CPH 100: Foundations for CPH

Risk Stratification

Irene Y. Chen

How can we make
Data 146 better for you?

Slides adapted from
David Sontag

% who say they think the impact of artificial intelligence (AI) on society over the next 20 years will be ...

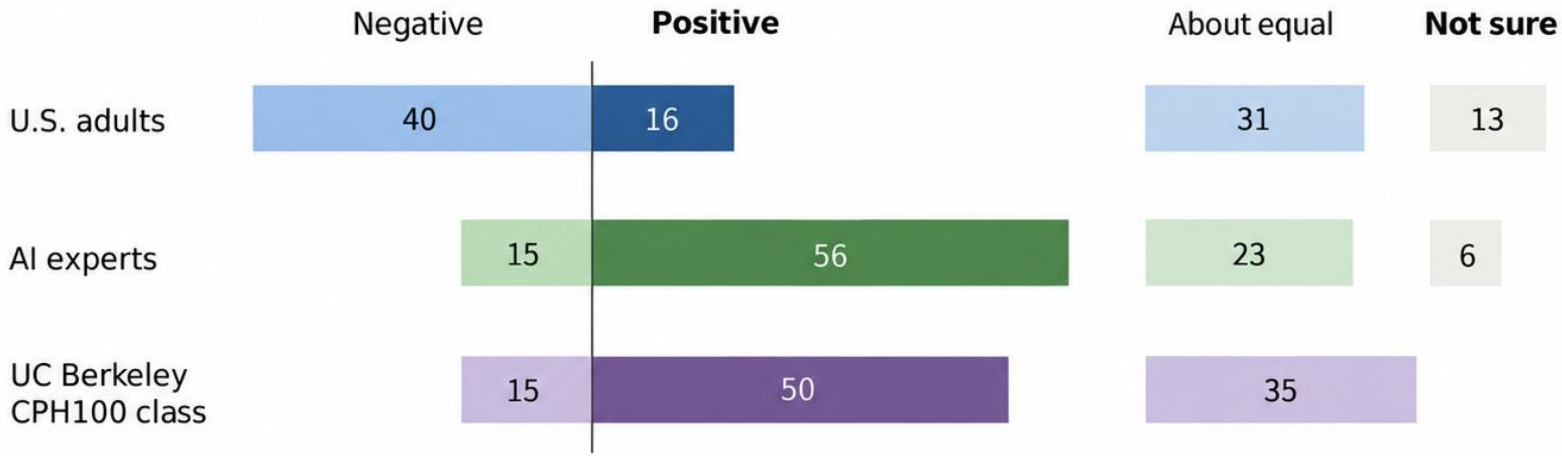

% who say they think the increased use of artificial intelligence (AI) is more likely to ...

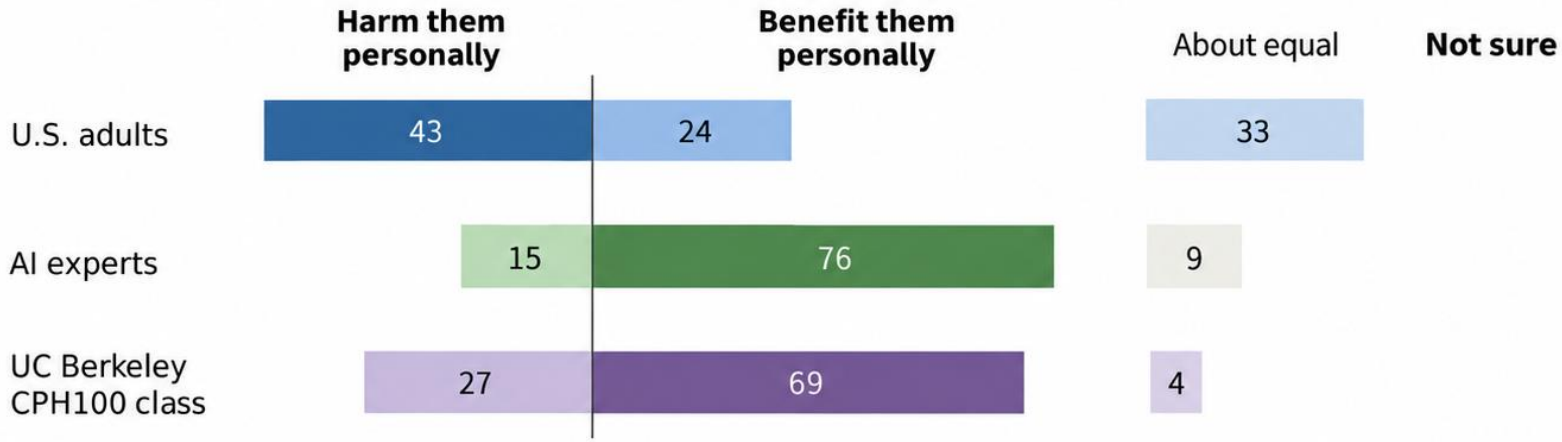

Source: Pew Research Center — U.S. adults, Feb. 17–23, 2026 (society) and Aug. 12–18, 2024 (personal); AI experts, Aug. 14–Oct. 31, 2024. UC Berkeley CPH100 class survey, Sept. 4, 2026 (n=26). Percentages rounded.

Outline

1. Introduction to risk stratification (15 mins)
2. Case study: Early detection of Type 2 diabetes (20 mins)
   - Encoding longitudinal structured health data
3. Framing as supervised learning problem (15 mins)
   - Deriving labels from EHR

How can we make Data 146 better for you?

Learning Objective: Formulate risk stratification as supervised learning problem

What *is* risk stratification?

- Separate a patient population into **high-risk** and **low-risk** of having an outcome
  - Predicting something in the future
- Coupled with **interventions** that target high-risk patients
- Goal is typically to reduce cost and improve patient outcomes

Examples of risk stratification

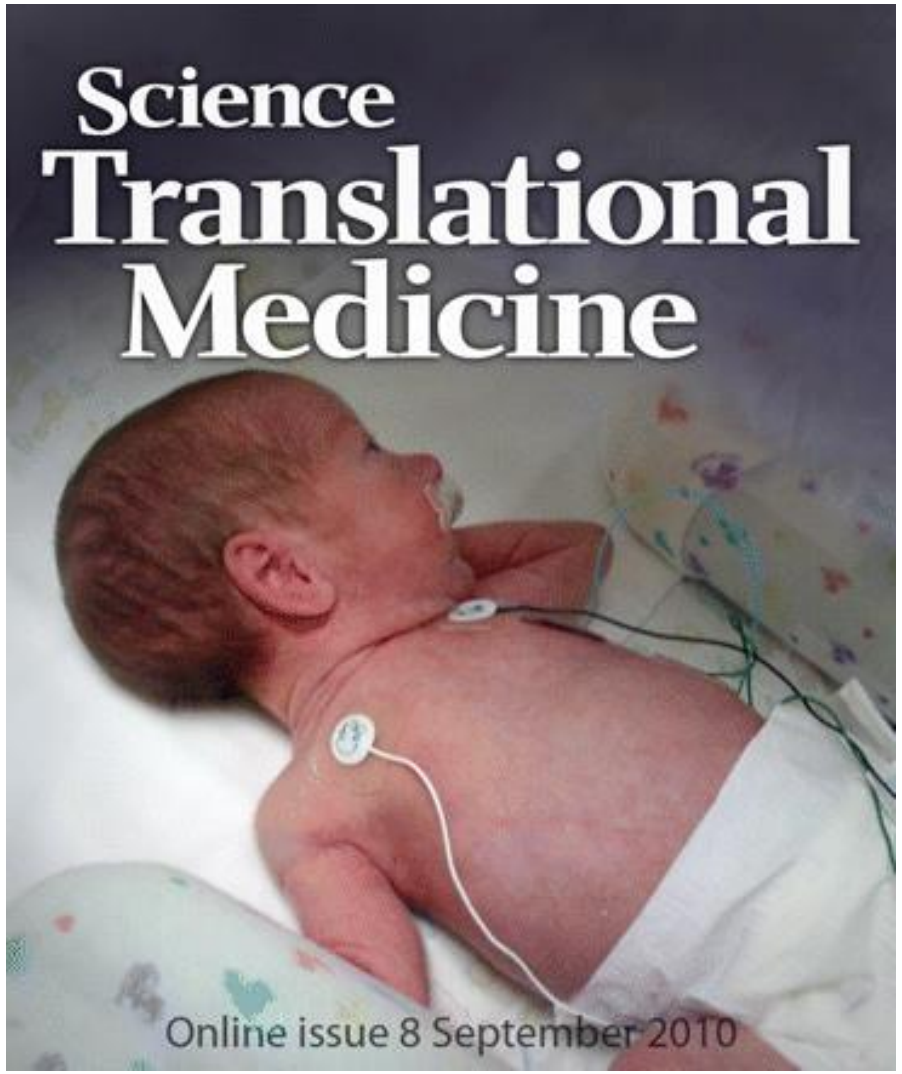

Preterm infant's risk of severe morbidity?

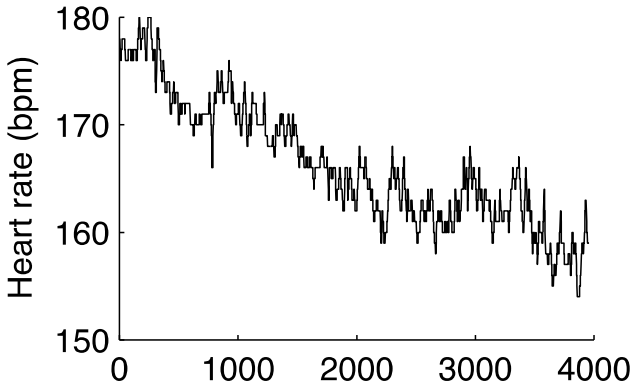

(Saria et al., Science Translational Medicine 2010)

Examples of risk stratification

Figure sources:
https://www.drmani.com/heart-attack/ (top)
https://www.emra.org/emresident/article/acute-mi-case-report/ (right)

Does this patient need to be admitted to the coronary-care unit?

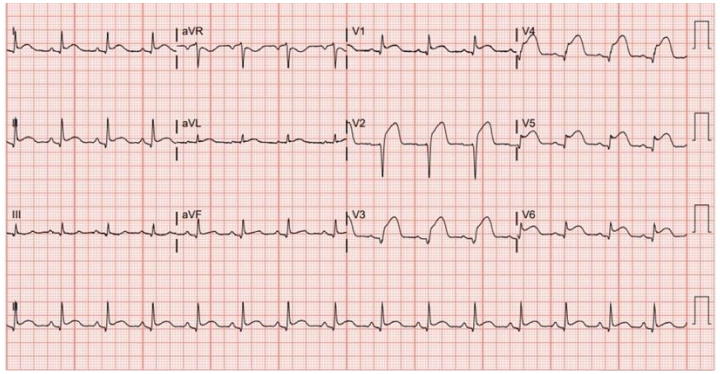

(Pozen et al., NEJM 1984)

Examples of risk stratification

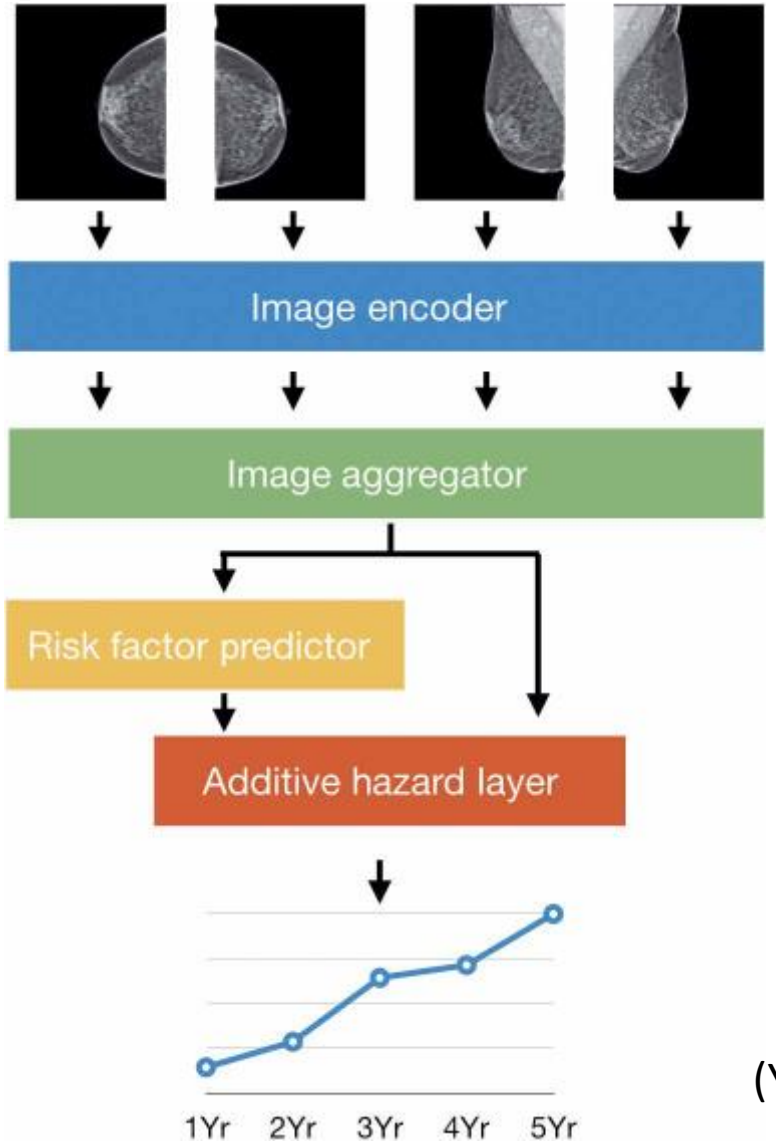

Will this woman develop breast cancer in the next 5 years?

(Yala et al., Science Translational Medicine 2021)

Examples of risk stratification

Irene, your genetics are associated with a **typical likelihood** of developing type 2 diabetes.

Based on data from 23andMe research participants, people of East Asian descent with genetics like yours have an estimated **30% chance** of developing type 2 diabetes at some point **between the ages of 32 (your current age) and 80.**

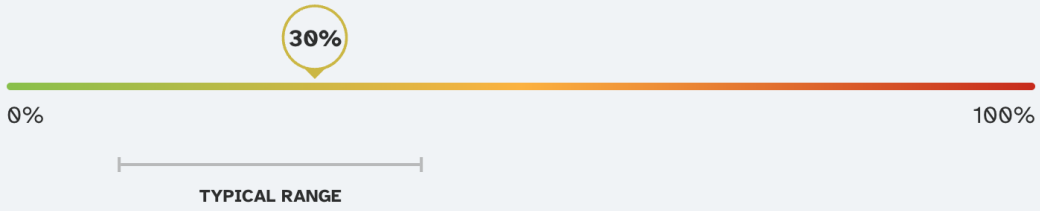

ETHNICITY	AUC VALUE
European	0.652
South Asian	0.603
Hispanic/Latino	0.638
East Asian	0.609
African	0.588

DNA
Deoxyribonucleic acid

Summary

This report is based on a statistical model that estimates the likelihood of developing type 2 diabetes by looking at genetic variants at 1,244 places in your DNA. We identified these variants and created this model using data from more than 1,110,000 23andMe research participants of European descent.

(Source: 23andme. https://permalinks.23andme.com/pdf/23_19-Type2Diabetes_March2019.pdf)

Differntial diagnosis?

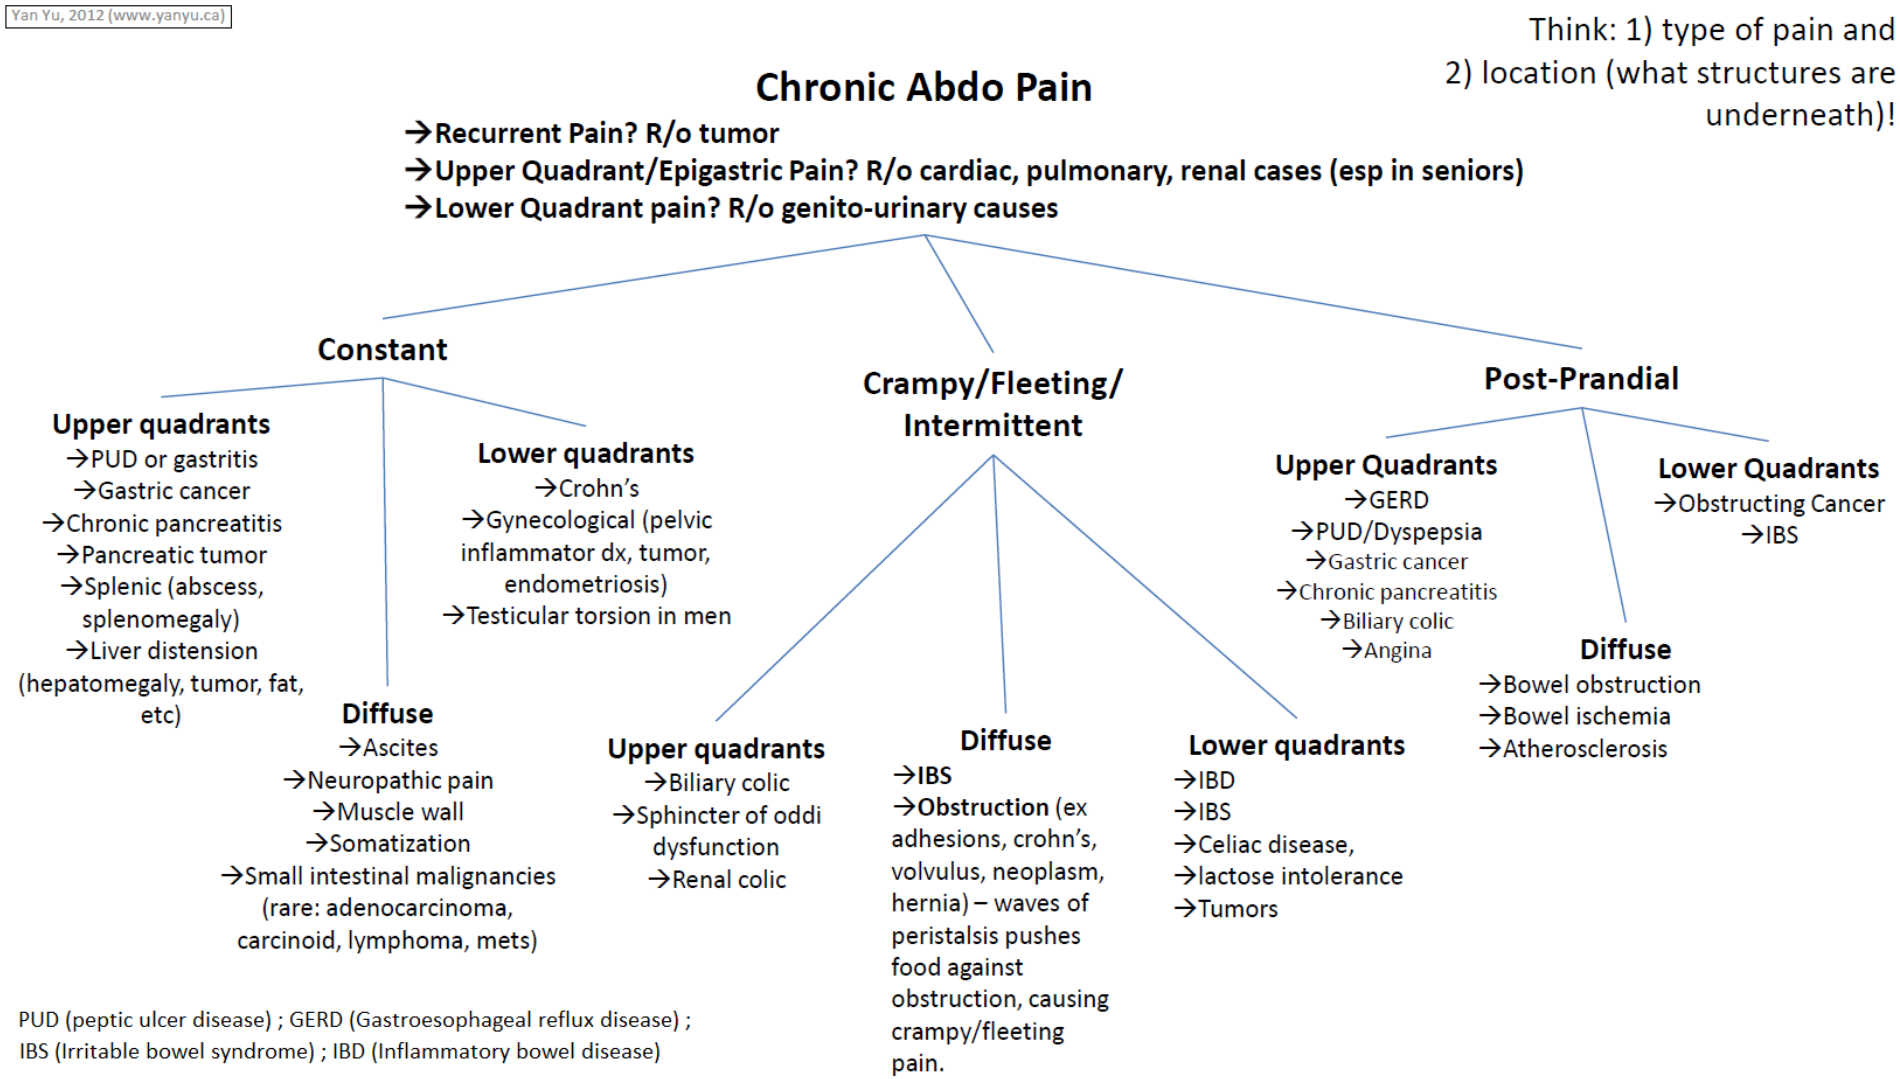

How does risk stratification differ from differential diagnosis?

How does risk stratification differ from differential diagnosis?

Differential diagnosis	Risk stratification
Usually iterative/active	Usually passive
Often considers a large set of conditions	Often just one condition
Has to consider rare conditions (needs hybrid knowledge/ML approaches)	Often focuses on settings where there is enough training data

Old vs. New

- Traditionally, risk stratification was based on simple scores using human-entered data

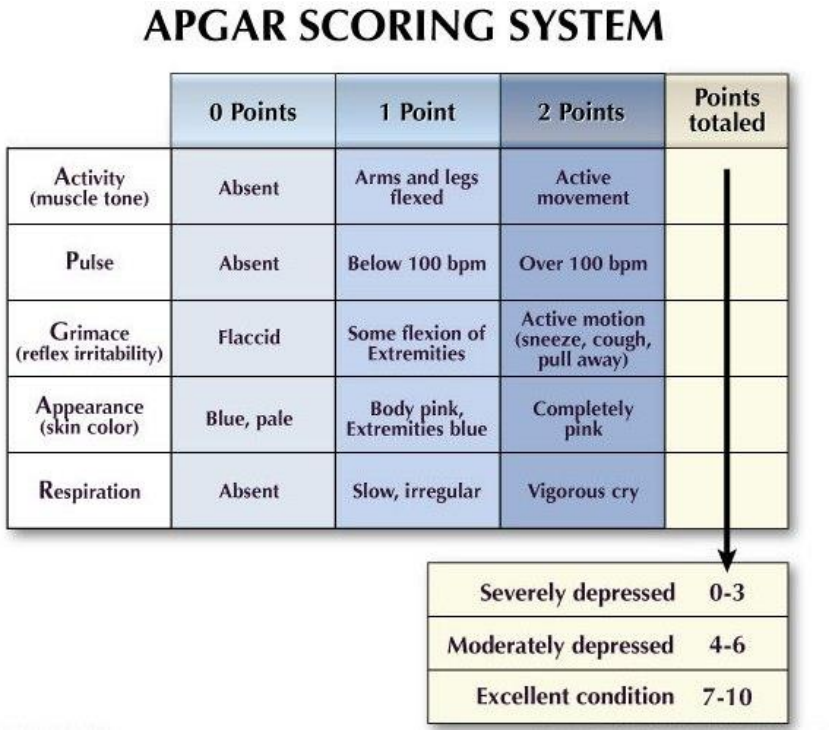

APGAR SCORING SYSTEM

	0 Points	1 Point	2 Points	Points totaled
Activity (muscle tone)	Absent	Arms and legs flexed	Active movement	
Pulse	Absent	Below 100 bpm	Over 100 bpm	
Grimace (reflex irritability)	Flaccid	Some flexion of Extremities	Active motion (sneeze, cough, pull away)	
Appearance (skin color)	Blue, pale	Body pink, Extremities blue	Completely pink	
Respiration	Absent	Slow, irregular	Vigorous cry	

Severely depressed	0-3
Moderately depressed	4-6
Excellent condition	7-10

Old vs. New

- Traditionally, risk stratification was based on simple scores using human-entered data
- Now, based on machine learning on high-dimensional data
  - Fits more easily into workflow
  - Higher accuracy
  - Quicker to derive (can special case)
- **But, ML approach comes with new challenges – to be discussed**

So, what do we need?

- Specification of prediction time / index date
- A way of encoding the data we have on the patient
  - CNN for images
  - Bag of words for text document
  - Longitudinal structured data…?
- A target, typically derived from the EHR
- Choice of appropriate supervised ML algorithm
  - Regression? Classification?

Outline for today's class

1. Introduction to risk stratification
2. **Case study: Early detection of Type 2 diabetes**
   - Encoding longitudinal structured health data
3. Framing as supervised learning problem
   - Deriving labels from EHR

[Razavian, Blecker, Schmidt, Smith-McLallen, Nigam, Sontag. *Big Data.* '16]

Type 2 Diabetes: A Major public health challenge

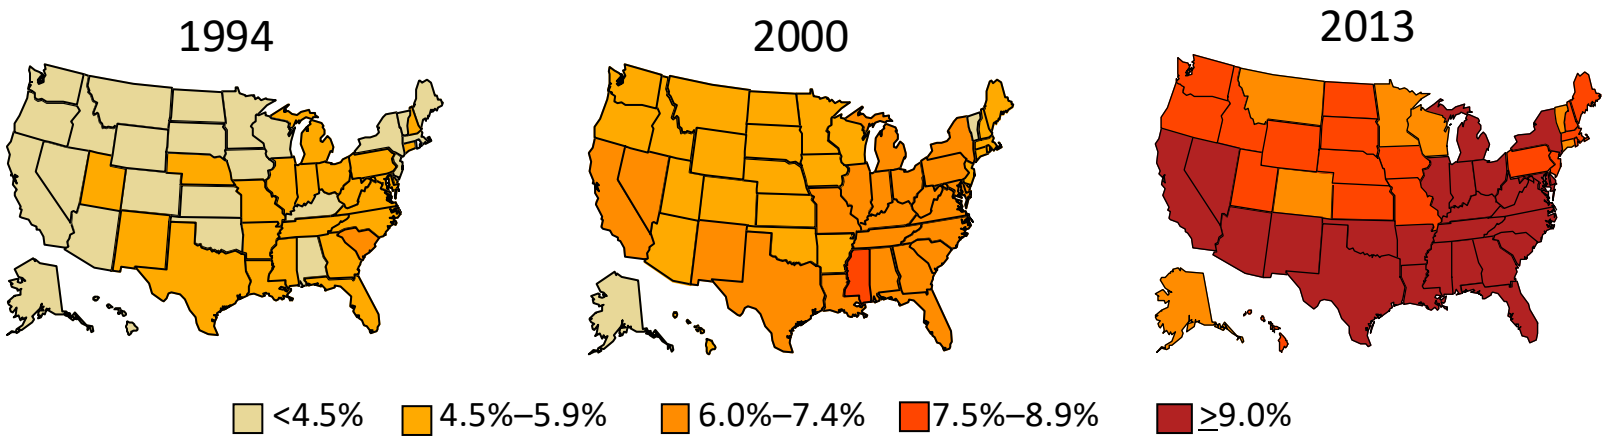

$245 billion: Total costs of diagnosed diabetes in the United States in 2012

$831 billion: Total fiscal year federal budget for healthcare in the United States in 2014

Type 2 Diabetes Can Be Prevented *

Requirement for successful large scale prevention program

1. Detect/reach truly at risk population
2. Improve the interventions
3. Lower the cost of intervention

* Diabetes Prevention Program Research Group. "Reduction in the incidence of type 2 diabetes with lifestyle intervention or metformin." The New England journal of medicine 346.6 (2002): 393.

Traditional Risk Prediction Models

- Successful Examples
  - ARIC
  - KORA
  - FRAMINGHAM
  - AUSDRISC
  - FINDRISC
  - San Antonio Model
- Easy to ask/measure in the office, or for patients to do online
- Simple model: can calculate scores by hand

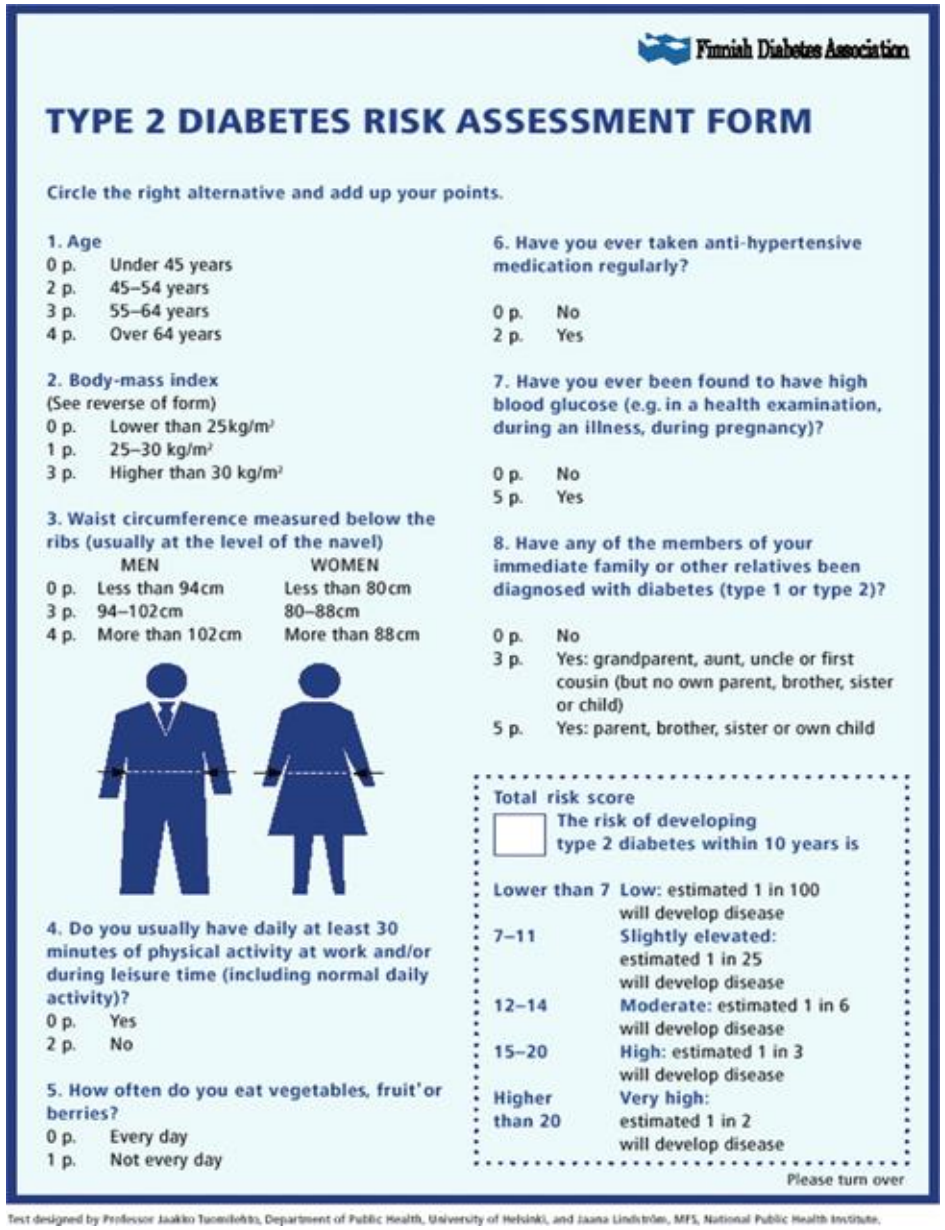

Finnish Diabetes Association

TYPE 2 DIABETES RISK ASSESSMENT FORM

Circle the right alternative and add up your points.

1. Age
0 p. Under 45 years
2 p. 45–54 years
3 p. 55–64 years
4 p. Over 64 years

2. Body-mass index
(See reverse of form)
0 p. Lower than 25kg/m²
1 p. 25–30 kg/m²
3 p. Higher than 30 kg/m²

3. Waist circumference measured below the ribs (usually at the level of the navel)

	MEN	WOMEN
0 p.	Less than 94cm	Less than 80cm
3 p.	94–102cm	80–88cm
4 p.	More than 102cm	More than 88cm

4. Do you usually have daily at least 30 minutes of physical activity at work and/or during leisure time (including normal daily activity)?
0 p. Yes
2 p. No

5. How often do you eat vegetables, fruit or berries?
0 p. Every day
1 p. Not every day

6. Have you ever taken anti-hypertensive medication regularly?
0 p. No
2 p. Yes

7. Have you ever been found to have high blood glucose (e.g. in a health examination, during an illness, during pregnancy)?
0 p. No
5 p. Yes

8. Have any of the members of your immediate family or other relatives been diagnosed with diabetes (type 1 or type 2)?
0 p. No
3 p. Yes: grandparent, aunt, uncle or first cousin (but no own parent, brother, sister or child)
5 p. Yes: parent, brother, sister or own child

Total risk score
The risk of developing type 2 diabetes within 10 years is

Lower than 7	Low: estimated 1 in 100 will develop disease
7–11	Slightly elevated: estimated 1 in 25 will develop disease
12–14	Moderate: estimated 1 in 6 will develop disease
15–20	High: estimated 1 in 3 will develop disease
Higher than 20	Very high: estimated 1 in 2 will develop disease

Please turn over

Test designed by Professor Jaakko Tuomilehto, Department of Public Health, University of Helsinki, and Jaana Lindström, MFS, National Public Health Institute.

Population-Level Risk Stratification

- Key idea: Use readily available administrative, utilization, and clinical data

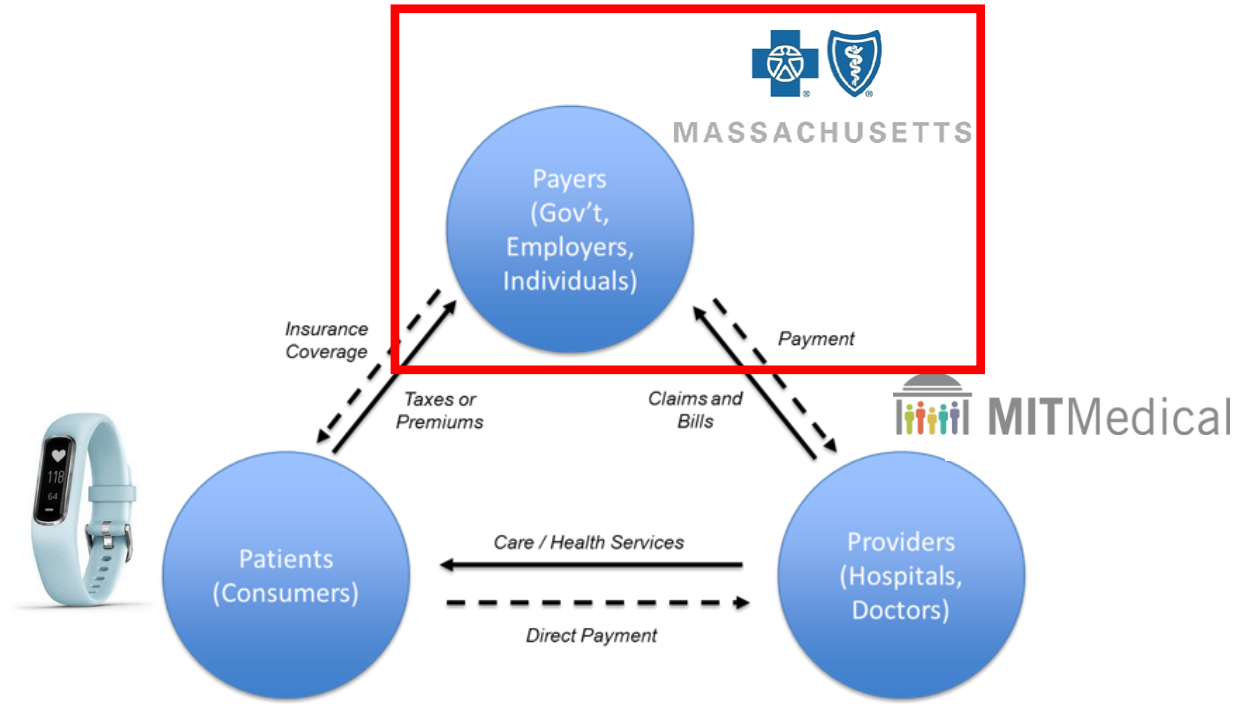

Source for figure: http://www.mahesh-vc.com/blog/understanding-whos-paying-for-what-in-the-healthcare-industry

Population-Level Risk Stratification

- Key idea: Use readily available administrative, utilization, and clinical data

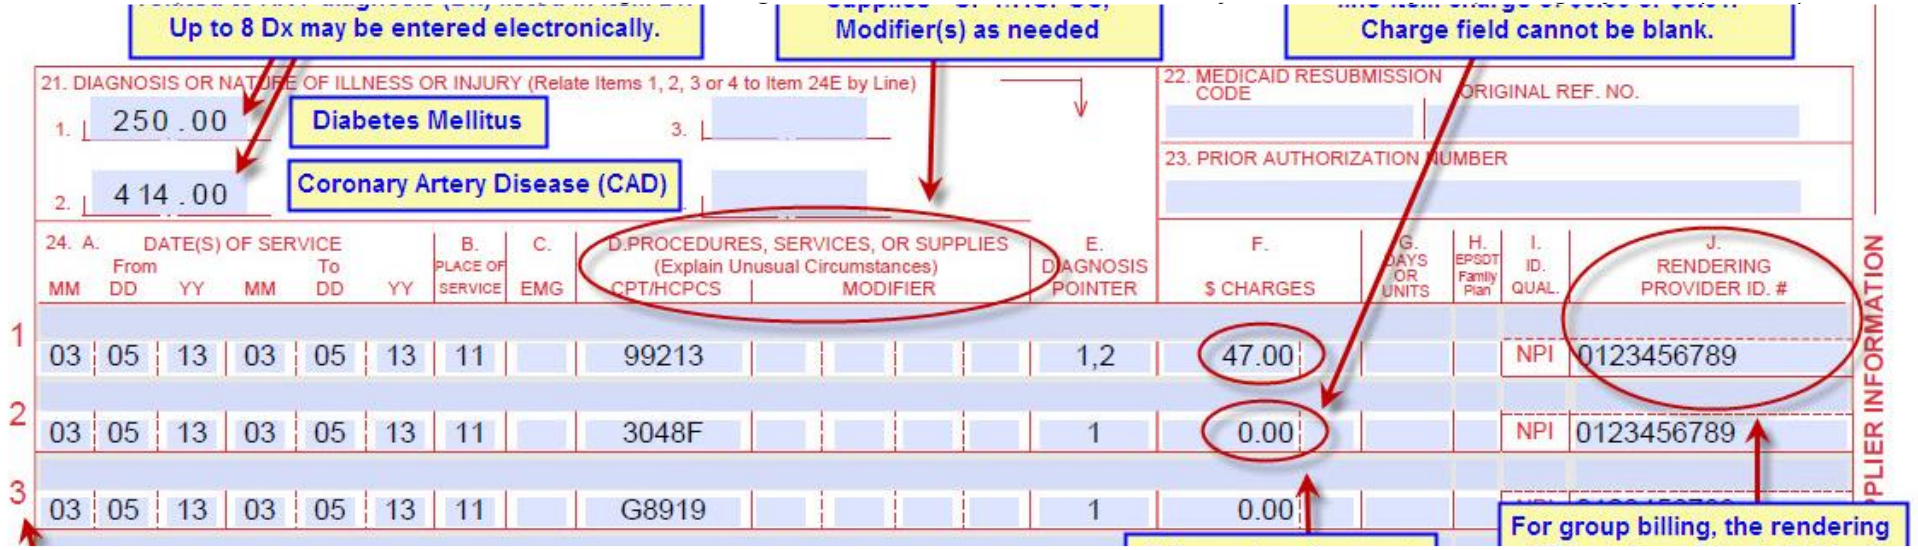

Population-Level Risk Stratification

- Key idea: Use readily available administrative, utilization, and clinical data
- Machine learning will find surrogates for risk factors that would otherwise be missing
- Perform risk stratification at the population level – millions of patients

A Data-Driven approach on Longitudinal Data

- Looking at individuals who got diabetes *today,* (compared to those who didn't)
  - Can we infer which variables in their record could have predicted their health outcome?

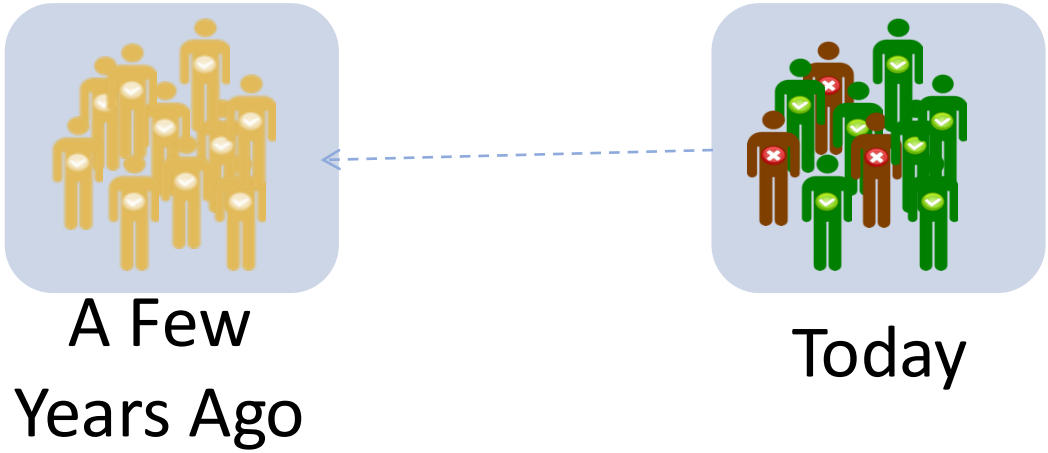

Administrative & Clinical Data

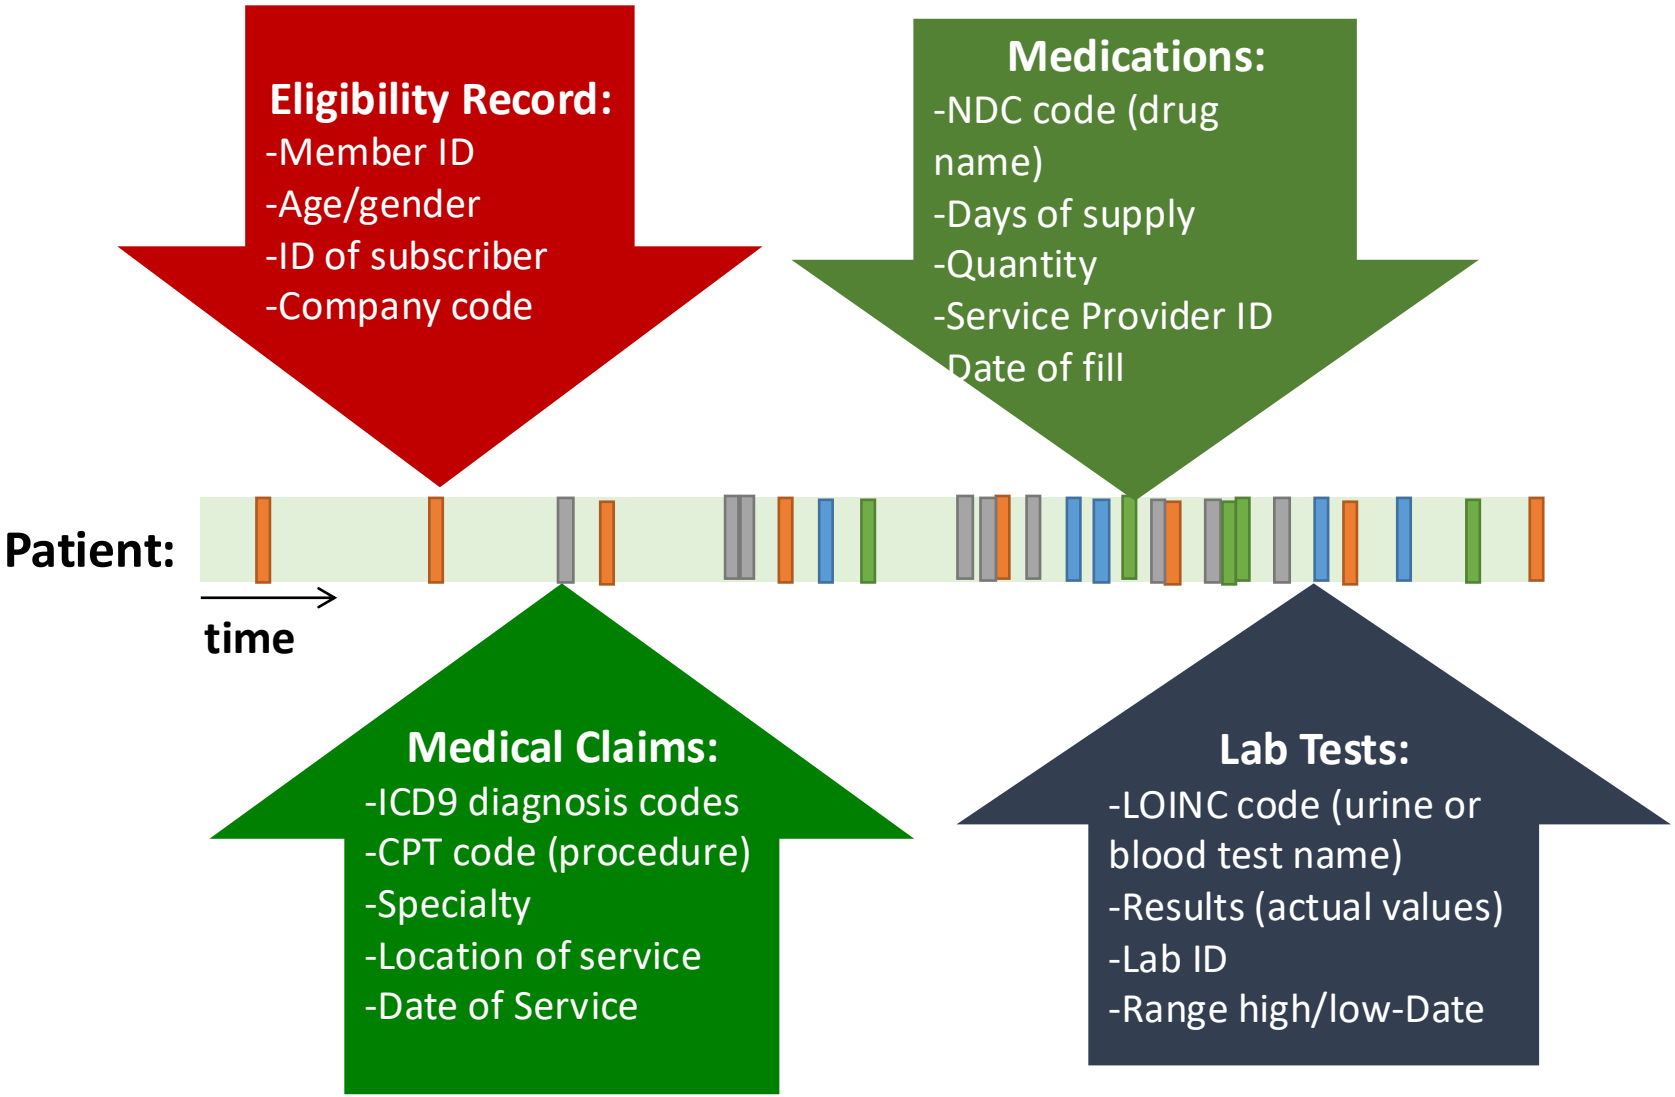

Top diagnosis codes

Disease	count
401.1 Benign hypertension	447017
272.4 Hyperlipidemia NEC/NOS	382030
401.9 Hypertension NOS	372477
250.00 DMII wo cmp nt st uncntr	339522
272.0 Pure hypercholesterolem	232671
272.2 Mixed hyperlipidemia	180015
V72.31 Routine gyn examination	178709
244.9 Hypothyroidism NOS	169829
780.79 Malaise and fatigue NEC	149797
V04.81 Vaccin for influenza	147858
724.2 Lumbago	137345
V76.12 Screen mammogram NEC	129445
V70.0 Routine medical exam	127848

Disease	count
530.81 Esophageal reflux	121064
427.31 Atrial fibrillation	113798
729.5 Pain in limb	112449
414.01 Crnry athrscl natve vssl	104478
285.9 Anemia NOS	103351
786.50 Chest pain NOS	91999
599.0 Urin tract infection NOS	87982
V58.69 Long-term use meds NEC	85544
496 Chr airway obstruct NEC	78585
477.9 Allergic rhinitis NOS	77963
414.00 Cor ath unsp vsl ntv/gft	75519

Disease	count
719.47 Joint pain-ankle	28648
300.4 Dysthymic disorder	28530
268.9 Vitamin D deficiency NOS	28455
V72.81 Preop cardiovsclr exam	27897
724.3 Sciatica	27604
787.91 Diarrhea	27424
V2.21 Supervis oth normal preg	27320
365.01 Opn angl brderln lo risk	26033
379.21 Vitreous degeneration	25592
424.1 Aortic valve disorder	25425
616.10 Vaginitis NOS	24736
702.19 Other sborheic keratosis	24453
380.4 Impacted cerumen	24046

Out of 135K patients who had laboratory data

Top lab test results

Lab test	
2160-0 Creatinine	1284737
3094-0 Urea nitrogen	1282344
2823-3 Potassium	1280812
2345-7 Glucose	1299897
1742-6 Alanine aminotransferase	1187809
1920-8 Aspartate aminotransferase	1187965
2885-2 Protein	1277338
1751-7 Albumin	1274166
2093-3 Cholesterol	1268269
2571-8 Triglyceride	1257751
13457-7 Cholesterol.in LDL	1241208
17861-6 Calcium	1165370
2951-2 Sodium	1167675

Lab test	
2085-9 Cholesterol.in HDL	1155666
718-7 Hemoglobin	1152726
4544-3 Hematocrit	1147893
9830-1 Cholesterol.total/Cholesterol.in HDL	1037730
33914-3 Glomerular filtration rate/1.73 sq M.predicted	561309
785-6 Erythrocyte mean corpuscular hemoglobin	1070832
6690-2 Leukocytes	1062980
789-8 Erythrocytes	1062445
787-2 Erythrocyte mean corpuscular volume	1063665

Lab test	
770-8 Neutrophils/100 leukocytes	952089
731-0 Lymphocytes	943918
704-7 Basophils	863448
711-2 Eosinophils	935710
5905-5 Monocytes/100 leukocytes	943764
706-2 Basophils/100 leukocytes	863435
751-8 Neutrophils	943232
742-7 Monocytes	942978
713-8 Eosinophils/100 leukocytes	933929
3016-3 Thyrotropin	891807
4548-4 Hemoglobin A1c/Hemoglobin.total	527062

Count of people who have the test result (ever)

There may be a varying amount of history per patient

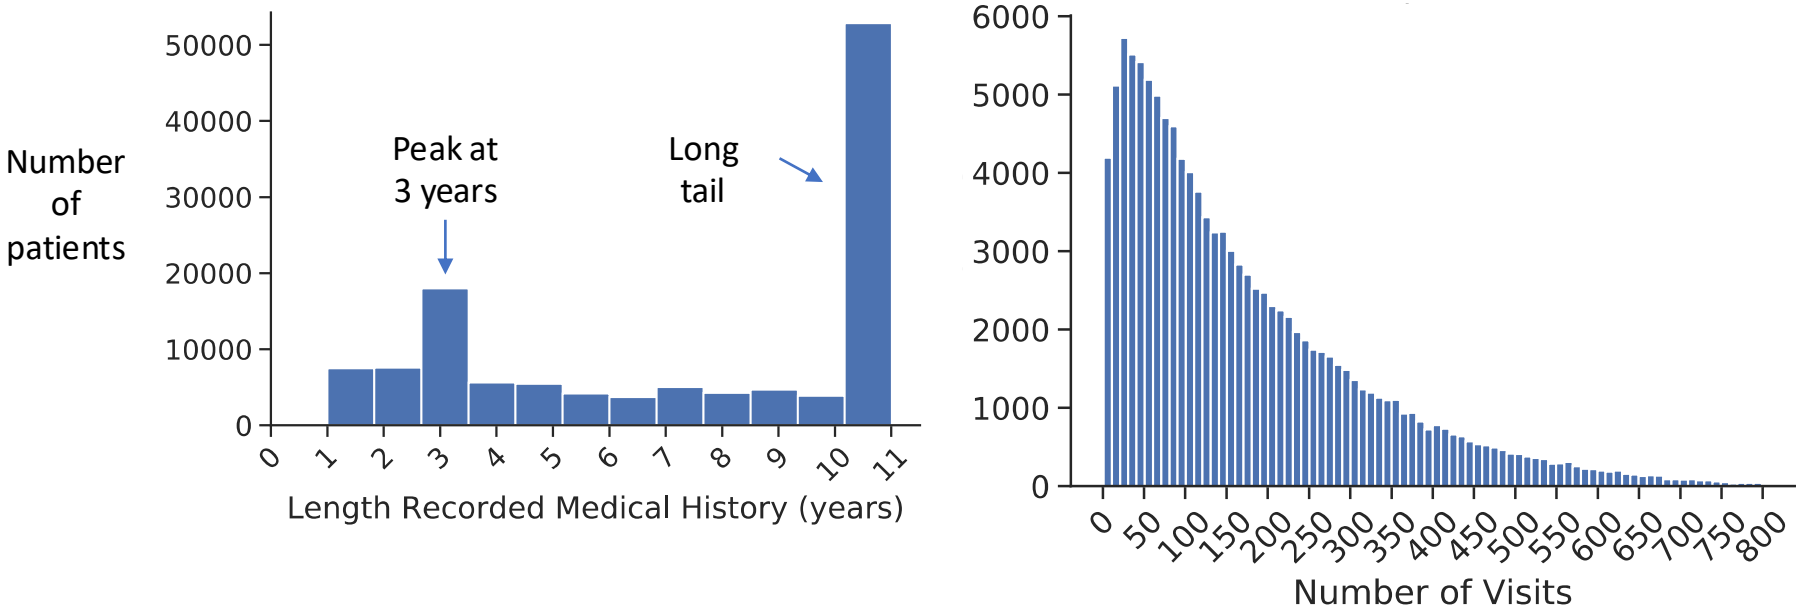

Kodialam et al., *Deep Contextual Clinical Prediction with Reverse Distillation*, AAAI 2021

Encoding the longitudinal health data

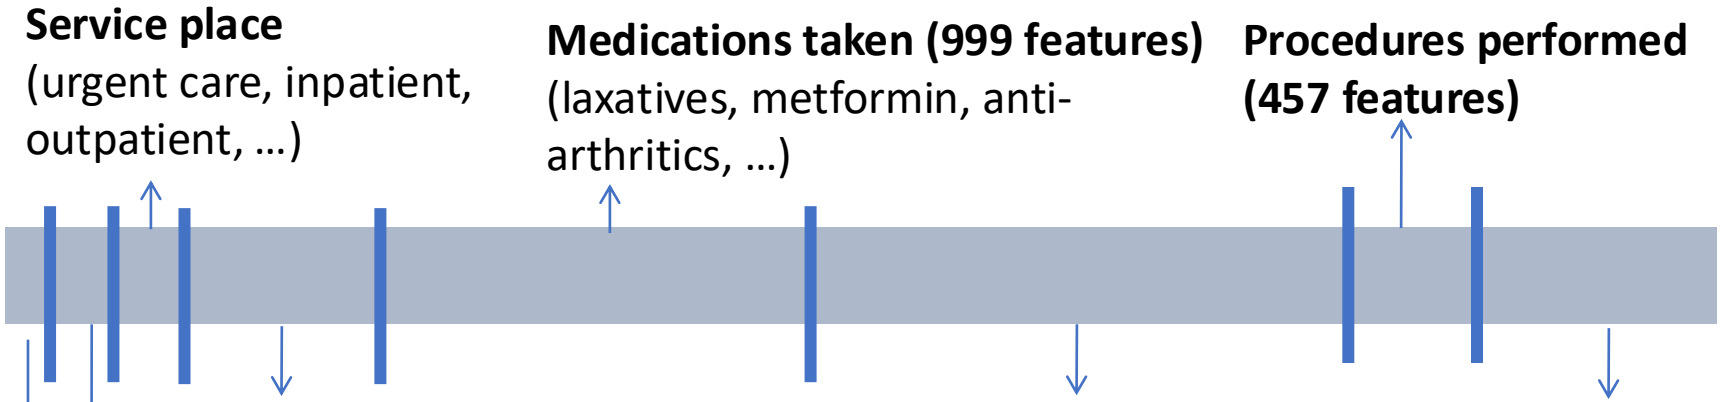

For the 1000 most frequent lab tests:

- Was the test ever administered?
- Was the result ever low?
- Was the result ever high?
- Was the result ever normal?
- Is the value increasing?
- Is the value decreasing?
- Is the value fluctuating?

Encoding the longitudinal health data

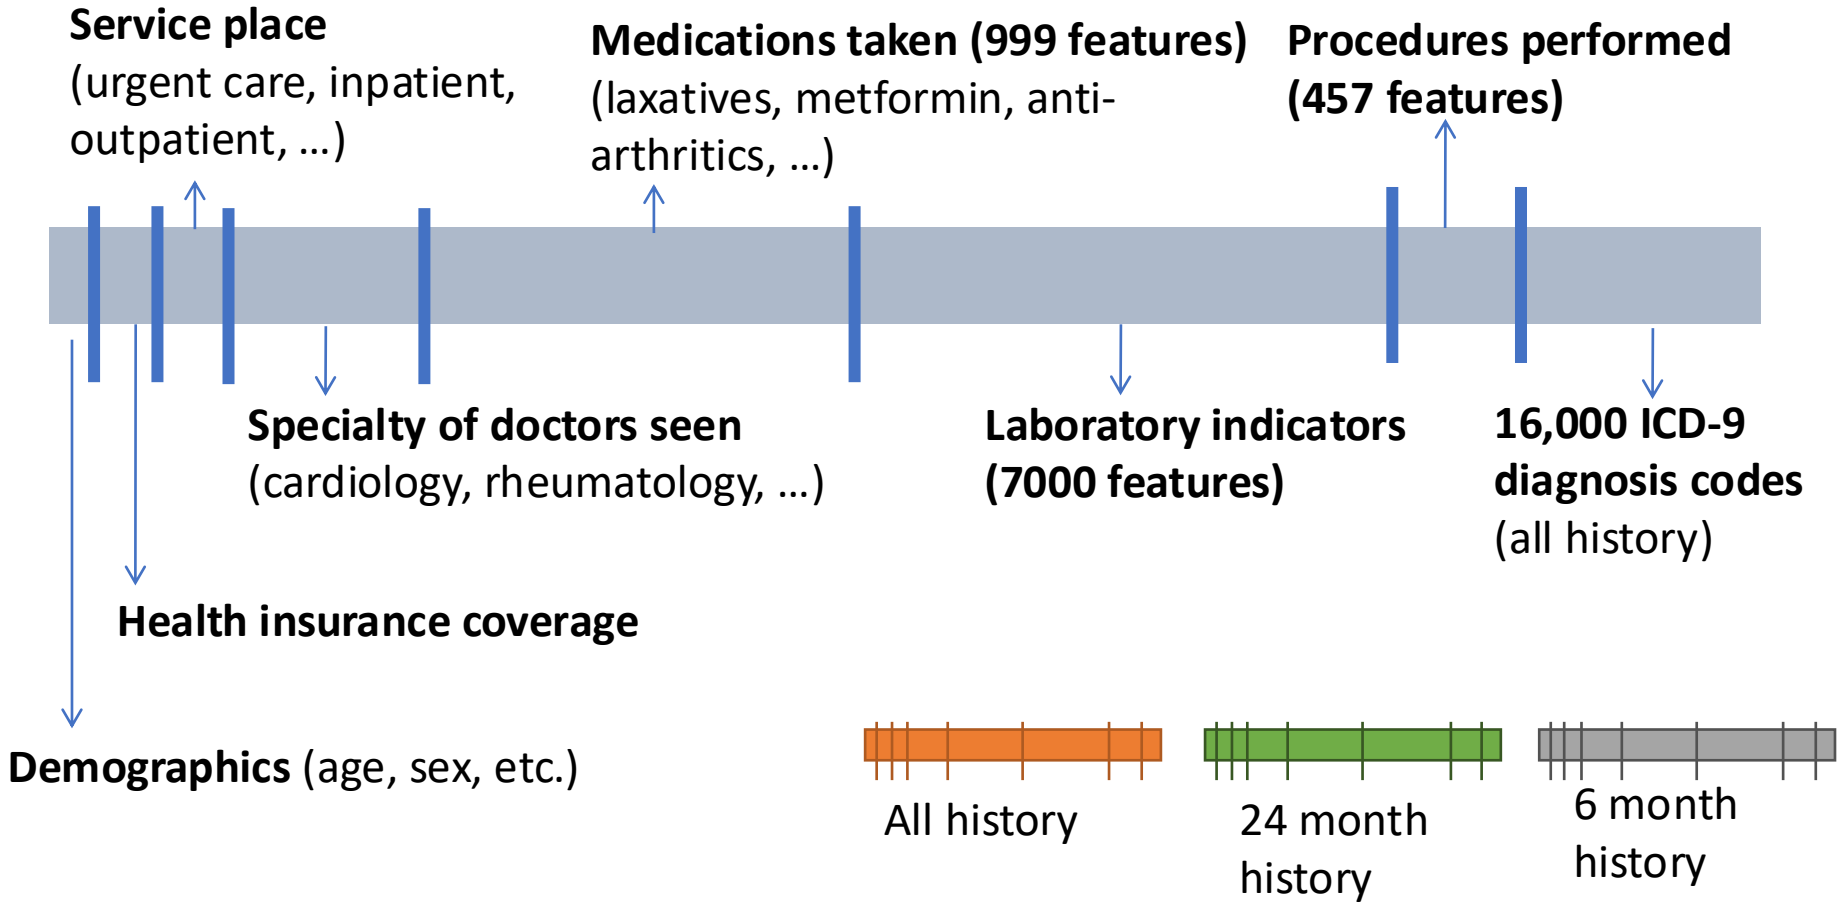

10s-100s of thousands of features

Encoding the longitudinal health data

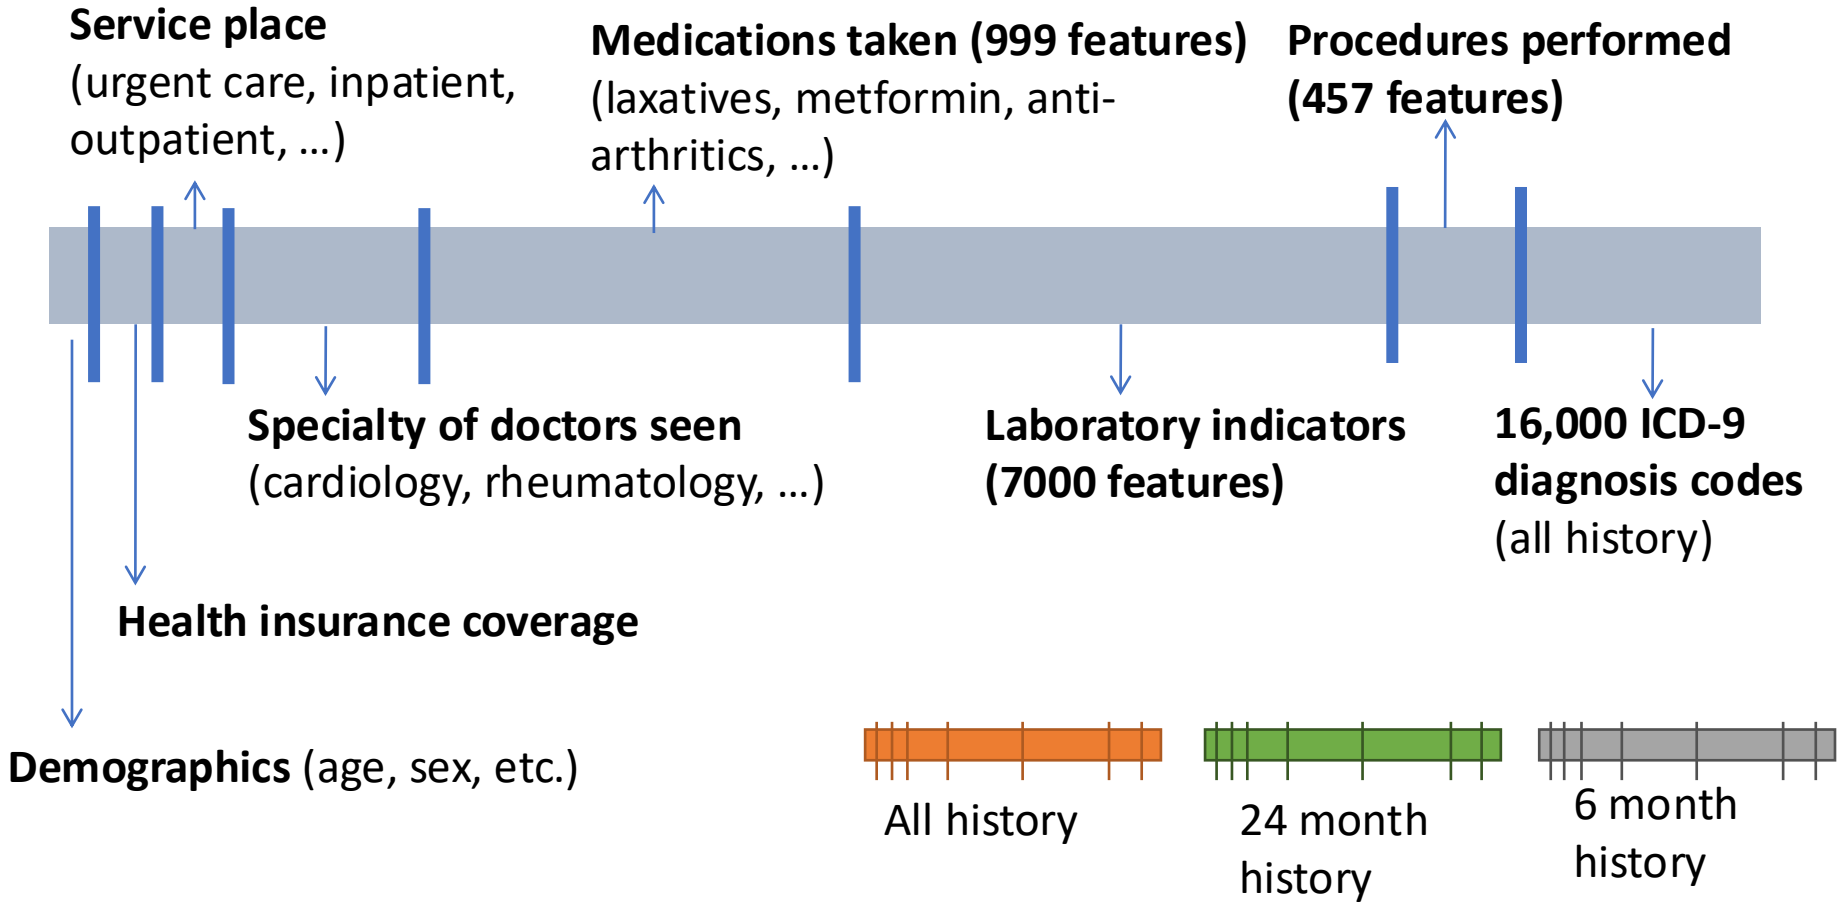

How does this deal with missing data? What are its limitations?

Outline for today's class

1. Introduction to risk stratification
2. Case study: Early detection of Type 2 diabetes
   - Encoding longitudinal structured health data
3. Framing as supervised learning problem
   - **Deriving labels from EHR**
   - L1-regularized logistic regression

Where do the labels come from?

Typical pipeline:

1. Manually label several patients' data by "chart review"
2. A) Come up with a simple rule to automatically derive label for all patients, **or**

   B) Use machine learning to get the labels themselves

Step 2: Example of a rule-based phenotype

Figure 1: Algorithm for identifying T2DM cases in the EMR.

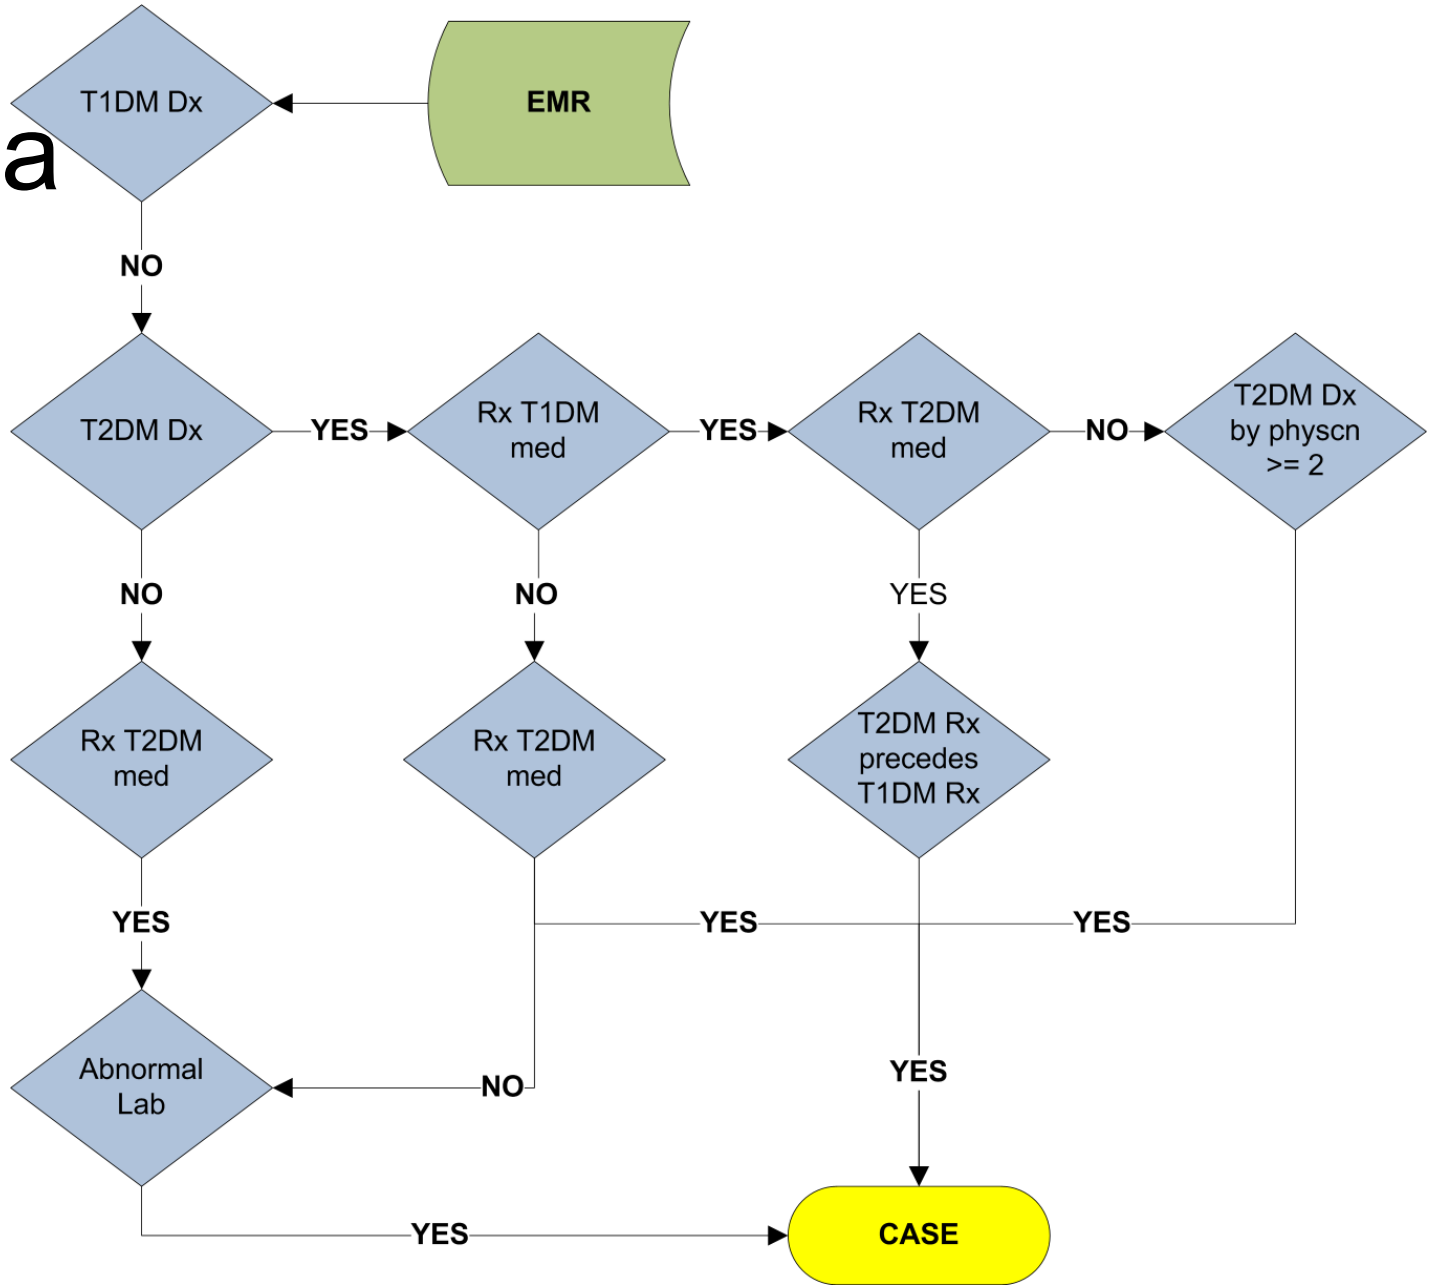

Source: https://phekb.org/sites/phenotype/files/T2DM-algorithm.pdf

Step 2: Example of a rule-based phenotype

Coverage of Different Diabetes Outcome Definitions on Claims Data

Condition	Percentage
Have 250.x diagnosis, or have been on diabetic medication, or have any HbA1c ≥ 6.5	**100 %**
Have been diagnosed 250.xx	**89.9 %**
Have been on diabetic medications	**15.0 %**
Have HbA1c values ≥ 6.5	**20.9 %**
Have 250.xx diagnosis on more than one distinct date	**40.0 %**
(Have 250.xx diagnosis, or have been on diabetic medication, or have any HbA1c ≥ 6.5) on more than one distinct date	**44.0 %**
(Have 250.xx diagnosis, or have been on diabetic medication, or have any HbA1c ≥ 6.5) on two dates separated by at least a week	**41.1 %**

The earliest date the rule triggers is defined as the date of diabetes diagnosis

Definition selected

[Razavian, Blecker, Schmidt, Smith-McLallen, Nigam, Sontag. *Big Data.* '16]

Repository of existing rule-based phenotypes

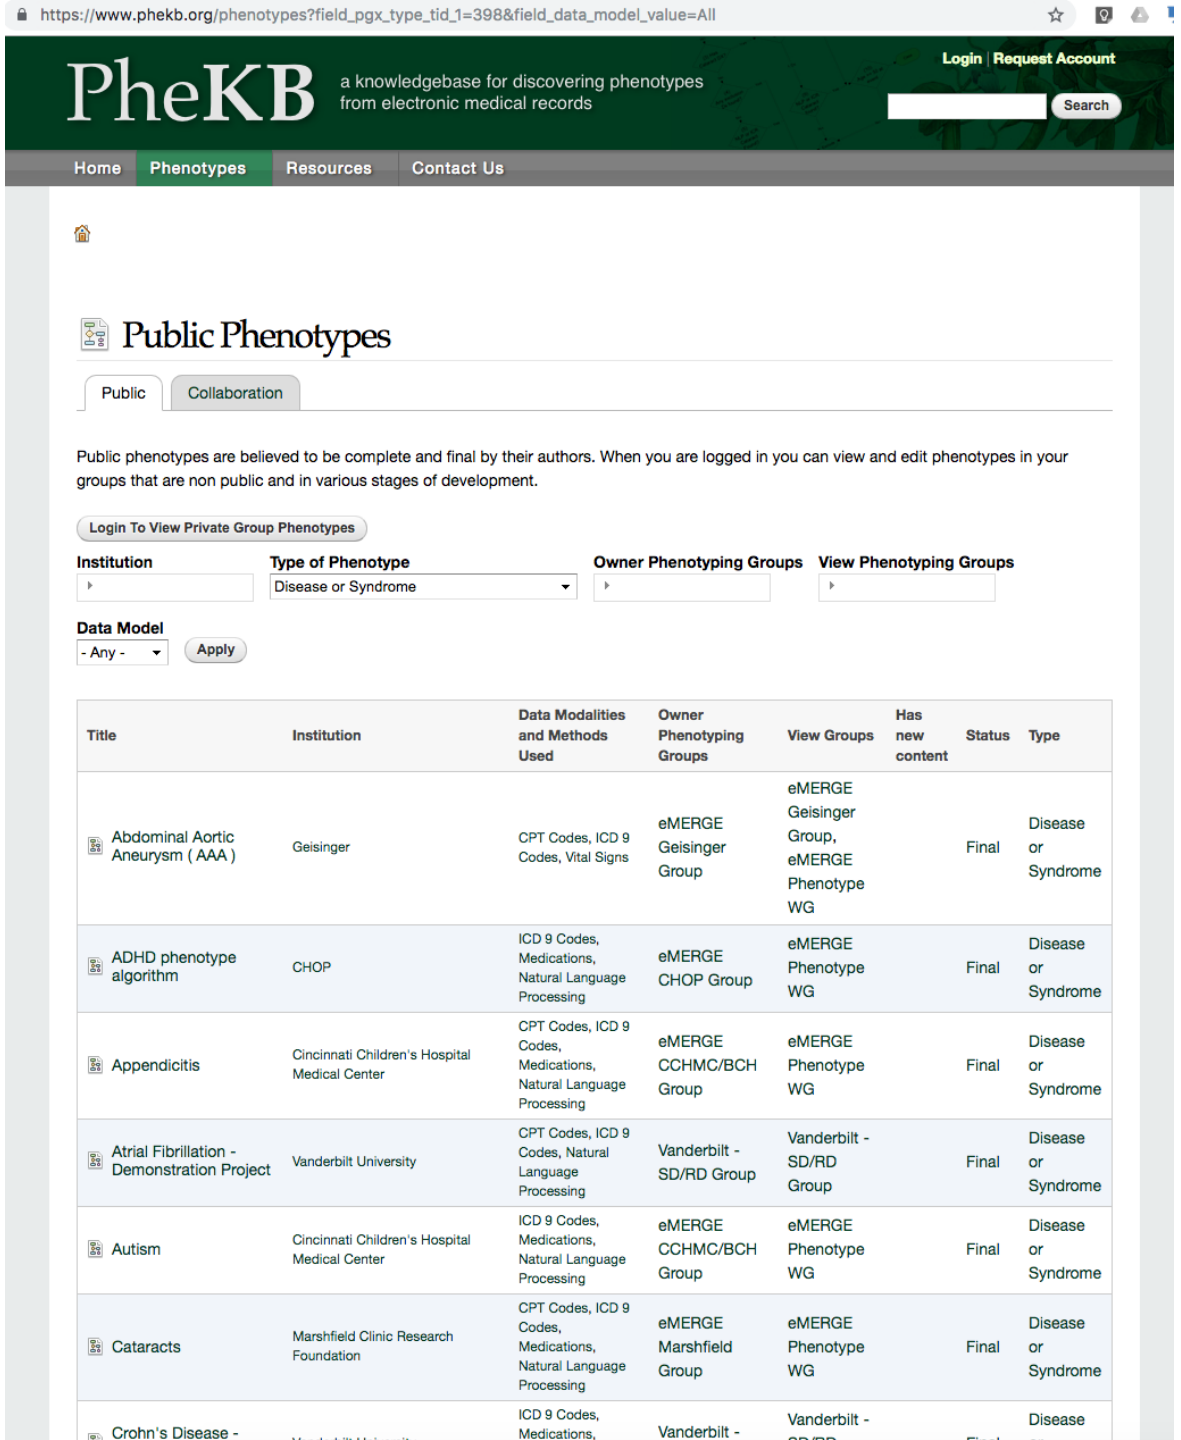

Framing for supervised machine learning

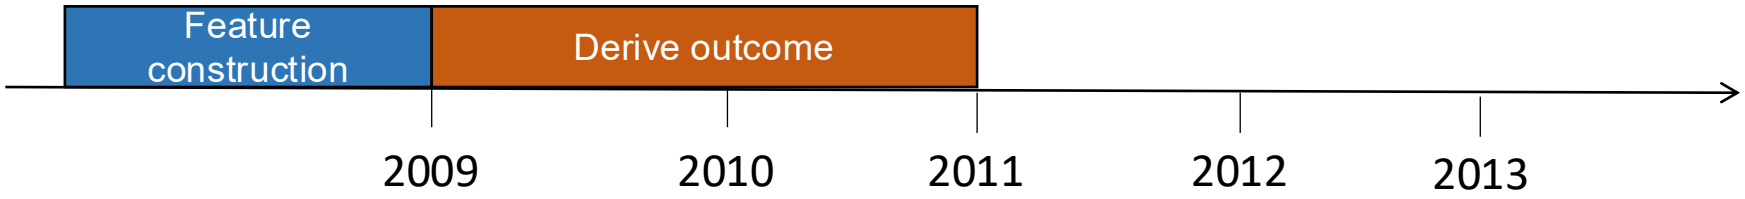

Exclusion criteria:

- Diabetes diagnosis (according to our rule) observed prior to January 1, 2009
- Less than 6 months of enrollment in feature construction window
- Member left health insurance prior to Jan. 1, 2011

What if someone is diagnosed with diabetes in 2012?

Why not model as “patient develops diabetes anytime after 2009”?

[Razavian, Blecker, Schmidt, Smith-McLallen, Nigam, Sontag. *Big Data.* ‘16]

Framing for supervised machine learning

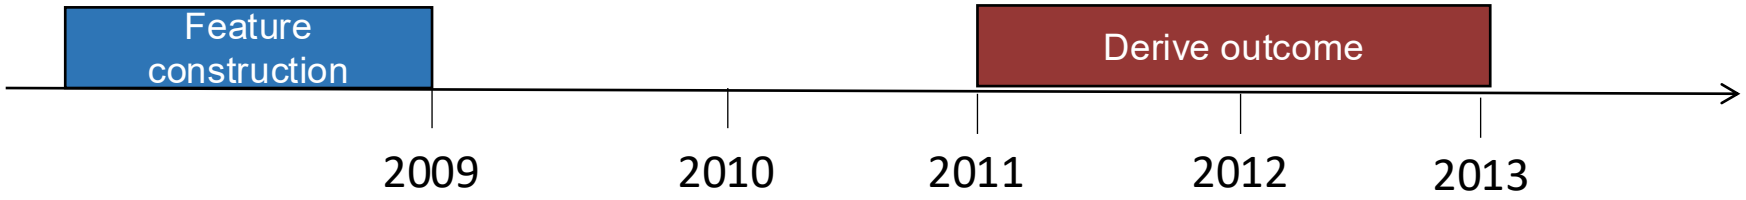

Exclusion criteria:

- Diabetes diagnosis (according to our rule) observed prior to January 1, ~~2009~~ 2011
- Less than 6 months of enrollment in feature construction window
- Member left health insurance prior to Jan. 1, ~~2011~~ 2013

[Razavian, Blecker, Schmidt, Smith-McLallen, Nigam, Sontag. *Big Data.* '16]

Framing for supervised machine learning

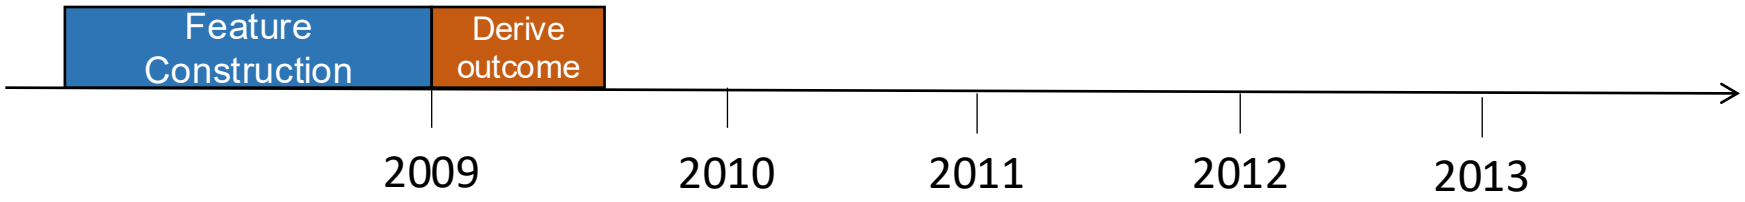

- Suppose we want to run the above model in August 2009. It may not have good performance due to *non-stationarity* in the data

- We now have data through 2021. Using a fixed prediction time / index date of Jan. 1, 2009 is ignoring most of the diabetes onsets!

Framing for supervised machine learning

- We can instead create ***many*** data points from each patient, using e.g. every month as an index date:

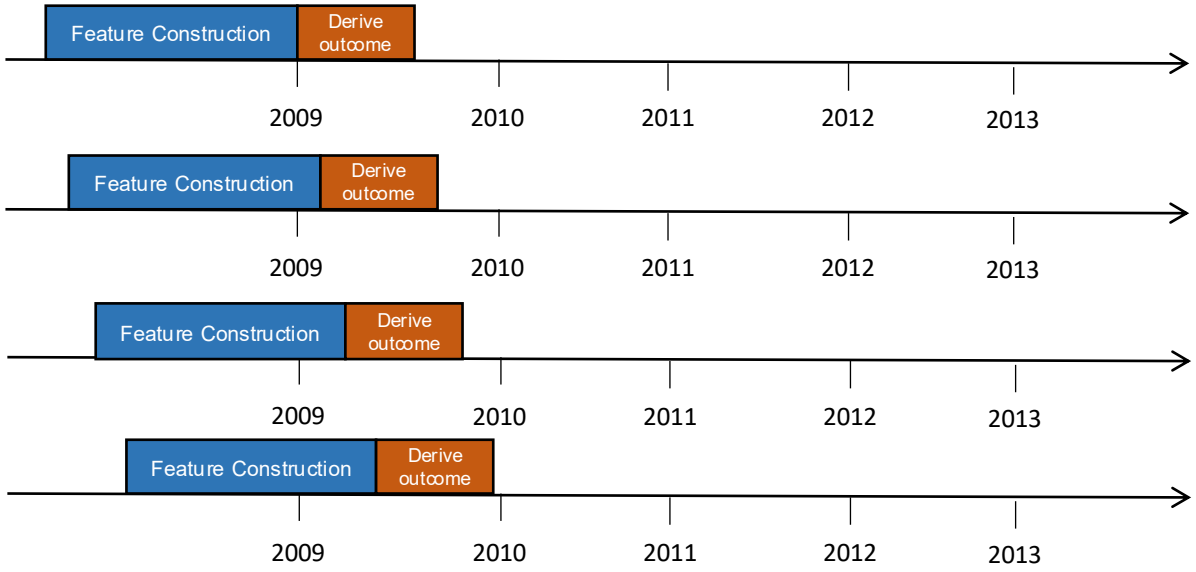

- **Important:** If multiple data points per patient, make sure each patient's data is in *only* train, validate, or test

Logistic regression with L1 regularization

- Penalizing the L1 norm of the weight vector leads to *sparse* (read: many 0's) solutions for *w.*

$$\min_w \sum_i \ell(x_i, y_i; w) + \lambda ||w||_1 \qquad ||\vec{w}||_1 = \sum_d |w_d|$$

instead of

$$\min_w \sum_i \ell(x_i, y_i; w) + \lambda ||w||_2^2 \qquad ||\vec{w}||_2^2 = \sum_d w_d^2$$

- Let's understand why…

Logistic regression with L1 regularization

- Penalizing the L1 norm of the weight vector leads to *sparse* (read: many 0's) solutions for *w.*

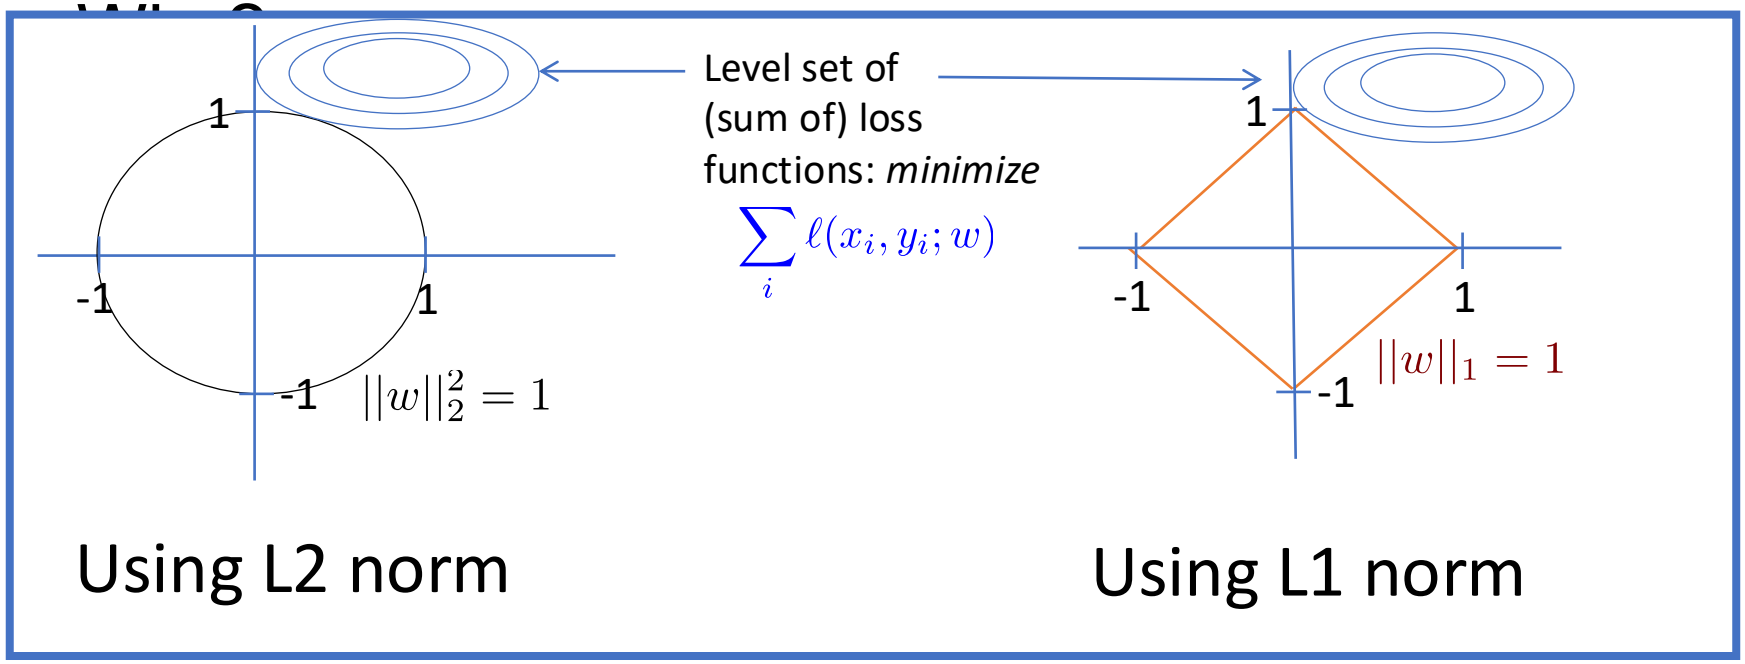

Big Data
Volume 3 Number 4, 2015
Mary Ann Liebert, Inc.
DOI: 10.1089/big.2015.0020

ORIGINAL ARTICLE

Population-Level Prediction of Type 2 Diabetes From Claims Data and Analysis of Risk Factors

Narges Razavian,[1] Saul Blecker,[2] Ann Marie Schmidt,[3] Aaron Smith-McLallen,[4] Somesh Nigam,[4] and David Sontag[1,*]

Sources of Clinical Data

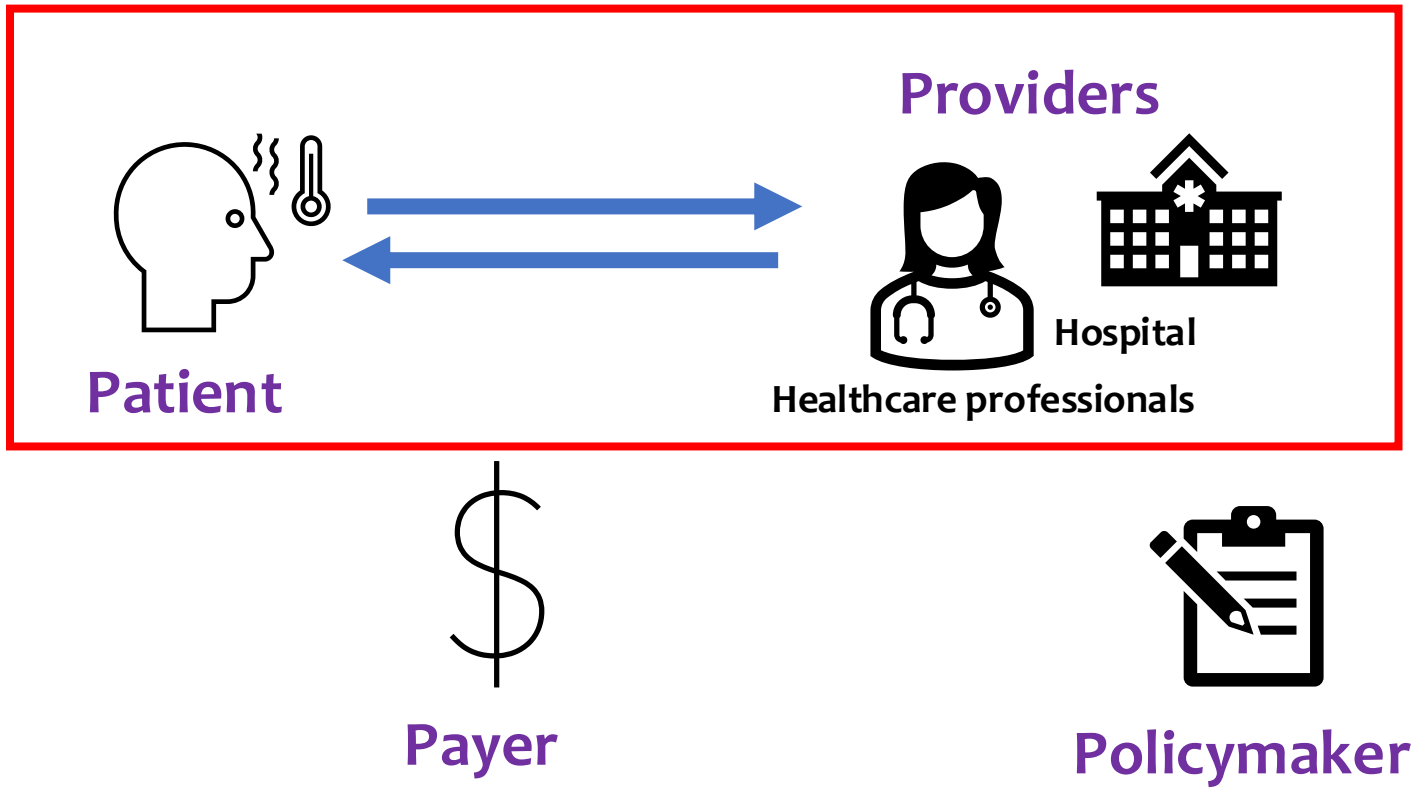

Electronic Health Records in the US

- Different hospitals use different EHR systems
  - Largest EHR systems in US
    1. EPIC
    2. Cerner
    3. Meditech

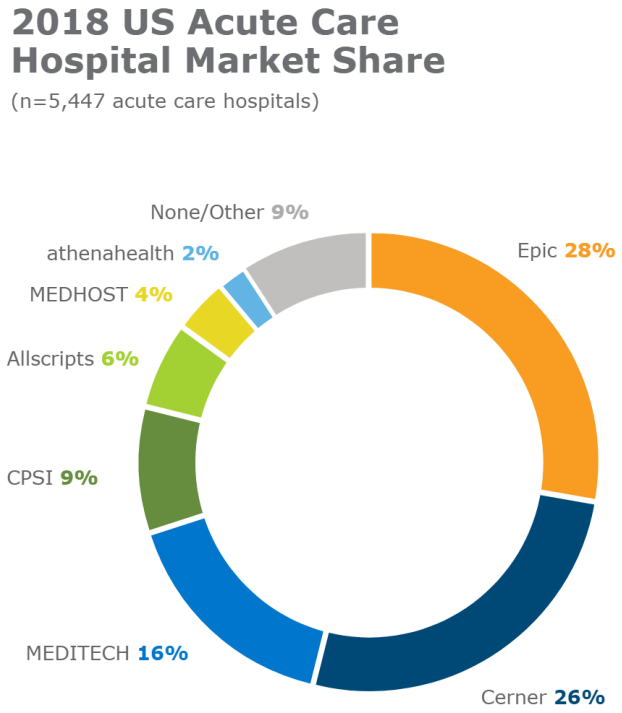

EPIC

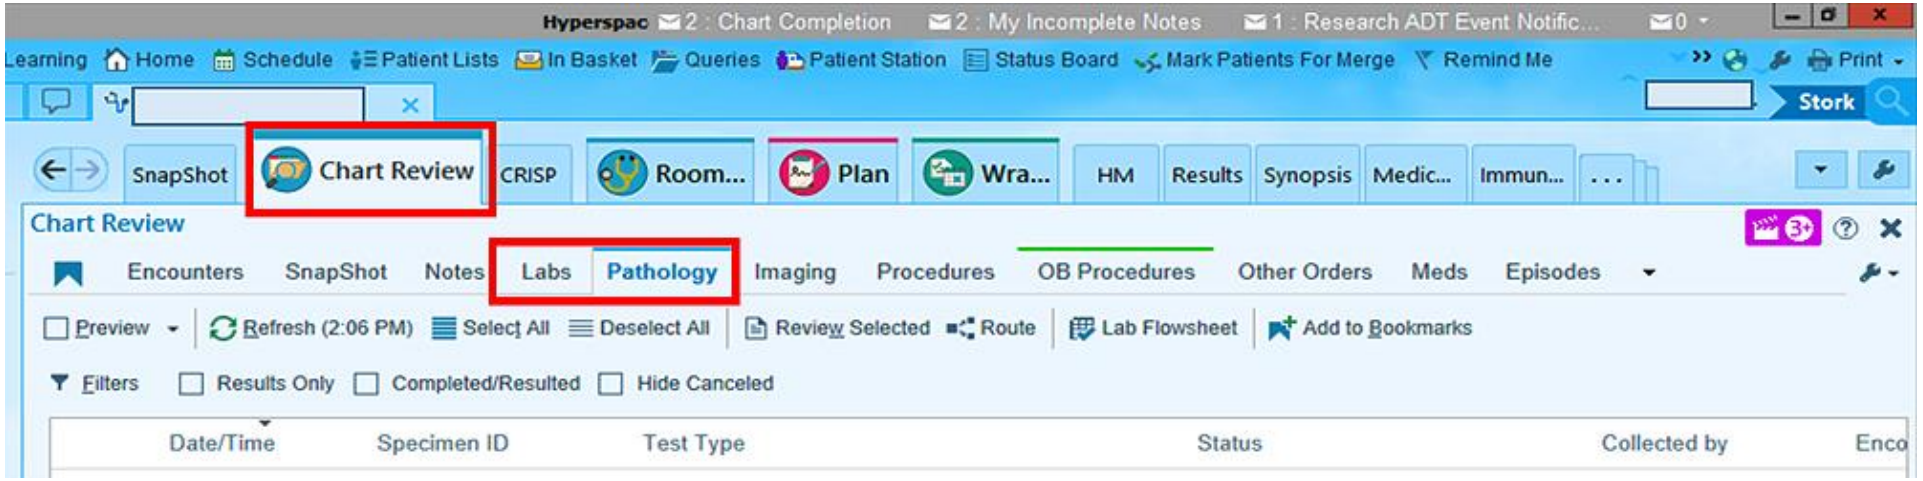

http://apps.pathology.jhu.edu/team-path-md/pathology-for-core-clinical-clerkships/how-to-find-pathology-results-and-reports-on-epic/

Types of data in healthcare

- History
  - Symptoms and their details, past medical/surgical history, medications, allergies, family history, etc.
- Physical exam
  - Height, weight, BMI, vital signs (temperature, blood pressure, heart rate, etc), tenderness, erythema (redness), etc.
- Labs
  - Complete blood count, serum electrolytes, urine culture, blood culture, etc.
- Imaging
  - Chest x-ray, CT scan, bone scan, MRI, ultrasound, etc.
- Pathology
  - Biopsy, surgical pathology
- Genetics
  - Germline testing, etc.

How do we evaluate models for risk stratification?

- Understand the population used for training
- Receiver-operator curve and precision-recall curve
- Interpreting linear models

“Table 1” – who did this study include?

Table 1. Subjects’ characteristics of the cohort included in training and validation

Characteristic	*Total population*	*Population with diabetes*
Average age (SD)	47.69 (17.1)	58.57 (13.3)
Female ratio	55%	51%
Average length of data in years (SD)	3.3 (1.0)	3.4 (1.0)
Hypertension (ICD9 401)	30.2%	62%
Hypercholesterolemia (ICD9 272.0)	18.7%	33.6%

SD, standard deviation.

But what about… Past hospitalizations? Number of years of historical data? Race/ethnicity?

[Razavian, Blecker, Schmidt, Smith-McLallen, Nigam, Sontag. *Big Data.* ‘16]

“Table 1”, better example

Table 1. Characteristics of 47 119 Hospitalized Patients

Characteristic	Finding[a]
Age, mean (SE), y	60.9 (18.15)
Female	23 952 (50.8)
Black/African American race	5258 (11.2)
Hispanic/Latino ethnicity	3667 (7.8)
Medicaid	8303 (17.6)
Heart failure in problem list	3630 (7.7)
Prior diagnosis of any heart failure	2985 (6.3)
Prior diagnosis of primary heart failure	615 (1.3)
Prior echocardiography	15 938 (33.8)
Loop diuretics	
Inpatient	6837 (14.5)
Outpatient	6427 (13.6)
ACE inhibitors or ARB	
Inpatient	13 166 (27.9)
Outpatient	14 797 (31.4)
β-Blockers	
Inpatient	19 748 (41.9)
Outpatient	14 870 (31.6)
Heart failure with β-blockers	
Inpatient	6310 (13.4)
Outpatient	8644 (18.4)
Blood pressure, mean (SE), mm Hg	
Systolic	123.3 (18.3)
Diastolic	67.8 (12.8)
Creatinine, mean (SE), mg/dL	1.01 (1.1)
Sodium, mean (SE), mEq/L	138.4 (3.7)
BNP, pg/mL	
<500	1721 (23.4)
500-999	878 (12.0)
1000-4999	2498 (34.0)
5000-9999	931 (12.7)
10 000-19 999	652 (8.9)
≥20 000	667 (9.1)
Blood pressure	
Any systolic	46 982 (99.7)
Any diastolic	46 982 (99.7)
Any creatinine	46 598 (98.9)
Any sodium	46 613 (98.9)
Any BNP	7347 (15.6)
Problem list	
Acute MI	952 (2.0)
Atherosclerosis	6147 (13.0)
Final discharge diagnosis of heart failure	
Any diagnosis	6549 (13.9)
Principal diagnosis	1214 (2.6)

[Blecker et al., Comparison of Approaches for Heart Failure Case Identification From Electronic Health Record Data, JAMA Cardiology 2016]

Receiver-operator characteristic curve

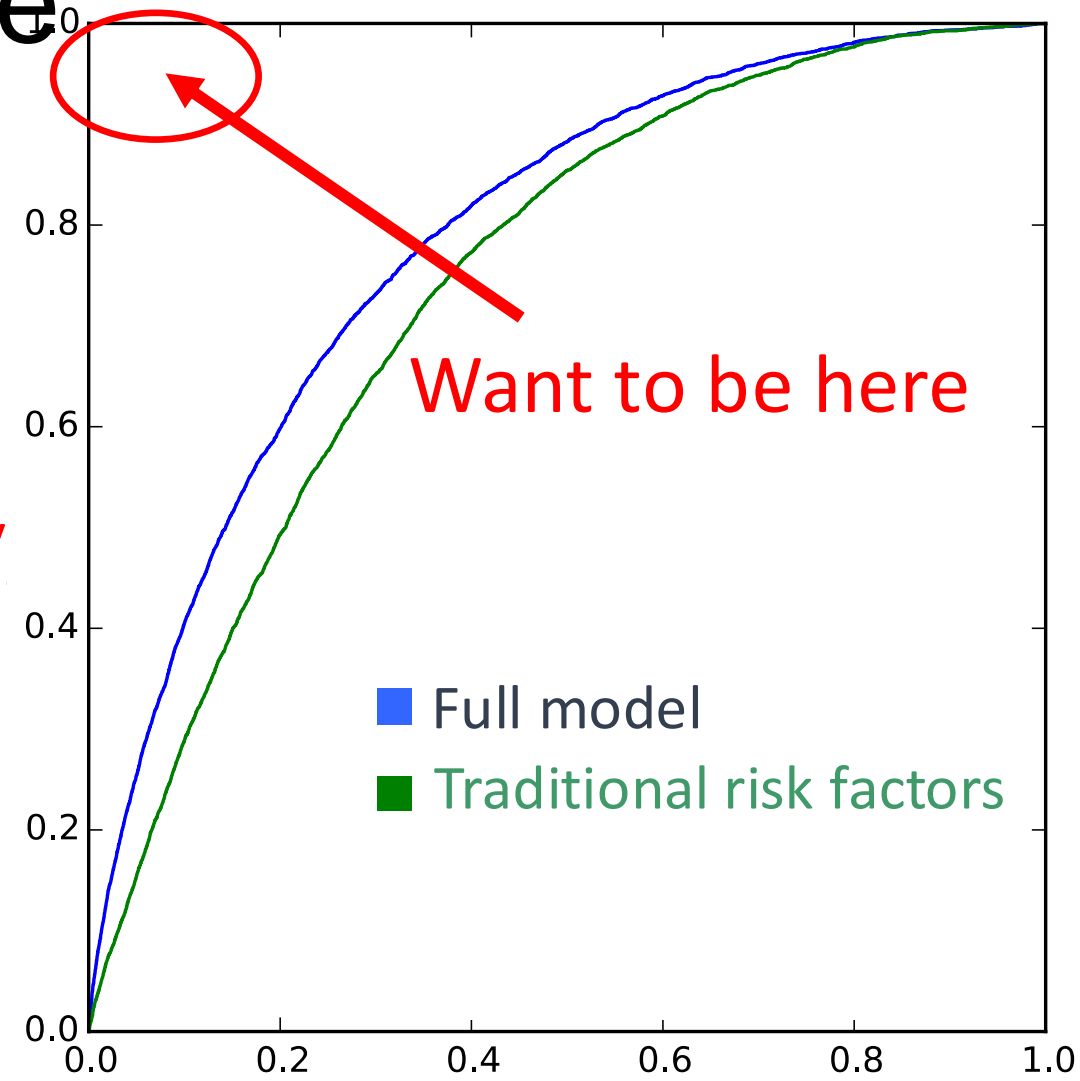

True positive rate (sensitivity or recall), *TP/P*

False positive rate (1-specificity), *FP/N*

Obtained by varying prediction threshold

Diabetes 1-year gap

Receiver-operator characteristic curve

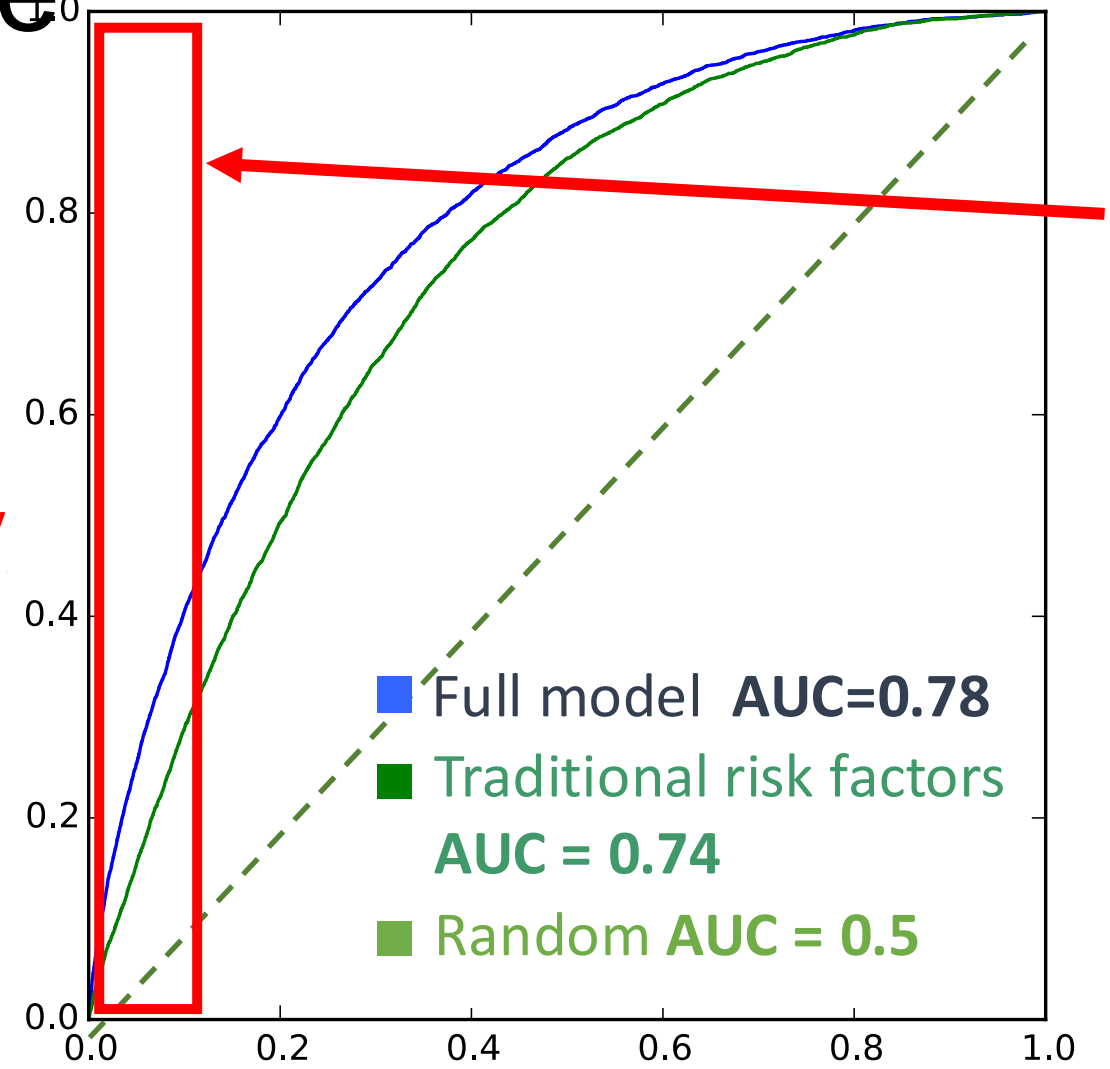

True positive rate (sensitivity or recall), *TP/P*

False positive rate (1-specificity), *FP/N*

Risk stratification usually focuses on just this region

(because of the cost of interventions)

Diabetes 1-year gap

Positive predictive value (PPV)

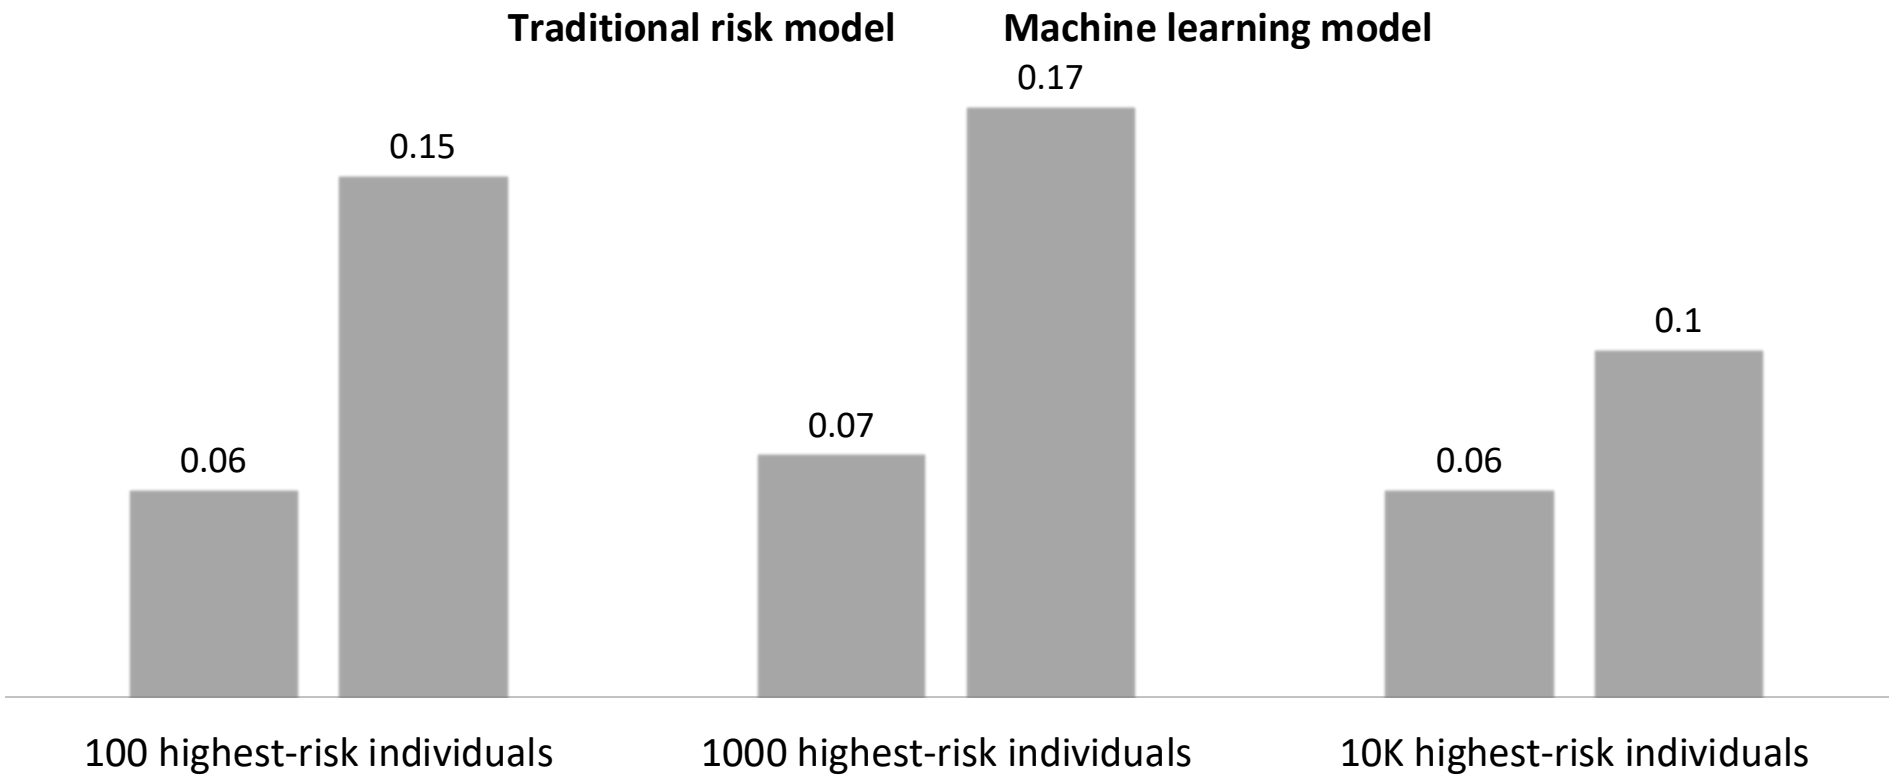

Predicting diabetes onset in 1-3 years

Precision-Recall Curve

- Instead of ROC, an alternative is the precision-recall curve
  - Particularly insightful with imbalanced datasets

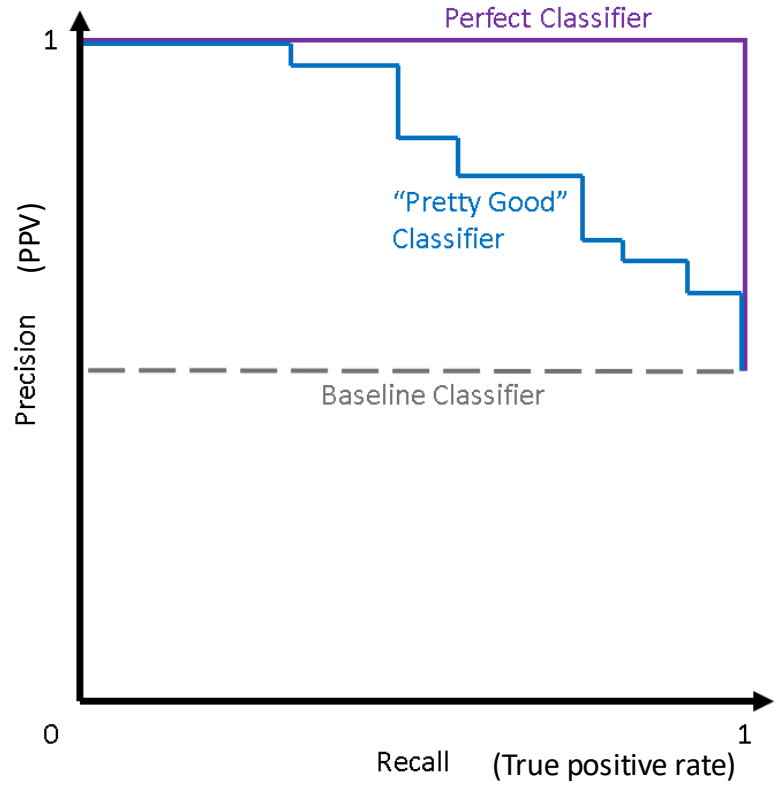

Figure from: https://medium.com/@douglaspsteen/precision-recall-curves-d32e5b290248

Logistic regression with L1 regularization

- 769 variables have non-zero weight. Look at most positive & most negative
- Positively weighted diagnosis codes include

  Pituitary dwarfism (253.3), Hepatomegaly(789.1), Chronic Hepatitis C (070.54), Hepatitis (573.3), Calcaneal Spur(726.73), Thyrotoxicosis without mention of goiter(242.90), Sinoatrial Node dysfunction(427.81), Acute frontal sinusitis (461.1), Hypertrophic and atrophic conditions of skin(701.9), Irregular menstruation(626.4), …

- Positively weighted laboratory features include

  Albumin/Globulin (Increasing -Entire history), Urea nitrogen/Creatinine -(high - Entire History), Specific gravity (Increasing, Past 2 years), Bilirubin (high -Past 2 years), …

[Razavian, Blecker, Schmidt, Smith-McLallen, Nigam, Sontag. *Big Data.* '16]

Interpreting high-dimensional linear models

- How do we interpret such high dimensional models…?
- A useful trick to build intuition: use higher value of λ (i.e., more regularization) than needed

$$\min_{w} \sum_{i} \ell(x_i, y_i; w) + \lambda ||w||_1$$

- What will the effect be?
- Intuition: often many predictive yet highly correlated features. Selects a representative set which still performs well

Features selected using model learned with more L1 regularization

History of Disease
Impaired Fasting Glucose (Code 790.21)
Abnormal Glucose NEC (790.29)
Hypertension (401)
Obstructive Sleep Apnea (327.23)
Obesity (278)
Abnormal Blood Chemistry (790.6)
Hyperlipidemia (272.4)
Shortness Of Breath (786.05)
Esophageal Reflux (530.81)

Top Lab Factors
Hemoglobin A1c /Hemoglobin.Total (High - past 2 years)
Glucose (High- Past 6 months)
Cholesterol.In VLDL (Increasing - Past 2 years)
Potassium (Low - Entire History)
Cholesterol.Total/Cholesterol.In HDL (High - Entire History)
Erythrocyte mean corpuscular hemoglobin concentration -(Low - Entire History)
Eosinophils (High - Entire History)
Glomerular filtration rate/1.73 sq M.Predicted (Low -Entire History)
Alanine aminotransferase (High Entire History)

[Razavian, Blecker, Schmidt, Smith-McLallen, Nigam, Sontag. *Big Data.* '16]

Interpreting high-dimensional linear models

- Model interpretation can be an important tool for identifying *errors* in your ML setup

- Illustrative examples:
  - Suppose a highly weighted positive feature is for “injection of aflibercept”, a treatment for diabetic macular edema. What could we infer?
  - Suppose we see many features for flu vaccines with high positive and negative weights. Looking up the NDC code for one of them, we see it is “influenza A virus A/Hong Kong/4801/2014 (H3N2) antigen 0.03 MG”. What could we infer?

Simple Models are Competitive in Healthcare Prediction

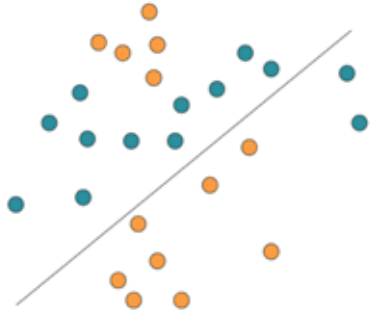

>

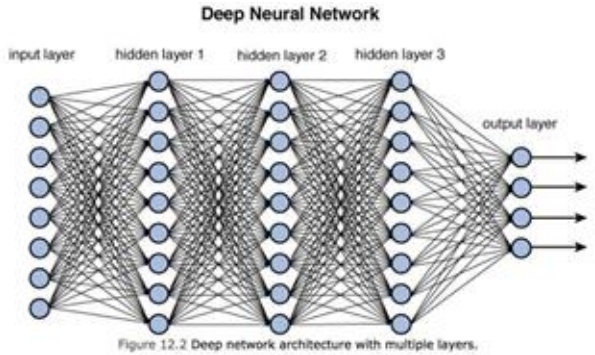

Figure 12.2 Deep network architecture with multiple layers.

- Razavian et al. (2015): **Type 2 Diabetes Onset**
- Ahmad et al. (2018): **Mortality Prediction**
- Rajkomar et al. (2018): **in-patient mortality and 30-day readmission**
- Beam et al. (2020): Surveyed 210 machine learning models for structured healthcare data: "**deep learning models do not perform better than traditional approaches like logistic regression in most cases.**"

Summary

- ✓ Introduction to risk stratification
- ✓ Case study: Early detection of Type 2 diabetes
- ✓ Framing as supervised learning problem

How can we make Data 146 better for you?

Next class: LLMs and healthcare

% who say they think the impact of artificial intelligence (AI) on society over the next 20 years will be ...

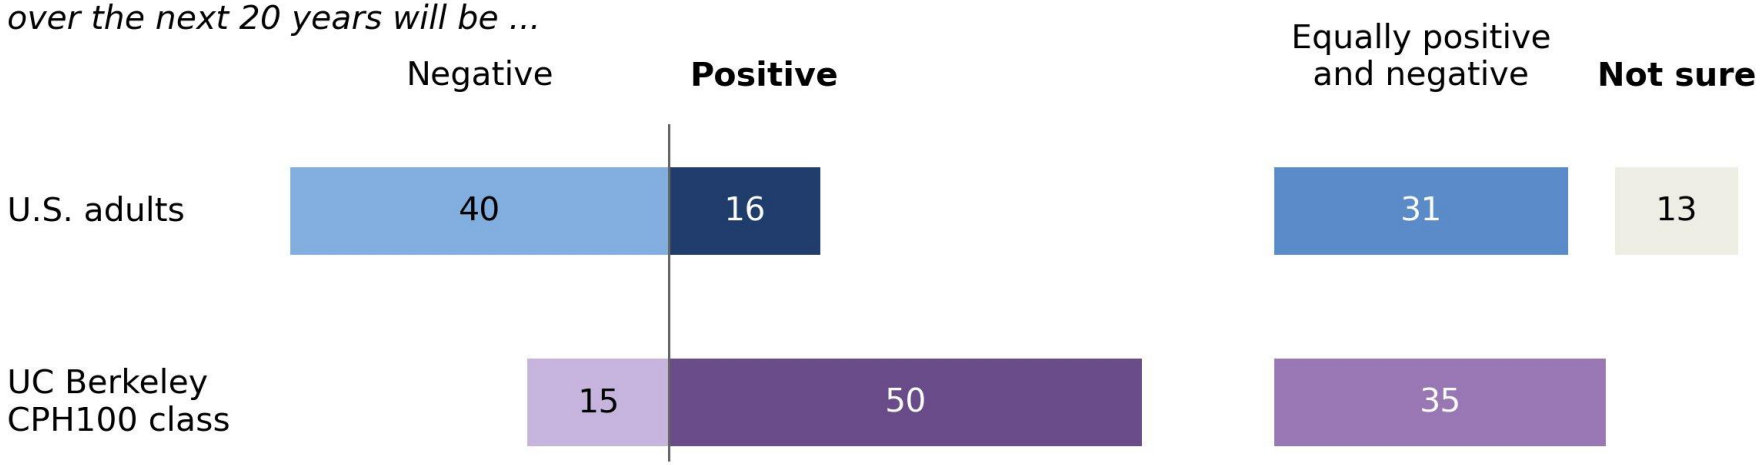

% who say they think the increased use of artificial intelligence (AI) is more likely to ...

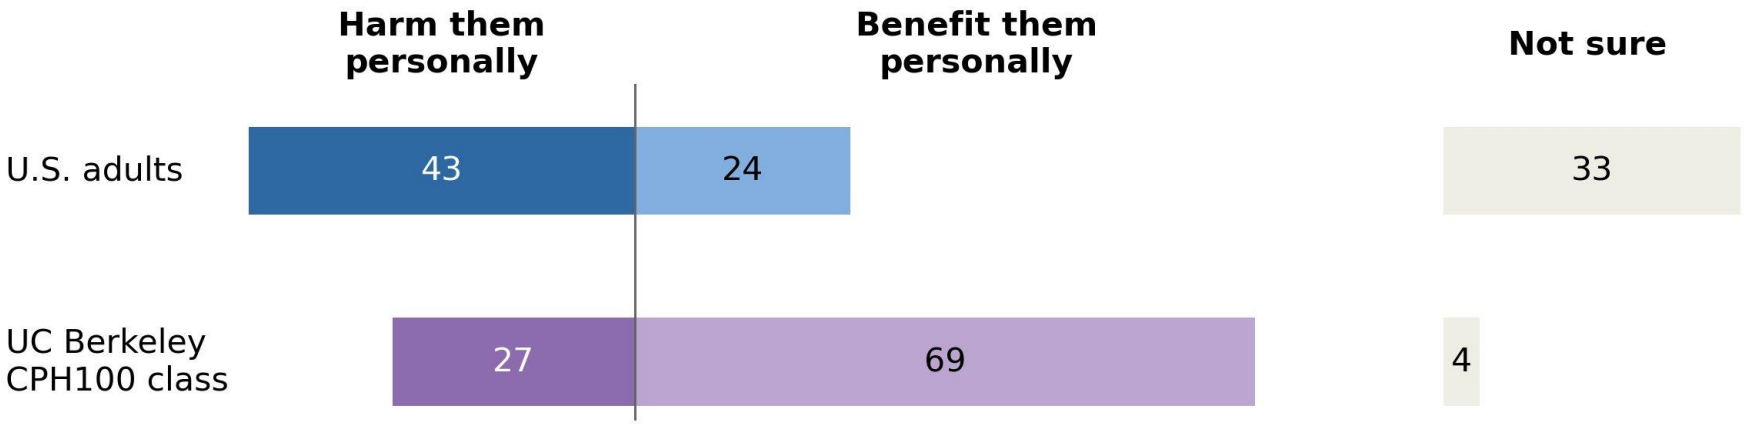

Source: U.S. adults — Pew Research Center (Feb. 17–23, 2026) for societal impact; Pew Research Center (Aug. 12–18, 2024) for personal impact. UC Berkeley CPH100 class survey, Sept. 4, 2026 (n=26). Percentages rounded.

AI experts more likely than the public to say AI will have a positive effect on the U.S. over next 20 years

% who say they think the impact of artificial intelligence (AI) on the U.S. over the next 20 years will be ...

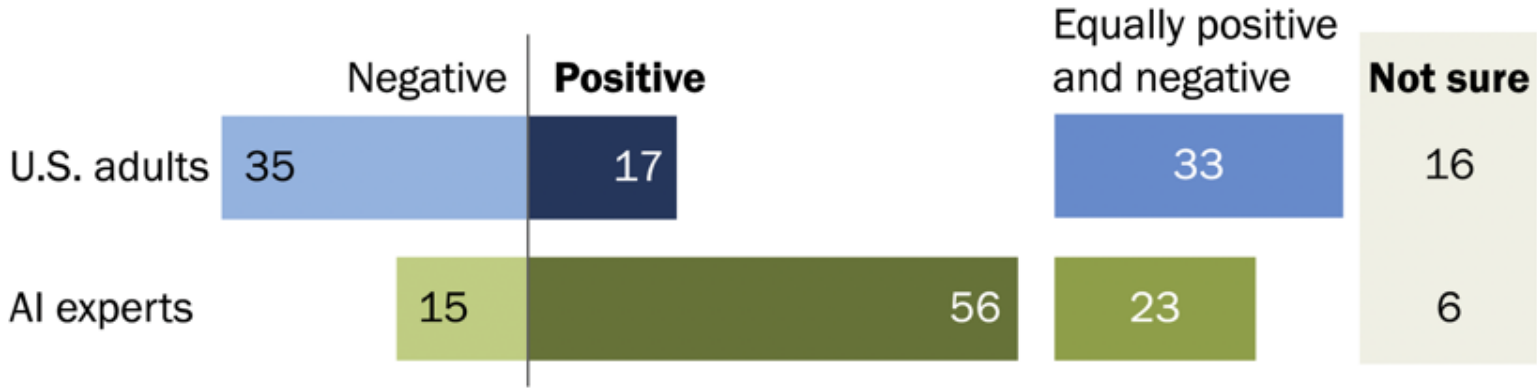

Note: "AI experts" refer to individuals whose work or research relates to AI. The AI experts surveyed are those who were authors or presenters at an AI-related conference in 2023 or 2024 and live in the U.S. Expert views are only representative of those who responded. For more details, refer to the methodology. "Very/somewhat positive" and "very/somewhat negative" are combined. Those who did not give an answer are not shown.
Source: Survey of U.S. adults conducted Aug. 12-18, 2024. Survey of AI experts conducted Aug. 14-Oct. 31, 2024.
"How the U.S. Public and AI Experts View Artificial Intelligence"

PEW RESEARCH CENTER

https://www.pewresearch.org/internet/2025/04/03/how-the-us-public-and-ai-experts-view-artificial-intelligence/

About three-quarters of AI experts think AI will likely benefit them personally, much higher than share of U.S. adults

% who say they think the increased use of artificial intelligence (AI) is more likely to ...

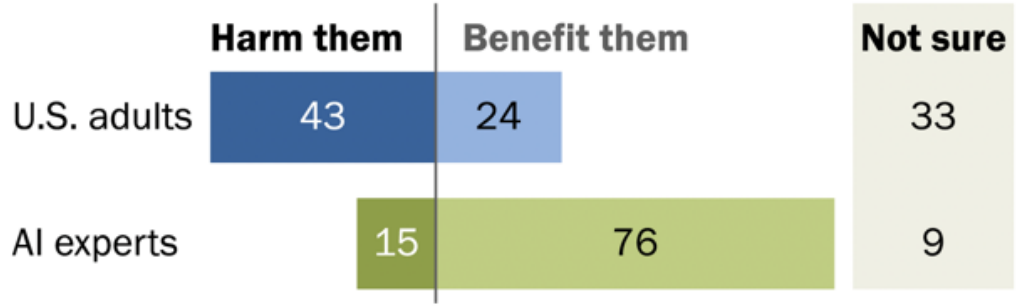

Note: "AI experts" refer to individuals whose work or research relates to AI. The AI experts surveyed are those who were authors or presenters at an AI-related conference in 2023 or 2024 and live in the U.S. Expert views are only representative of those who responded. For more details, refer to the methodology. For full question wording, refer to the topline. Those who did not give an answer are not shown.
Source: Survey of U.S. adults conducted Aug. 12-18, 2024. Survey of AI experts conducted Aug. 14-Oct. 31, 2024.
"How the U.S. Public and AI Experts View Artificial Intelligence"

PEW RESEARCH CENTER

https://www.pewresearch.org/internet/2025/04/03/how-the-us-public-and-ai-experts-view-artificial-intelligence/