

Machine Learning for Healthcare

HST.956, 6.S897

Lecture 1: What makes healthcare unique?

Prof. David Sontag & Pete Szolovits



Need 2 scribes for today's lecture

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 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. **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



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

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. The navigation menu includes links for MEDIA ROOM, SUPPORT, CAREER, SOLUTIONS, EVENTS, KNOWLEDGE, and ABOUT. The Truven Health Analytics logo is visible, along with the text "an IBM Company". The main content area features a purple header with "Life Sciences" and a breadcrumb trail: Home » Life Sciences » Data & Tools » MarketScan Databases. Below this is a large image of a hand holding a smartphone displaying a data visualization. A sidebar on the left lists navigation options: Market Knowledge, Real World Evidence, Stakeholder Management, Data & Tools (highlighted), MarketScan Databases, Treatment Pathways, Inpatient/Outpatient View, PULSE, and Heartbeat Profiler. The main content area contains the 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." There are also several small icons and partial text visible on the right side of the page.

“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



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>	

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

NDC 0015-3479-11


Standardization



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

Industry interest in ML & healthcare



Google DeepMind

Home AlphaGo DQN Health Press

DeepMind Health

CLINICIAN-LED TECHNOLOGY



PathAI


What we do About us Careers Pathologists Partner with Us

Pathology Evolved.

Advanced learning toward faster, more accurate diagnosis of disease.

IBM Watson for Oncology

Get oncologists the assistance they need to make more informed treatment decisions. Watson for Oncology analyzes a patient's medical information against a vast array of data and expertise to provide evidence-based treatment options.



BAYLABS

Better heart health reimaged through artificial intelligence.

Who We Are

Bay Labs combines deep learning, a type of artificial intelligence, with cardiovascular imaging to help in the diagnosis and management of heart disease.


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



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 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:

Predicted 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:

Contextual auto-complete

- 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 Cancell

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>

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. *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, joint IMES & CSAIL
 - PhD MIT, then 5.5 years as professor at NYU
 - Leads clinical machine learning 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

- Willie Boag (teaching assistant)
 - PhD student with Pete Szolovits
 - Research in clinical NLP
 - Master's thesis on quantifying racial disparities in end-of-life care
- Irene Chen (teaching assistant)
 - PhD student with David Sontag
 - Research in fairness in ML, and modeling disease progression
 - Before PhD, worked for 2 years at Dropbox
- **Office hours Monday 1pm, 32-G 9th floor lobby**



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
- **Because of space limitations, no listeners or auditors will be permitted**

Logistics

- Course website:
<https://mlhc19mit.github.io/>
- All announcements made via Piazza – make sure you are signed up for it!
- **Recitation (optional): Fridays, starting next week (details TBD)**
- Grading:
 - 40% homework (6 problem sets)
 - 40% course project
 - 20% participation (scribing, MLHC community consulting, reading responses, and in-class discussion)

Homework (tentative)

- **PS0 (due Monday): human subjects training & data use agreements**
- PS1: Predicting mortality in ICUs using labs and clinical text
- PS2: Risk stratification using health insurance claims
- PS3: Clinical natural language processing
- PS4: Physiological time-series
- PS5: Causal inference (theory)
- PS6: Inferring the effect opioid prescription on addiction

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/>