

Machine Learning for Healthcare

HST.956, 6.S897

Lecture 1: What makes healthcare unique?

Prof. David Sontag & Pete Szolovits



The Problem

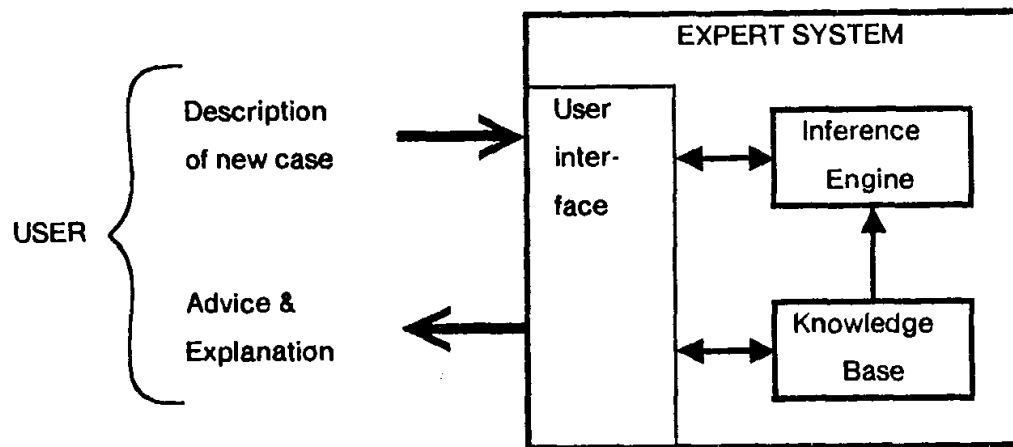
- Cost of health care expenditures in the US are over \$3 trillion, and rising
- Despite having some of the best clinicians in the world, chronic conditions are
 - Often diagnosed late
 - Often inappropriately managed
- Medical errors are pervasive

Outline for today's class

1. **Brief history of AI and ML in healthcare**
2. *Why now?*
3. Examples of how ML will transform healthcare
4. *What is unique* about ML in healthcare?
5. Overview of class syllabus

1970's: MYCIN expert system

- 1970's (Stanford): MYCIN expert system for identifying bacteria causing severe infections
- Proposed a good therapy in ~69% of cases. Better than infectious disease experts



Dialogue interface

I am ready

** THIS IS A 26 YEAR OLD MALE PATIENT

My understanding is:

The age of the patient is 26

The sex of the patient is male

** FIVE DAYS AGO, HE HAD RESPIRATORY-TRACT SYMPTOMS

What is his name?

** JO

My understanding is:

The name of the patient is Jo

Respiratory-tract is one of the symptoms that the patient had

** A COUPLE OF DAYS BEFORE THE ADMISSION, HE HAD A MALAISE

Please give me the date of admission

** MARCH 12, 1979

My understanding is:

The patient was admitted at the hospital 3 days ago

Malaise is one of the symptoms that the patient had 5 days ago

FIGURE 33-1 Short sample dialogue. The physician's inputs appear in capital letters after the double asterisks.

FIGURE 1-1 Major parts of an expert system. Arrows indicate information flow.

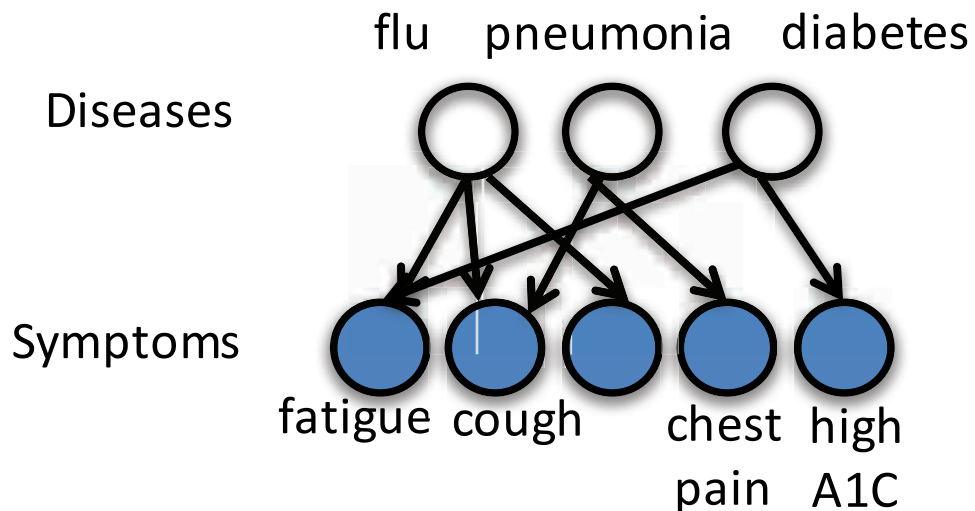
1980's: INTERNIST-1/QMR model

- 1980's (Univ. of Pittsburgh): INTERNIST-1/Quick Medical Reference
- Diagnosis for internal medicine

Probabilistic model relating:

570 binary disease variables
4,075 binary symptom variables
45,470 directed edges

Elicited from doctors:
15 person-years of work

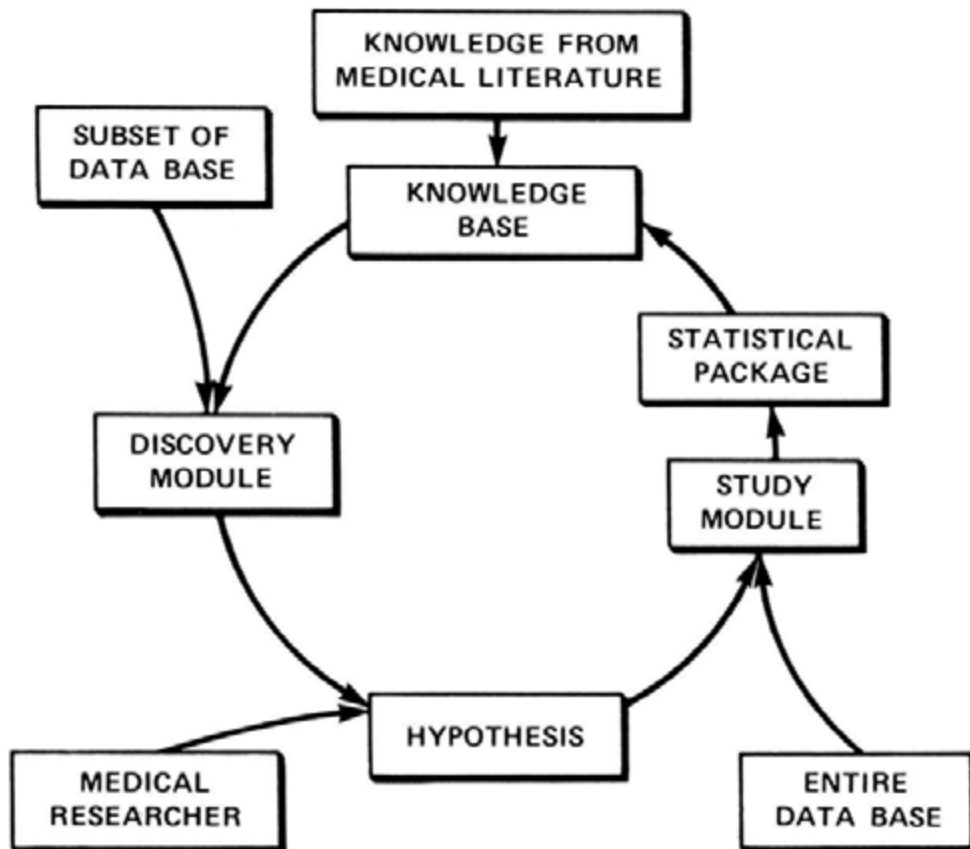


Led to advances in ML & AI
(Bayesian networks, approximate inference)

- Problems:**
1. Clinicians entered symptoms *manually*
 2. Difficult to maintain, difficult to generalize

1980's: automating medical discovery

RX PROJECT: AUTOMATED KNOWLEDGE ACQUISITION



Discovers that prednisone
elevates cholesterol
(Annals of Internal Medicine, '86)

[Robert Blum, "Discovery, Confirmation and Incorporation of Causal Relationships from a Large Time-Oriented Clinical Data Base: The RX Project". Dept. of Computer Science, Stanford. 1981]

1990's: neural networks in medicine

- Neural networks with clinical data took off in 1990, with 88 new studies that year
- Small number of features (inputs)
- Data often collected by chart review

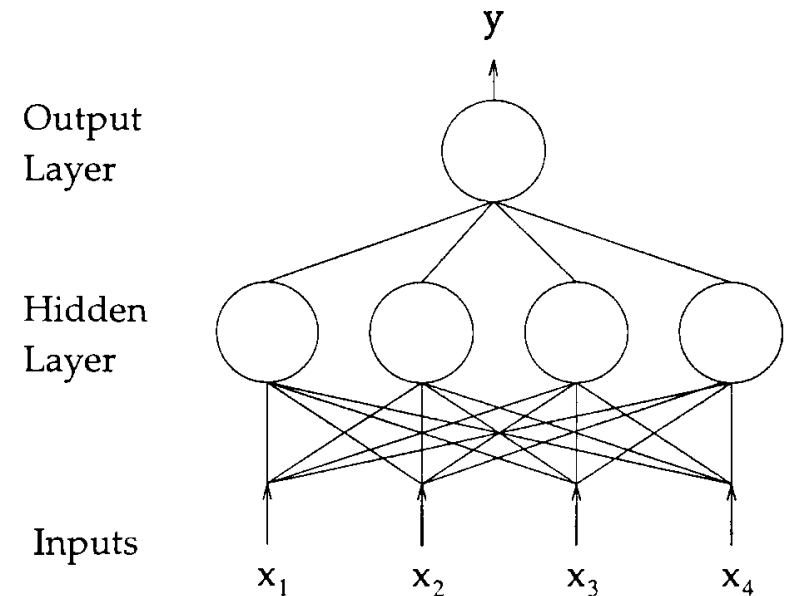


FIGURE 2. A multilayer perceptron. This is a two-layer perceptron with four inputs, four hidden units, and one output unit.

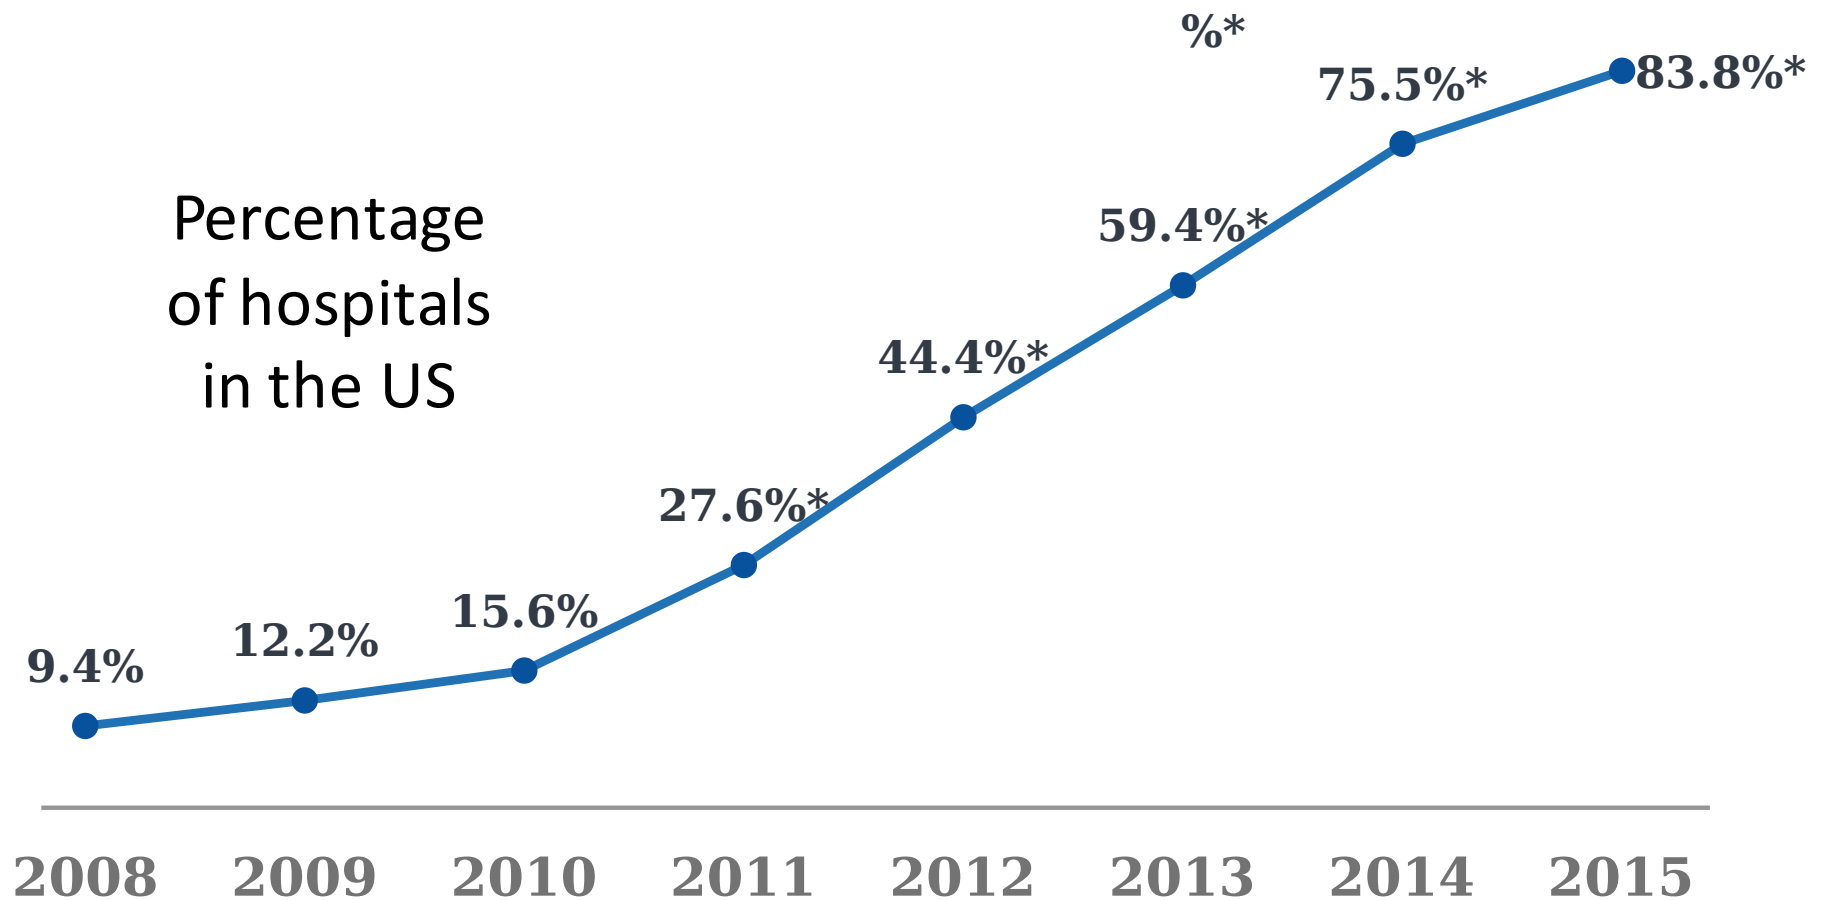
- Problems:**
1. Did not fit well into clinical workflow
 2. Hard to get enough training data
 3. Poor generalization to new places

[Penny & Frost, Neural Networks in Clinical Medicine. Med Decis Making, 1996]

Outline for today's class

1. Brief history of AI and ML in healthcare
2. **Why now?**
3. Examples of how ML will transform healthcare
4. What is *unique* about ML in healthcare?
5. Overview of class syllabus

The Opportunity: Adoption of Electronic Health Records (EHR) has increased 9x in US since 2008



Courtesy of Health and Human Services. Image is in the public domain.

[Henry et al., ONC Data Brief, May 2016]

Large datasets



Laboratory for
Computational
Physiology

De-identified
health data from
~40K critical care
patients

Demographics,
vital signs,
laboratory tests,
medications,
notes, ...

If you use MIMIC data or code in your work, please cite the following publication:

MIMIC-III, a freely accessible critical care database. Johnson AEW, Pollard TJ, Shen L, Lehman L, Feng M, Ghassemi M, Moody B, Szolovits P, Celi LA, and Mark RG. *Scientific Data* (2016). DOI: 10.1038/sdata.2016.35. Available from: <http://www.nature.com/articles/sdata201635>

Large datasets

President Obama's initiative to create a 1 million person research cohort

Core data set:

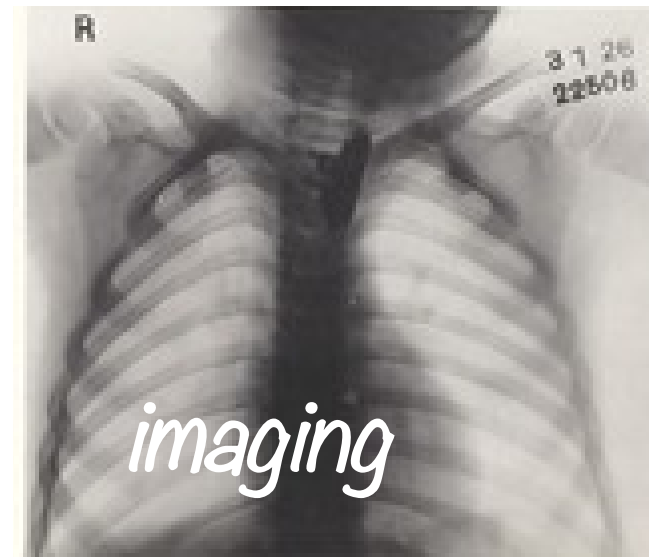
- Baseline health exam
- Clinical data derived from electronic health records (EHRs)
- Healthcare claims
- Laboratory data

[Precision Medicine Initiative (PMI) working Group Report, Sept. 17 2015]

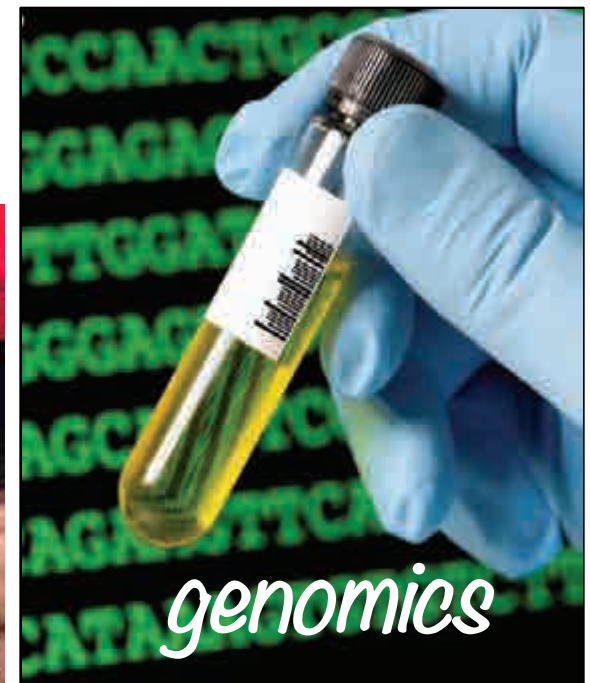
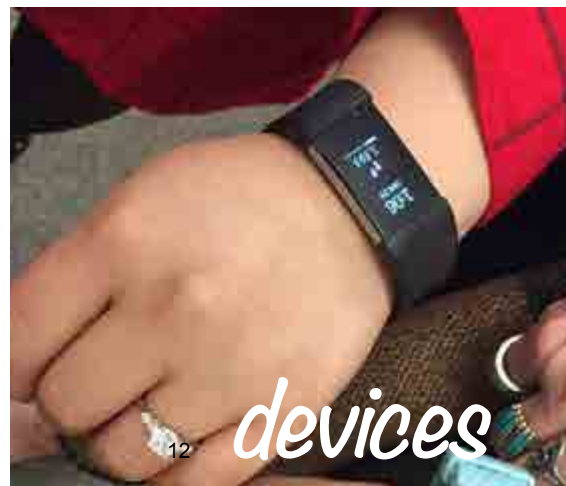
Diversity of digital health data



proteomics



social media



Standardization

- Diagnosis codes: ICD-9 and ICD-10 (International Classification of Diseases)

...

ICD-9 codes 290–319: mental disorders

ICD-9 codes 320–359: diseases of the nervous system

ICD-9 codes 360–389: diseases of the sense organs

ICD-9 codes 390–459: diseases of the circulatory system

ICD-9 codes 460–519: diseases of the respiratory system

ICD-9 codes 520–579: diseases of the digestive system

ICD-9 codes 580–629: diseases of the genitourinary system

ICD-9 codes 630–679: complications of pregnancy, childbirth,

...

...

https://en.wikipedia.org/wiki/List_of_ICD-9_codes

<https://blog.curemd.com/the-most-bizarre-icd-10-codes-infographic/>

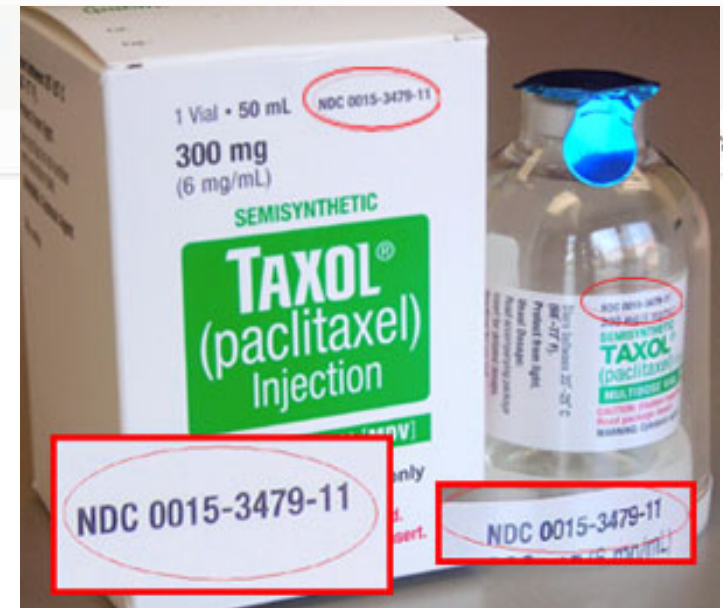
Standardization

- Diagnosis codes: ICD-9 and ICD-10 (International Classification of Diseases)
- Laboratory tests: LOINC codes
- Pharmacy: National Drug Codes (NDCs)
- Unified Medical Language System (UMLS): millions of medical concepts

LOINC
From Regenstrief

1 / 5

| LOINC | LongName |
|----------------|---|
| <u>27353-2</u> | Glucose mean value [Mass/volume] in Blood Estimated from glycated hemoglobin |
| <u>2352-3</u> | Glucose in CSF/Glucose plas |
| <u>49689-3</u> | Glucose tolerance [Interpretation] in Serum or Plasma Narrative—post 100 g glucose PO |
| <u>49688-5</u> | |
| <u>72650-5</u> | |



Standardization

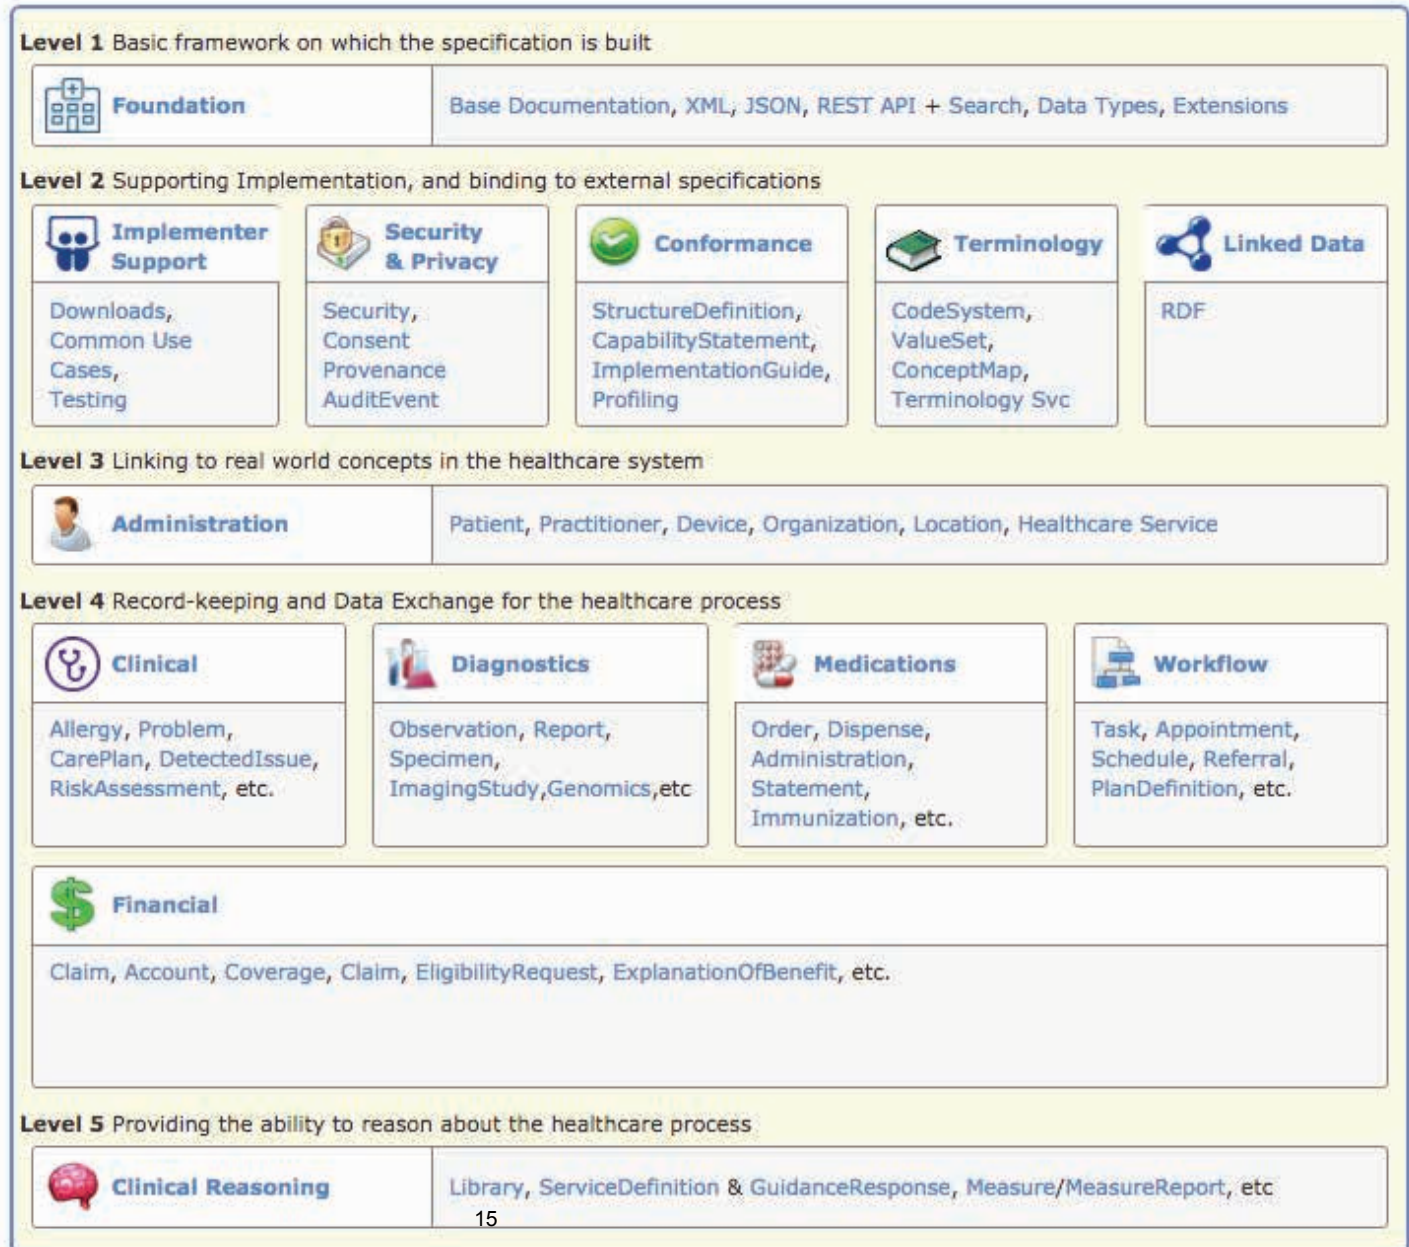
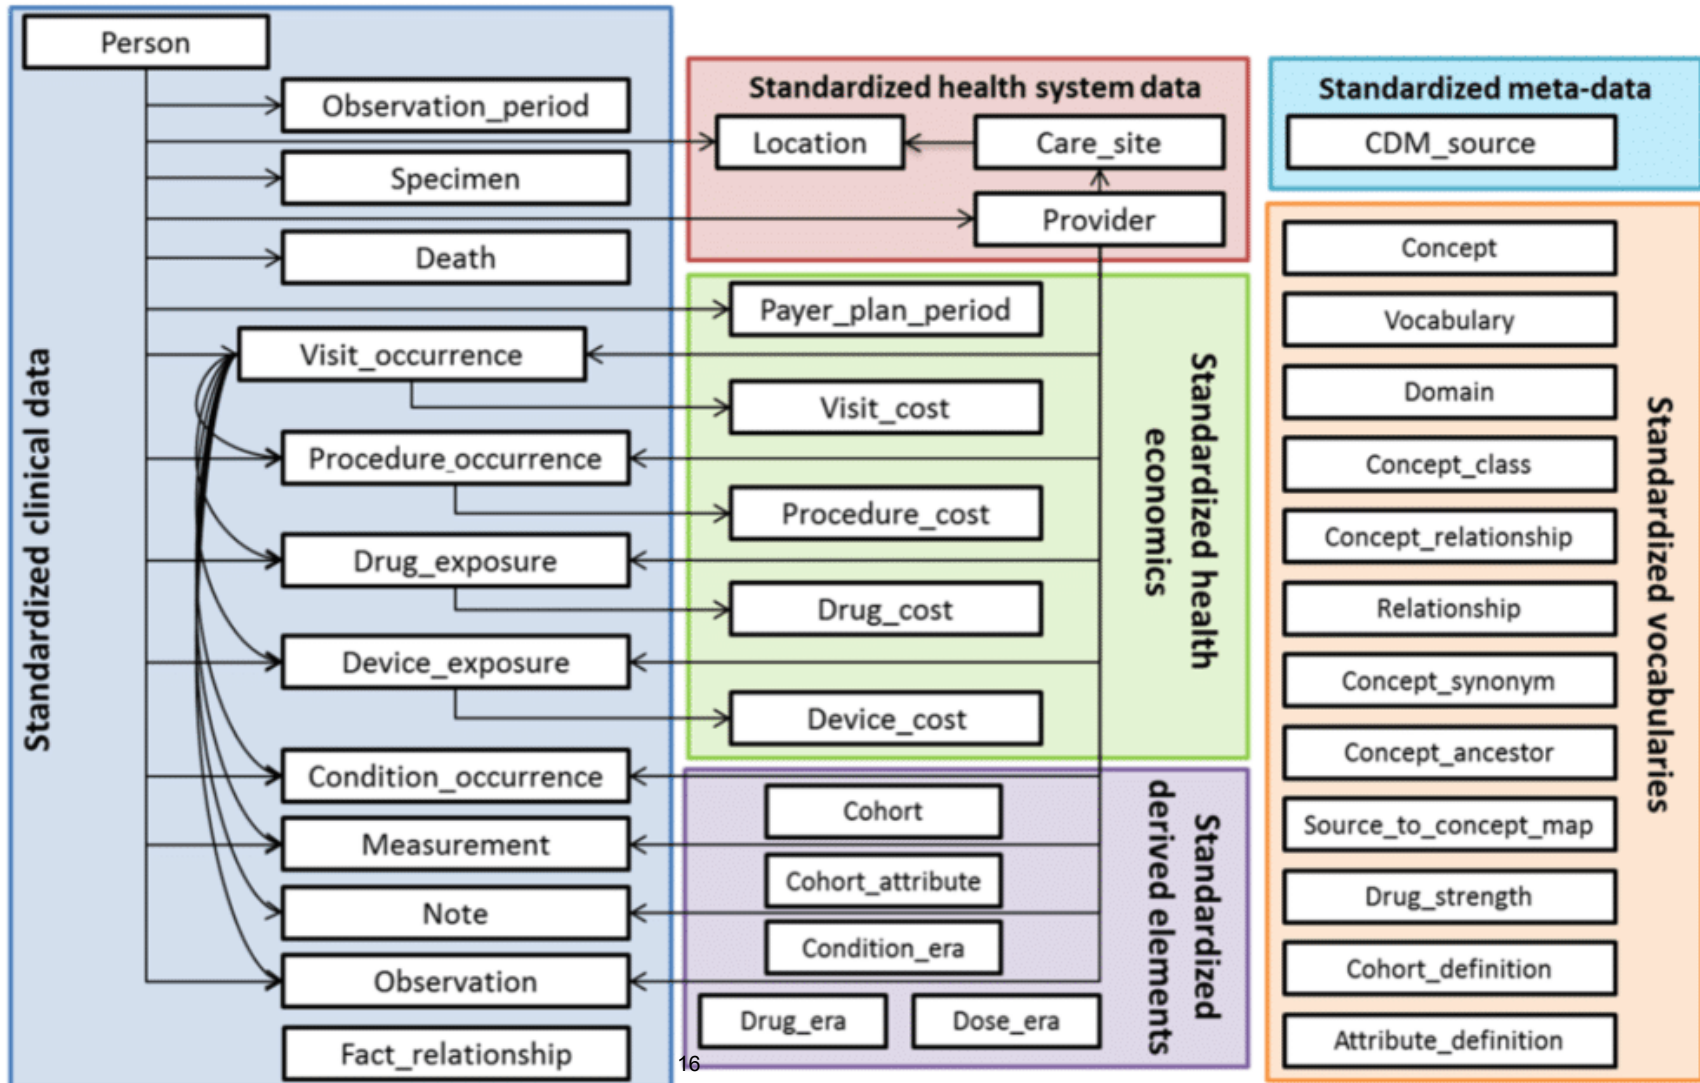


Image is in the public domain.

Standardization



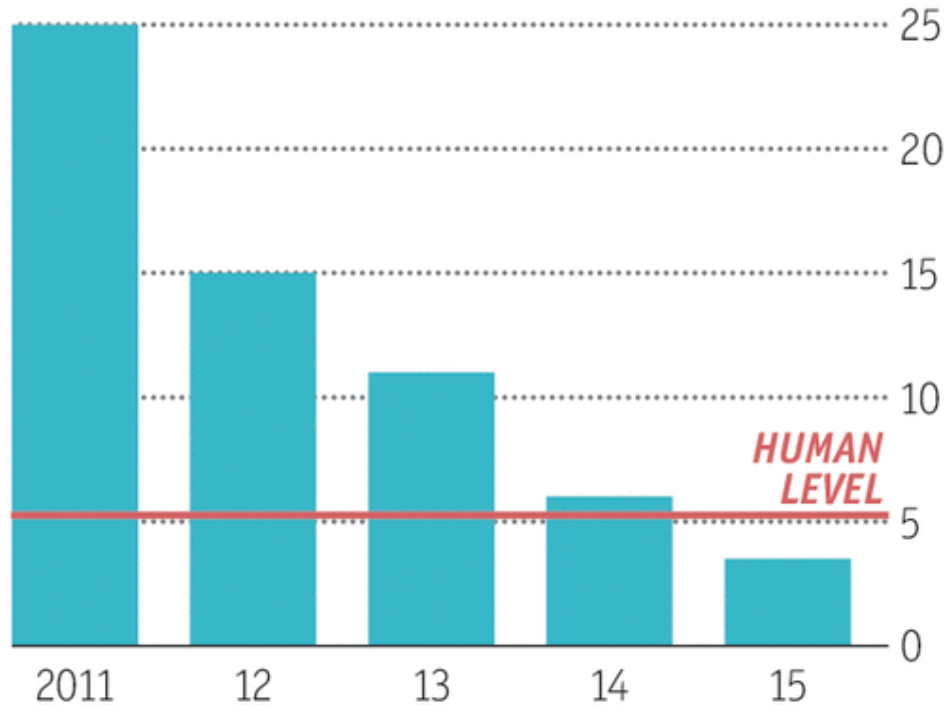
OMOP
Common
Data
Model v5.0



Breakthroughs in machine learning

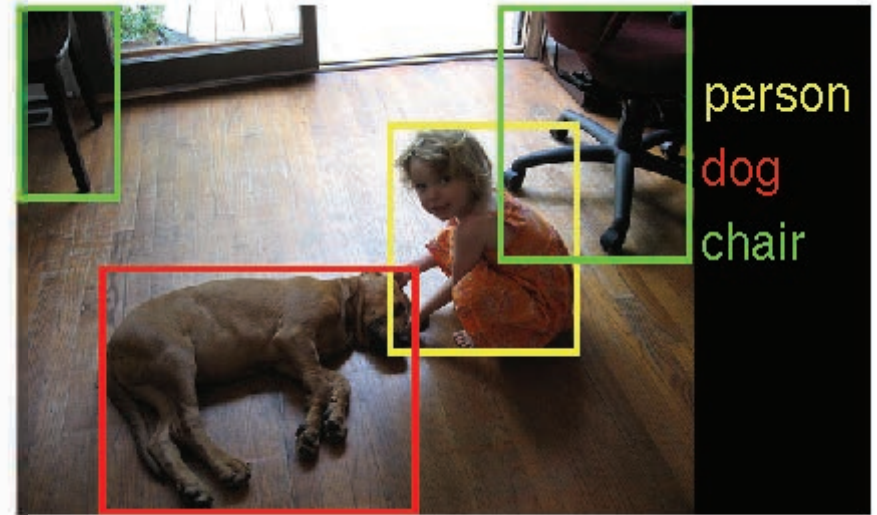
Ever cleverer

Error rates on ImageNet Visual Recognition Challenge, %



Sources: ImageNet; Stanford Vision Lab

Economist.com



Why now?

- Big data
- Algorithmic advances
- Open-source software

Breakthroughs in machine learning

- Major advances in ML & AI
 - Learning with high-dimensional features (e.g., l1-regularization)
 - Semi-supervised and unsupervised learning
 - Modern deep learning techniques (e.g. convnets, variants of SGD)
- Democratization of machine learning
 - High quality open-source software, such as Python's scikit-learn, TensorFlow, Torch, Theano

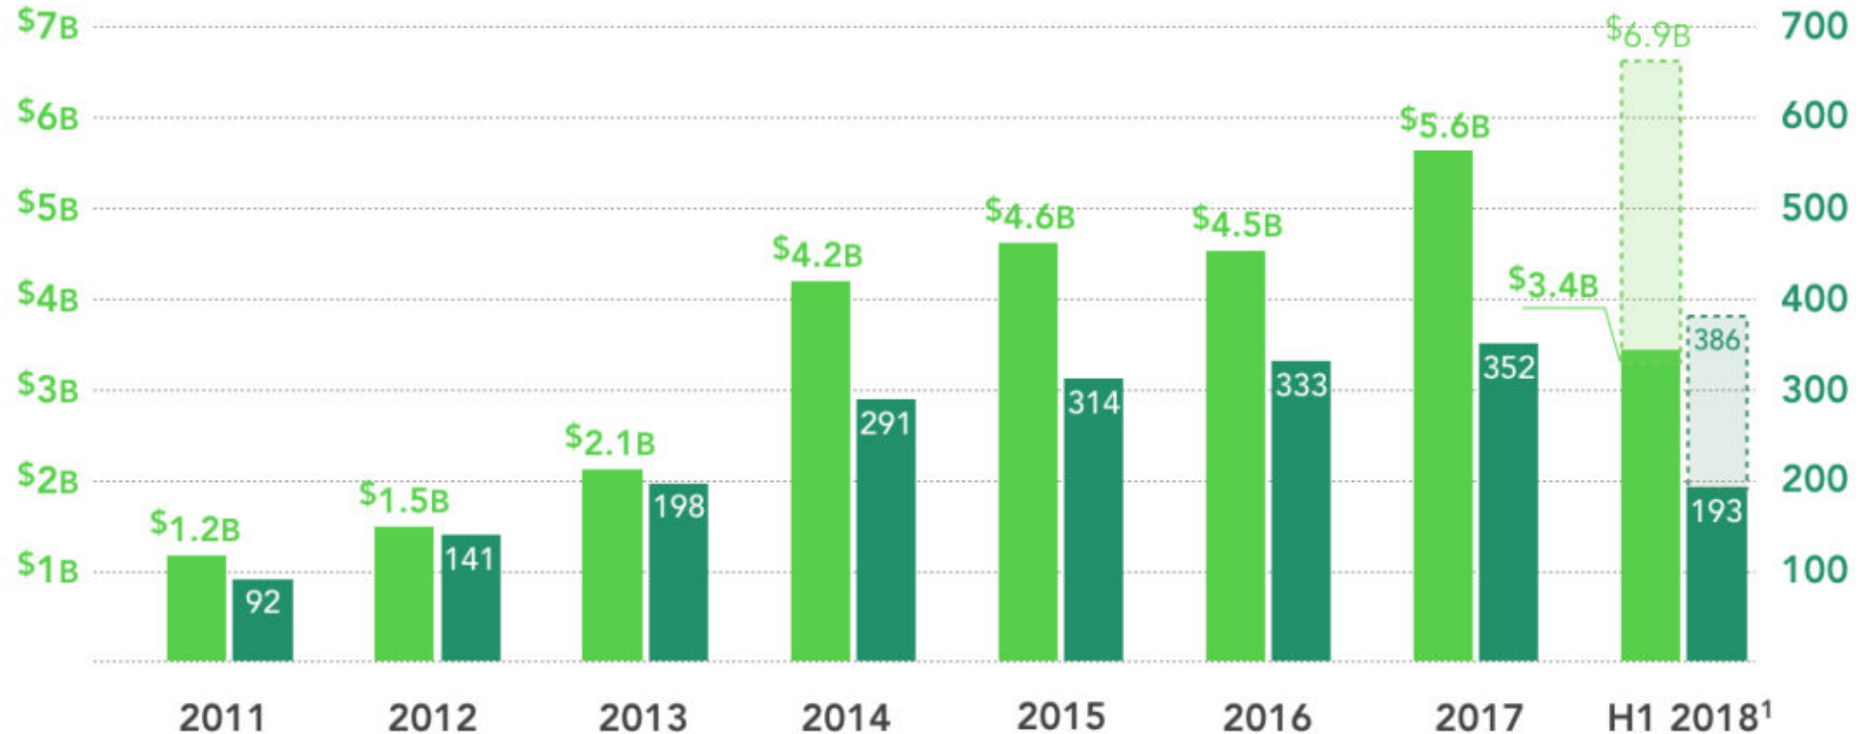
DIGITAL HEALTH FUNDING

2011-H1 2018



TOTAL VENTURE FUNDING

OF DEALS



AVERAGE DEAL SIZE



Source: Rock Health Funding Database

1: Shaded portion shows projections for entire year of 2018, assuming current funding pace continues.

Note: Only includes U.S. deals >\$2M; data through June 30, 2018

© Rock Health. All rights reserved. This content is excluded from our Creative Commons license.

For more information, see <https://ocw.mit.edu/help/faq-fair-use/>



106 STARTUPS TRANSFORMING HEALTHCARE WITH AI



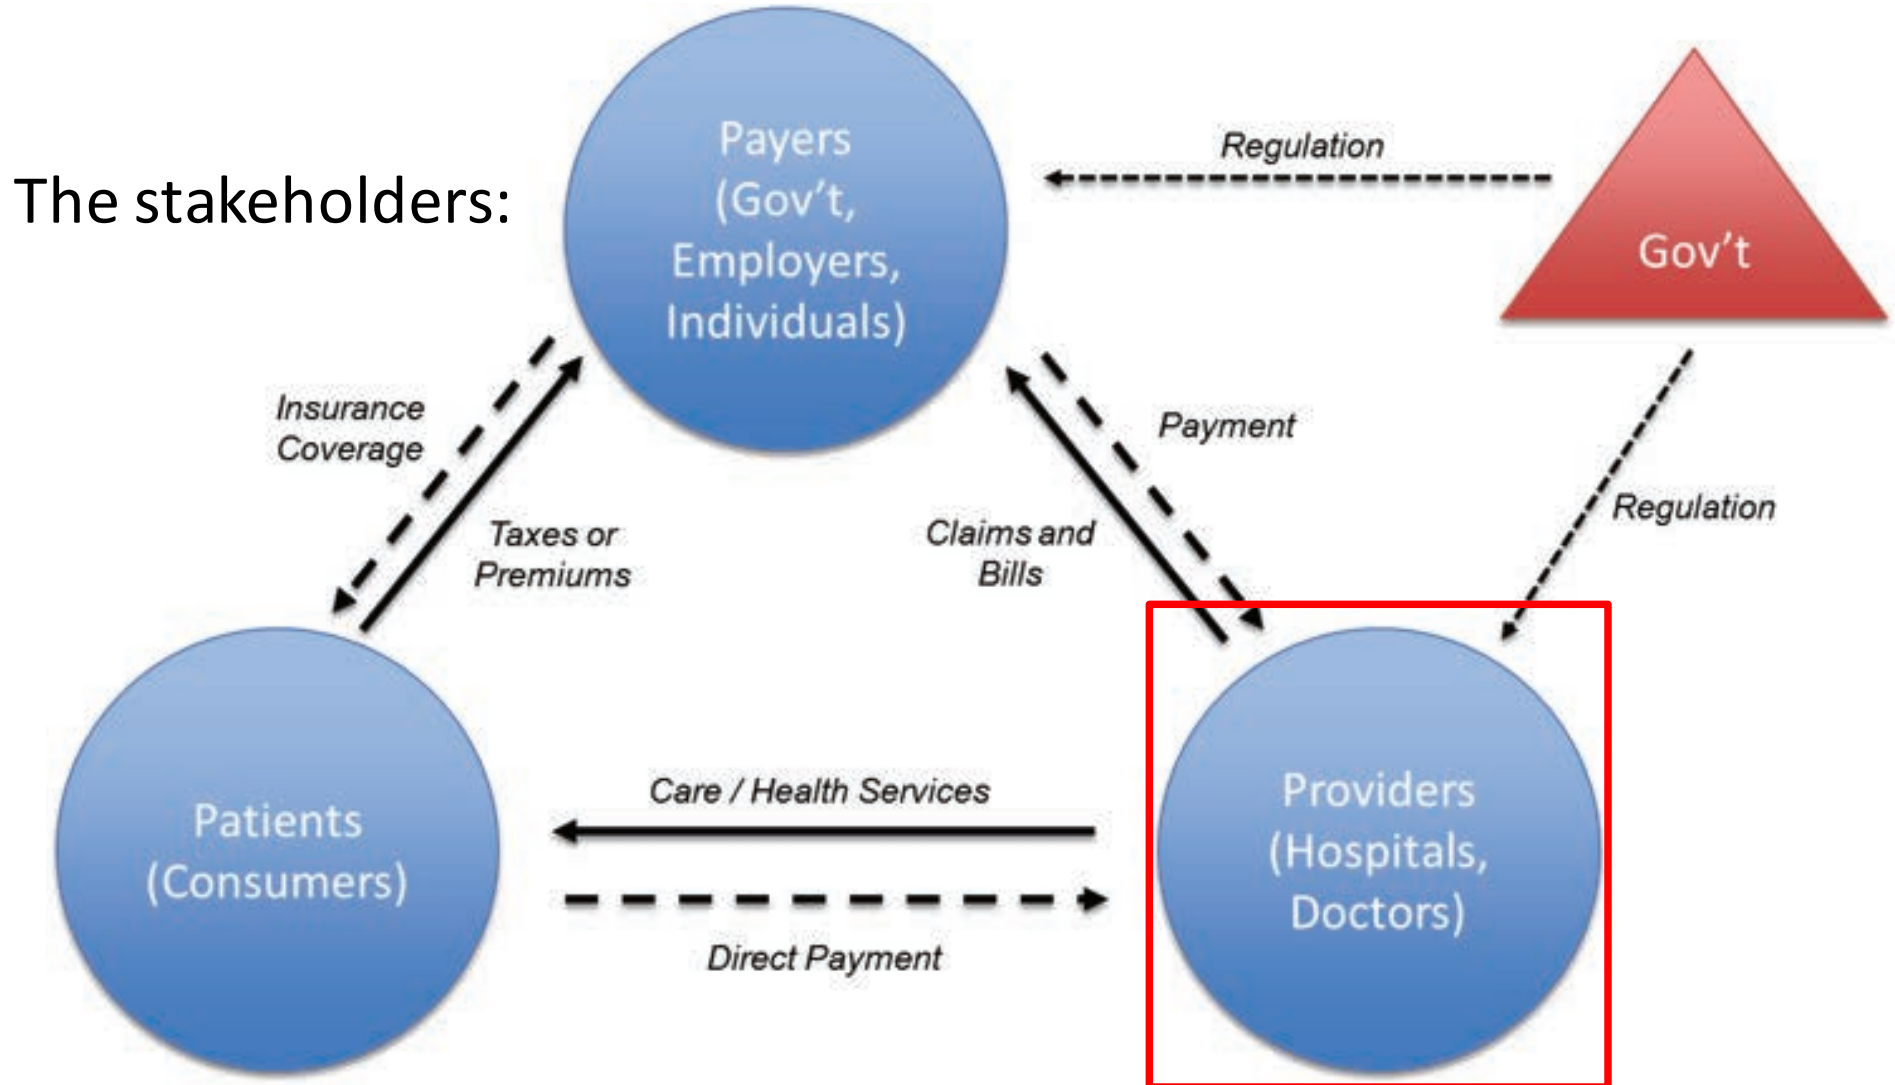
Industry interest in ML & healthcare

- Major acquisitions to get big data for ML:
 - Merge (\$1 billion purchase by IBM, 2015)
medical imaging
 - Truven Health Analytics (\$2.6 billion purchase by IBM, 2016)
health insurance claims
 - Flatiron Health (\$1.9 billion purchase by Roche, 2018)
electronic health records (oncology)

Outline for today's class

1. Brief history of AI and ML in healthcare
2. *Why now?*
3. **Examples of how ML will transform healthcare**
4. *What is unique* about ML in healthcare?
5. Overview of class syllabus

ML will transform every aspect of healthcare



© source unknown. All rights reserved. This content is excluded from our Creative Commons license. For more information, see <https://ocw.mit.edu/help/faq-fair-use/>

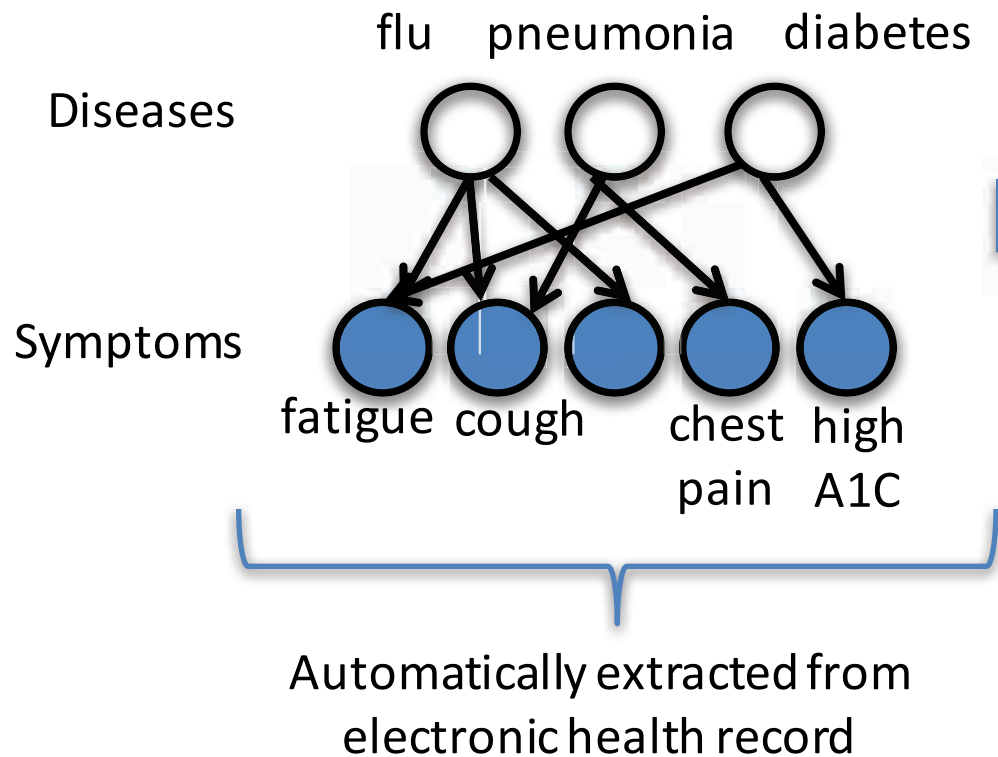


Emergency Department:

- Limited resources
- Time sensitive
- Critical decisions

What will the ER of the future be like?

Behind-the-scenes reasoning about the patient's conditions (current and future)



- Better triage
- Faster diagnosis
- Early detection of adverse events
- Prevent medical errors

What will the ER of the future be like?

Propagating best practices

The ED Dashboard decision support algorithms have determined that this patient may be eligible for the Atrius Cellulitis pathway. Please choose from the following options:

Enroll in pathway

Decline

You can include a comment for the reviewers: *Mandatory if Declining*

Below are links to the pathway and/or other supporting documents:

[Atrius Cellulitis Pathway](#)

What will the ER of the future be like?

Anticipating the clinicians' needs

- Psych Order Set

To be drawn immediately Add-on

Laboratory

CBC + Diff

+ Chem-7

+ Serum Tox

+ Urine Tox

Order

- Chest Pain Order Set

To be drawn immediately Add-on

Initial

Place IV (saline lock);
flush per protocol

Continuous Cardiac monitoring

Continuous Pulse oximetry

EKG (pick 1)

Indication: Chest Pain

Indication: Dyspnea

Laboratory

CBC + Diff

+ Chem-7

Troponin

Aspirin (pick 1)

Aspirin 324 mg PO chewed

Aspirin 243 mg PO chewed

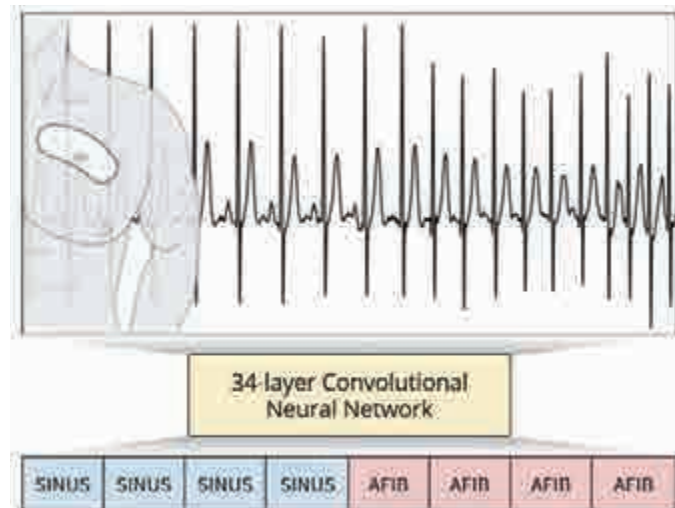
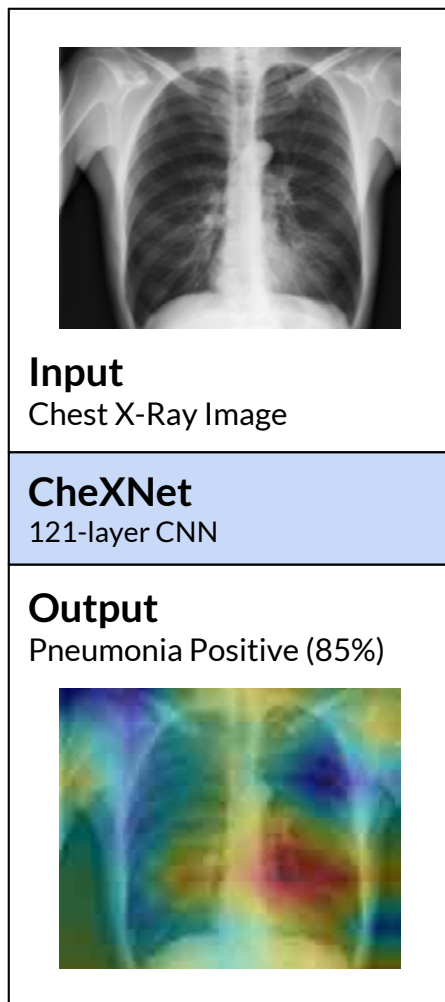
Aspirin taken before arrival

Imaging

XR Chest PA & Lateral

What will the ER of the future be like?

Reducing the need for specialist consults



Arrhythmia?

© Rajpurkar et al. All rights reserved. This content is excluded from our Creative Commons license. For more information, see <https://ocw.mit.edu/help/faq-fair-use/>

Figure sources: Rajpurkar et al., arXiv:1711.05225'17

28 Rajpurkar et al., arXiv:1707.01836, '17

What will the ER of the future be like?

Automated documentation and billing

KERMIT, F [69 / M]

Temp 99 HR 102 BP 150/70 RR 24 O2sat 99%

69 y/o M Patient with severe intermittent RUQ pain. Began soon after eating. Also is a heavy drinker.

Chief Complaints:

- RUQ abdominal pain
- Allergic reaction
- L Knee pain
- Rectal pain
- Right sided abdominal pain

Transfer
MCI

Triage note

Test

ED

Hornig, Steven

Logout

Options

- My Patients
- Overview
- Team 1
- Team 2
- Team 3
- Core+Red
- Farr-8
- Role Dash
- Steves Research
- Register
- Big Screen

KERMIT, F [69 / M]

Temp 99 HR 102 BP 150/70 RR 24 O2sat 99%

69 y/o M Patient with severe intermittent RUQ pain. Began soon after eating. Also is a heavy drinker.

Chief Complaints: a

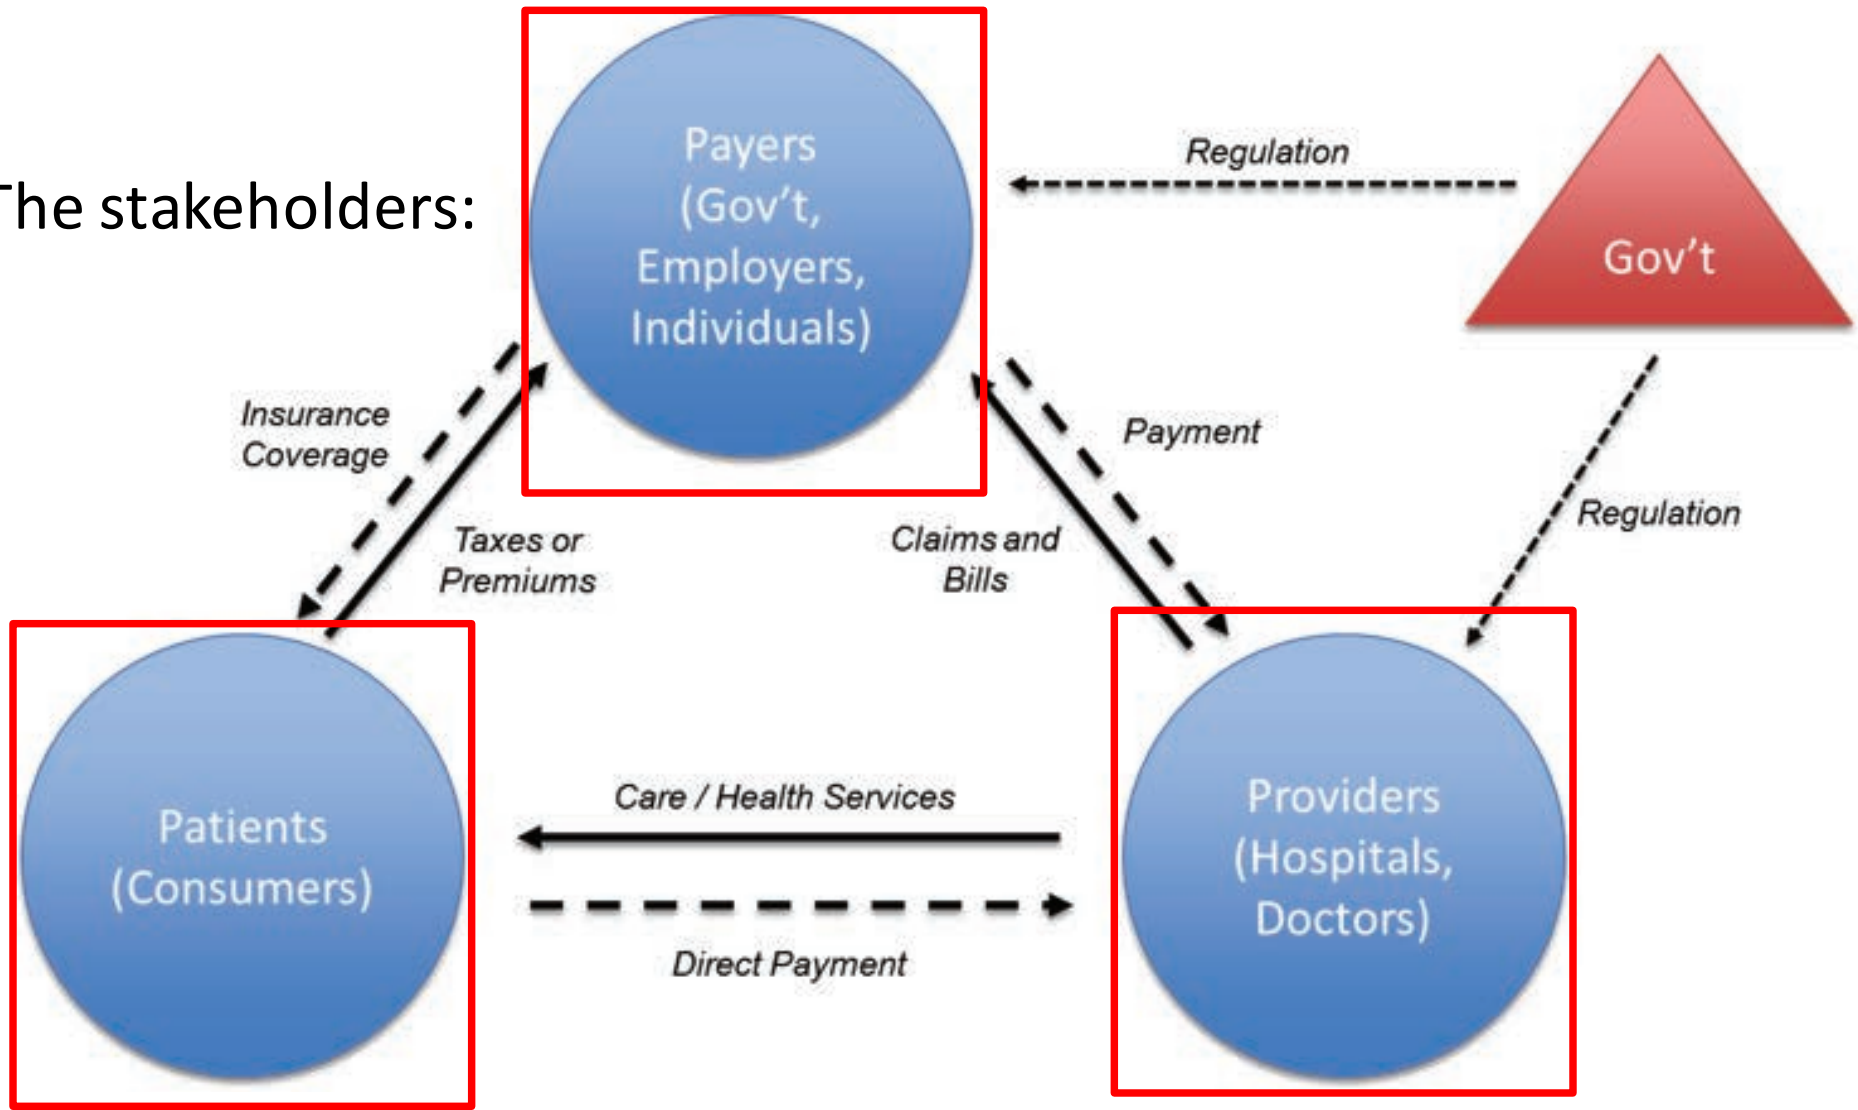
- RIGHT UPPER QUADRANT PAIN
- RUQ ABDOMINAL PAIN
- RUQ PAIN
- ALLERGIC REACTION
- L KNEE PAIN
- RECTAL PAIN
- RIGHT SIDED ABD PAIN
- RIGHT SIDED ABDOMINAL PAIN
- L WRIST PAIN
- RIGHT SIDED CHEST PAIN
- TESTICULAR PAIN
- KNEE PAIN
- ELBOW PAIN
- RIB PAIN
- L ELBOW PAIN
- HAND PAIN
- VAGINAL PAIN

Predicted chief complaints

Contextual auto-complete

ML will transform every aspect of healthcare

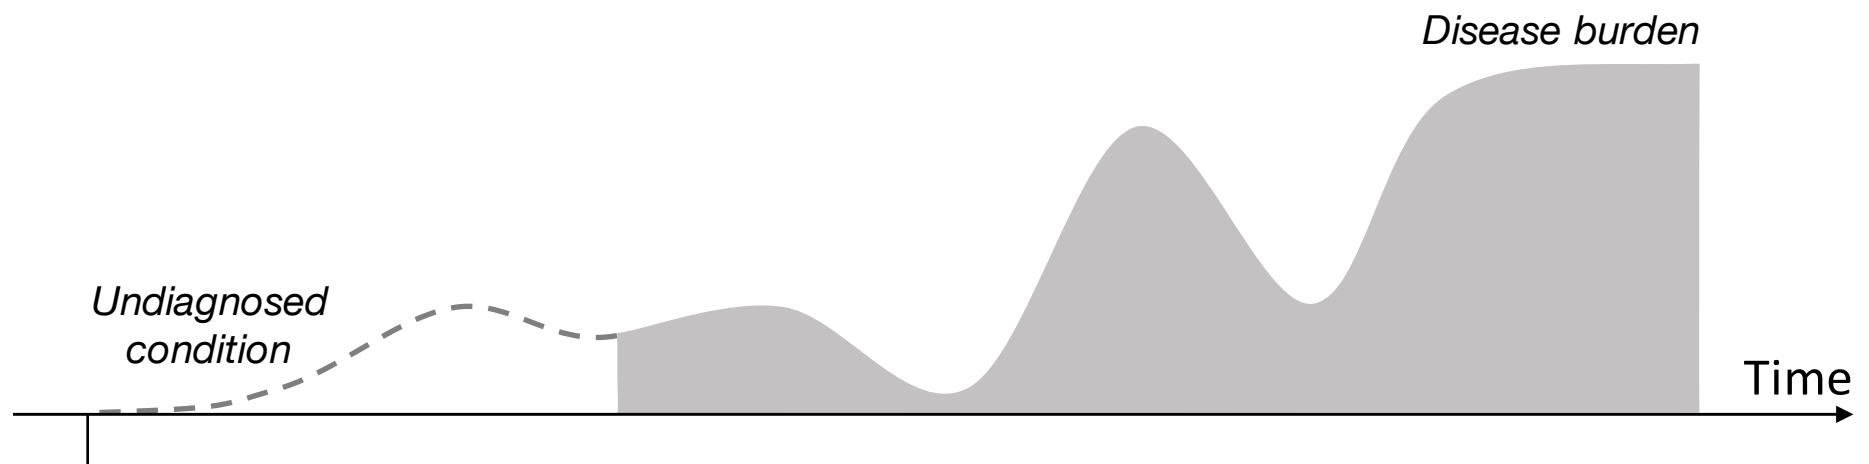
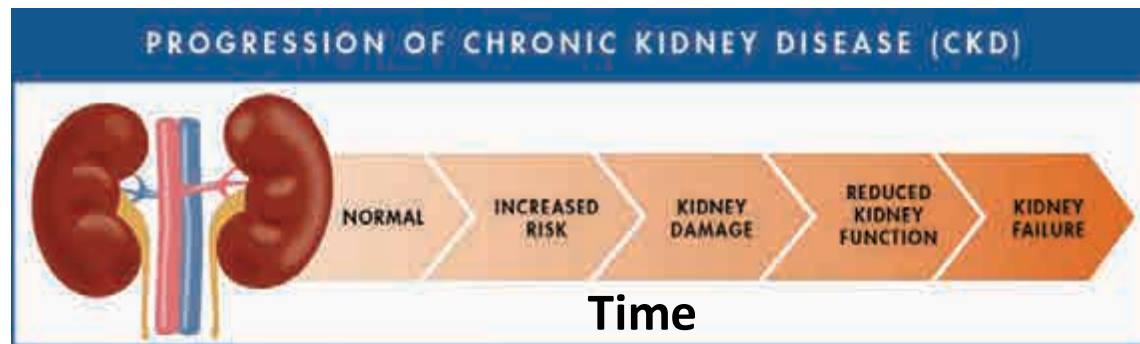
The stakeholders:



© source unknown. All rights reserved. This content is excluded from our Creative Commons license. For more information, see <https://ocw.mit.edu/help/faq-fair-use/>

What is the future of how we treat chronic disease?

- Predicting a patient's future disease progression



Courtesy of the CDC. Image is in the public domain.

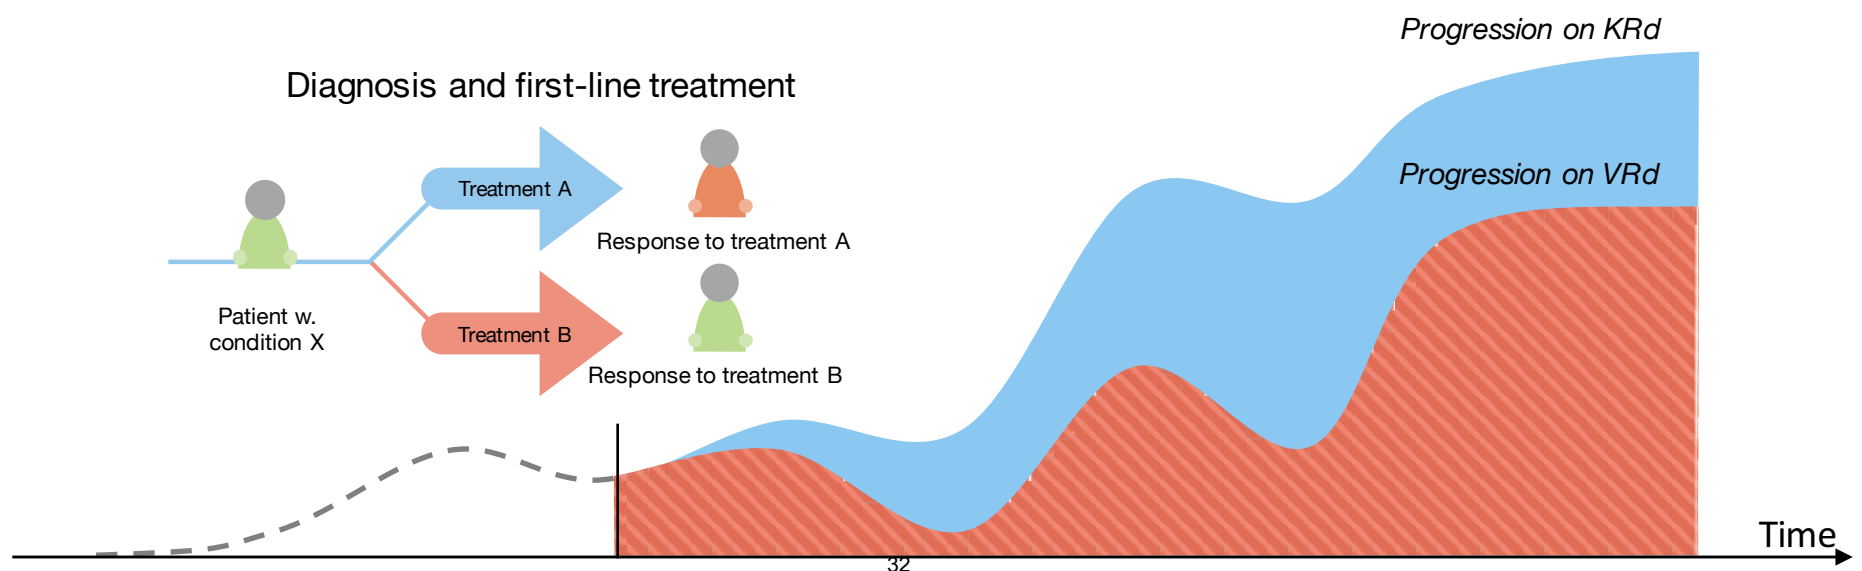
Figure credit: <https://www.cdc.gov/kidneydisease/prevention-risk.html>

What is the future of how we treat chronic disease?

- Predicting a patient's future disease progression
- Precision medicine

Choosing first line therapy in multiple myeloma

A) KRd: carfilzomib-lenalidomide-dexamethasone, **B) VRd:** bortezomib-lenalidomide-dexamethasone



What is the future of how we treat chronic disease?

- Early diagnosis, e.g. of diabetes, Alzheimer's, cancer
- Continuous monitoring and coaching, e.g. for the elderly, diabetes, psychiatric disease
- Discovery of new disease subtypes; design of new drugs; better targeted clinical trials

Outline for today's class

1. Brief history of AI and ML in healthcare
2. *Why now?*
3. Examples of how ML will transform healthcare
4. **What is *unique* about ML in healthcare?**
5. Overview of class syllabus

What makes healthcare different?

- Life or death decisions
 - Need **robust** algorithms
 - Checks and balances built into ML deployment
 - (Also arises in other applications of AI such as autonomous driving)
 - Need **fair** and **accountable** algorithms
- Many questions are about unsupervised learning
 - Discovering disease subtypes, or answering question such as “characterize the types of people that are highly likely to be readmitted to the hospital”?
- Many of the questions we want to answer are *causal*
 - Naïve use of supervised machine learning is insufficient

What makes healthcare different?

- Very little labeled data
 - Motivates semi-supervised learning algorithms
- Sometimes small numbers of samples (e.g., a rare disease)
 - Learn as much as possible from other data (e.g. healthy patients)
 - Model the problem carefully
- Lots of missing data, varying time intervals, censored labels

What makes healthcare different?

- Difficulty of de-identifying data
 - Need for data sharing agreements and sensitivity
- Difficulty of deploying ML
 - Commercial electronic health record software is difficult to modify
 - Data is often in silos; everyone recognizes need for interoperability, but slow progress
 - Careful testing and iteration is needed

Goals for the semester

- Intuition for working with healthcare data
- How to set up as machine learning problems
- Understand which learning algorithms are likely to be useful and when
- Appreciate subtleties in safely & robustly applying ML in healthcare
- Set the research agenda for the next decade

6.S897/HST.956 vs 6.874

- Our class will focus on **clinical data** and its use to improve health care
- For reasons of time & scope, we will not go deep into applications in the life sciences
 - For this, we recommend taking **6.874**
Computational Systems Biology: Deep Learning in the Life Sciences

MIT OpenCourseWare

<https://ocw.mit.edu>

6.S897 / HST.956 Machine Learning for Healthcare

Spring 2019

For information about citing these materials or our Terms of Use, visit: <https://ocw.mit.edu/terms>