

Machine Learning for Healthcare

6.871, HST.956

Lecture 1: What makes healthcare unique?

David Sontag



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. Interlude: Student & faculty introductions
3. Why *now*? What has changed?
4. Examples of how ML will transform healthcare
5. What is *unique* about ML in healthcare?
6. Course logistics & 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

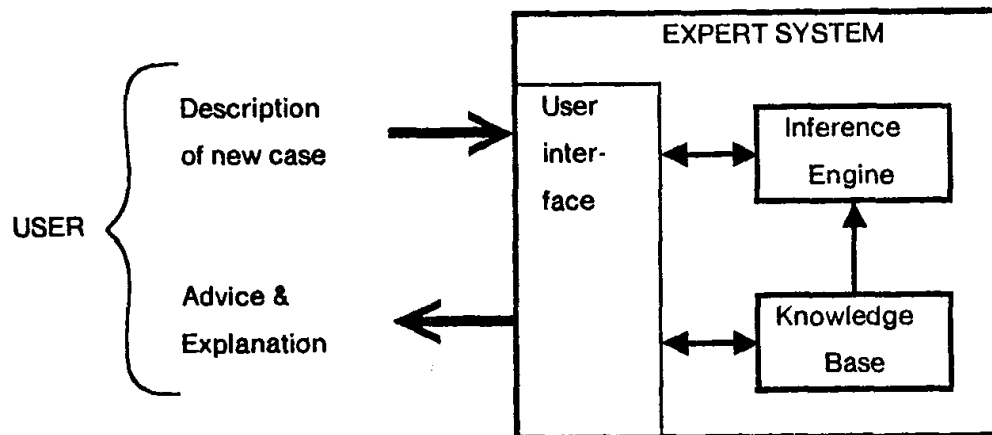


FIGURE 1-1 Major parts of an expert system. Arrows indicate information flow.

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

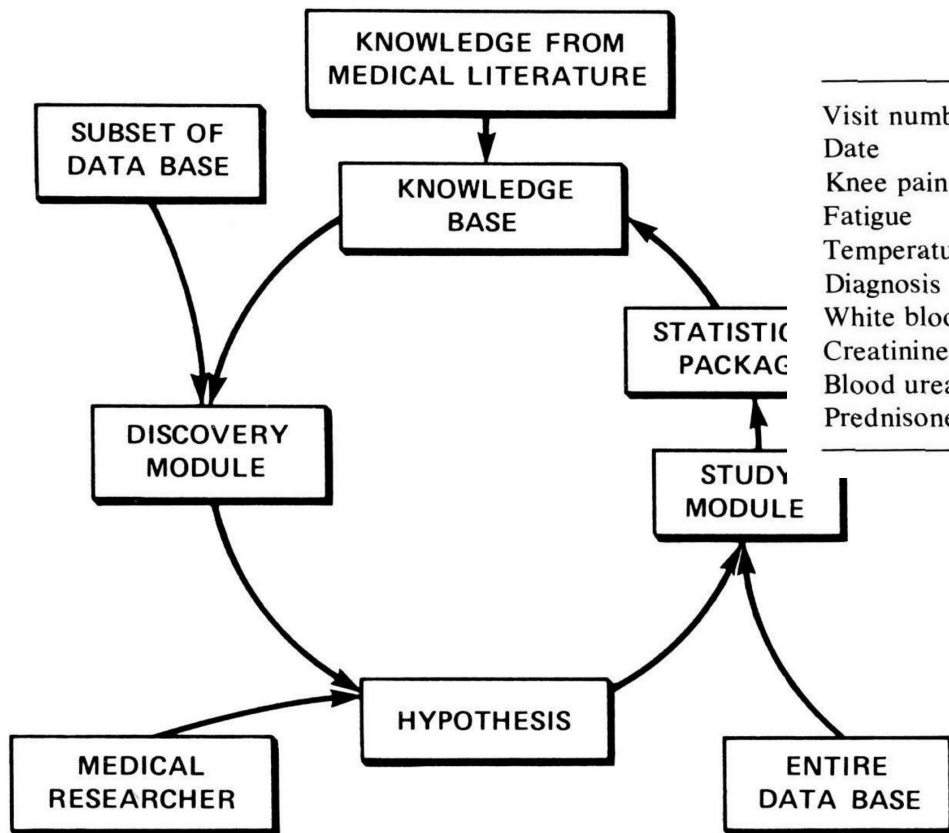


TABLE 1

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 ⁴	57	20	60	9-15-2	0.6	80	75
Vasculitis ²	404	403	73	8-5-1	8.0	94	—
Myocardial infarction ⁶	351	331	89	20-10-10-1	1.1	97	84
Myocardial infarction ⁸	356	350	87	20-10-10-1	1.1	97	94
Low back pain ¹¹	100	100	25	50-48-2	0.2	90	90
Cancer outcome ¹³	5,169	3,102	—	54-40-1	1.4	0.779	0.776
Psychiatric length of stay ¹⁷	957	106	73	48-400-4	0.2	74	76
Intensive care outcome ²³	284	138	91	27-18-1	0.5	0.82	0.82
Skin tumor ²¹	150	100	80	18	—	80	90
Evoked potentials ³⁵	100	67	52	14-4-3	3.8	77	77
Head injury ⁴⁷	500	500	50	6-3-3	20	66	77
Psychiatric outcome ⁵⁴	289	92	60	41-10-1	0.7	79	—
Tumor classification ⁵⁵	53	6	38	8-9-3	1.4	99	88
Dementia ⁵⁷	75	18	19	80-10-7-7	0.6	61	—
Pulmonary embolism ⁵⁹	607	606	69	50-4-1	2.9	0.82	0.83
Heart disease ⁶²	460	230	54	35-16-8-2	3	83	84
Thyroid function ⁶²	3,600	1,800	93	21-16-8-3	22	98	93
Breast cancer ⁶²	350	175	66	9-4-4-2	10	97	96
Diabetes ⁶²	384	192	65	8-4-4-2	12	77	75
Myocardial infarction ⁶³	2,856	1,429	56	291-1	9.8	85	—
Hepatitis ⁶⁵	39	42	38	4-4-3	3.3	74	79
Psychiatric admission ⁷⁶	319	339	85	53-1-1	6.0	91	—
Cardiac length of stay ⁸³	713	696	73	15-12-1	3.5	0.70	—
Anti-cancer agents ⁸⁹	127	141	25	60-7-6	1.5	91	86
Ovarian cancer ⁹¹	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. Interlude: Student & faculty introductions**
3. Why *now*? What has changed?
4. Examples of how ML will transform healthcare
5. What is *unique* about ML in healthcare?
6. Course logistics & 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
- Madhur Nayan (instructor)
 - Surgeon at MGH, Fellow in Urologic Oncology at Harvard
 - MD '12 McGill, PhD '17 Univ. of Toronto



Course staff – teaching assistants

- Intae Moon

- PhD student in EECS advised by Alexander Gusev (Dana-Farber)
- Research on diagnosing cancers of unknown primary origin using genomics, survival analysis



- Zeshan Hussain

- PhD student in EECS advised by David Sontag
- Research on disease progression modeling, causal inference, deep generative models




Student intros

- We have a diverse set of students from MIT, Harvard, and local hospitals – let's start to get to know each other!

Outline for today's class

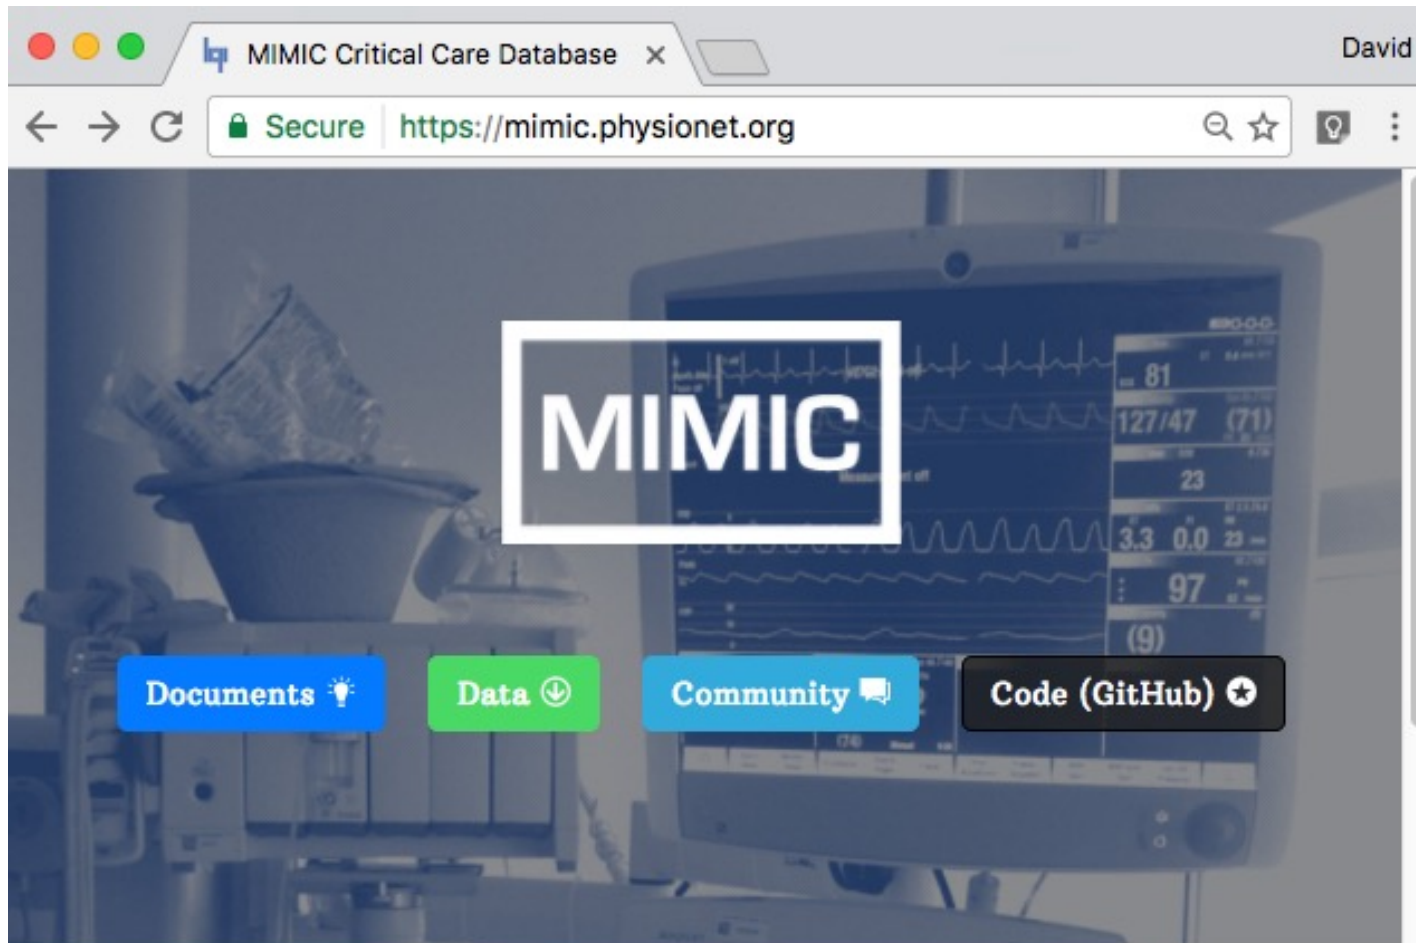
1. Brief history of AI and ML in healthcare
2. Interlude: Student & faculty introductions
3. **Why *now*? What has changed?**
4. Examples of how ML will transform healthcare
5. What is *unique* about ML in healthcare?
6. Course logistics & syllabus

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
~60K critical care
patients

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

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

Standardization

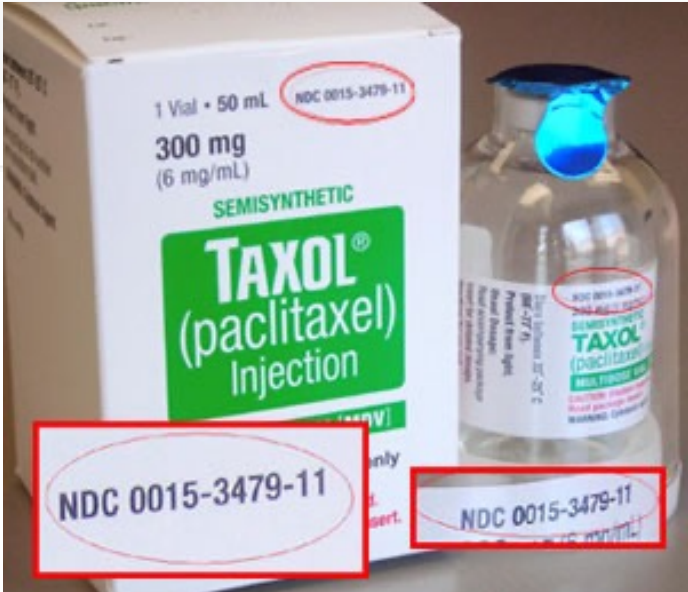
- Diagnosis codes: ICD-9 and ICD-10 (International Classification of Diseases)
- Laboratory tests: LOINC codes
- Pharmacy: National Drug Codes (NDCs)
- Unified Medical Language System (UMLS): millions of medical concepts

LOINC[®] From Regenstrief

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>	



1 Vial • 50 mL
300 mg
(6 mg/mL)
SEMISYNTHETIC
TAXOL[®]
(paclitaxel)
Injection

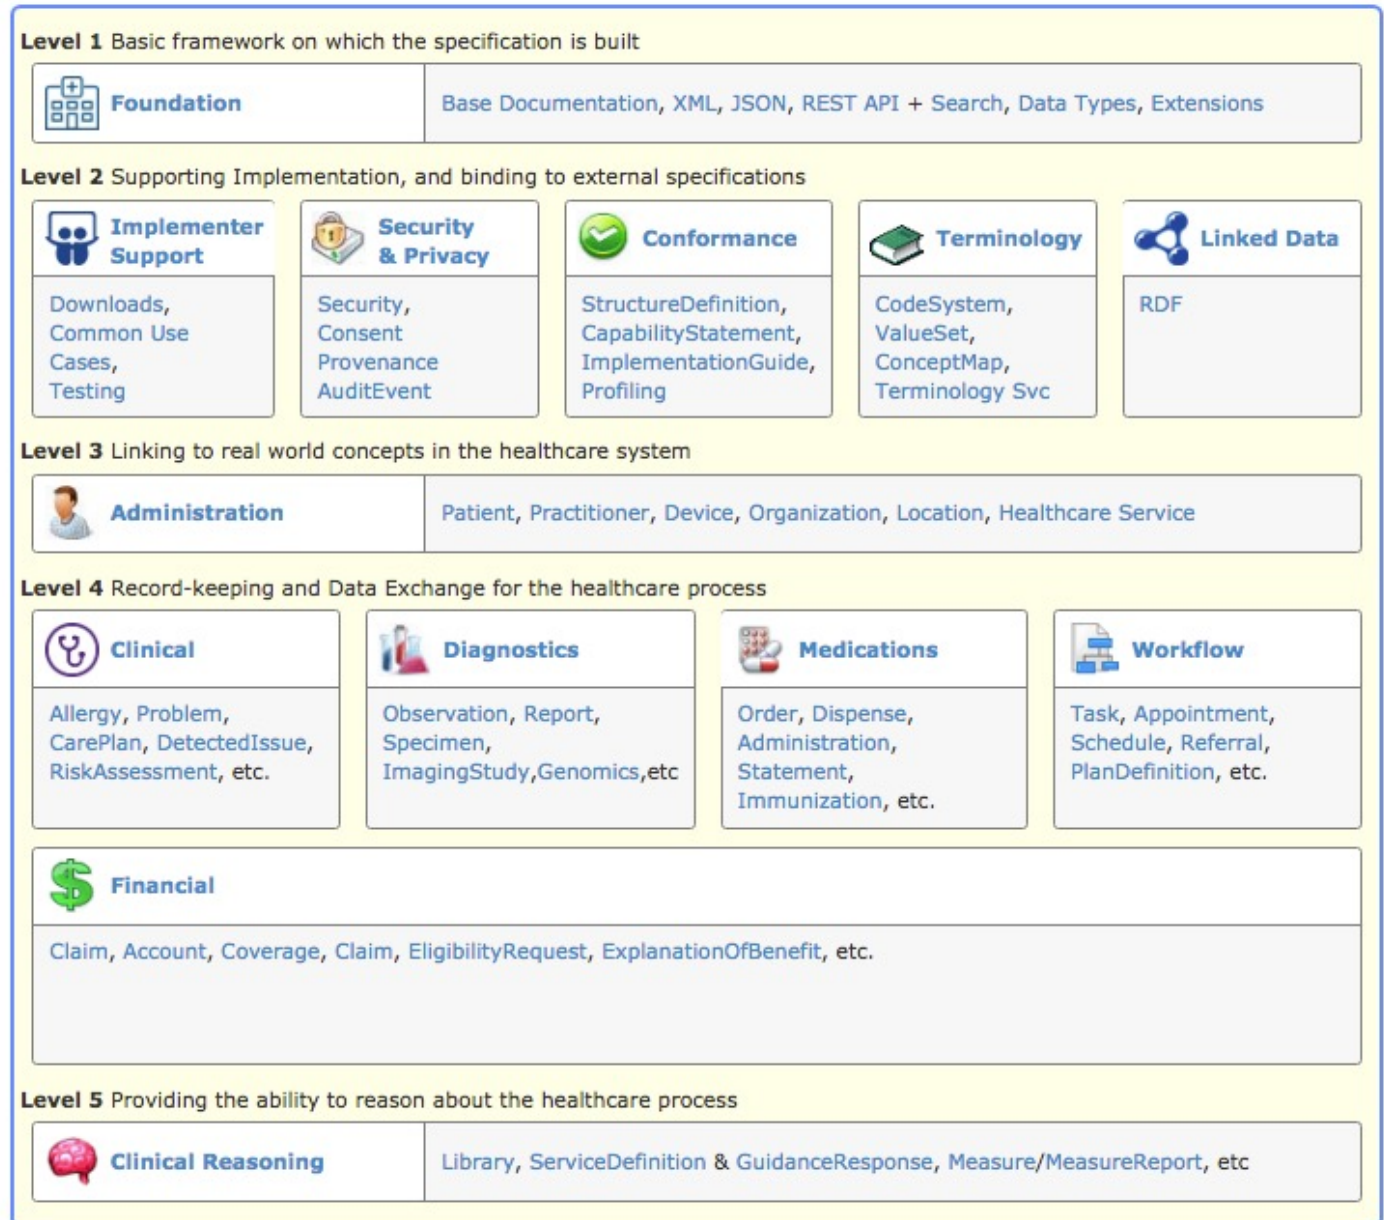
NDC 0015-3479-11

NDC 0015-3479-11

NDC 0015-3479-11

[http://oplinc.com/newsletter/index_May08.htm]

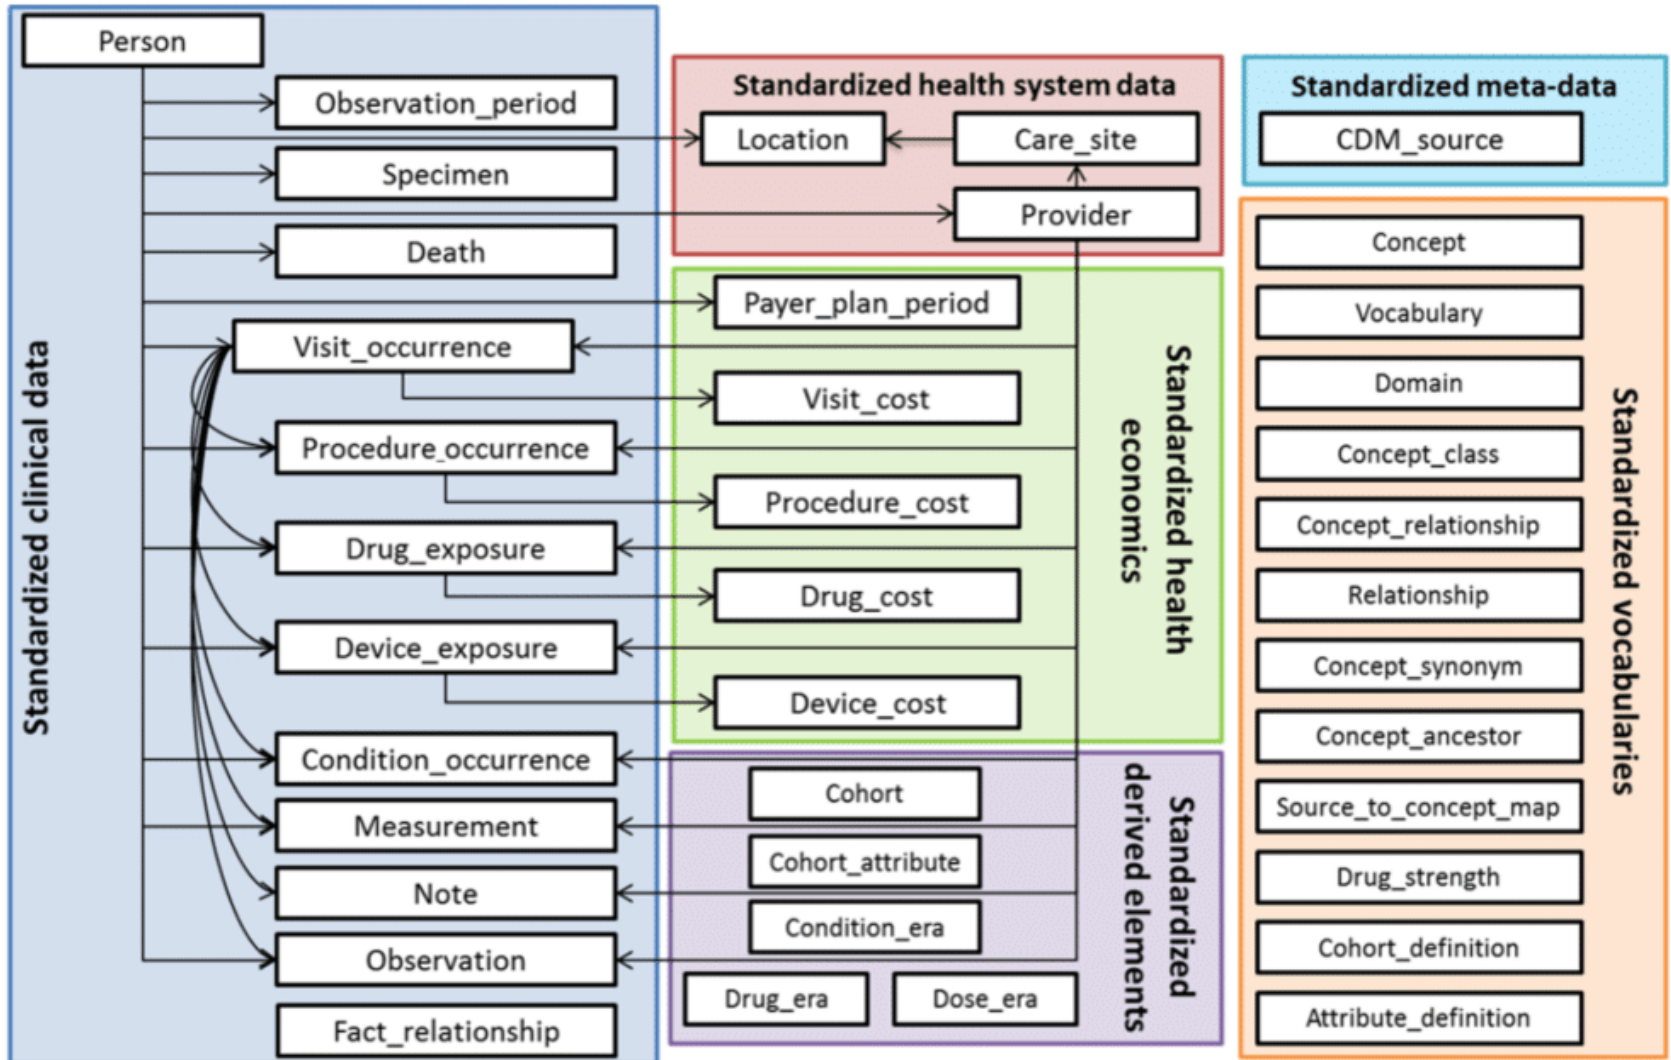
Standardization



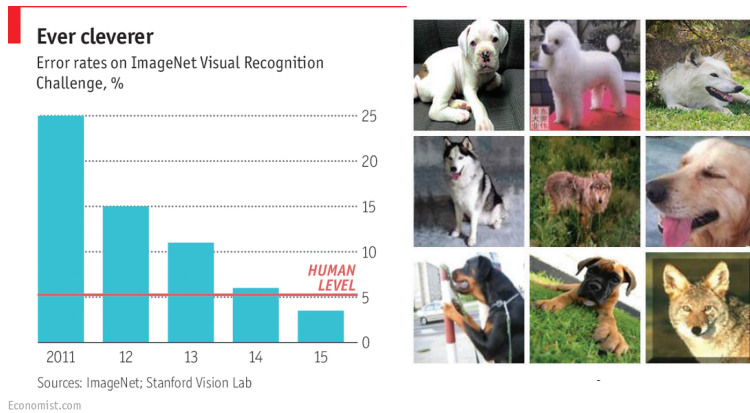
Standardization



OMOP
Common
Data
Model v5.0



Breakthroughs in machine learning



Object recognition
(deep neural networks)

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 24, 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)



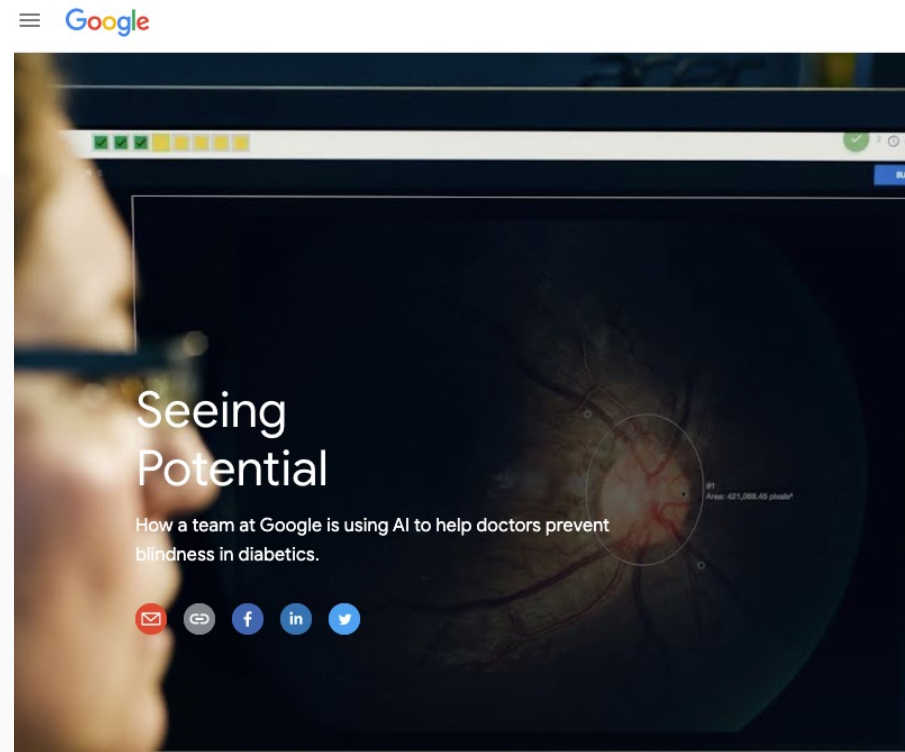
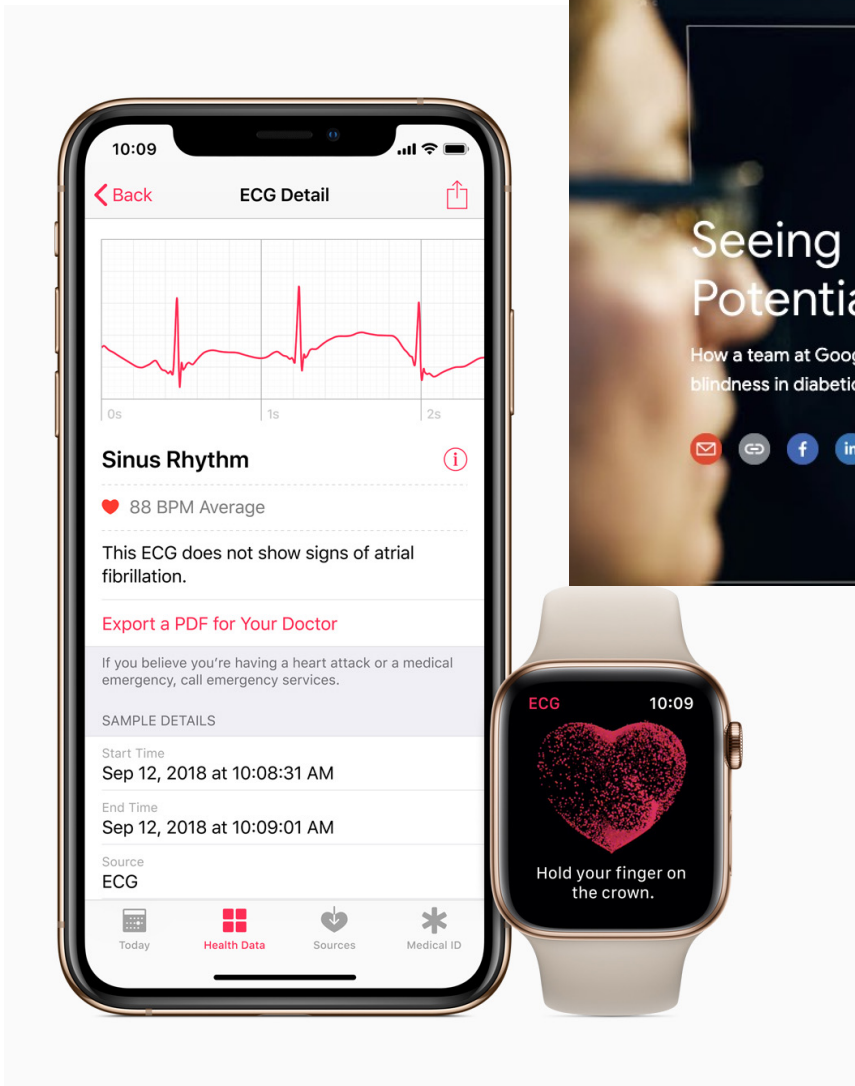
AlphaGo
(reinforcement learning)



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

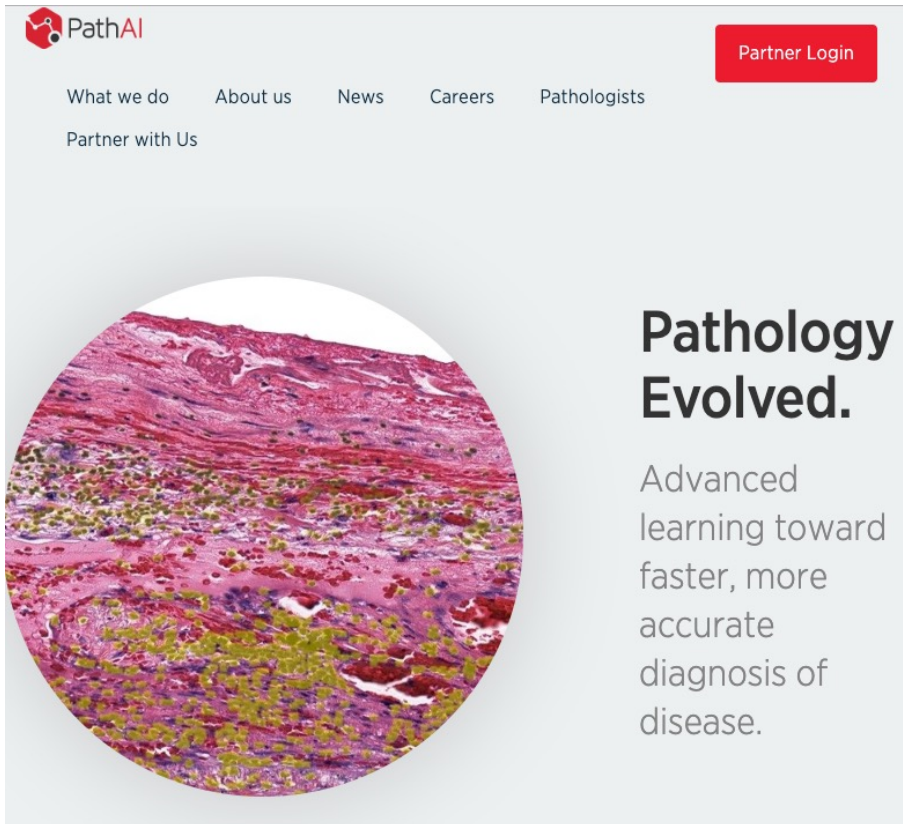
Generating realistic data
(GANs, VAEs)

Tech industry interest in health care



The image is a screenshot of the Amazon Comprehend Medical website. The top navigation bar includes the 'aws' logo, a search icon, and a menu icon. Below the navigation bar, the text 'Amazon Comprehend' is followed by a dropdown arrow. The main content area features the heading 'Amazon Comprehend Medical' in large white text. Below the heading, the text reads: 'Extract information from unstructured medical text accurately and quickly' and 'No machine learning experience required'. A yellow button with the text 'Get started with Amazon Comprehend Medical' is positioned below the text. At the bottom of the page, there is a 'TECH TALK' section with the heading 'AI-Powered Health Data Masking' and the subtext: '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

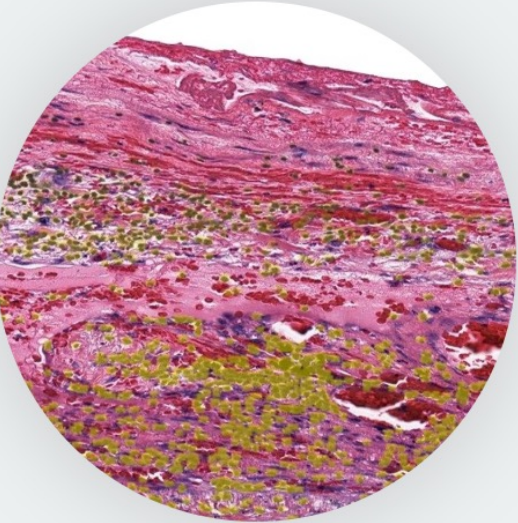


PathAI

Partner Login

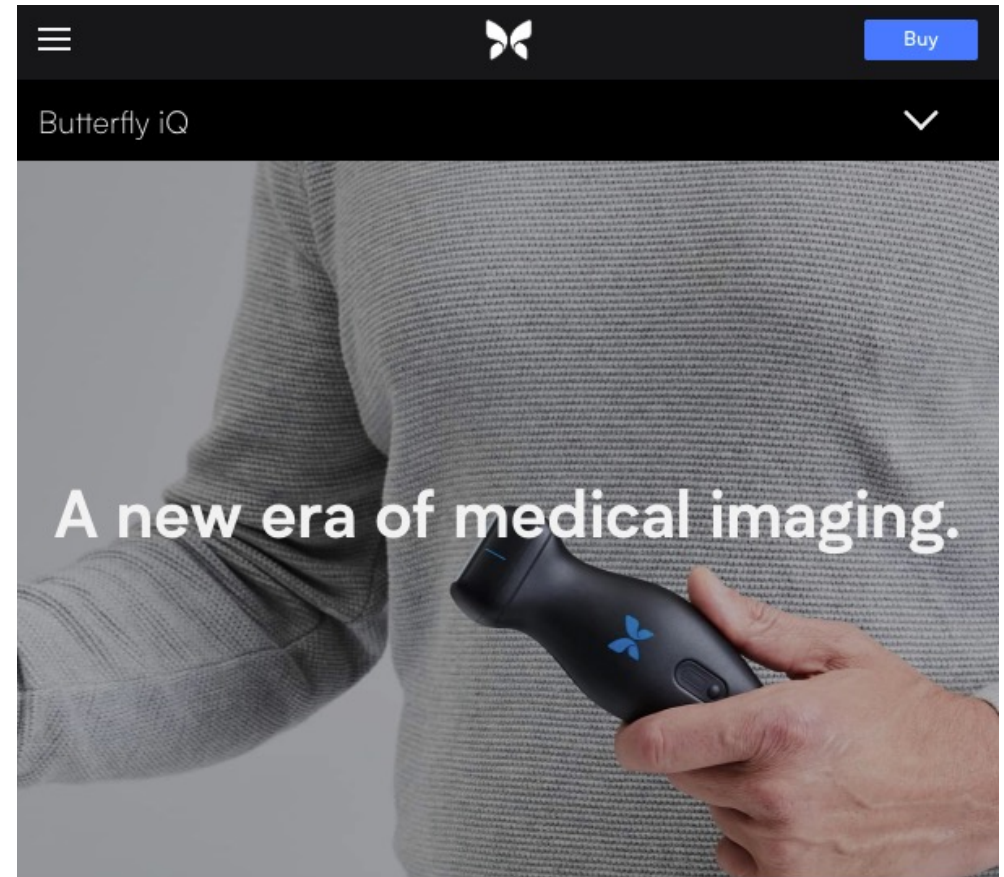
What we do About us News Careers Pathologists

Partner with Us




Pathology Evolved.

Advanced learning toward faster, more accurate diagnosis of disease.



Butterfly iQ

Buy

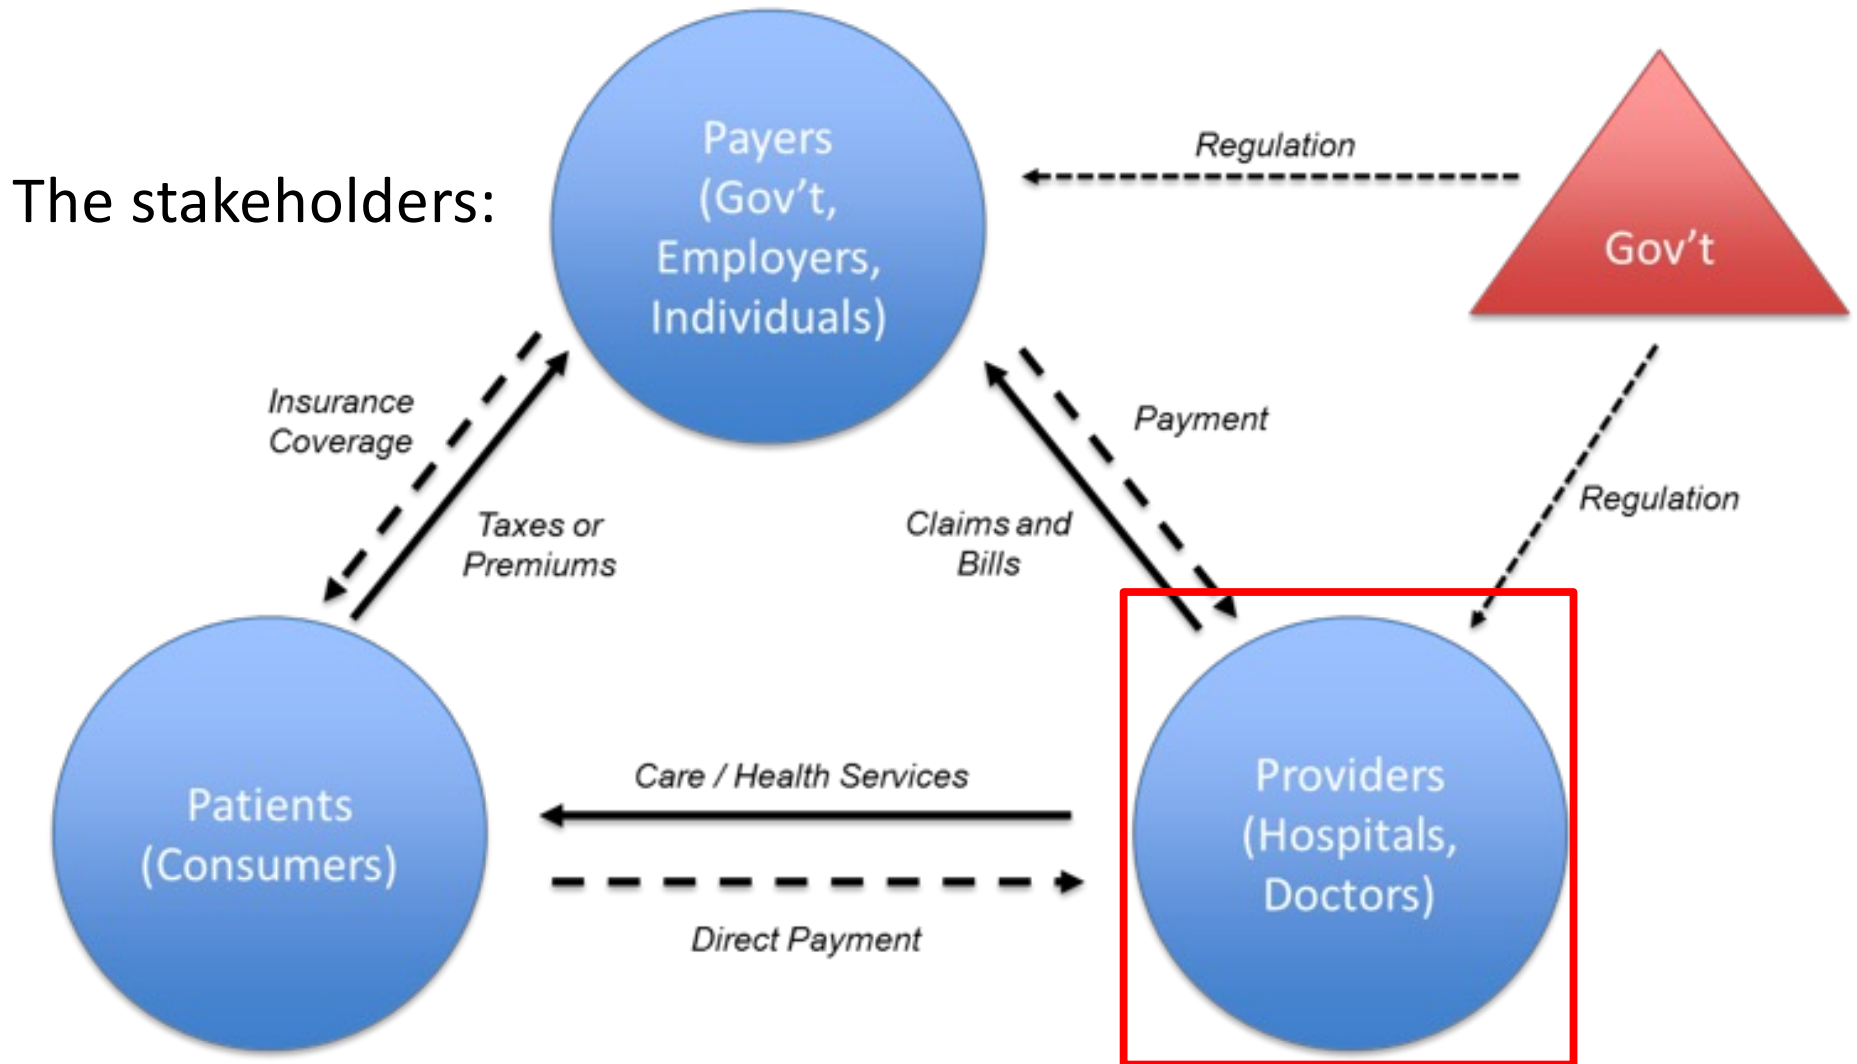


A new era of medical imaging.

Outline for today's class

1. Brief history of AI and ML in healthcare
2. Interlude: Student & faculty introductions
3. *Why now?* What has changed?
4. **Examples of how ML will transform healthcare**
5. *What is unique* about ML in healthcare?
6. Course logistics & 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


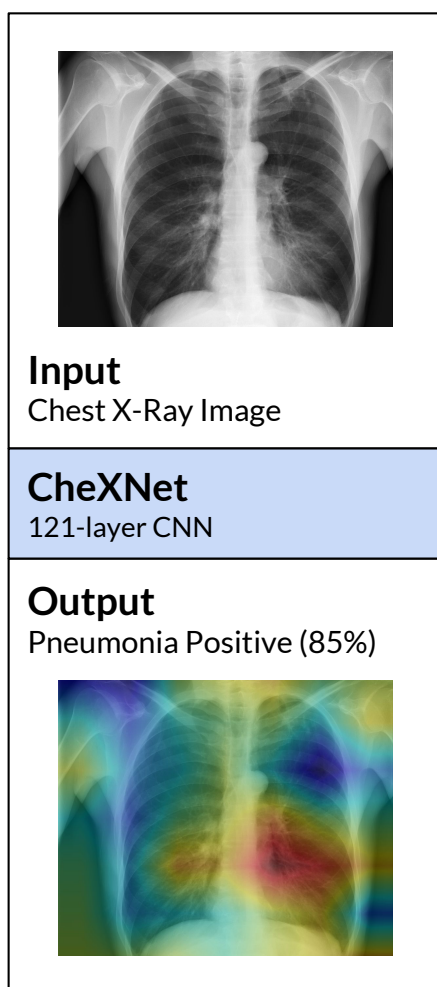
Aspirin taken before arrival

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

Triage note

Predicted chief complaints

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

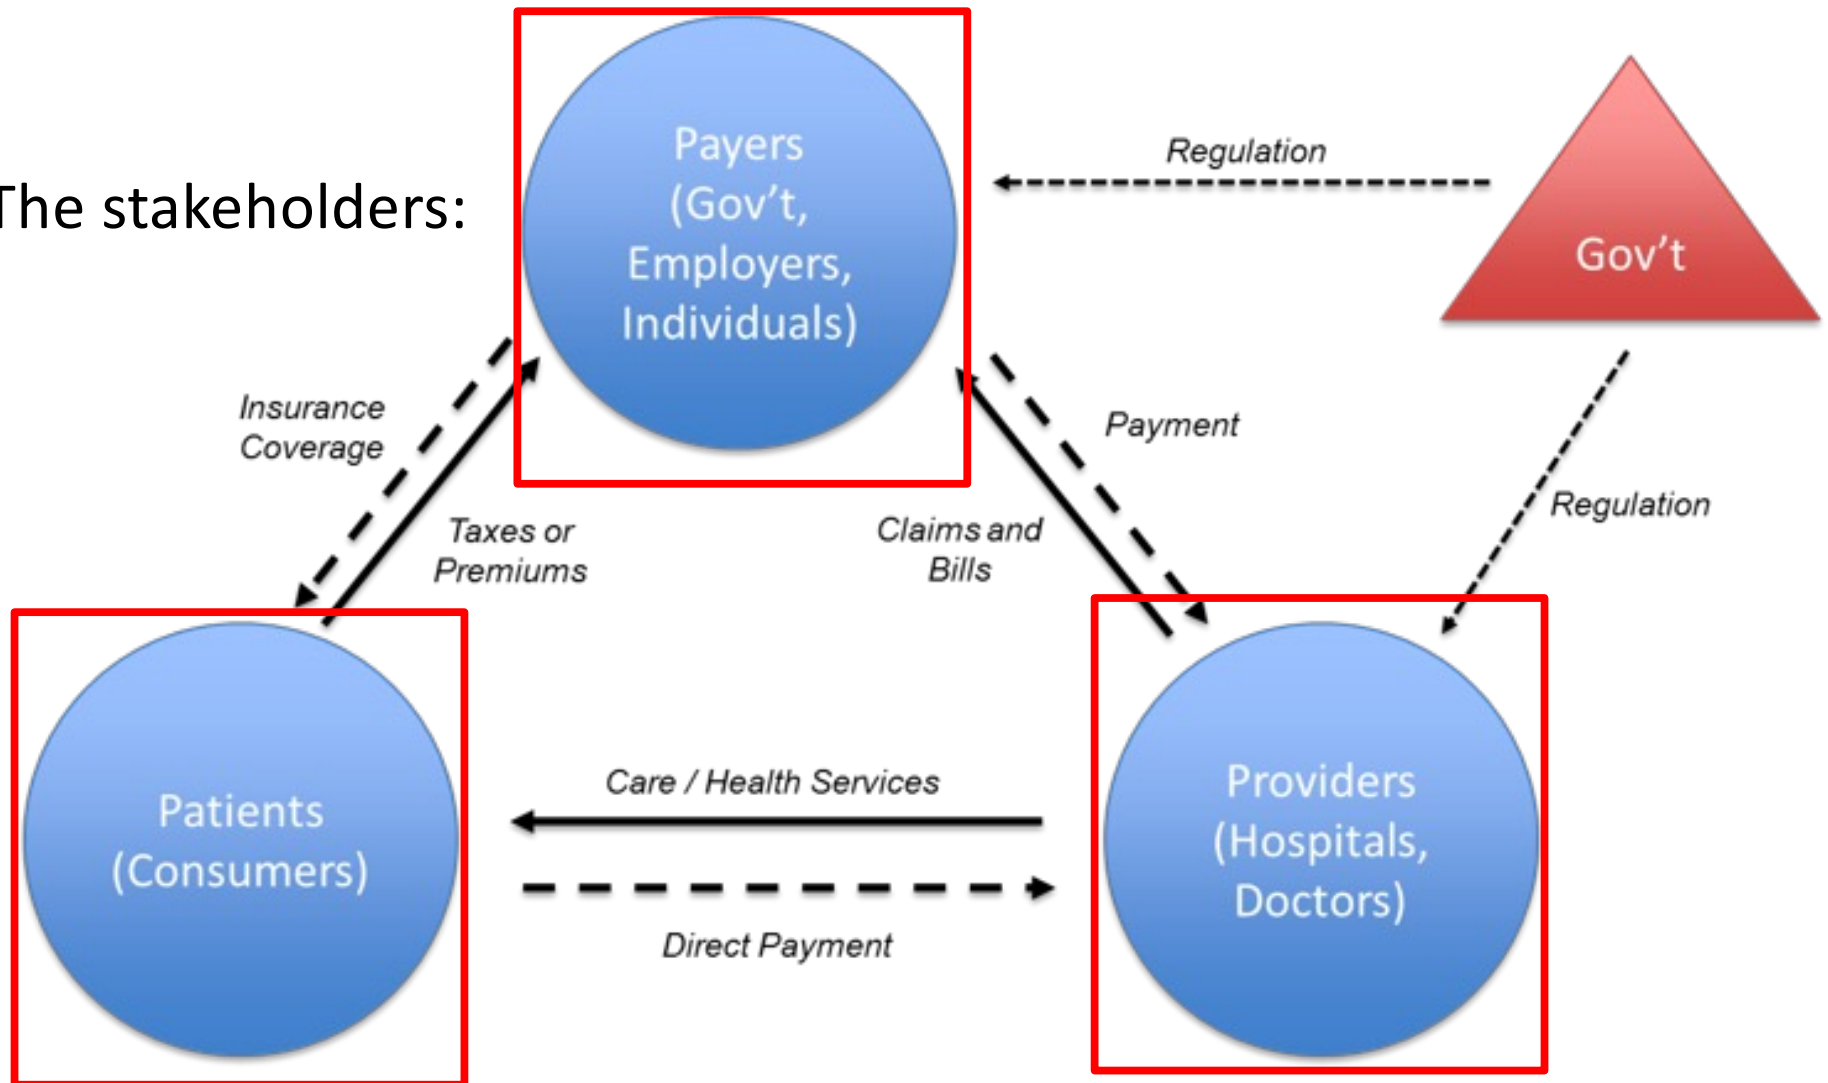
RIGHT UPPER QUADRANT PAIN
RUQ ABDOMINAL PAIN
RUQ PAIN
ALLERGIC REACTION
L KNEE PAIN
RECTAL PAIN
RIGHT SIDED ABD PAIN
RIGHT SIDED ABDOMINAL PAIN
L WRIST PAIN
RIGHT SIDED CHEST PAIN
TESTICULAR PAIN
KNEE PAIN
ELBOW PAIN
RIB PAIN
L ELBOW PAIN
HAND PAIN
VAGINAL PAIN

Enter Cancel

Contextual auto-complete

ML will transform every aspect of healthcare

The stakeholders:


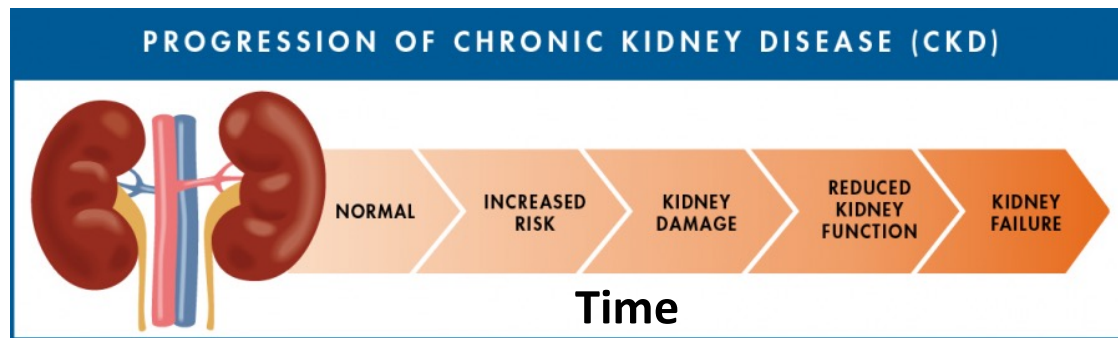


Source for figure:

<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




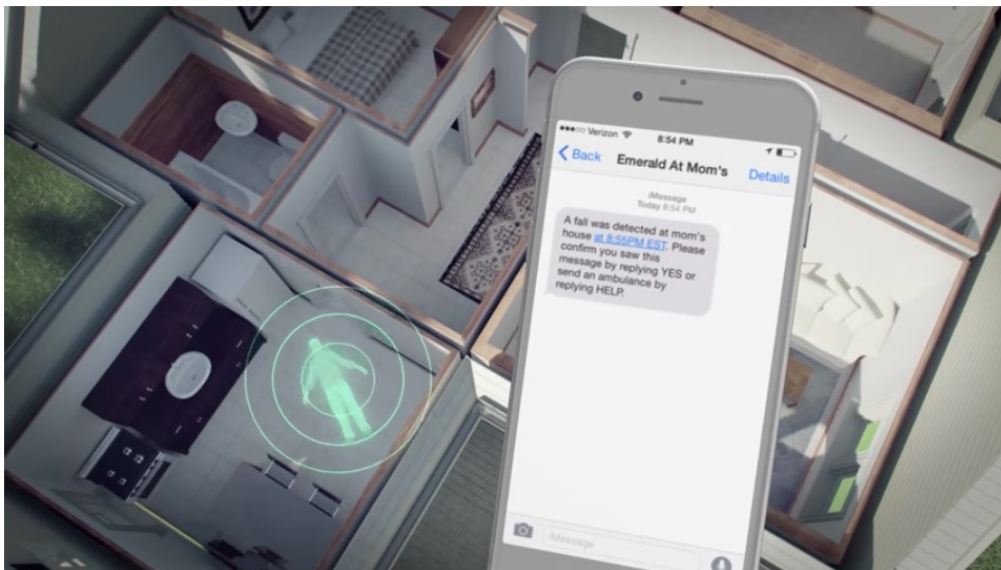


Figure sources: NIH,
https://www.roche.com/research_and_development/what_we_are_working_on/oncology/liquid-biopsy.htm

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

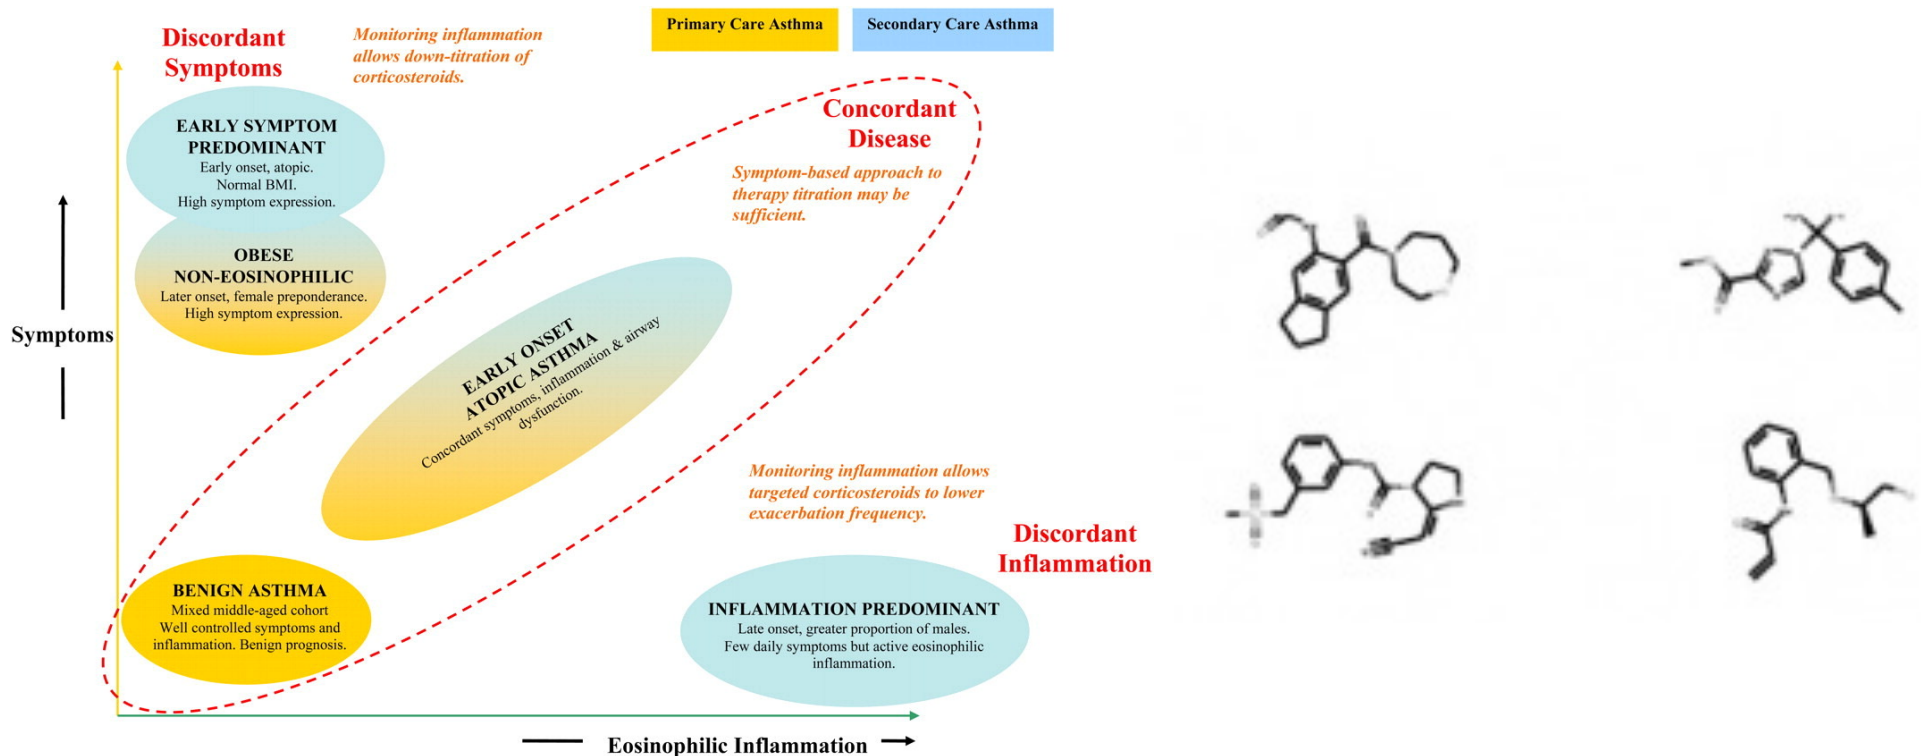


Figure sources: Haldar et al., Am J Respir Crit Care Med, 2008

<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. Interlude: Student & faculty introductions
3. *Why now?* What has changed?
4. Examples of how ML will transform healthcare
5. **What is *unique* about ML in healthcare?**
6. Course logistics & 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. Interlude: Student & faculty introductions
3. *Why now?* What has changed?
4. Examples of how ML will transform healthcare
5. What is *unique* about ML in healthcare?
6. **Course logistics & syllabus**

Prerequisites

- Previous undergraduate-level ML (e.g. 6.036):
 - Machine learning methodology (e.g. generalization, cross-validation)
 - Supervised machine learning techniques (e.g. linear and logistic regression, neural networks)
 - Loss functions, regularization, and optimization (e.g. stochastic gradient descent)
 - Statistical modeling (e.g. Gaussian mixture models)
- Python

Logistics

- Course website:
<https://mlhcmits.github.io/>
- All announcements made via Canvas
- Use Piazza for Q&A with staff and each other
- Recitation (required): Fridays 3-4pm in 4-270
(starts this week)
- Office hours TBD

Grading

- 40% course project
- 35% homework (5 problem sets; both theory & practice)
- 20% final exam (date to be scheduled by registrar)
- 5% participation – note: class attendance is required*

* Exceptions will be made for quarantine/isolation.

This week's assignments/readings

- PS0 (due **Weds 2/2, 11:59pm**): human subjects training & MIMIC data use agreement
- Project rankings (due **Fri 2/4, 1pm**)
- Reading response (due **Fri 2/4, 1pm**)

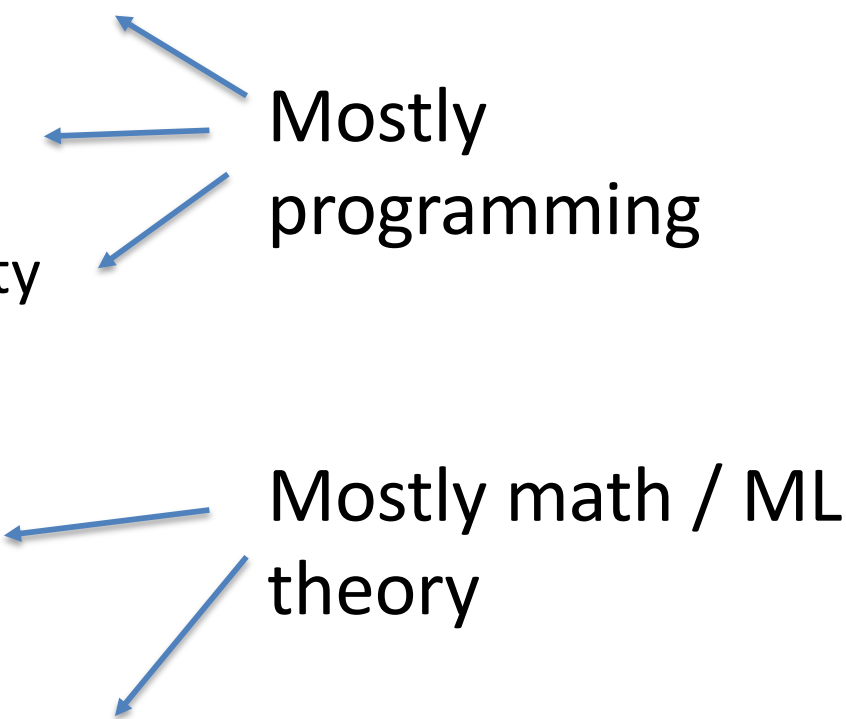
AI in Health and Medicine – required reading

Pranav Rajpurkar, Emma Chen, Oishi Banerjee & Eric J. Topol
Nature Medicine, 2022

Machine Learning in Medicine – optional reading

Alvin Rajkomar, Jeffrey Dean, Isaac Kohane
New England Journal of Medicine, 2019

Future assignments (dates approximate)

- PS1: EMRs, physiological data, risk modeling
 - Released 2/7, Due 2/16
 - PS2: Clinical NLP
 - Released 2/16, Due 2/25
 - PS3: Imaging, interpretability
 - Released 2/25, Due 3/9
 - ***Project checkpoint 1, 3/2***
 - PS4: Causal inference
 - Released 3/9, Due 3/18
 - ***Project checkpoint 2, 4/6***
 - PS5: Dataset shift, learning with imperfect data
 - Released 4/6, due 4/20
 - ***Project reports due, 5/4***
- Mostly programming
- Mostly math / ML theory
- 

Course project

- Teams of 3-5 students
- Each project will have one or more clinicians involved as mentors and/or students
- Project ranking form sent out later today
- Project poster presentations **May 5, 9:30am-12pm** in 34-401 (no recitation 5/6)