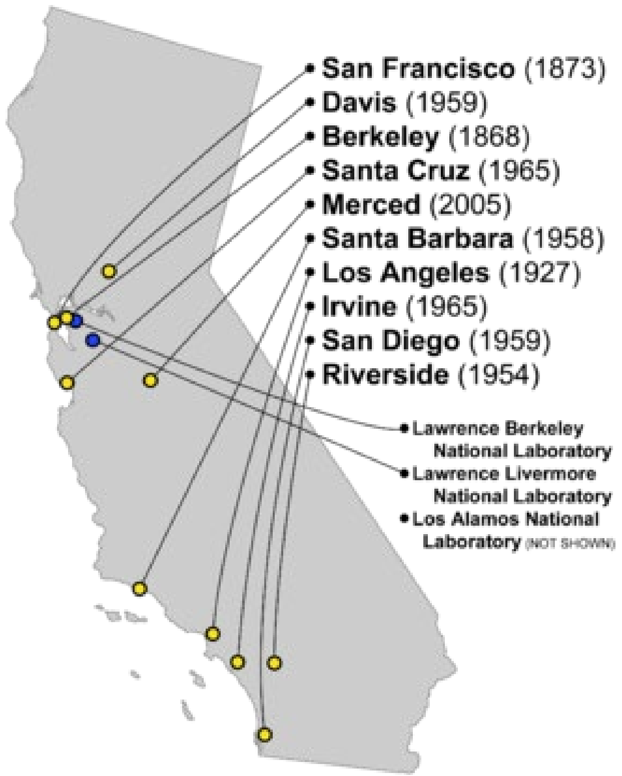

The University of California has an incredible view of the medical system

- Combined EHR data from UCSF, UCLA, UC Irvine, UC Davis, UC San Diego, and UC Riverside
  - 15 million patients treated over the past 15 years
- Central database built using OMOP (not Epic) as a data backend
  - Structured data from 2012 to the present day
  - **6.7 million patients with “modern” data**
  - 205M encounters, 563M procedures, 747M med orders, 680M diagnosis codes, 2B lab tests and vital signs
  - “From Tylenol to CAR-T cells…”
  - OSHPD data, pathology and radiology text elements, death index
  - Claims data from our self-funded plans now included
  - Continually harmonizing elements
- Quality and performance dashboards

Real world data: Comparing actually treated patients with original pivotal studies

- AbbVie Humira
  - RA: PREMIER 2 year double-blind, active comparator-controlled, multicenter study, N=799 (542 treated with Humira)
  - RA: DE019 4-week, randomized, double-blind, placebo-controlled, multicenter study, N=619
  - IBD: CLASSIC-I 4-week, randomized, double-blind, placebo-controlled, multicenter study, N=299
  - IBD: CHARM 56-week randomized, double-blind, placebo-controlled, multicenter study, N=854
  - University of California: **10,698 patients treated** with 59,163 prescriptions/orders/doses
- Celgene Revlimid
  - Randomized, multicenter, open-label, three-arm trial, N = 1,623
  - University of California: **5,152 patients treated** with 91,383 prescriptions/orders/doses
- Regeneron Praluent
  - ODYSSEY LONG TERM double-blind, placebo-controlled study, N=2,341
  - University of California: **1,341 patients treated** with 10,133 prescriptions/orders/doses

Predicting the future state of a patient with Rheumatoid Arthritis

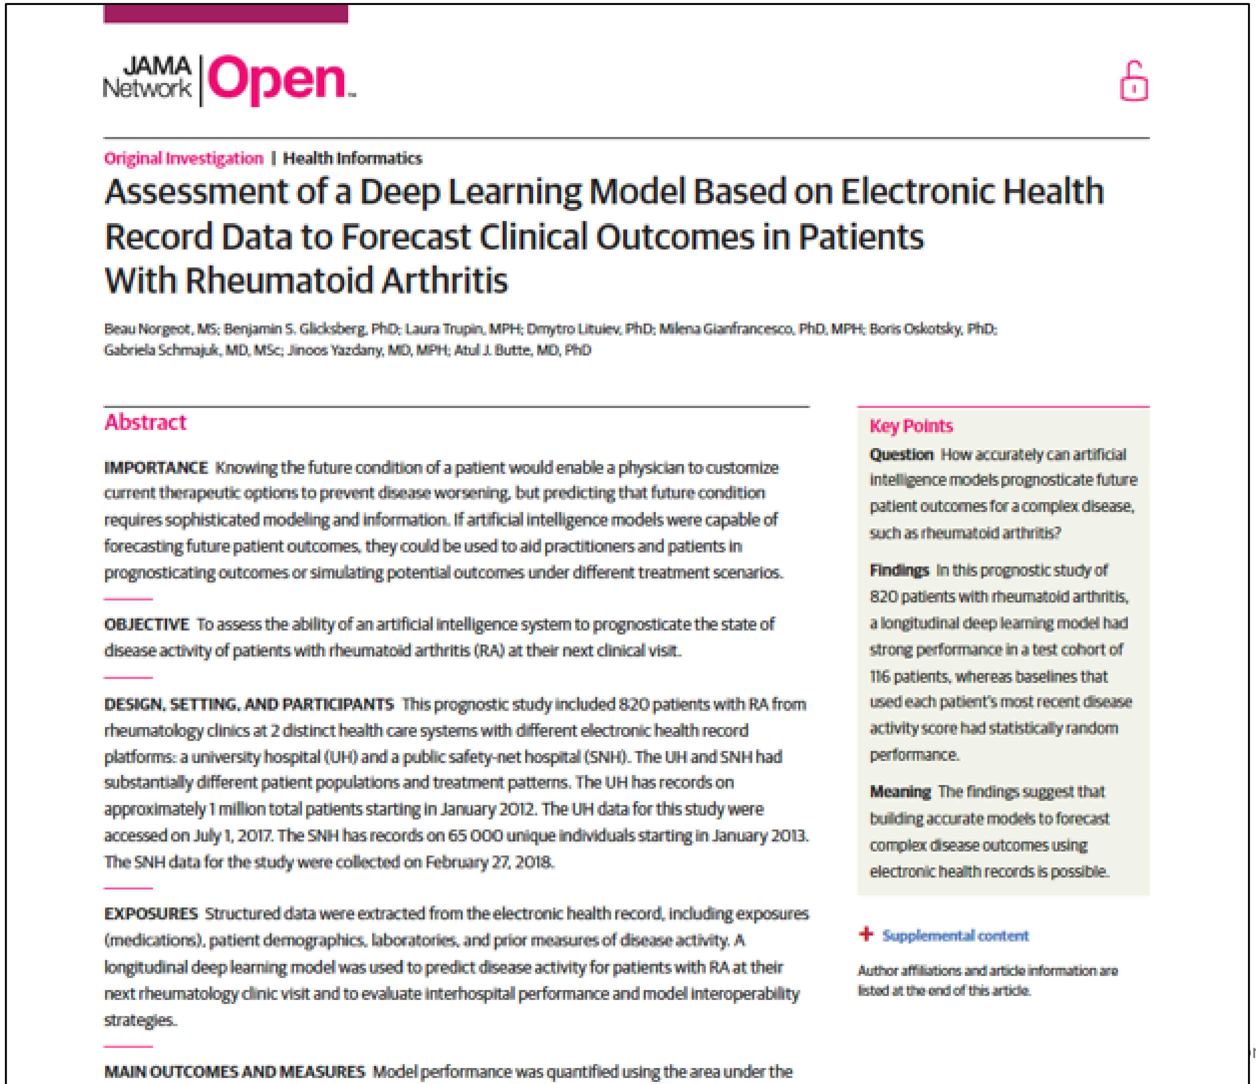
JAMA Network | Open

Original Investigation | Health Informatics

Assessment of a Deep Learning Model Based on Electronic Health Record Data to Forecast Clinical Outcomes in Patients With Rheumatoid Arthritis

Beau Norgeot, MS; Benjamin S. Glicksberg, PhD; Laura Trupin, MPH; Dmytro Lituiev, PhD; Milena Gianfrancesco, PhD, MPH; Boris Oskotsky, PhD; Gabriela Schmajuk, MD, MSc; Jinoos Yazdany, MD, MPH; Atul J. Butte, MD, PhD

Abstract

IMPORTANCE Knowing the future condition of a patient would enable a physician to customize current therapeutic options to prevent disease worsening, but predicting that future condition requires sophisticated modeling and information. If artificial intelligence models were capable of forecasting future patient outcomes, they could be used to aid practitioners and patients in prognosticating outcomes or simulating potential outcomes under different treatment scenarios.

OBJECTIVE To assess the ability of an artificial intelligence system to prognosticate the state of disease activity of patients with rheumatoid arthritis (RA) at their next clinical visit.

DESIGN, SETTING, AND PARTICIPANTS This prognostic study included 820 patients with RA from rheumatology clinics at 2 distinct health care systems with different electronic health record platforms: a university hospital (UH) and a public safety-net hospital (SNH). The UH and SNH had substantially different patient populations and treatment patterns. The UH has records on approximately 1 million total patients starting in January 2012. The UH data for this study were accessed on July 1, 2017. The SNH has records on 65 000 unique individuals starting in January 2013. The SNH data for the study were collected on February 27, 2018.

EXPOSURES Structured data were extracted from the electronic health record, including exposures (medications), patient demographics, laboratories, and prior measures of disease activity. A longitudinal deep learning model was used to predict disease activity for patients with RA at their next rheumatology clinic visit and to evaluate interhospital performance and model interoperability strategies.

MAIN OUTCOMES AND MEASURES Model performance was quantified using the area under the

Key Points

Question How accurately can artificial intelligence models prognosticate future patient outcomes for a complex disease, such as rheumatoid arthritis?

Findings In this prognostic study of 820 patients with rheumatoid arthritis, a longitudinal deep learning model had strong performance in a test cohort of 116 patients, whereas baselines that used each patient's most recent disease activity score had statistically random performance.

Meaning The findings suggest that building accurate models to forecast complex disease outcomes using electronic health records is possible.

+ Supplemental content

Author affiliations and article information are listed at the end of this article.

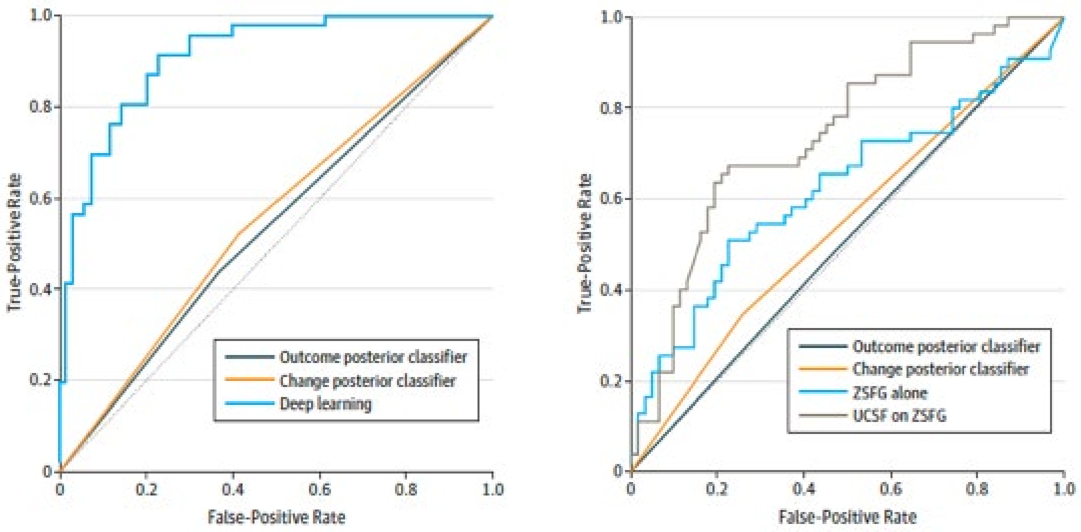

Beau Norgeot

bit.ly/jamaRA

Enabling UC researchers and patients to go beyond… machine learning in a safe, respectful, fair, equitable way in medicine


npj | Digital Me

ARTICLE OPEN

Scalable and
records

Alvin Rajkomar[1,2], Eyal Oren[1], Kai Chen[1], Andrew M. Dai[1], Nissan Hajaj[1], Michaela Hardt[1], Peter J. Liu[1], Xiaobing Liu[1], Jake Marcus[1], Mimi Sun[1], Patrik Sundberg[1], Hector Yee[1], Kun Zhang[1], Yi Zhang[1], Gerardo Flores[1], Gavin E. Duggan[1], Jamie Irvine[1], Quoc Le[1], Kurt Litsch[1], Alexander Mossin[1], Justin Tansuwan[1], De Wang[1], James Wexler[1], Jimbo Wilson[1], Dana Ludwig[2], Samuel L. Volchenboum Katherine Chou[1], Michael Pearson[1], Srinivasan Madabushi[1], Nigam H. Shah[4], Atul J. Butte[2], Michael D. Howell[1], Claire Cui[1], Greg S. Corrado[1] and Jeffrey Dean[1]

Predictive modeling with electronic health record (EHR) data is anticipated to drive personalized medicine and improve healthcare quality. Constructing predictive statistical models typically requires extraction of curated predictor variables from normalized EHR data, a labor-intensive process that discards the vast majority of information in each patient's record. We propose a representation of patients' entire raw EHR records based on the Fast Healthcare Interoperability Resources (FHIR) format. We demonstrate that deep learning methods using this representation are capable of accurately predicting multiple medical events from multiple centers without site-specific data harmonization. We validated our approach using de-identified EHR data from two US academic medical centers with 216,221 adult patients hospitalized for at least 24 h. In the sequential format we propose, this volume of EHR data unrolled into a total of 46,864,534,945 data points, including clinical notes. Deep learning models achieved high accuracy for tasks such as predicting: in-hospital mortality (area under the receiver operator curve [AUROC] across sites 0.93–0.94), 30-day unplanned readmission (AUROC 0.75–0.76), prolonged length of stay (AUROC 0.85–0.86), and all of a patient's final discharge diagnoses (frequency-weighted AUROC 0.90). These models outperformed traditional, clinically-used predictive models in all cases. We believe that this approach can be used to create accurate and scalable predictions for a variety of clinical scenarios. In a case study of a particular prediction, we demonstrate that neural networks can be used to identify relevant information from the patient's chart.

npj Digital Medicine (2018)1:18 ; doi:10.1038/s41746-018-0029-1

INTRODUCTION

The promise of digital medicine stems in part from the hope that, by digitizing health data, we might more easily leverage computer information systems to understand and improve care. In fact, routinely collected patient healthcare data are now approaching the genomic scale in volume and complexity.[1] Unfortunately, most of this information is not yet used in the sorts of predictive statistical models clinicians might use to improve care delivery. It is widely suspected that use of such efforts, if successful, could provide major benefits not only for patient safety and quality but also in reducing healthcare costs.[2–6]

In spite of the richness and potential of available data, scaling the development of predictive models is difficult because, for traditional predictive modeling techniques, each outcome to be predicted requires the creation of a custom dataset with specific variables.[7] It is widely held that 80% of the effort in an analytic model is preprocessing, merging, customizing, and cleaning datasets,[8,9] not analyzing them for insights. This profoundly limits the scalability of predictive models.

Another challenge is that the number of potential predictor variables in the electronic health record (EHR) may easily number in the thousands, particularly if free-text notes from doctors, nurses, and other providers are included. Traditional modelin approaches have dealt with this complexity simply by choosing very limited number of commonly collected variables to consider This is problematic because the resulting models may produ imprecise predictions: false-positive predictions can overwhel physicians, nurses, and other providers with false alarms ar concomitant alert fatigue,[10] which the Joint Commission identifie as a national patient safety priority in 2014.[11] False-negativ predictions can miss significant numbers of clinically importa events, leading to poor clinical outcomes.[11,12] Incorporating th entire EHR, including clinicians' free-text notes, offers some hop of overcoming these shortcomings but is unwieldy for mo predictive modeling techniques.

Recent developments in deep learning and artificial neur networks may allow us to address many of these challenges ar unlock the information in the EHR. Deep learning emerged as th preferred machine learning approach in machine percepti problems ranging from computer vision to speech recognitio but has more recently proven useful in natural languag processing, sequence prediction, and mixed modality da settings.[13–17] These systems are known for their ability to hand large volumes of relatively messy data, including errors in labe

[1]Google Inc, Mountain View, CA, USA; [2]University of California, San Francisco, San Francisco, CA, USA; [3]University of Chicago Medicine, Chicago, IL, USA and [4]Stanford Universi Stanford, CA, USA
Correspondence: Alvin Rajkomar (alvinrajkomar@google.com)
These authors contributed equally: Alvin Rajkomar, Eyal Oren

Received: 26 January 2018 Revised: 14 March 2018 Accepted: 26 March 2018
Published online: 08 May 2018

Published in partnership with the Scripps Translational Science Institute

npj nature partner journals

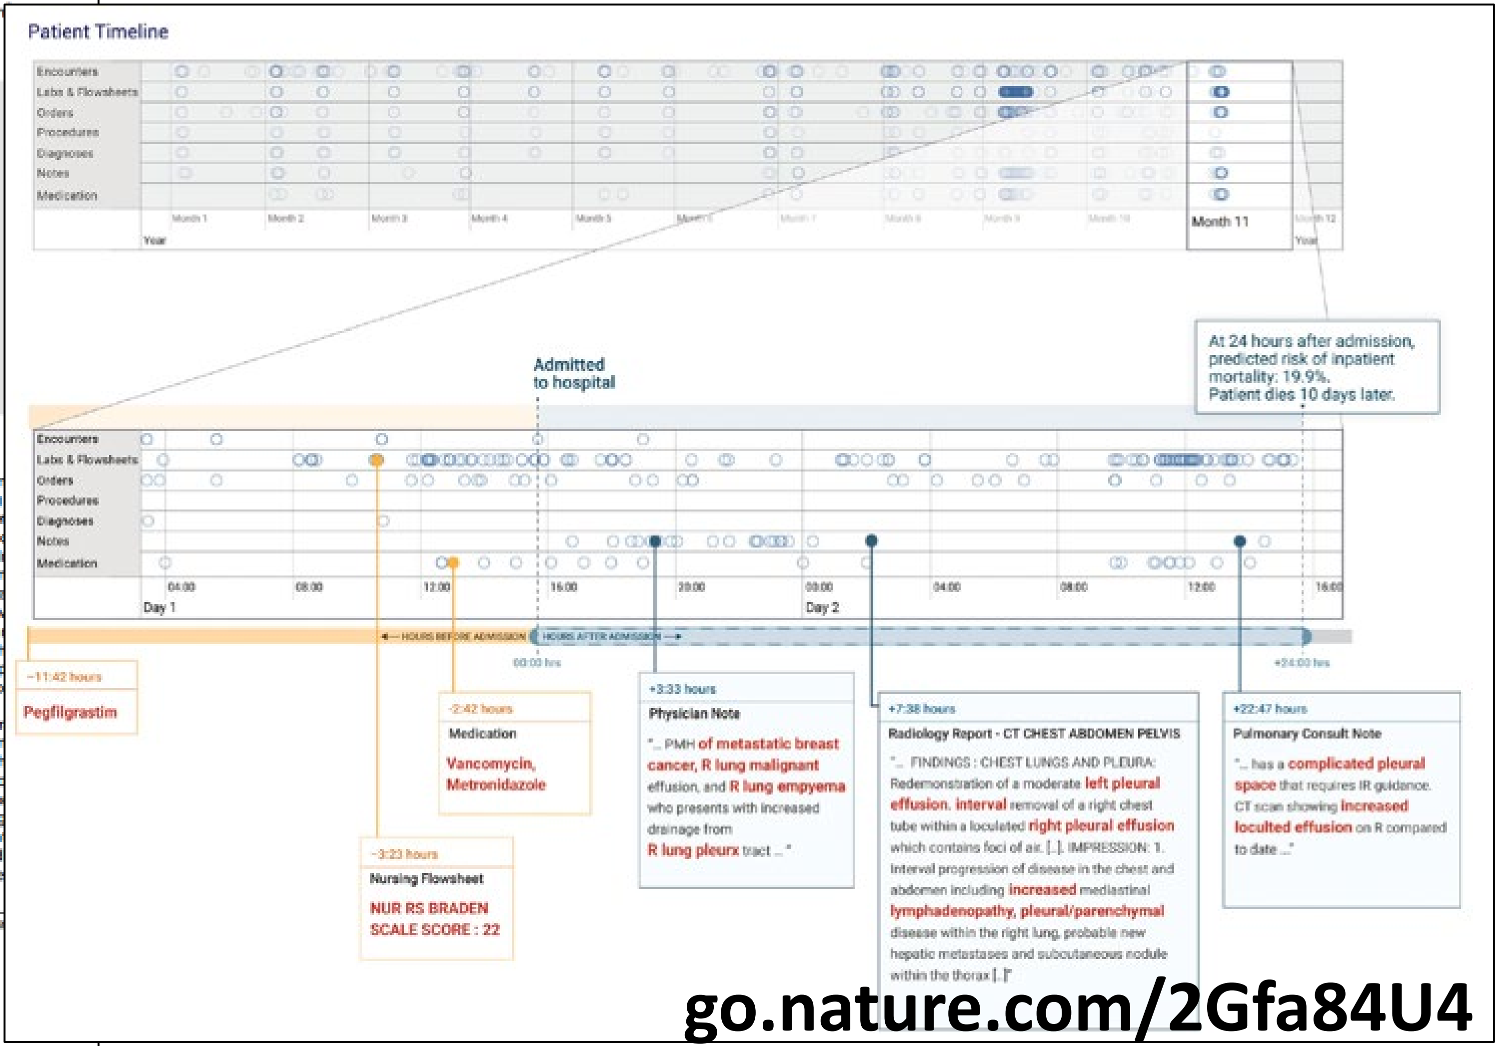

go.nature.com/2Gfa84U4

/staff & contractors by the European Medicines Agency

A New Deep Learning Healthcare System

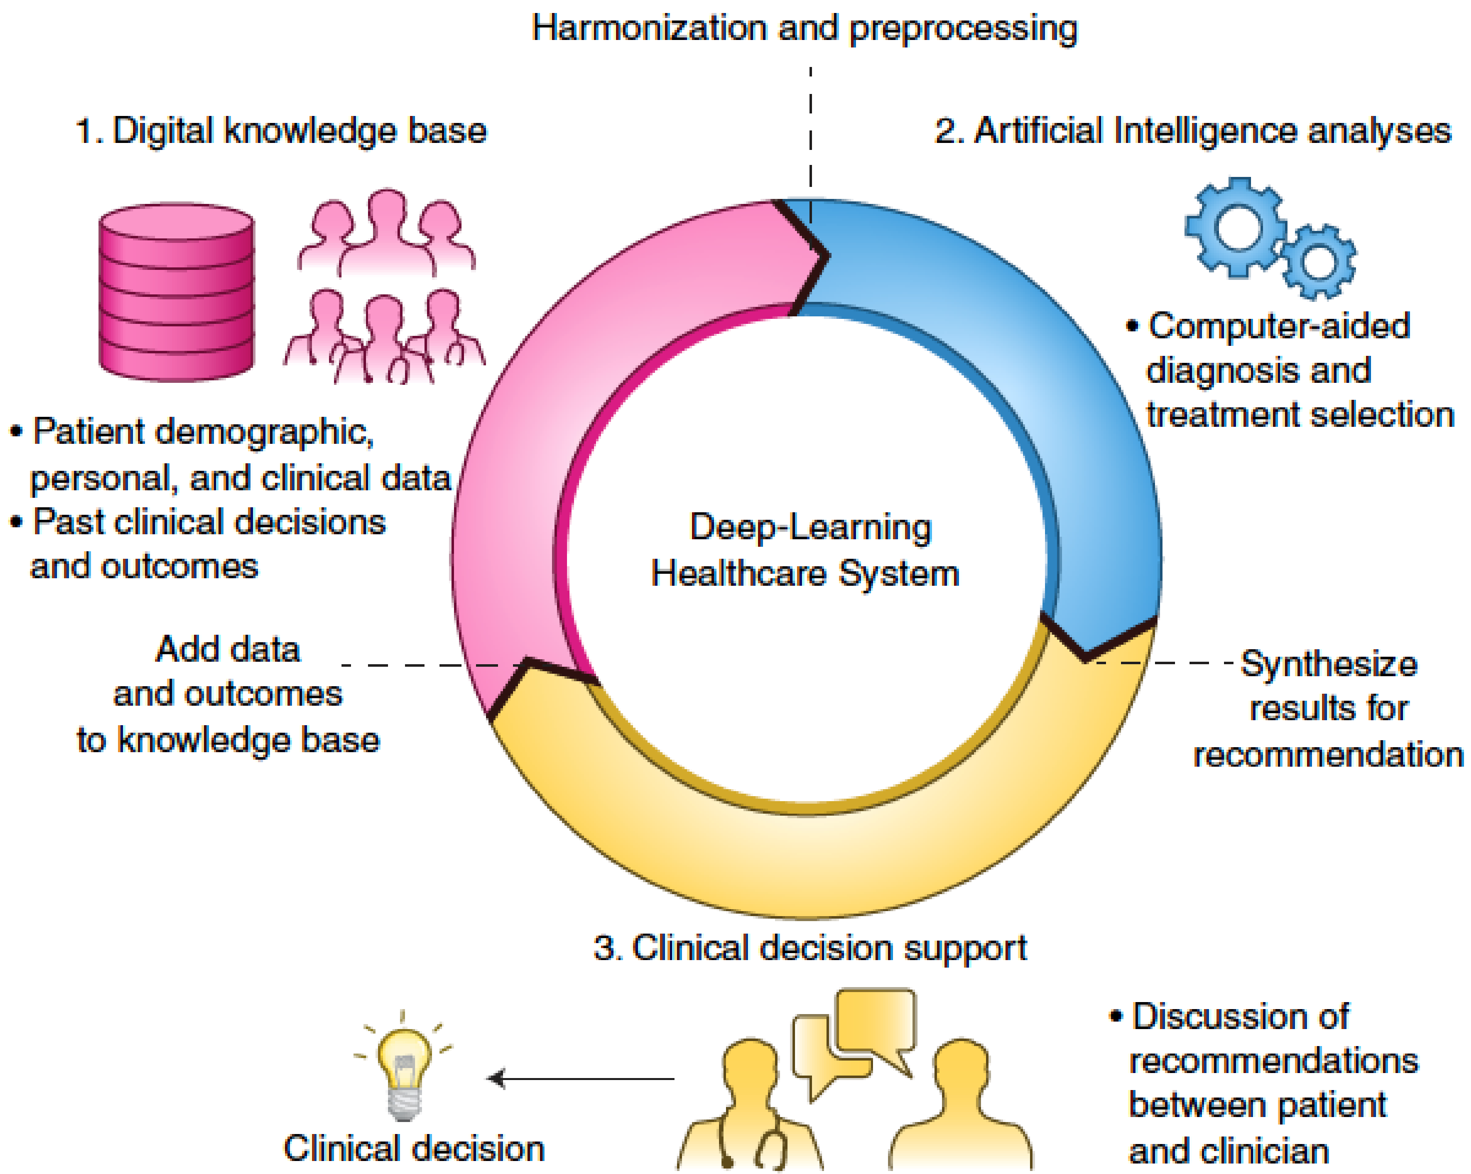

Beau Norgeot

bit.ly/DLHCare

Must be aware of biases in training data and methods

MIT Technology Review

Artificial intelligence Oct 25

A biased medical algorithm favored
people for health-care programs

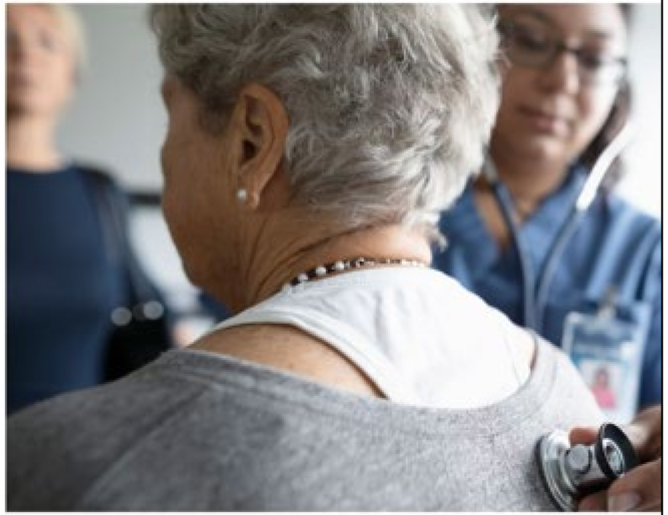

A study has highlighted the risks inherent in using historical dat
learning algorithms to make predictions.

The news: An algorithm that many US health providers use to
will most need extra medical care privileged white patients ove
according to researchers at UC Berkeley, whose study was pu
Effectively, it bumped whites up the queue for special treatments for complex
conditions like kidney problems or diabetes.

The study: The researchers dug through almost 50,000 records from a large,
undisclosed academic hospital. They found that white patients were given hig

How Biased is GPT-3?

Despite its impressive performance, the world's newest language model reflects societal biases in gender, race, and religion

Catherine Yeo Follow
Jun 3 · 4 min read ★

Last week, O
model that b
known as th
175 billion p
billion paran

GPT-3
that hu
article
GPT-3

Its release w
Learners", a
that this pap
broader soc
Representa

NEWS · 24 JANUARY 2020

The battle for ethical AI at the world's biggest machine-learning conference

Bias and the prospect of societal harm increasingly plague artificial-intelligence research – but it's not clear who should be on the lookout for these problems.

Elizabeth Gibney

PDF version

RELATED ARTICLES

How one conference embraced diversity

Can a major AI conference shed its reputation for hosting sexist behaviour?

Bias detectives: the researchers striving to make

Minimum information about clinical artificial intelligence modeling: the MI-CLAIM checklist

Here we present the MI-CLAIM checklist, a tool intended to improve transparent reporting of AI algorithms in medicine.

Beau Norgeot, Giorgio Quer, Brett K. Beaulieu-Jones, Ali Torkamani, Raquel Dias, Milena Gianfrancesco, Rima Arnaout, Isaac S. Kohane, Suchi Saria, Eric Topol, Ziad Obermeyer, Bin Yu and Atul J. Butte

The application of artificial intelligence (AI) in medicine is an old idea[1–3], but methods for this in the past involved programming computers with patterns or rules ascertained from human experts, which resulted in deterministic, rules-based systems. The study of AI in medicine has grown tremendously in the past few years due to increasingly available datasets from medical practice, including clinical images, genetics, and electronic health records, as well as the maturity of methods that use data to teach computers[4–6]. The use of data labeled by clinical experts to train machine, probabilistic, and statistical models is called 'supervised machine learning'. Successful uses of these new machine-learning approaches include targeted real-time early-warning systems for adverse events[7], the detection of diabetic retinopathy[8], the classification of pathology and other images, the prediction of the near-term future state of patients with rheumatoid arthritis[9], patient discharge disposition[10], and more.

go.nature.com/36Y8bGW

Before paper submission			
Study design (Part 1)	**Completed: page number**		**Notes if not completed**
The clinical problem in which the model will be employed is clearly detailed in the paper.	□		
The research question is clearly stated.	□		
The characteristics of the cohorts (training and test sets) are detailed in the text.	□		
The cohorts (training and test sets) are shown to be representative of real-world clinical settings.	□		
The state-of-the-art solution used as a baseline for comparison has been identified and detailed.	□		
Data and optimization (Parts 2, 3)	**Completed: page number**		**Notes if not completed**
The origin of the data is described and the original format is detailed in the paper.	□		
Transformations of the data before it is applied to the proposed model are described.	□		
The independence between training and test sets has been proven in the paper.	□		
Details on the models that were evaluated and the code developed to select the best model are provided.	□		
Is the input data type structured or unstructured?	□ Structured □ Unstructured		
Model performance (Part 4)	**Completed: page number**		**Notes if not completed**
The primary metric selected to evaluate algorithm performance (e.g., AUC, F-score, etc.), including the justification for selection, has been clearly stated.	□		
The primary metric selected to evaluate the clinical utility of the model (e.g., PPV, NNT, etc.), including the justification for selection, has been clearly stated.	□		
The performance comparison between baseline and proposed model is presented with the appropriate statistical significance.	□		
Model examination (Part 5)	**Completed: page number**		**Notes if not completed**
Examination technique 1[a]	□		
Examination technique 2[a]	□		
A discussion of the relevance of the examination results with respect to model/algorithm performance is presented.	□		
A discussion of the feasibility and significance of model interpretability at the case level if examination methods are uninterpretable is presented.	□		
A discussion of the reliability and robustness of the model as the underlying data distribution shifts is included.	□		
Reproducibility (Part 6): choose appropriate tier of transparency			**Notes**
Tier 1: complete sharing of the code		□	
Tier 2: allow a third party to evaluate the code for accuracy/fairness; share the results of this evaluation		□	
Tier 3: release of a virtual machine (binary) for running the code on new data without sharing its details		□	
Tier 4: no sharing		□	

PPV, positive predictive value; NNT, numbers needed to treat.

[a]Common examination approaches based on study type: for studies involving exclusively structured data, coefficients and sensitivity analysis are often appropriate; for studies involving unstructured data in the domains of image analysis or natural language processing, saliency maps (or equivalents) and sensitivity analyses are often appropriate.

go.nature.com/36Y8bGW

VIEWPOINT

The Case for Algorithmic Stewardship for Artificial Intelligence and Machine Learning Technologies

Stephanie Eaneff, MSP
Berkeley Institute for Data Science, University of California, Berkeley; and Bakar Computational Health Sciences Institute, University of California, San Francisco.

Ziad Obermeyer, MD
University of California, Berkeley School of Public Health.

Atul J. Butte, MD, PhD
Bakar Computational Health Sciences Institute, University of California, San Francisco; and Center for Data-Driven Insights and Innovation, University of California Health, Oakland.

The first manual on hospital administration, published in 1808, described a hospital steward as "an individual who [is] honest and above reproach," with duties including the purchasing and management of hospital materials.[1] Today, a steward's job can be seen as ensuring the safe and effective use of clinical resources. The Joint Commission, for instance, requires antimicrobial stewardship programs to support appropriate antimicrobial use, including by monitoring antibiotic prescribing and resistance patterns.

A similar approach to "algorithmic stewardship" is now warranted. Algorithms, or computer-implementable instructions to perform specific tasks, are available for clinical use, including complex artificial intelligence (AI) and machine learning (ML) algorithms and simple rule-based algorithms. More than 50 AI/ML algorithms have been cleared by the US Food and Drug Administration[2] for uses that include identifying intracranial hemorrhage from brain computed tomographic scans[3] and detecting seizures in real time.[4] Algorithms are also used to inform clinical operations, such as predicting which patients will health systems must also develop oversight frameworks to ensure that algorithms are used safely, effectively, and fairly. Such efforts should focus particularly on complex and predictive algorithms that necessitate additional layers of quality control. Health systems that use predictive algorithms to provide clinical care or support operations should designate a person or group responsible for algorithmic stewardship. This group should be advised by clinicians who are familiar with the language of data, patients, bioethicists, scientists, and safety and regulatory organizations. In this Viewpoint, drawing from best practices from other areas of clinical practice, several key considerations for emerging algorithmic stewardship programs are identified.

Create and Maintain an Algorithm Inventory

Health systems should inventory all predictive algorithms currently in use, with a particular emphasis on understanding the exact outcome being predicted and the decisions made on the basis of those predictions. This is particularly important because recent work has shown that algorithms can reach enormous scale

bit.ly/algostew

Figure. Existing and Proposed Processes and Tools to Ensure Appropriate Use of Drugs for Algorithmic Stewardship Efforts

	Existing processes and tools	Proposed processes and tools for algorithmic stewardship
1 Clinical trials	Phase 1, 2, and 3 trials	Assess safety, efficacy, and fairness (potentially via clinical trials)
2 Scale-up and early adoption	Hospital formulary	Algorithm inventory
3 Postmarket use and evaluation	Medication use evaluations	Algorithm use evaluations
4 Ongoing oversight	Antimicrobial steward role	Algorithmic steward role

bit.ly/algostew

go.nature.com/32qDBT9

npj Digital Medicine

www.nature.com/npjdigitalmed

COMMENT OPEN

Check for updates

Characteristics and challenges of the clinical pipeline of digital therapeutics

Nisarg A. Patel[1,2]✉ and Atul J. Butte[2]

In this Comment, we characterize the current pipeline of digital therapeutics and offer a clinical perspective into the advantages, challenges, and barriers to implementation of this treatment modality for patient care, which we hope will inform future regulatory policy, prescribing decisions, and scope of real-world evidence collection.

npj Digital Medicine (2020)3:159 ; https://doi.org/10.1038/s41746-020-00370-8

Digital therapeutics (DTx), defined by the Digital Therapeutics Alliance as "evidence-based therapeutic interventions driven by high-quality software programs to prevent, manage, or treat a medical disorder or disease[1]", have emerged as a new therapeutic modality for the prevention, management, or treatment of chronic, behavior-modifiable disease. Akin to biopharmaceuticals and medical devices, DTx undergo review and are cleared or approved by the U.S. Food and Drug Administration (FDA) and are either available over-the-counter or prescribed by physicians. As of this writing, the FDA has cleared or approved multiple DTx on the basis of superiority trial data, such as Welldoc's BlueStar for the management of Type II diabetes in 2010 and Pear Therapeutics'

If the DTx is prescription-only, physicians must currently complete and send a patient enrollment to the DTx manufacturer, who then onboards the patient directly onto the DTx via a mobile app store access code. Patients interact with the DTx over a predetermined period of time, per the drug's label, directly from a personal smartphone or tablet; clinicians may also be able to monitor their patients' progress and input-related information, such as drug screen results and appointment compliance, on a dedicated web dashboard[9].

DTx may serve as a useful complement, and in certain cases, replacement, to biomedical therapeutics. DTx to date have largely targeted neurological and psychiatric conditions with significant

Company	Product	Indication(s)	Status	Commercial partner	Investment partner
Pear Therapeutics	reSET	Substance use disorder[25,26]	Marketed	Sandoz[a][27]	Novartis
	reSET-O	Opioid use disorder[28]	Marketed	Sandoz[a]	
	Somryst	Chronic insomnia[29]	Marketed		
	Pear-004	Schizophrenia	Marketed[b] Pivotal		
	Pear-006	Multiple sclerosis	Discovery	Novartis[30]	
	Unspecified	Gastrointestinal conditions	Discovery	Ironwood Pharmaceuticals[31]	
Welldoc	Bluestar	Type I diabetes, Type II diabetes[32]	Marketed	Astellas[33]	
Akili Interactive	Endeavor	ADHD[34]	Marketed	Shionogi[35]	Amgen, Merck
	AKL-T02	Autism spectrum disorder[36]	Pilot		
	AKL-T03	Major depressive disorder	Pilot		
Nightware	Nightware	Post-traumatic stress disorder	Marketed		
Click Therapeutics	CT-152	Major depressive disorder	Pivotal	Otsuka[37]	Sanofi, Hikma
Cognoa	Autism Diagnostic	Autism spectrum disorder	Pivotal	EVERSANA	
	Autism Therapeutic	Autism spectrum disorder[38]	Feasibility		
Biofourmis	BiovitalsHF V1	Heart failure	Marketed	Novartis[39]	
	BiovitalsHF V2	Heart failure[40]	Pivotal		
	BF140	Pain[41]	Pilot	Chugai[42]	
Propeller Health	Propeller	Asthma[43], chronic obstructive pulmonary disease[44]	Marketed	AstraZeneca[45], GlaxoSmithKline, Novartis, Orion, Boehringer Ingelheim	
AppliedVR	EaseVRx	Chronic pain[46]	Efficacy		
	RelieVRx	Acute postoperative pain[47]	Efficacy		
	AnxietyVRx	Generalized anxiety	Discovery		
Happify Health	Happify	Multiple sclerosis-associated depression and anxiety	Discovery	Sanofi[48]	

go.nature.com/32qDBT9

U.S. FOOD & DRUG ADMINISTRATION