

# Machine Learning for Healthcare

## 6.871, HST.956

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

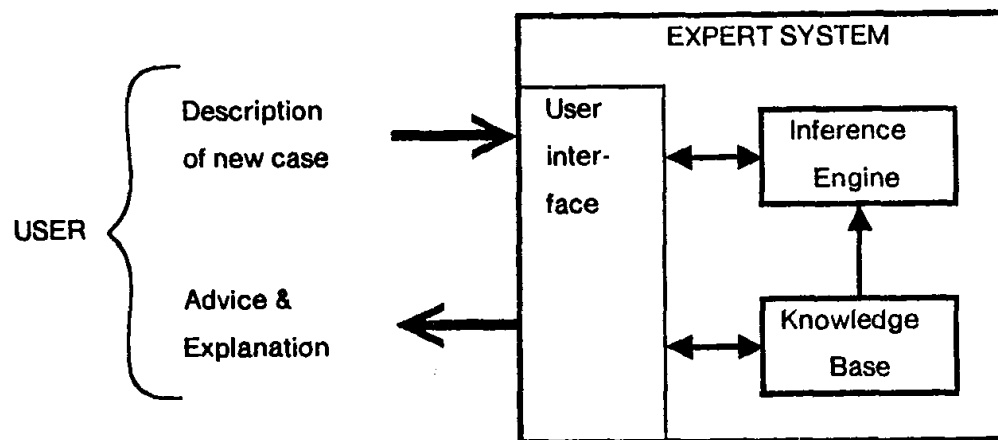


FIGURE 1-1 Major parts of an expert system. Arrows indicate information flow.

## 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.

# 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

[Miller et al., '86, Shwe et al., '91]

# 1980's: automating medical discovery

RX PROJECT: AUTOMATED KNOWLEDGE ACQUISITION

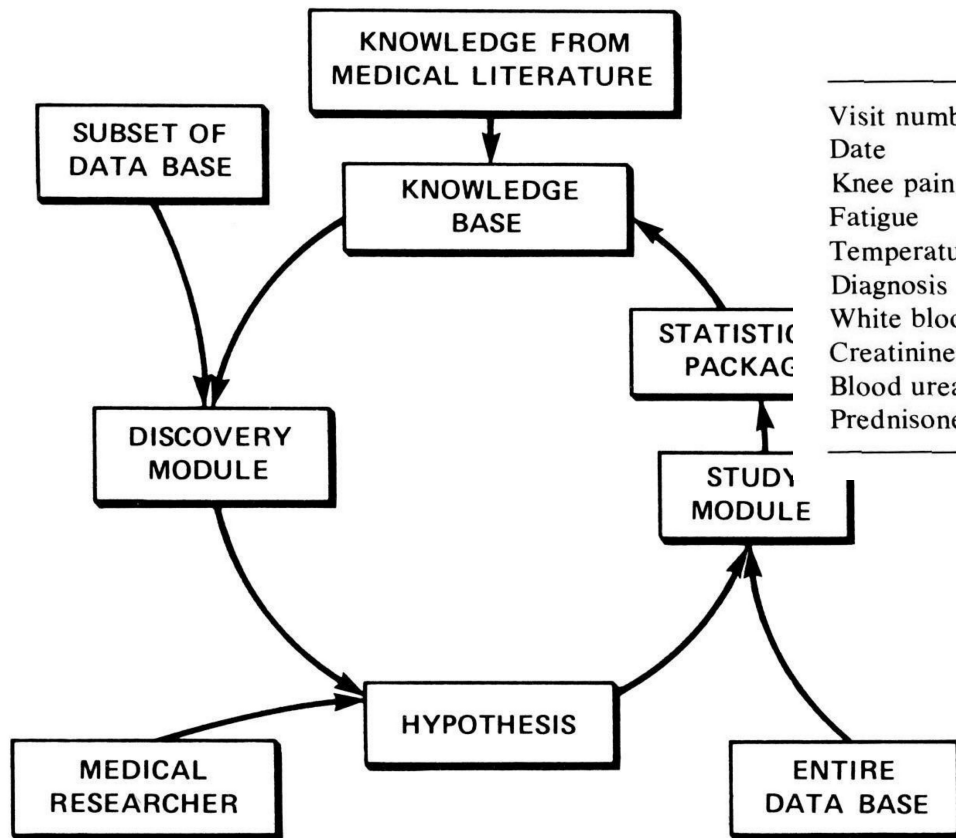


TABLE 1

HYPOTHETICAL TIME-ORIENTED RECORD FOR ONE PATIENT

Visit number	1	2	3
Date	January 17, 79	June 23, 79	July 1, 79
Knee pain	Severe	Mild	Mild
Fatigue	Moderate	—	Moderate
Temperature	38.5	37.5	36.9
Diagnosis	Systemic lupus		
White blood count	3500	4700	4300
Creatinine clearance	45	—	65
Blood urea nitrogen	36	33	—
Prednisone	30	25	20

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

**Table 1** • 25 Neural Network Studies in Medical Decision Making\*

Subject	No. of Examples		P†	Network	D‡	Accuracy§	
	Training	Test				Neural	Other
Breast cancer <sup>4</sup>	57	20	60	9-15-2	0.6	80	75
Vasculitis <sup>2</sup>	404	403	73	8-5-1	8.0	94	—
Myocardial infarction <sup>6</sup>	351	331	89	20-10-10-1	1.1	97	<del>84</del>
Myocardial infarction <sup>8</sup>	356	350	87	20-10-10-1	1.1	97	<del>84</del>
Low back pain <sup>11</sup>	100	100	25	50-48-2	0.2	90	90
Cancer outcome <sup>13</sup>	5,169	3,102	—	54-40-1	1.4	0.779	0.776
Psychiatric length of stay <sup>17</sup>	957	106	73	48-400-4	0.2	74	76
Intensive care outcome <sup>23</sup>	284	138	91	27-18-1	0.5	0.82	0.82
Skin tumor <sup>21</sup>	150	100	80	18	—	80	90
Evoked potentials <sup>35</sup>	100	67	52	14-4-3	3.8	77	77
Head injury <sup>47</sup>	500	500	50	6-3-3	20	66	77
Psychiatric outcome <sup>54</sup>	289	92	60	41-10-1	0.7	79	—
Tumor classification <sup>55</sup>	53	6	38	8-9-3	1.4	99	<del>88</del>
Dementia <sup>57</sup>	75	18	19	80-10-7-7	0.6	61	—
Pulmonary embolism <sup>59</sup>	607	606	69	50-4-1	2.9	0.82	0.83
Heart disease <sup>62</sup>	460	230	54	35-16-8-2	3	83	<del>84</del>
Thyroid function <sup>62</sup>	3,600	1,800	93	21-16-8-3	22	98	93
Breast cancer <sup>62</sup>	350	175	66	9-4-4-2	10	97	96
Diabetes <sup>62</sup>	384	192	65	8-4-4-2	12	77	75
Myocardial infarction <sup>63</sup>	2,856	1,429	56	291-1	9.8	85	—
Hepatitis <sup>65</sup>	39	42	38	4-4-3	3.3	74	79
Psychiatric admission <sup>76</sup>	319	339	85	53-1-1	6.0	91	—
Cardiac length of stay <sup>83</sup>	713	696	73	15-12-1	3.5	0.70	—
Anti-cancer agents <sup>89</sup>	127	141	25	60-7-6	1.5	91	86
Ovarian cancer <sup>91</sup>	75	98	—	6-6-2	2.6	84	81
MEDIAN VALUE	350	175	71	20	2.8		

\*For reference citations, see the reference list

†P = prior probability of most prevalent category.

‡D = ratio of training examples to weights per output.


§A single integer in the accuracy column denotes percentage overall classification rate and a single real number between 0 and 1 indicates the AUROCC value. Neural = accuracy of neural net, Other = accuracy of best other method.



# Outline for today's class

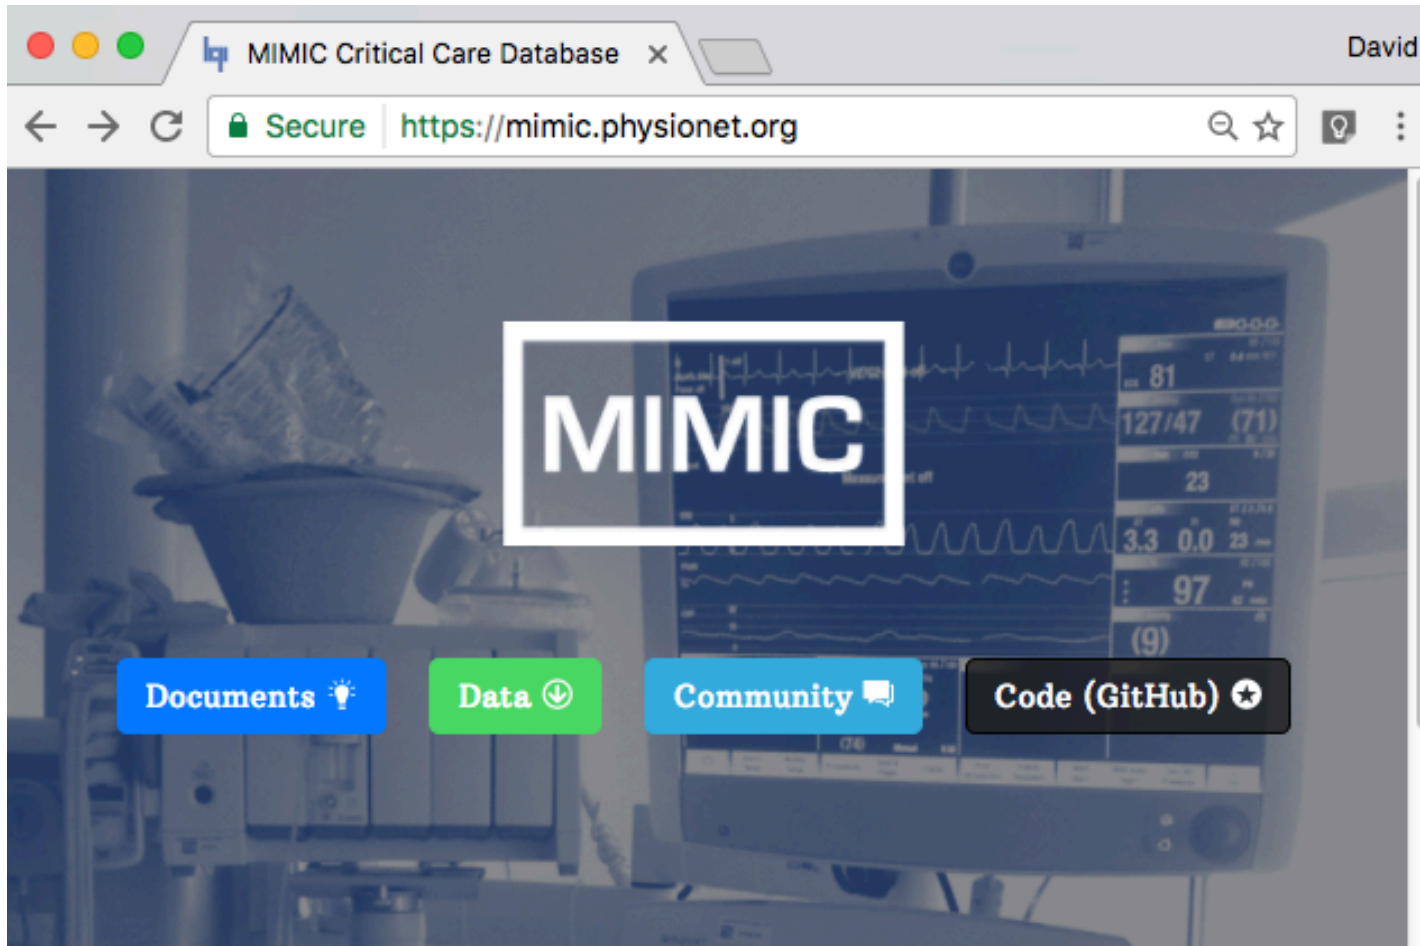
1. Brief history of AI and ML in healthcare
2. **Why now?**
3. Examples of how ML will transform healthcare
4. What is *unique* about ML in healthcare?
5. Overview of class syllabus

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



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

# Large datasets



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](https://doi.org/10.1038/sdata.2016.35). Available from: <http://www.nature.com/articles/sdata201635>



Laboratory for  
Computational  
Physiology

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

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

# Large datasets

The screenshot shows the Truven Health Analytics website. The browser address bar displays the URL: [truvenhealth.com/markets/life-sciences/products/data-tools/marketscan-databases](https://truvenhealth.com/markets/life-sciences/products/data-tools/marketscan-databases). The navigation menu includes links for MEDIA ROOM, SUPPORT, and CAREER. The Truven Health Analytics logo is visible, along with the text "an IBM Company". The main navigation bar contains links for SOLUTIONS, EVENTS, KNOWLEDGE, and AB. The page title is "Life Sciences" and the breadcrumb trail is "Home » Life Sciences » Data & Tools » MarketScan Databases". The main content area features a large heading "Putting Research Data Into Your Hands with the MarketScan Databases" and a sub-heading "The Family of MarketScan® Research Databases is the largest of its kind in the industry, with data on nearly 230 million unique patients since 1995." A sidebar on the left lists various data tools and services, with "Data & Tools" selected. Other visible elements include a "PULSE" button and a "Heartbeat Profiler" link.

← → ↻ ⓘ [truvenhealth.com/markets/life-sciences/products/data-tools/marketscan-databases](https://truvenhealth.com/markets/life-sciences/products/data-tools/marketscan-databases)

MEDIA ROOM | SUPPORT | CAREER

**TRUVEN**  
HEALTH ANALYTICS™  
an IBM Company

SOLUTIONS | EVENTS | KNOWLEDGE | AB

**Life Sciences** [Home](#) » [Life Sciences](#) » [Data & Tools](#) » [MarketScan Databases](#)

Market Knowledge

Real World Evidence

Stakeholder Management

**Data & Tools**

[MarketScan Databases](#)

Treatment Pathways

Inpatient/Outpatient View

PULSE

Heartbeat Profiler

## Putting Research Data Into Your Hands with the MarketScan Databases

The Family of MarketScan® Research Databases is the largest of its kind in the industry, with data on nearly 230 million unique patients since 1995.

Mar  
Bibliog

W

“Data on nearly 230 million unique patients since 1995”

\$\$\$

# 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



# 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://en.wikipedia.org/wiki/List_of_ICD-9_codes)]



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

# Standardization

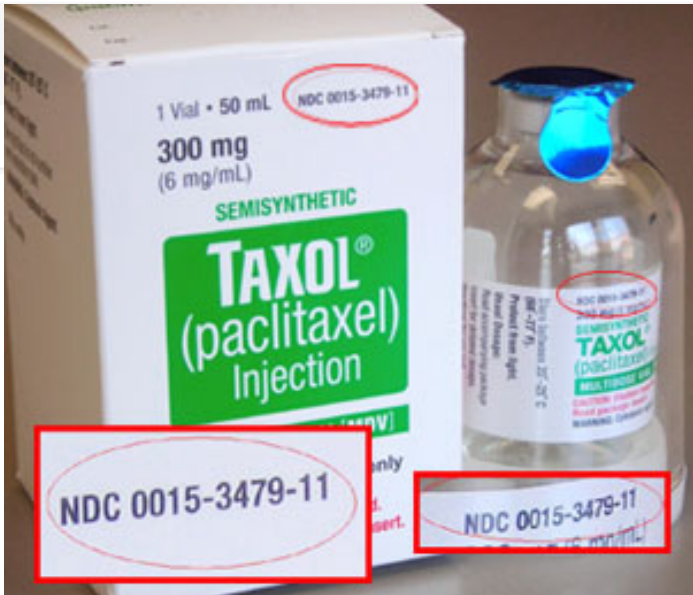
- Diagnosis codes: ICD-9 and ICD-10 (International Classification of Diseases)
- Laboratory tests: LOINC codes
- Pharmacy: National Drug Codes (NDCs)
- Unified Medical Language System (UMLS): millions of medical concepts

**LOINC**<sup>®</sup> From Regenstrief

glucose

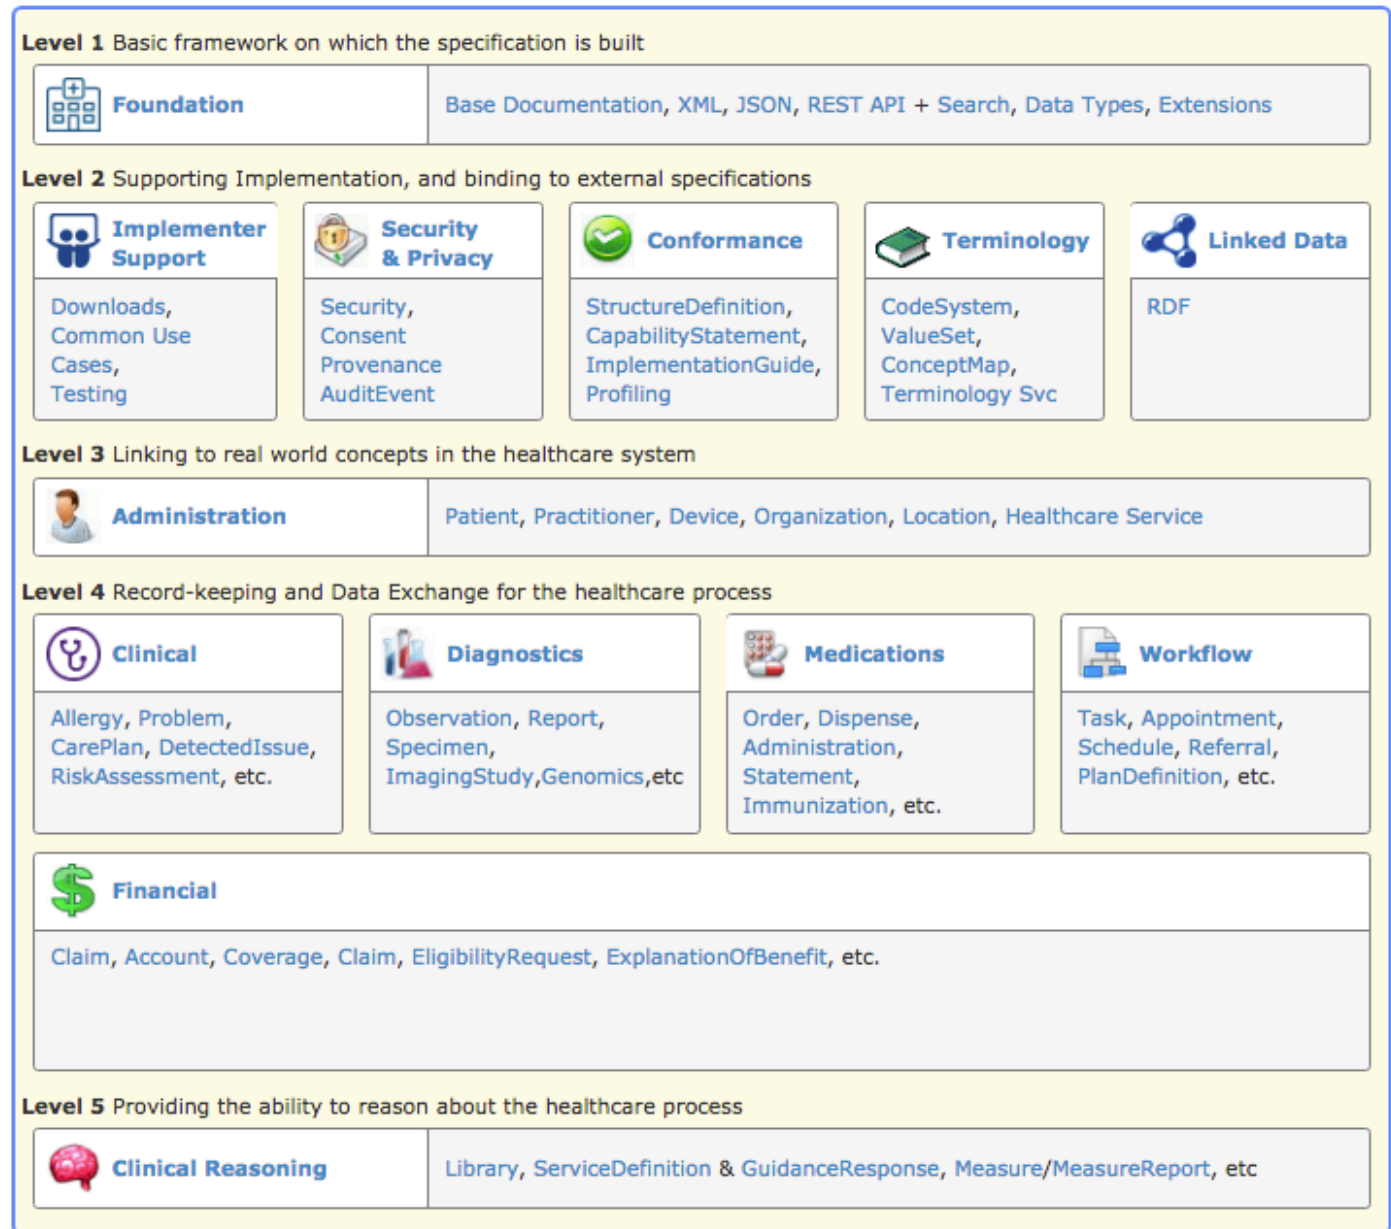
1 / 5

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>	





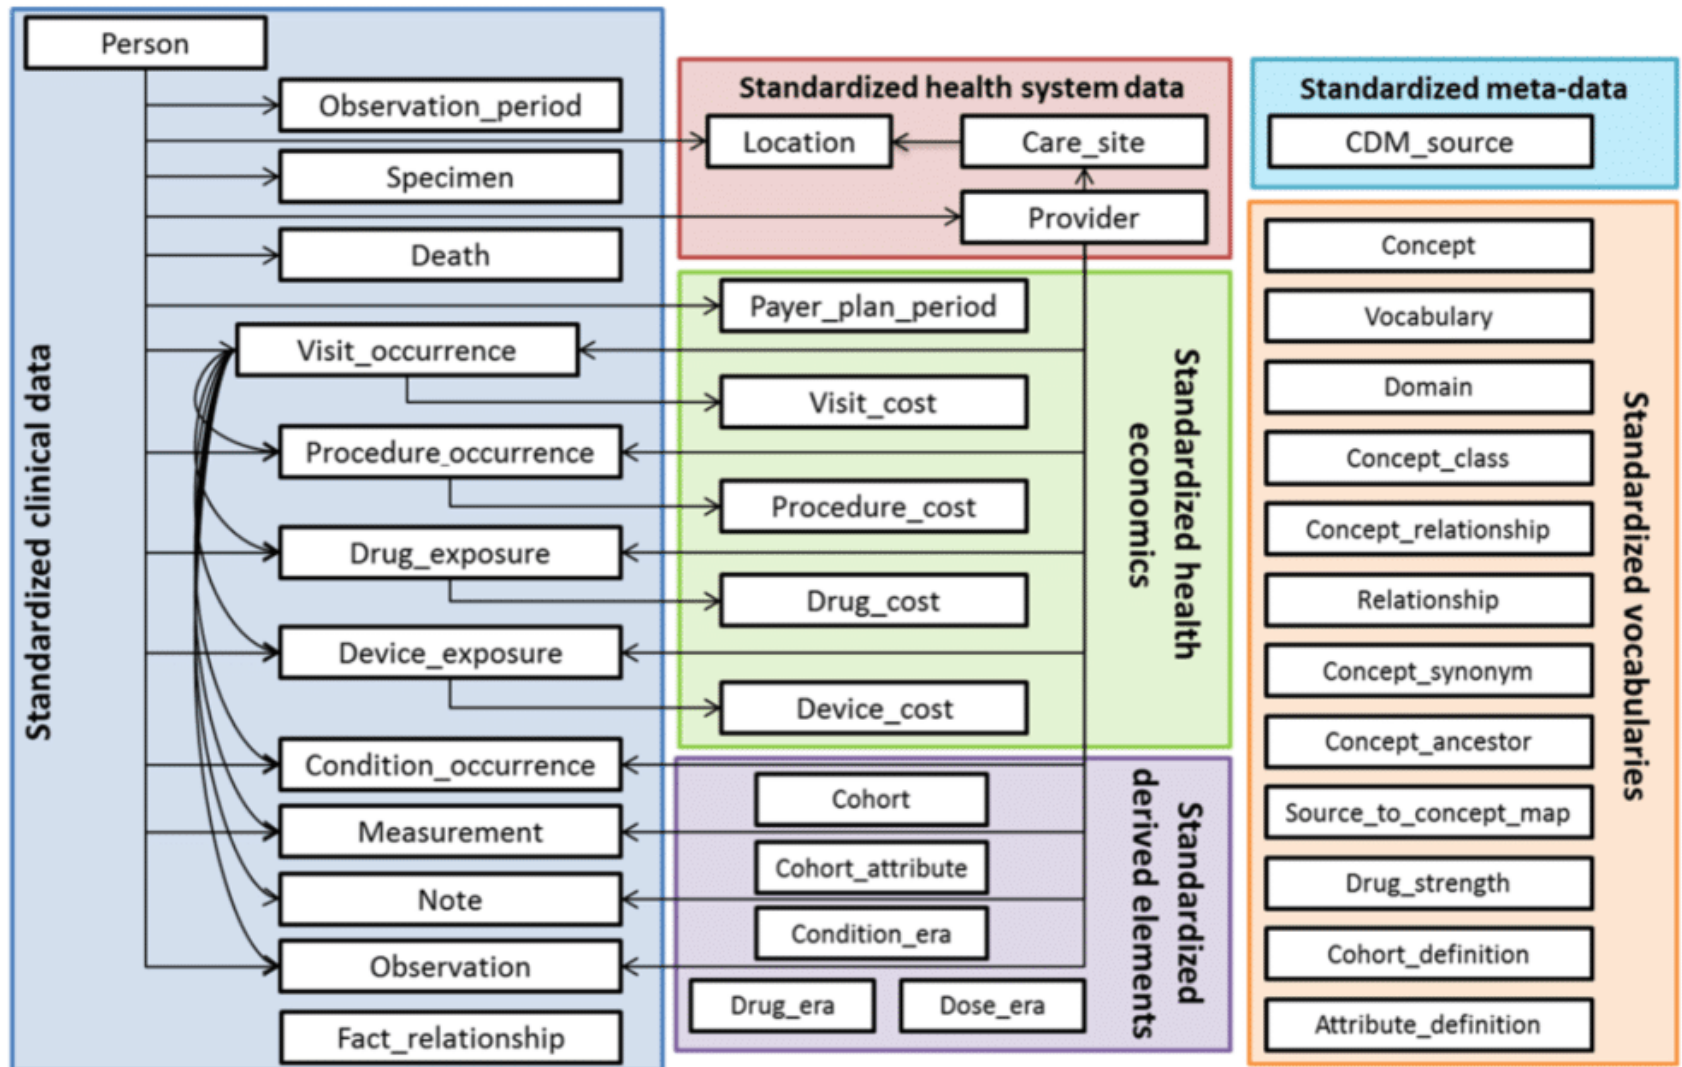
# Standardization



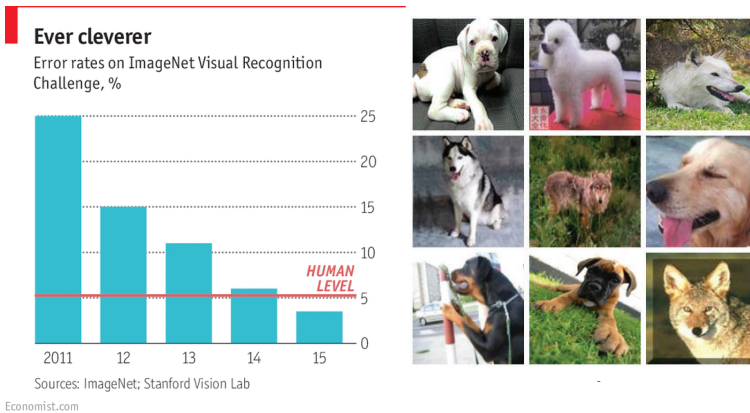
# Standardization



OMOP  
Common  
Data  
Model v5.0



# Breakthroughs in machine learning



Object recognition  
(deep neural networks)



AlphaGo  
(reinforcement learning)

**EXAMPLE**  
The 2008 Summer Olympics torch relay was run from March 24 until August 8, 2008, prior to the 2008 Summer Olympics, with the theme of "one world, one dream". Plans for the relay were announced on April 26, 2007, in Beijing, China. The relay, also called by the organizers as the "Journey of Harmony", lasted 129 days and carried the torch 137,000 km (85,000 mi) – the longest distance of any Olympic torch relay since the tradition was started ahead of the 1936 Summer Olympics.

After being lit at the birthplace of the Olympic Games in Olympia, Greece on March 26, the torch traveled to the Panathinaiko Stadium in Athens, and then to Beijing, arriving on March 31. From Beijing, the torch was following a route passing through six continents. The torch has visited cities along the Silk Road, symbolizing ancient links between China and the rest of the world. The relay also included an ascent with the flame to the top of Mount Everest on the border of Nepal and Tibet, China from the Chinese side, which was closed specially for the event.

Q. What was the theme?  
A: "one world, one dream".

Q. What was the length of the race?  
A: 137,000 km

Q. Was it larger than previous ones?  
A: No

Text comprehension  
(language models)

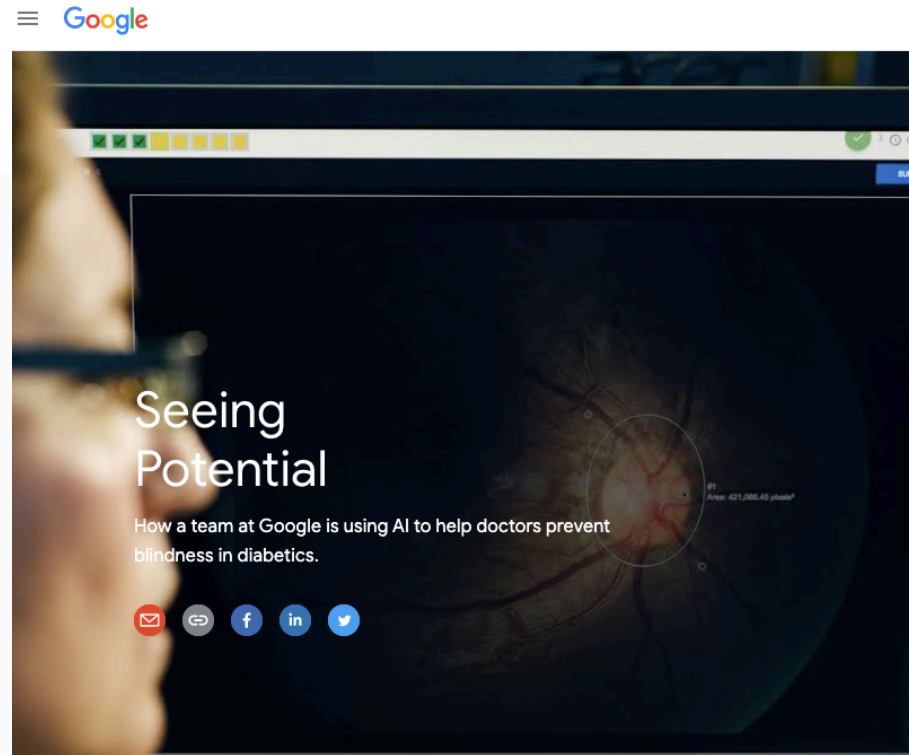
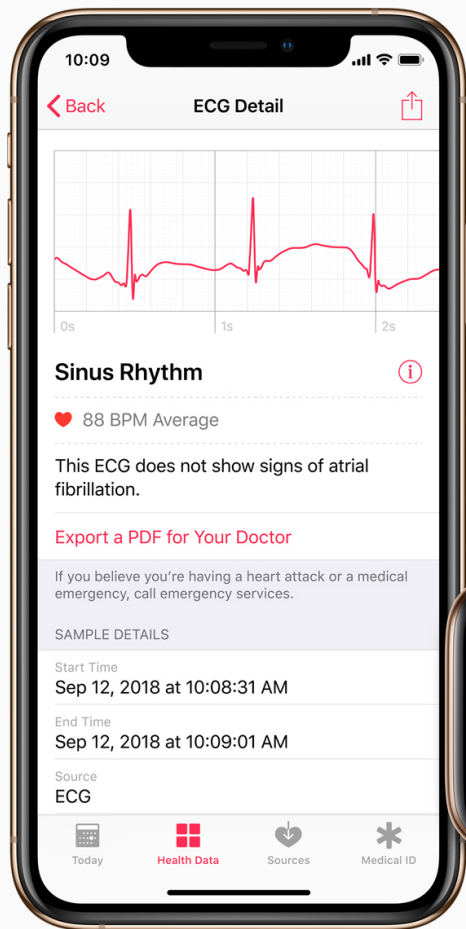


Figure 5: Linear interpolation in latent space between real images.

Generating realistic data  
(GANs, VAEs)

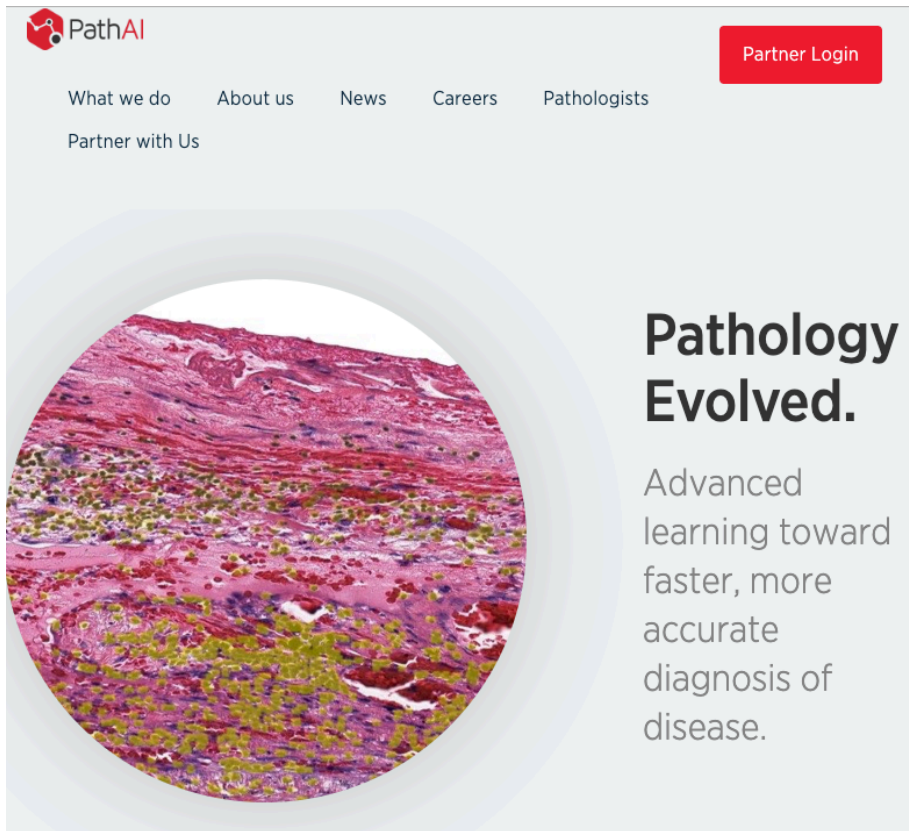


# Tech industry interest in health care



An advertisement for Amazon Comprehend Medical. The top shows the Amazon logo and search and menu icons. Below the logo is the text "Amazon Comprehend" with a dropdown arrow. The main heading is "Amazon Comprehend Medical". Below the heading is the text "Extract information from unstructured medical text accurately and quickly" and "No machine learning experience required". A yellow button contains the text "Get started with Amazon Comprehend Medical". At the bottom, there is a "TECH TALK" section with the heading "AI-Powered Health Data Masking" and the text "Learn how to use a pre-built solution from AWS to identify and mask health data in images or text."

# Tech industry interest in health care

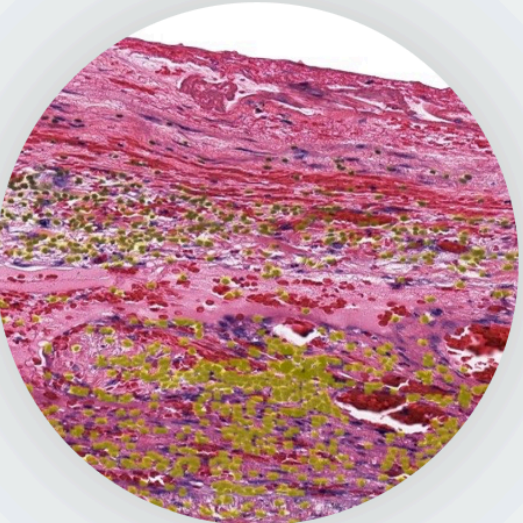


PathAI

Partner Login

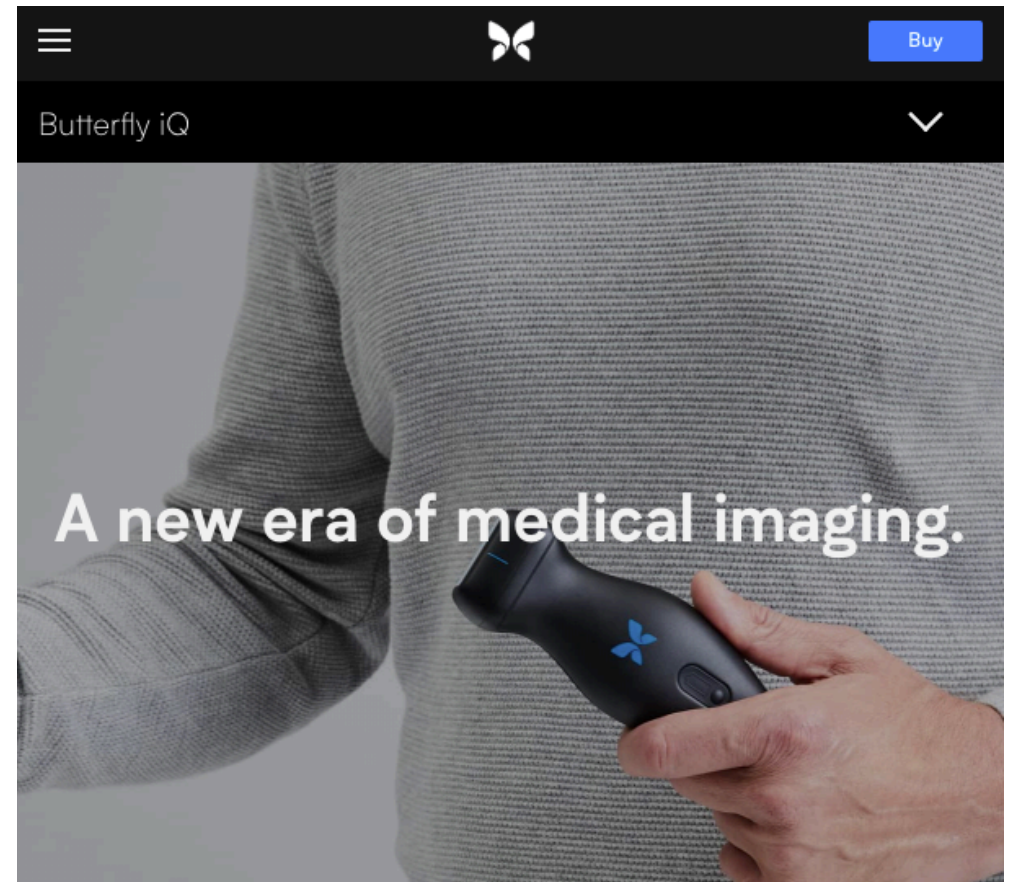
What we do About us News Careers Pathologists

Partner with Us




## Pathology Evolved.

Advanced learning toward faster, more accurate diagnosis of disease.



Butterfly iQ

Buy



## A new era of medical imaging.



# 106 STARTUPS TRANSFORMING HEALTHCARE WITH AI



istock.com/hilch

# Tech/pharma interest in health care

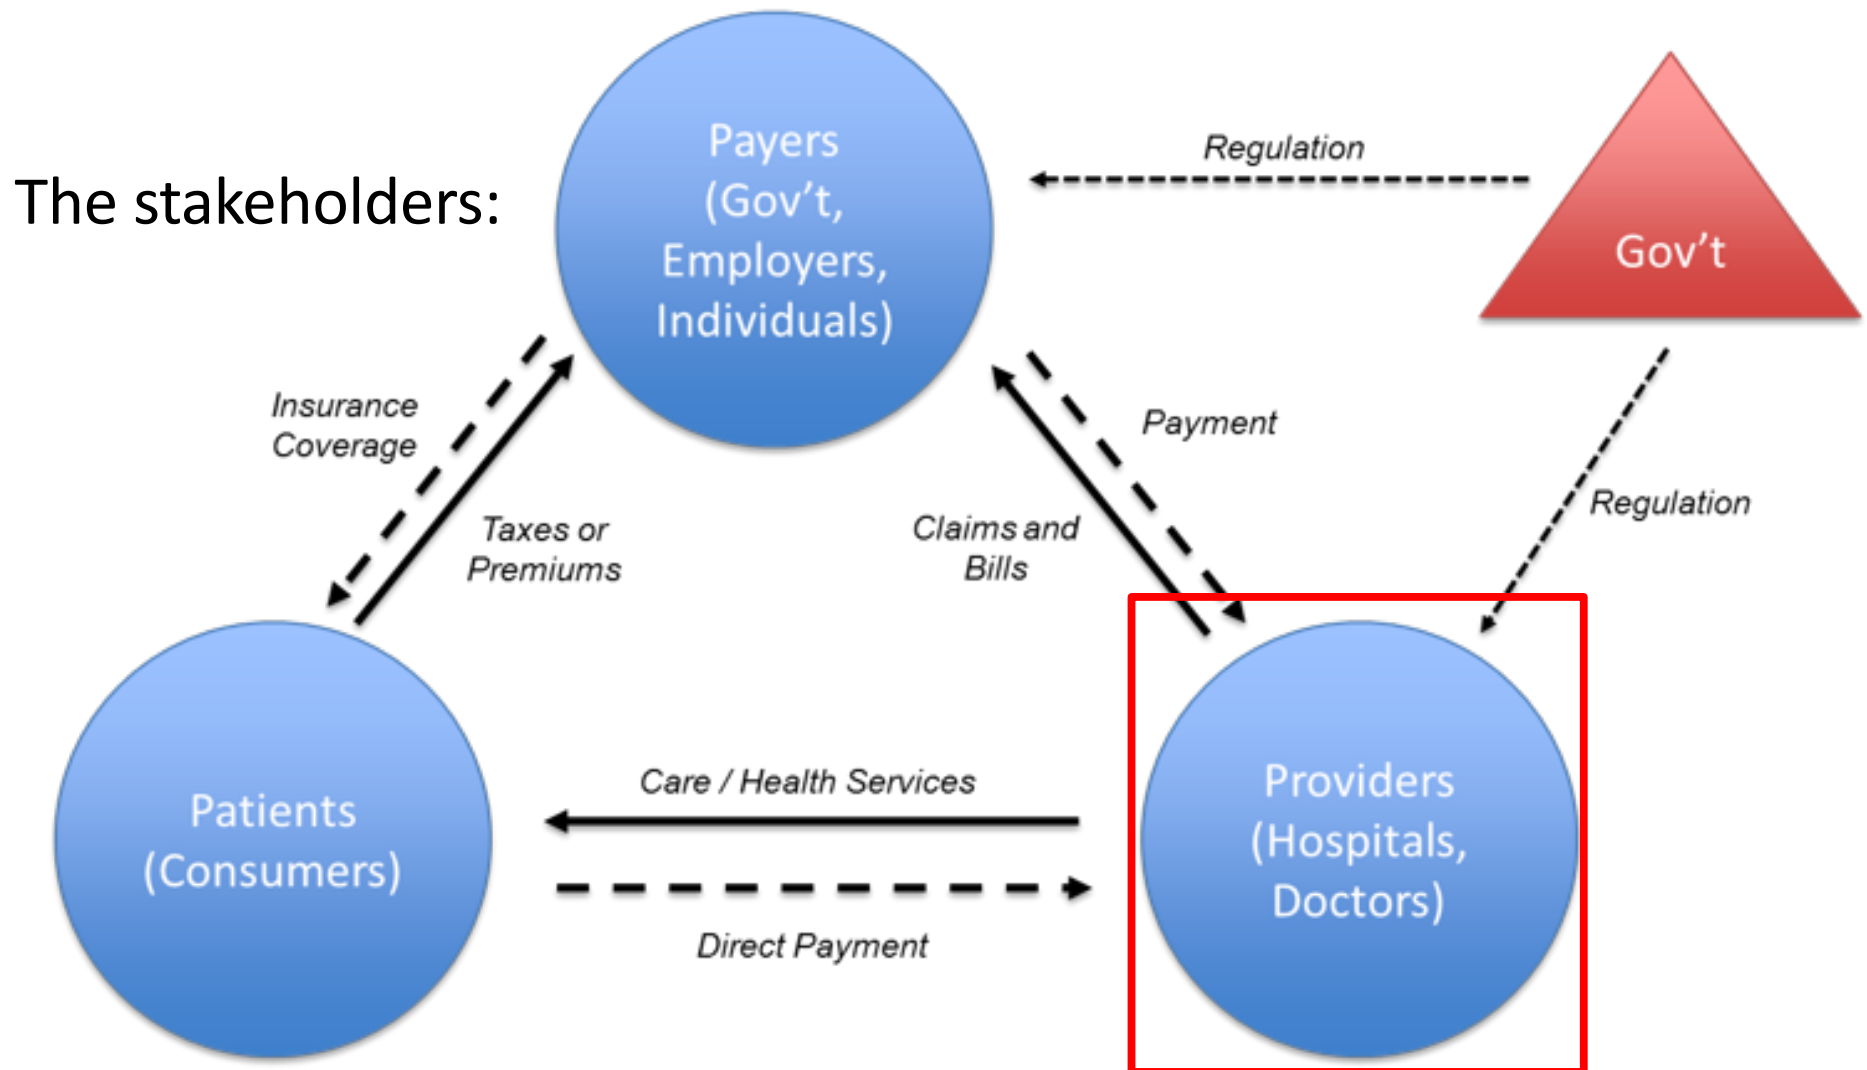
- Major acquisitions to get big data for ML:
  - Merge (\$1 billion purchase by IBM, 2015)  
***medical imaging***
  - Truven Health Analytics (\$2.6 billion purchase by IBM, 2016)  
***health insurance claims***
  - Flatiron Health (\$1.9 billion purchase by Roche, 2018)  
***electronic health records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

<http://www.mahesh-vc.com/blog/understanding-whos-paying-for-what-in-the-healthcare-industry>




## 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


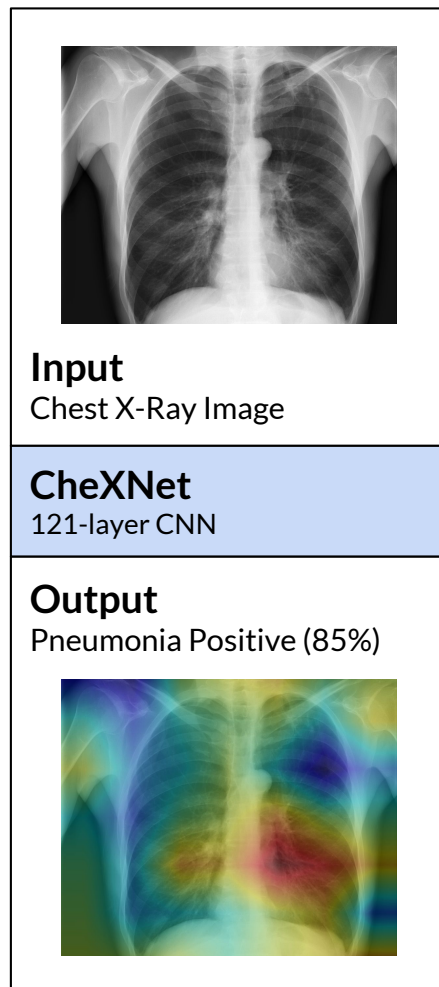
Aspirin taken before arrival

**Imaging**

XR Chest PA & Lateral

# What will the ER of the future be like?

Reducing the need for specialist consults



Arrhythmia?

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

Enter Cancel

**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

Enter Canc

**Triage note**

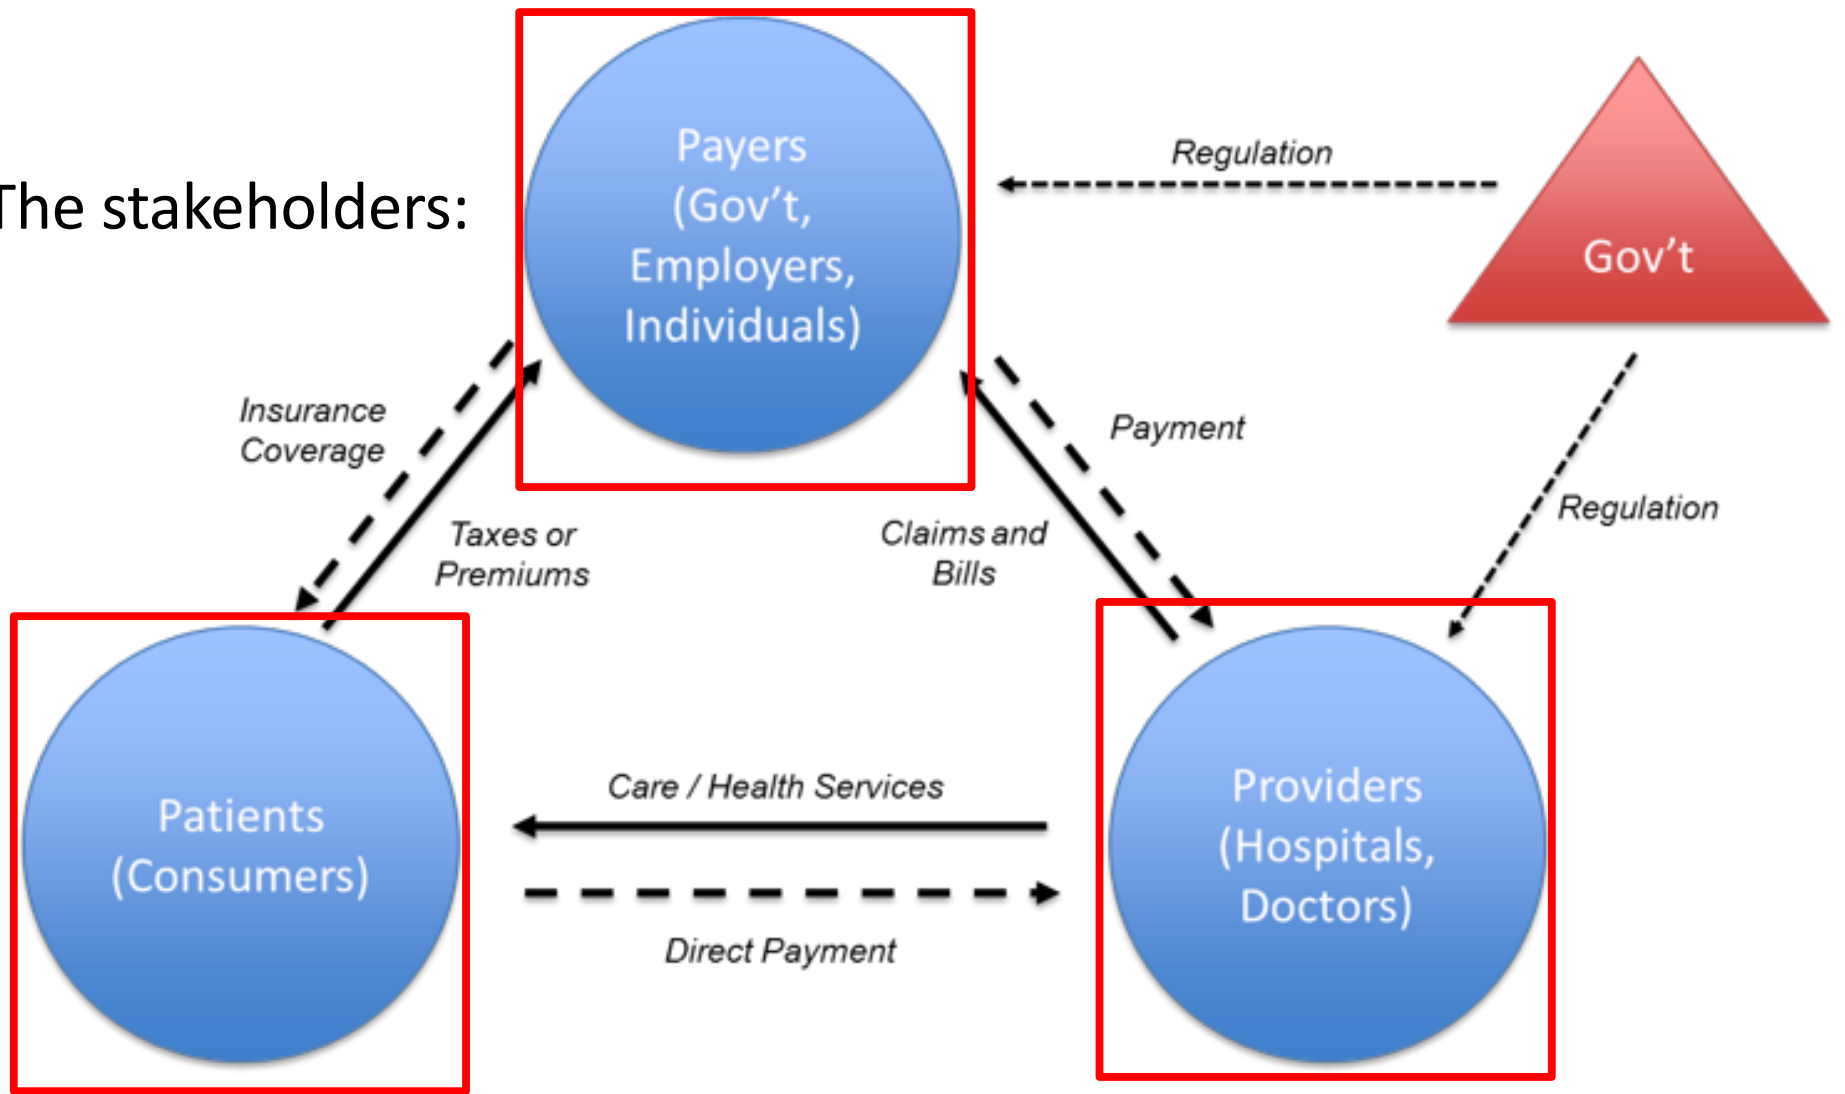
**Predicted chief complaints**

**Contextual auto-complete**



# ML will transform every aspect of healthcare

The stakeholders:


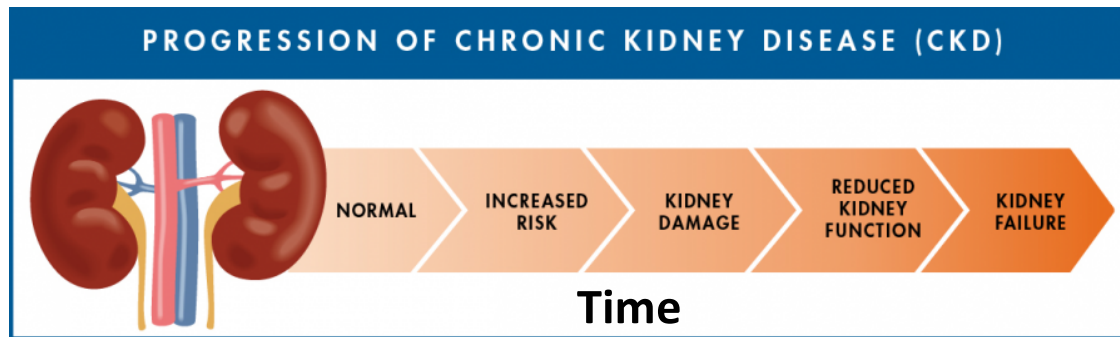


Source for figure:

<http://www.mahesh-vc.com/blog/understanding-whos-paying-for-what-in-the-healthcare-industry>

# What is the future of how we treat chronic disease?

- Predicting a patient's future disease progression




# 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




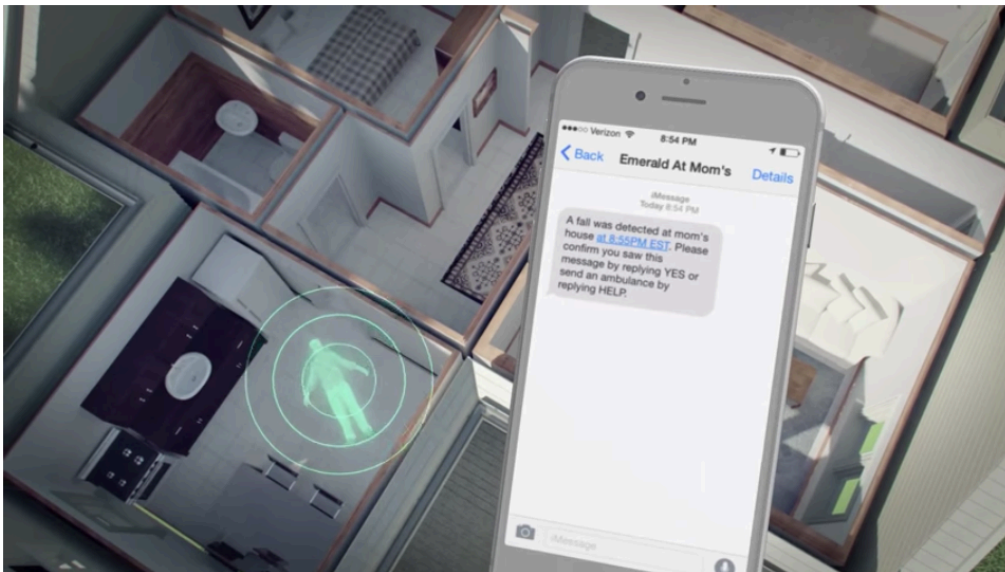


Figure sources: NIH,  
[https://www.roche.com/research\\_and\\_development/what\\_we\\_are\\_working\\_on/oncology/liquid-biopsy.htm](https://www.roche.com/research_and_development/what_we_are_working_on/oncology/liquid-biopsy.htm)

# What is the future of how we treat chronic disease?

- Continuous monitoring and coaching, e.g. for the elderly, diabetes, psychiatric disease



# What is the future of how we treat chronic disease?

- Discovery of new disease subtypes; design of new drugs; better targeted clinical trials

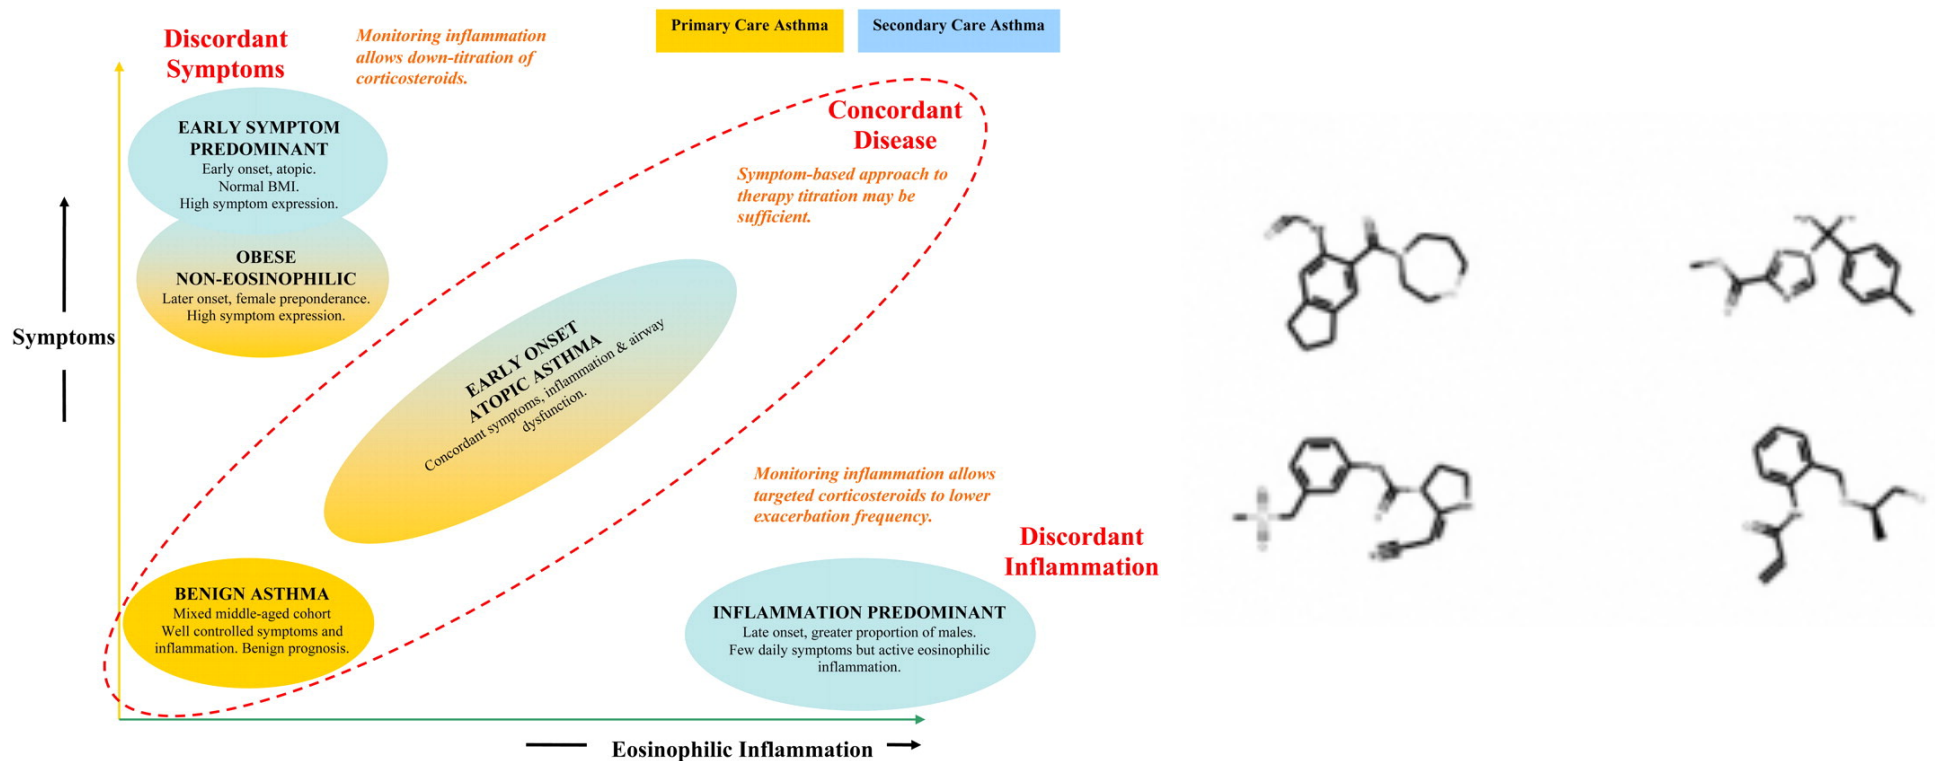


Figure sources: Haldar et al., Am J Respir Crit Care Med, 2008

<http://news.mit.edu/2018/automating-molecule-design-speed-drug-development-0706>

# 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



Recent breakthroughs in AI depended on *lots* of labeled data!

## Loud and clear

Speech-recognition word-error rate, selected benchmarks, %




The Switchboard corpus is a collection of recorded telephone conversations widely used to train and test speech-recognition systems

Sources: Microsoft; research papers

## Ever cleverer

Error rates on ImageNet Visual Recognition Challenge, %



Sources: ImageNet; Stanford Vision Lab

Economist.com

# 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

# 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**


# Course staff

- David Sontag (instructor)
  - Associate Professor in EECS (course 6) and part of CSAIL and IMES
  - PhD '10, then 5 years as professor at NYU
  - Leads clinical machine learning research group
- Peter Szolovits (instructor)
  - Professor in EECS, associate faculty in IMES
  - Researching AI in medicine since 1975 (!)
  - Leads clinical decision making group in CSAIL



# Course staff

- **Monica Agrawal (teaching assistant)**
  - PhD student with David Sontag
  - Research in clinical NLP
  - Undergrad at Stanford (with Jure Leskovec) and intern at Flatiron Health
- **Matthew McDermott (teaching assistant)**
  - PhD student with Pete Szolovits
  - Research on deep learning in machine learning for health care
  - Executive committee for new CHIL conference, organizer of NeurIPS ML4H workshop
- **Office hours Monday 12:30pm, location TBA**



# Prerequisites & Enrollment

- **Must submit pre-req quiz (on course website) by 11:59PM EST today**
- We assume previous undergraduate-level ML, and comfort with:
  - Machine learning methodology (e.g. generalization, cross-validation)
  - Supervised machine learning techniques (e.g. support vector machines, neural networks)
  - Optimization for ML (e.g. stochastic gradient descent)
  - Statistical modeling (e.g. Gaussian mixture models)
  - Python
- **Listeners and auditors by permission only; please fill out pre-req quiz**



# Logistics

- Course website:  
<https://mlhcmit.github.io/>
- All announcements made via Piazza – make sure you are signed up for it!
- **Recitation (optional): Fridays, starting next week (details TBD)**
- Grading:
  - 35% homework (~5 problem sets; both theory & practice)
  - 20% quiz (early/mid April)
  - 35% course project
  - 10% participation (scribing, reading responses, and in-class discussion)

# Upcoming problem sets

- PS0 (due Mon 5pm): human subjects training & MIMIC data use agreement
- PS1: Predicting mortality in ICUs using labs and clinical text
  - Released late next week
  - Due Feb 24 (tentative)

# Draft Schedule

- **Module 1: Overview of clinical care & data** (3 lectures)
- **Module 2: Using ML for risk stratification and diagnosis** (9 lectures)
  - Supervised learning with noisy, biased, or censored labels
  - Interpretability; Methods for detecting dataset shift; Fairness; Uncertainty
- **Module 3: Suggesting treatments** (4 lectures)
  - Causal inference; Off-policy reinforcement learning
- **Module 4: Understanding disease and its progression** (3 lectures)
  - Unsupervised learning on censored time series with substantial missing data
  - Discovery of disease subtypes; Precision medicine
- **Module 5: Human factors** (3 lectures)
  - Differential diagnosis; Utility-theoretic trade-offs
  - Automating clinical workflows
  - Translating technology into the clinic

# 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**  
<https://mit6874.github.io/>