

# Application of Machine Learning to Cardiac Imaging

Rahul Deo, MD PHD  
Lead Investigator, Brigham and Women's  
Hospital  
One Brave Idea

Adjunct Associate Professor,  
University of California San Francisco

Member of Faculty  
Harvard Medical School



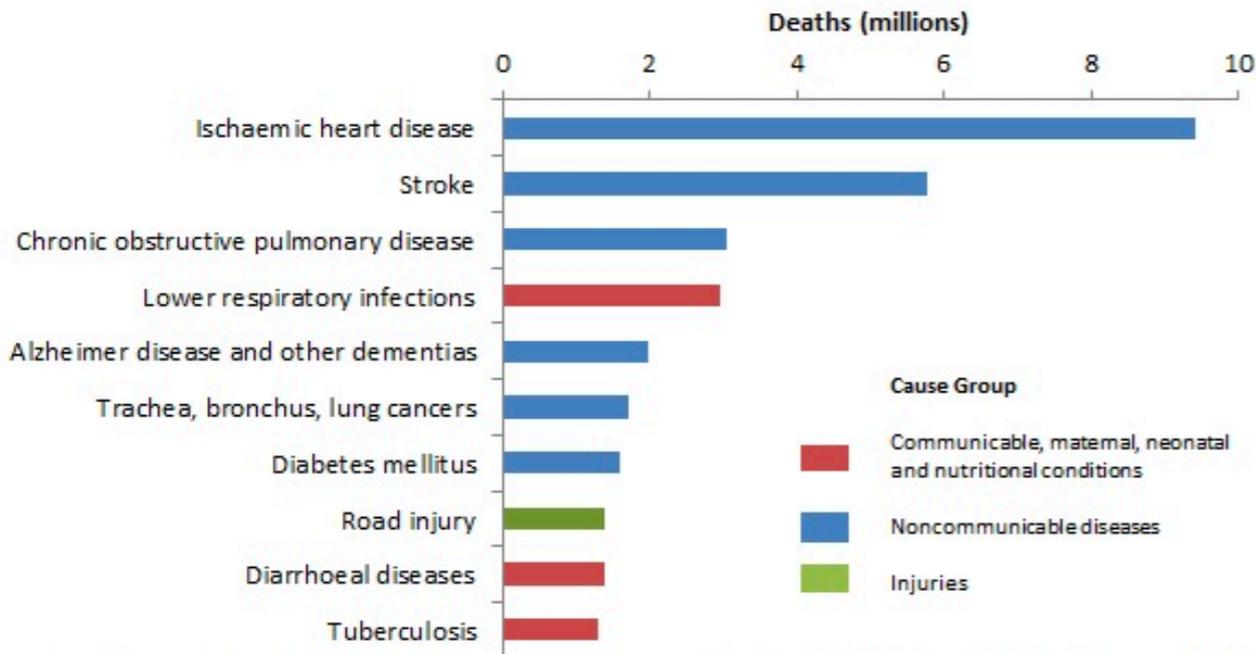
# Outline

1. Introduction to cardiac structure and function
2. Major types of cardiac diagnostics and how they are used
3. Where's the data?
4. Basic computer vision topics relevant to cardiac imaging
5. A fully automated pipeline for echocardiogram interpretation
6. Rethinking the future of automated interpretation: lessons from the ECG
7. What about the biology?

# **A Brief Introduction to Cardiac Structure and Function**

# Coronary heart disease (CHD) is the leading global cause of death

Top 10 global causes of deaths, 2016

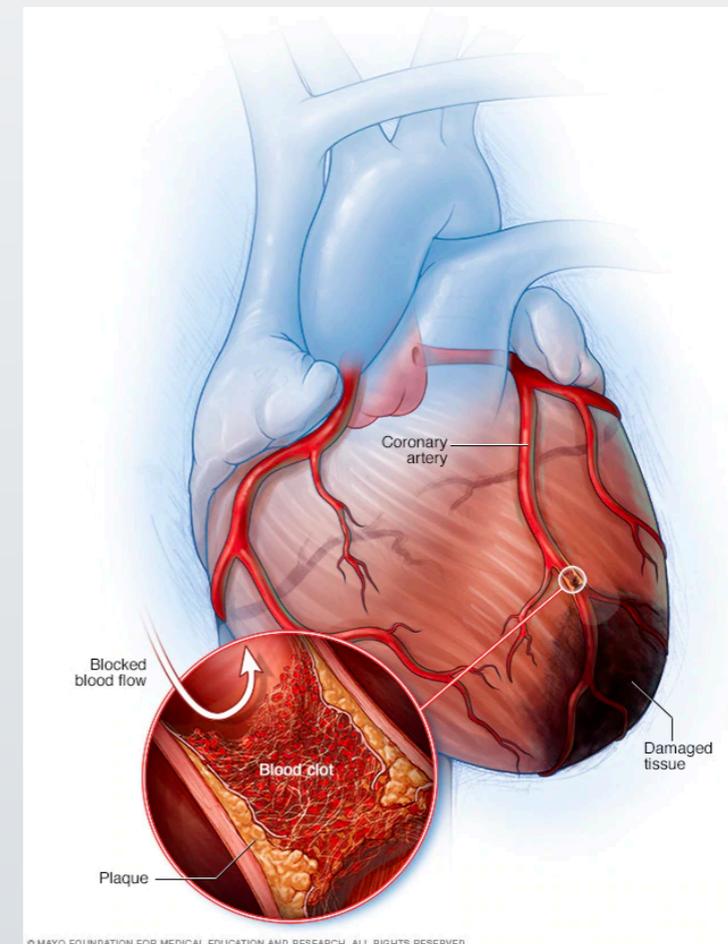


Source: Global Health Estimates 2016: Deaths by Cause, Age, Sex, by Country and by Region, 2000-2016. Geneva, World Health Organization; 2018.

CHD is the leading cause of death in both developed and developing countries.

Lancet, 2018

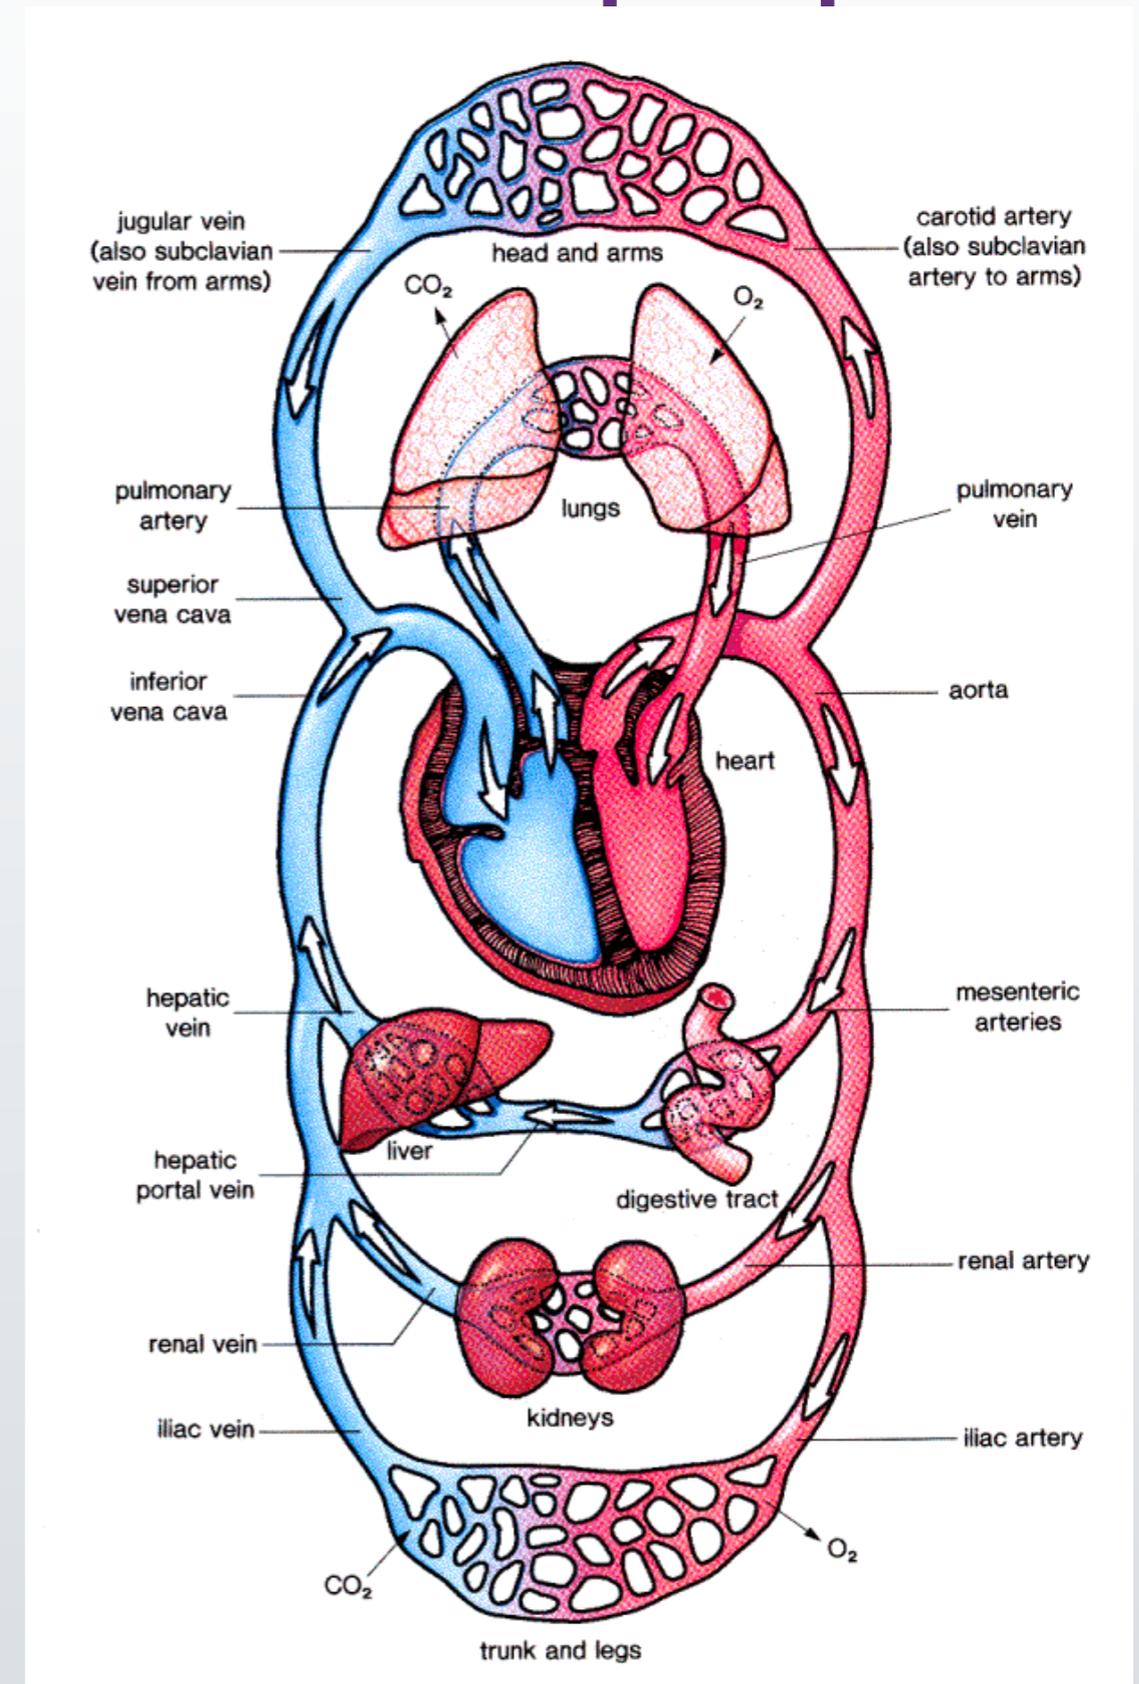
	All-age deaths (thousands)		Age-standardised death rate (per 100 000)	
	2017	Percentage change, 2007-17	2017	Percentage change, 2007-17
<b>All causes</b>	55 945.7 (55 356.4 to 56 516.7)	9.3% (8.2 to 10.2)*	737.7 (729.9 to 745.4)	-14.2% (-15.0 to -13.5)*
<b>Communicable, maternal, neonatal, and nutritional diseases</b>	10 389.9 (10 004.0 to 10 975.9)	-22.2% (-24.0 to -20.0)*	143.8 (138.4 to 151.6)	-31.8% (-33.3 to -30.1)*
<b>Non-communicable diseases</b>	41 071.1 (40 470.9 to 41 548.9)	22.7% (21.5 to 23.9)*	536.1 (528.4 to 542.2)	-7.9% (-8.8 to -7.0)*
<b>Neoplasms</b>	9 556.2 (9 395.7 to 9 692.3)	25.4% (23.9 to 27.0)*	121.2 (119.1 to 122.9)	-4.4% (-5.6 to -3.3)*
<b>Cardiovascular diseases</b>	17 790.9 (17 527.1 to 18 042.7)	21.1% (19.7 to 22.6)*	233.1 (229.7 to 236.4)	-10.3% (-11.4 to -9.3)*



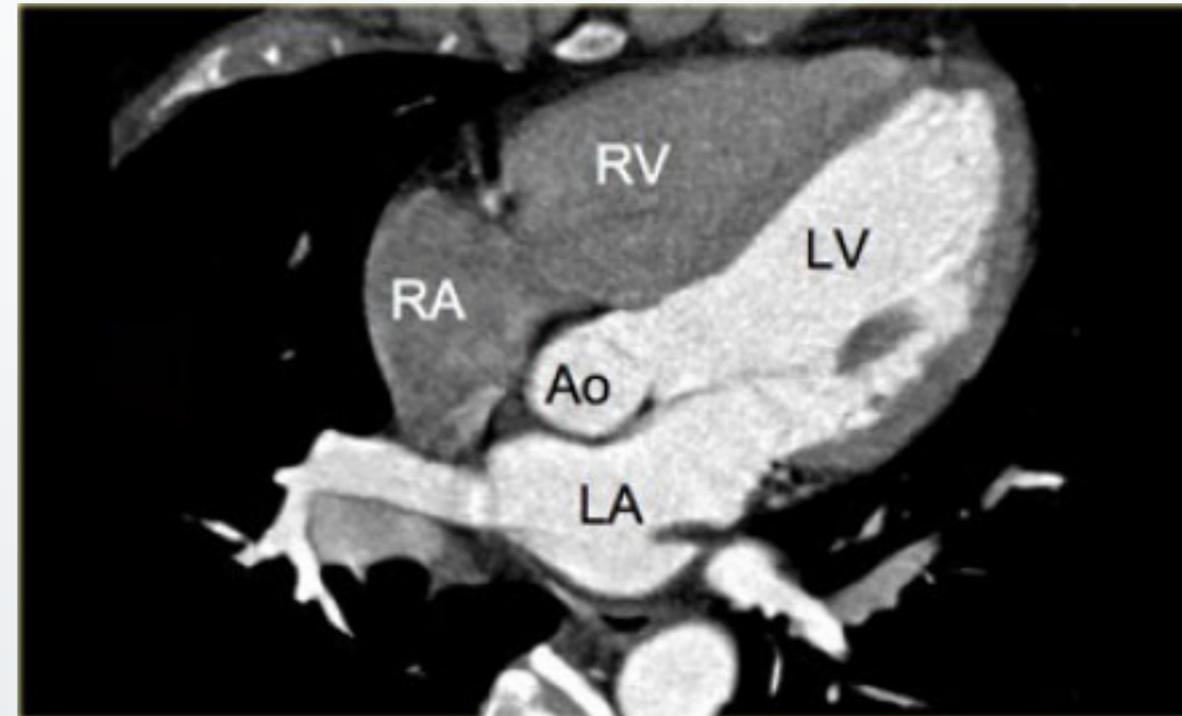
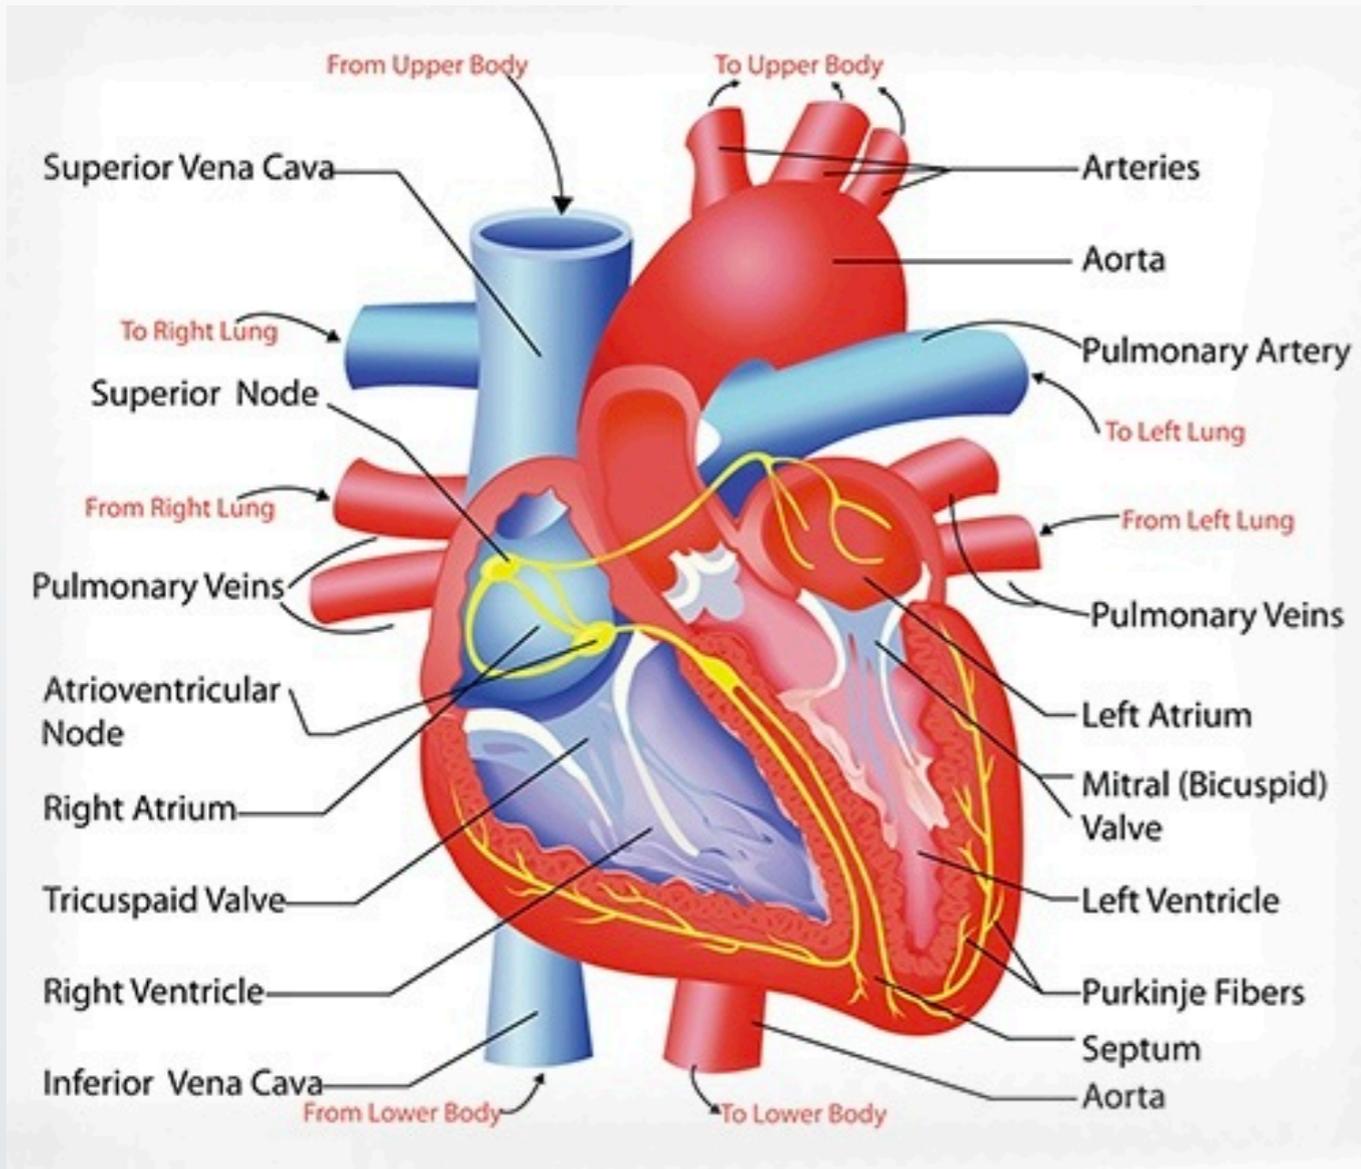
© MAYO FOUNDATION FOR MEDICAL EDUCATION AND RESEARCH. ALL RIGHTS RESERVED.

# The heart's primary function is as a pump

1. The heart must deliver oxygenated blood throughout the circulatory system
2. Blood supplies tissues with oxygen for ATP production, delivers and receives signaling molecules, and removes waste
3. The heart pumps ~5L of blood per minute, which can expand to 20-35L per minute during exercise
4. The rhythmic function of the heart results in >2 billion heart beats in a typical lifetime



# The structure of the heart



**4 chambers:** RA, RV, LA, LV

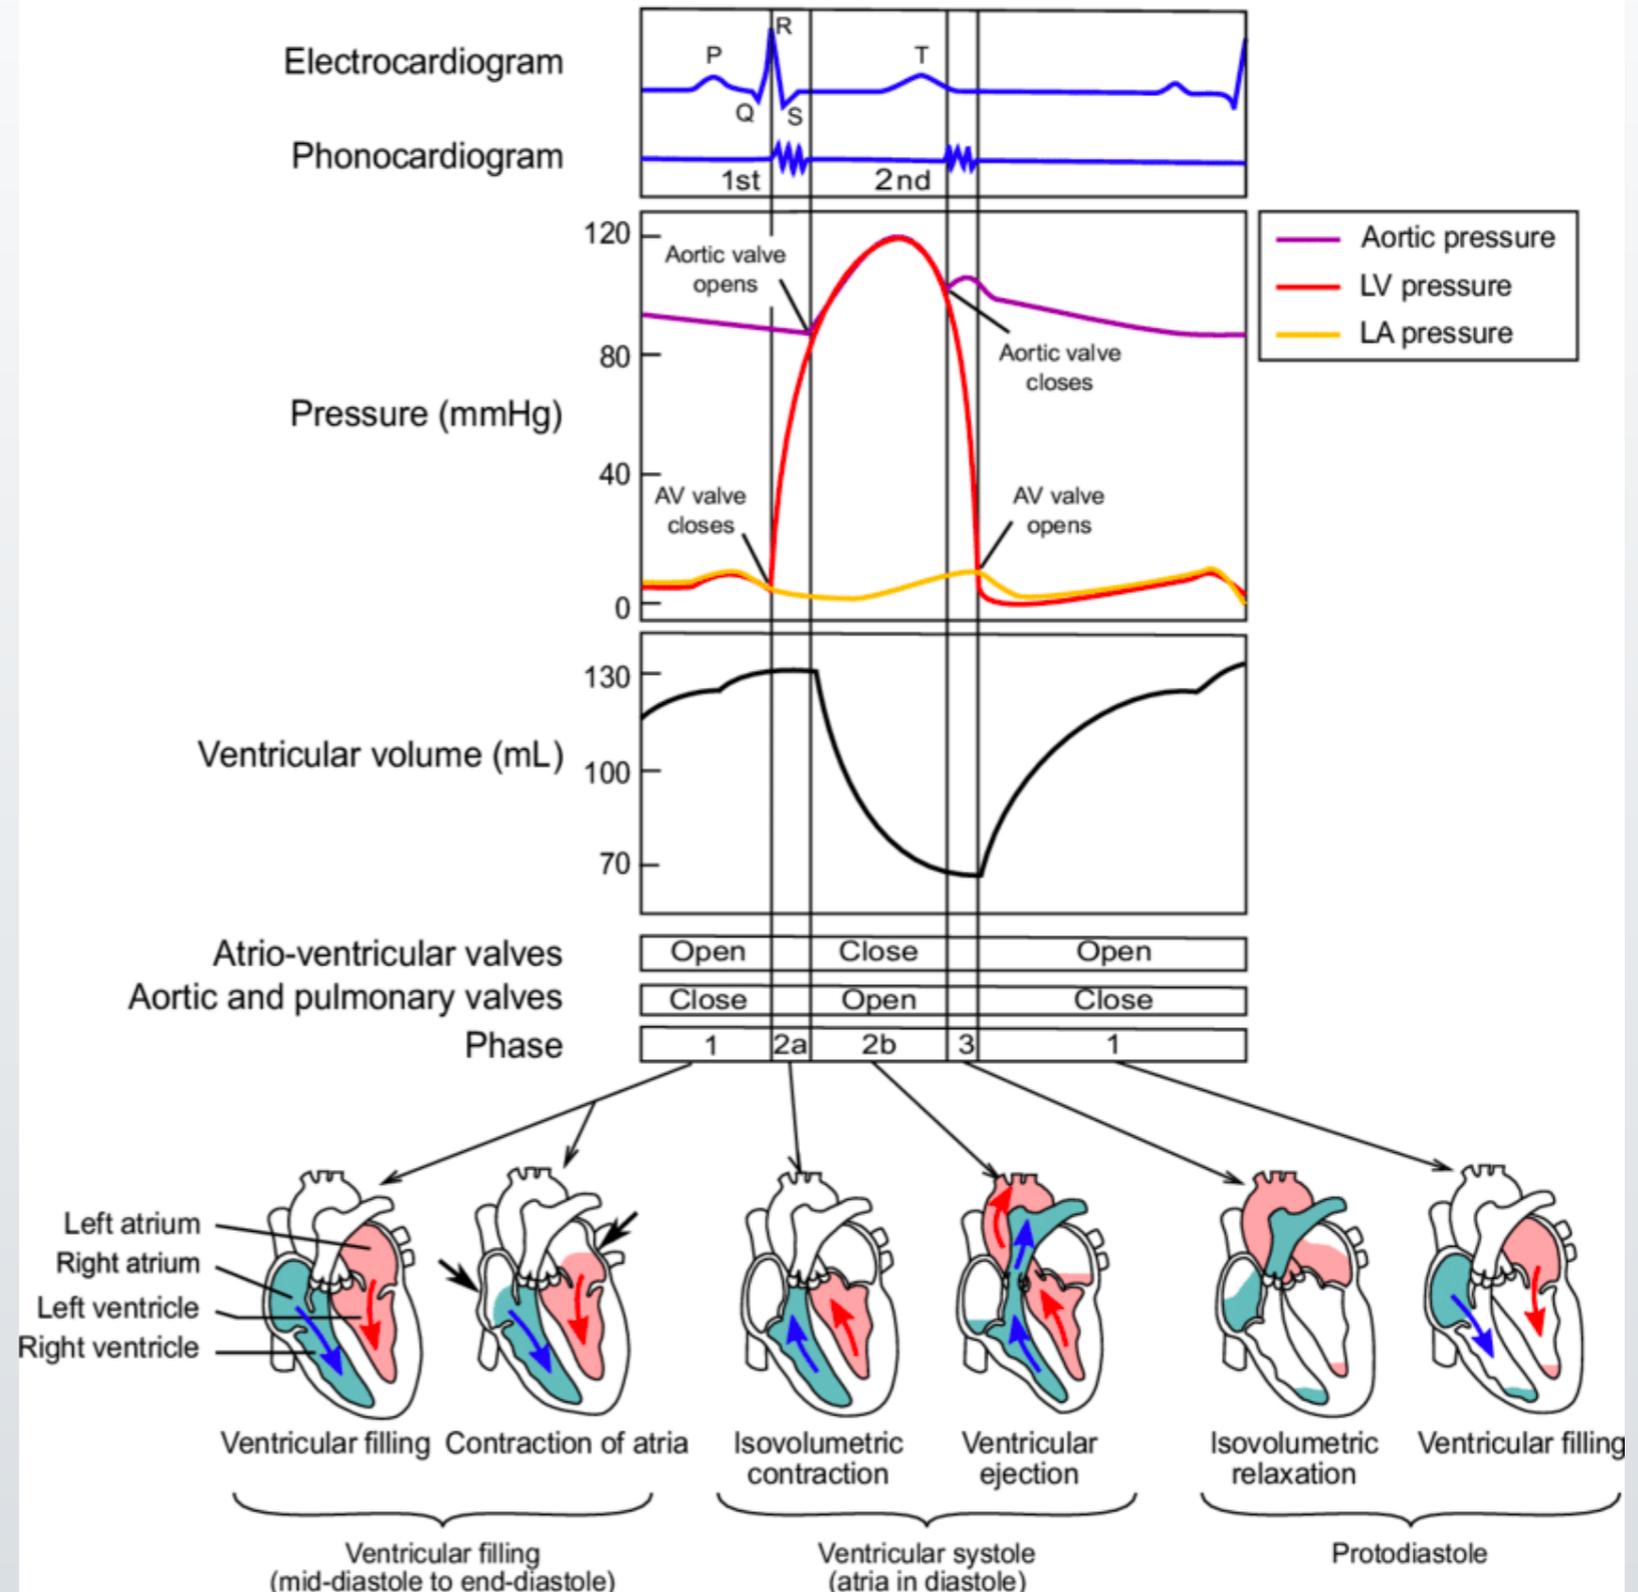
**4 valves:** TV, PV, MR, AV

**2 circulations in series:**

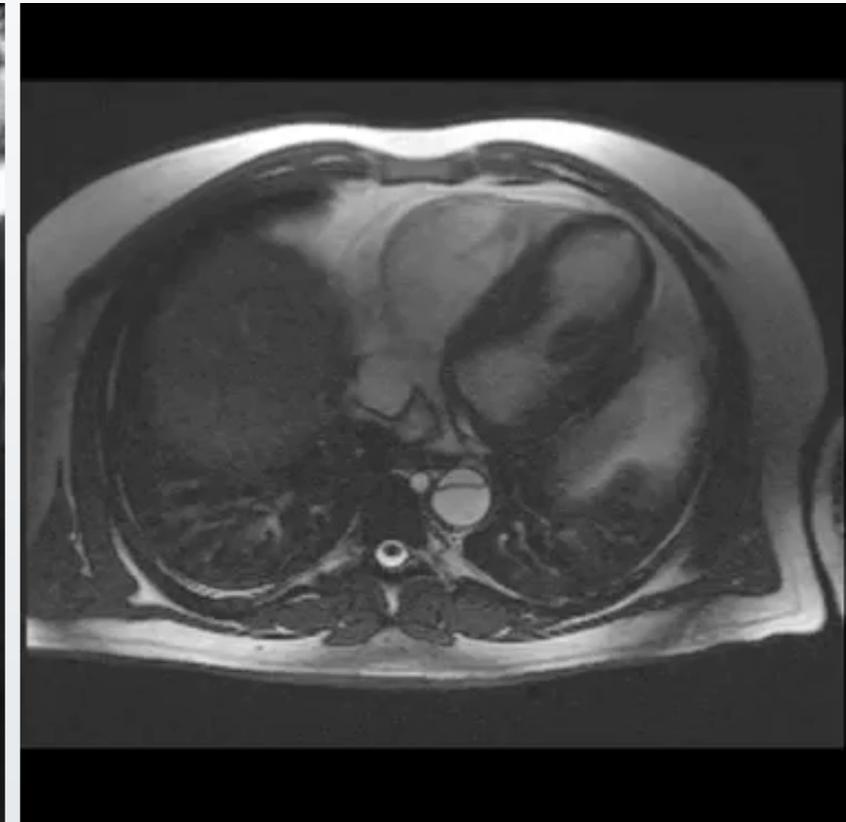
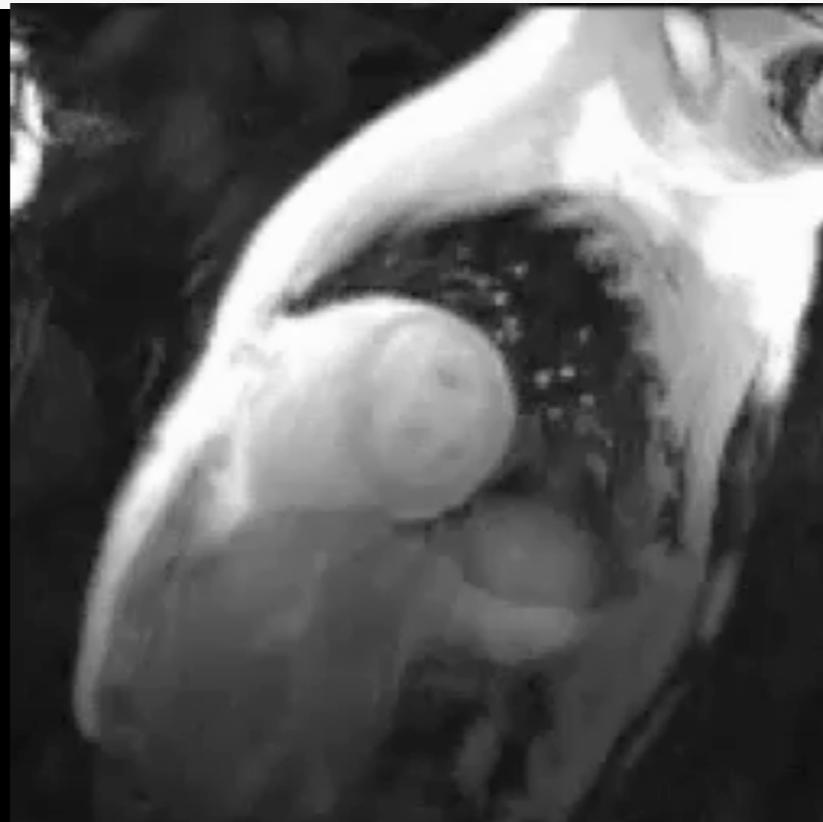
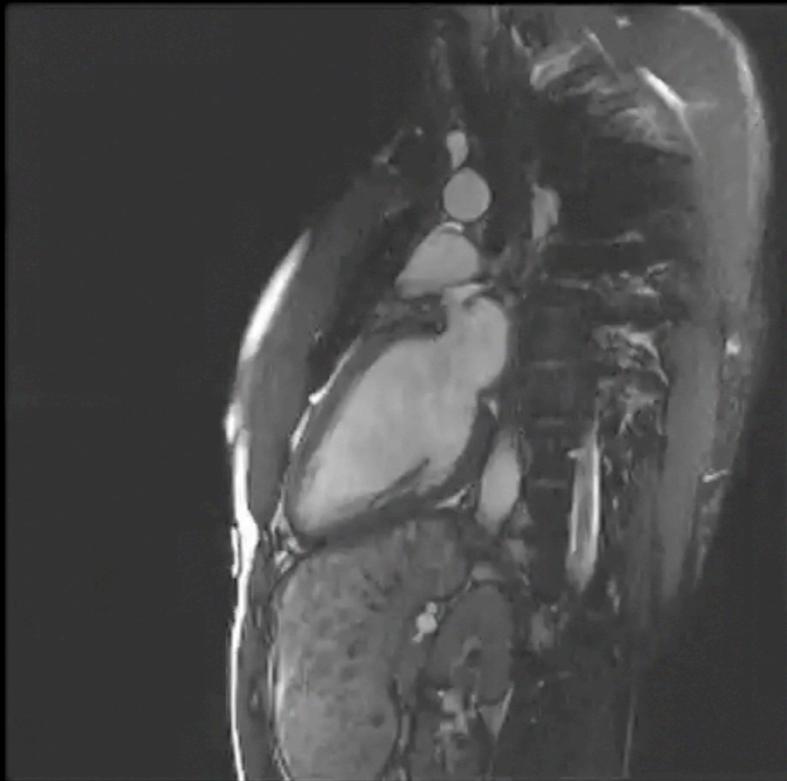
pulmonary and systemic

# The Cardiac Cycle: Synchronized Electrical and Mechanical Activation

1. The **Wiggers diagram** aligns mechanical and electrical events (and heart sounds)
2. The heart alternates between periods of relaxation and filling (**diastole**) and periods of contraction and ejection of blood (**systole**)



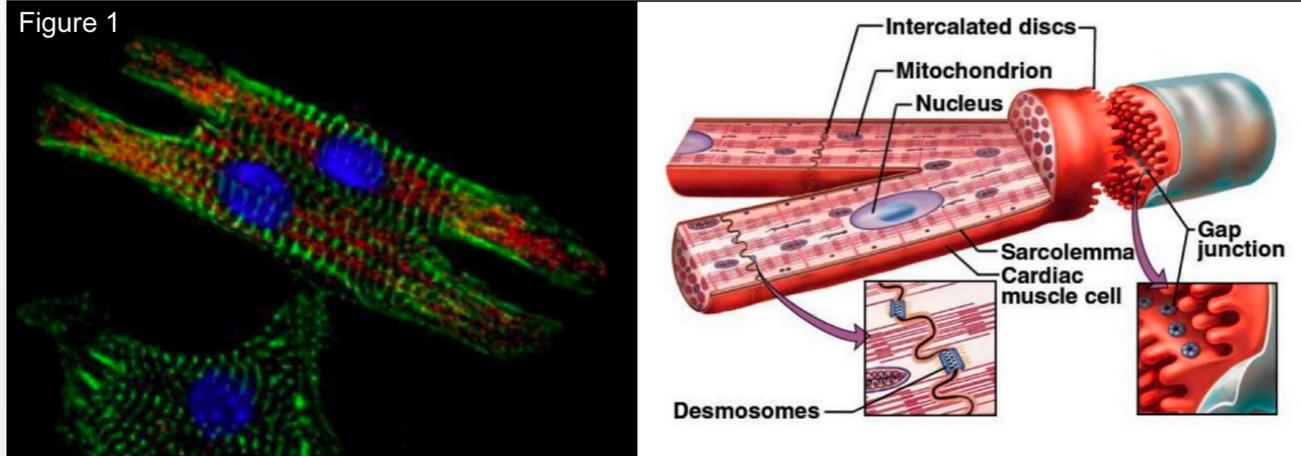
# Visualizing the Heart in Motion



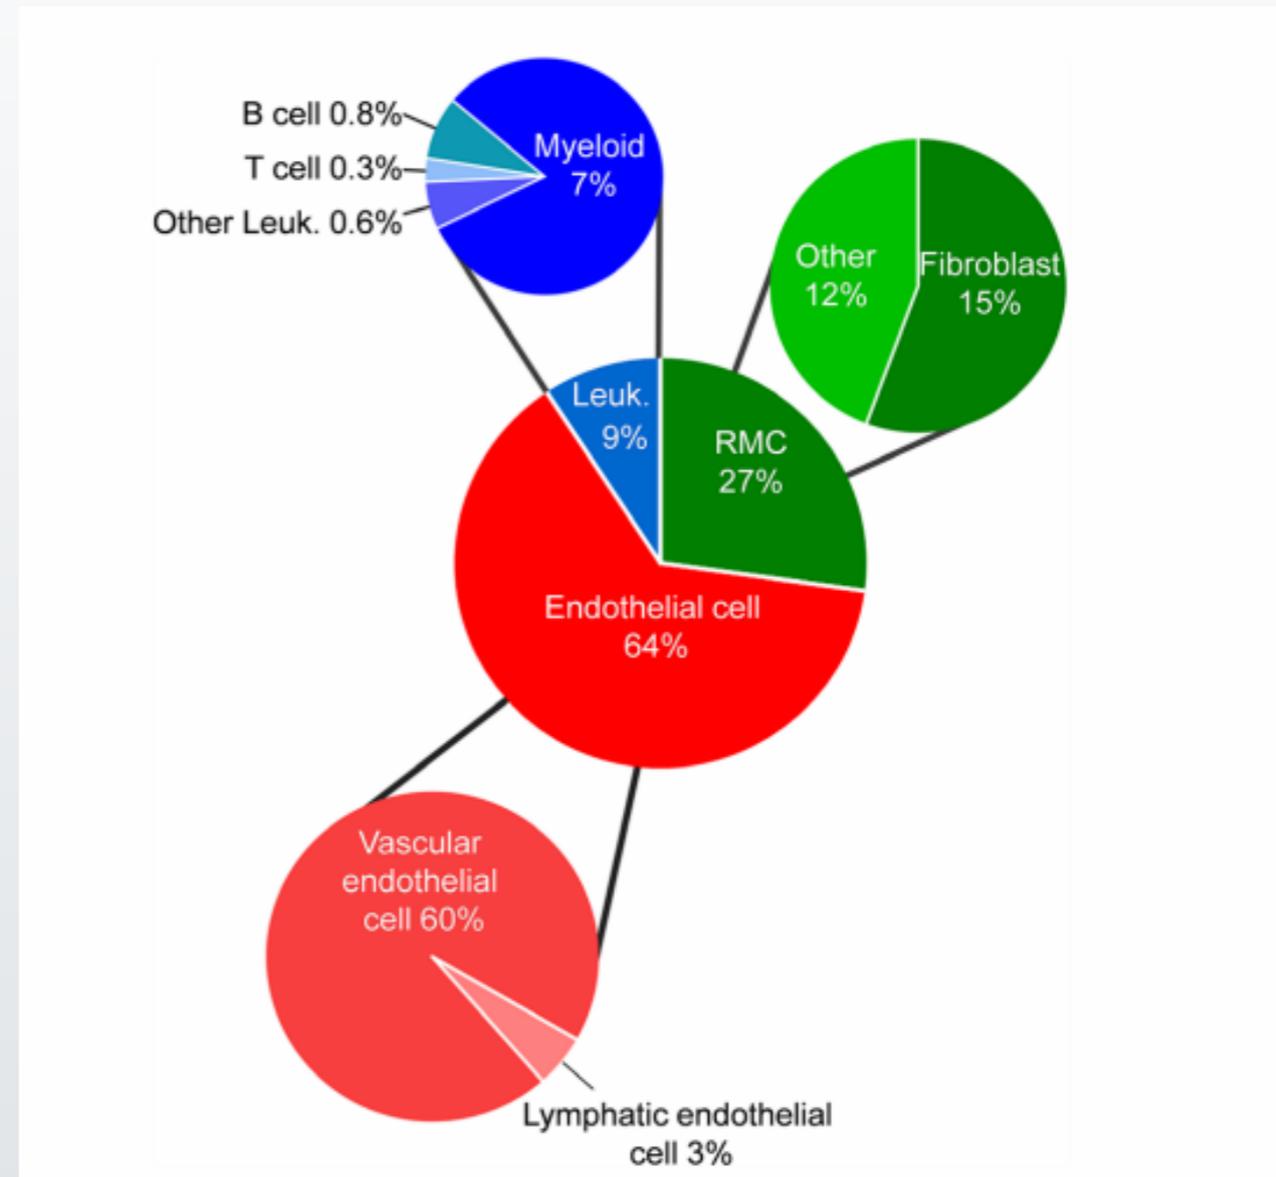
# Diseases of the heart are organized into abnormalities of contractile function, coronary blood supply, circulatory flow, or heart rhythm

Abnormality	Disease Names	Presentation	Treatment
<b>Contractile function</b>	Heart failure	Shortness of breath, fluid buildup in legs	medications, ventricular assist device, transplant
<b>Coronary blood supply</b>	Coronary artery disease, myocardial infarction	Chest pain, shortness of breath	angioplasty/stenting; coronary artery bypass grafting
<b>Circulatory flow</b>	Aortic stenosis/regurgitation, mitral stenosis/regurgitation,	Light headedness, shortness of breath, fainting	valve replacement, valve repair
<b>Heart rhythm</b>	Atrial fibrillation/flutter, ventricular tachycardia, sick sinus syndrome	palpitations, fainting, cardiac arrest	ablation, implantable defibrillator, pacemaker

# The heart is a complex multicellular organ



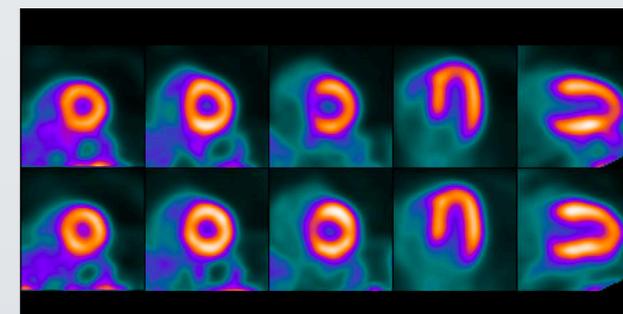
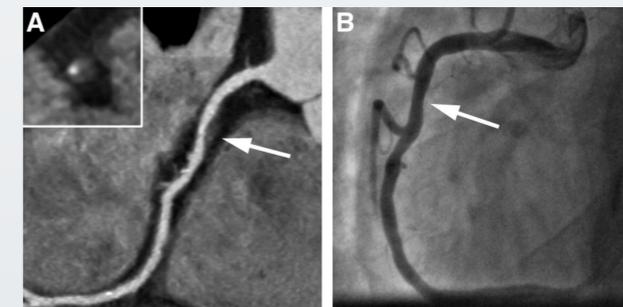
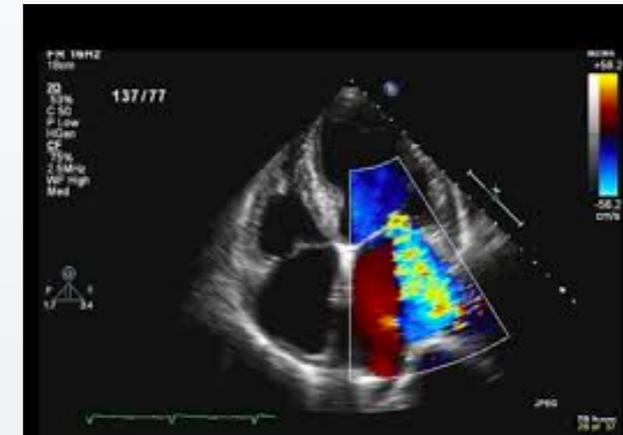
1. The cardiomyocyte is the primary excitable and contractile cell in the heart
2. But ... **only 31% of cells in the heart are cardiomyocytes**
3. Cardiac function and disease arises from the interplay of a broad group of cells
4. Other cell types: endothelial, fibroblast, leukocytes



# Cardiac Imaging in Medical Decision Making

# Cardiac imaging plays a critical role in diagnosis (and definitions of disease)

Modality	Cost	Approach	Diagnostic Utility
Electrocardiogram (ECG)	\$	Voltage differences	Myocardial infarction
Echocardiography	\$\$-\$\$\$	Ultrasound (sound waves); Doppler shift	Quantitation of cardiac structure and function, heart failure, valvular disease, pulmonary hypertension
MRI	\$\$\$\$	Magnetic resonance (volumetric reconstruction)	Quantitation of cardiac structure and function, heart failure, valvular disease
Angiography (Fluoroscopy and Computed Tomography)	\$\$\$\$	X-ray (volumetric reconstruction for CT)	Epicardial coronary artery disease
SPECT/PET	\$\$\$\$	Radionuclide tracer	Coronary artery disease (inferred); microvascular disease
Intracardiac pressure transducers	\$\$\$	Pressure transducer	Heart failure, valvular disease, pulmonary hypertension



Many cardiac diseases are defined (for better or worse) as departures from normal anatomic/physiologic values

# Cardiac decisions are often (but not always) guided by inputs from imaging

Disease	Decision	Inputs
Heart failure	Decision to implant a defibrillator to prevent sudden death	Symptoms + ejection fraction of the heart <35%
Coronary artery disease	Angioplasty and stenting of a coronary artery	Symptoms + stenosis > 70%
Aortic stenosis	Valve replacement	Symptoms + valve area + enlargement of the heart
Atrial fibrillation	Decision to start anticoagulation to prevent stroke	Age, sex, other diagnoses
Myocardial infarction	Decision to start aspirin and a statin to prevent a future heart attack	A risk model based on age, sex, lab values, blood pressure, diabetes

1. Information content of imaging can be very high ... but decisions are based on **historical patient populations followed through time** with the relevant disease
2. Risk model and decision analysis is dictated by **what data are available** for these historical populations
3. Imaging is available for patient populations for which it is a part of the **accepted management plan** ... but is unlikely to be found for other diseases given cost

# Imaging Modalities and Data

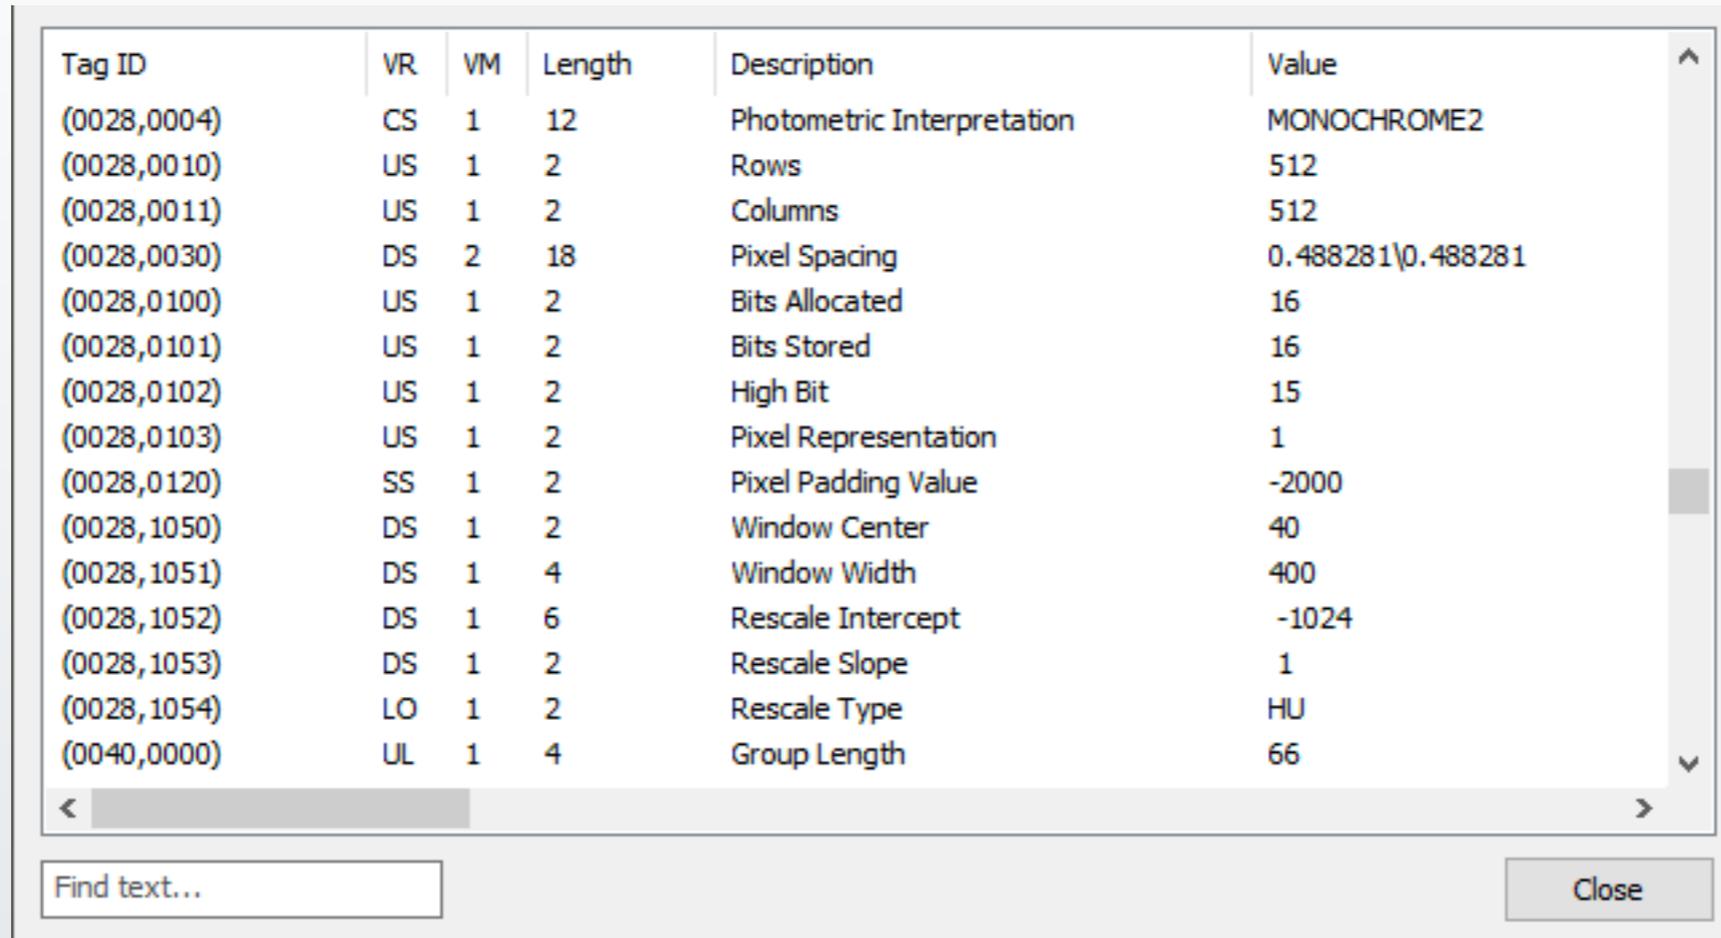
# How Medical Imaging Data Are Stored

1. **DICOM (Digital Imaging and Communications Standard)** is the international standard to transmit, store, retrieve, print, process, and display imaging information
2. Image/video files are stored in DICOM format, and combine a compressed image with a **DICOM “header”** which includes characteristics of the image
3. Open access libraries like **GDCM, pydicom** facilitate compressing/uncompressing; reading and editing header
4. Osirix Lite provides a free DICOM viewer

# Where can I get access to data?

1. Most imaging data is housed in **data archives** (increasingly “vendor neutral”)
2. **Access is often highly limited:**
  1. Some images have burned in pixels with patient names, dates of birth
  2. Scalable solutions for download and de-identification are not always available (these would facilitate changing vendors)
  3. Some systems have monetized their imaging data
3. Labels (e.g. diagnoses, measurements) are often **stored separately in the electronic health record**
4. Scale of data (at BWVH) - clearly relates to cost of the study as well as perceived utility:
  1. Electrocardiograms: 30 million ECGs
  2. Echocardiography: 300,000-500,000 studies
  3. Cardiac PET: 8000 studies

# Example of a DICOM header



The image shows a screenshot of a DICOM header viewer window. The window contains a table with the following columns: Tag ID, VR, VM, Length, Description, and Value. The table lists various DICOM header fields and their corresponding values. At the bottom of the window, there is a search box labeled 'Find text...' and a 'Close' button.

Tag ID	VR	VM	Length	Description	Value
(0028,0004)	CS	1	12	Photometric Interpretation	MONOCHROME2
(0028,0010)	US	1	2	Rows	512
(0028,0011)	US	1	2	Columns	512
(0028,0030)	DS	2	18	Pixel Spacing	0.488281\0.488281
(0028,0100)	US	1	2	Bits Allocated	16
(0028,0101)	US	1	2	Bits Stored	16
(0028,0102)	US	1	2	High Bit	15
(0028,0103)	US	1	2	Pixel Representation	1
(0028,0120)	SS	1	2	Pixel Padding Value	-2000
(0028,1050)	DS	1	2	Window Center	40
(0028,1051)	DS	1	4	Window Width	400
(0028,1052)	DS	1	6	Rescale Intercept	-1024
(0028,1053)	DS	1	2	Rescale Slope	1
(0028,1054)	LO	1	2	Rescale Type	HU
(0040,0000)	UL	1	4	Group Length	66

Unfortunately, there can be some instrument to instrument variability in how some fields are represented.

# Characteristics of Cardiac Imaging Data

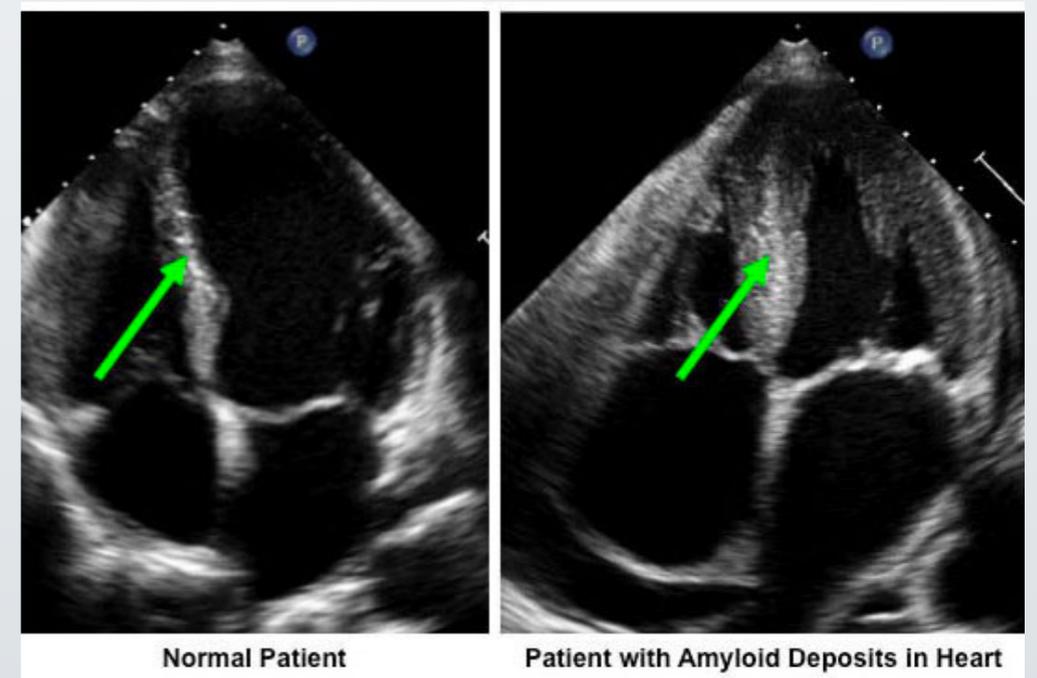
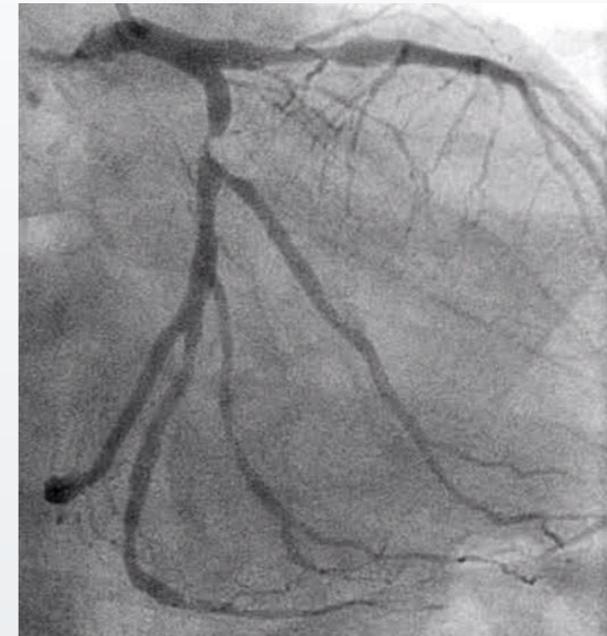
1. **Compression:** lossy vs. lossless
2. **Spatial resolution:** number of pixels; pixel dimensions
3. **Sampling frequency (temporal resolution):** very high for various ultrasound modalities and coronary angiography, moderate for CT scanners
4. **Coronary artery velocity** is 10-65 mm/seconds
5. **“Gating”:** electrocardiogram information can be coupled with imaging information to average images across corresponding portions of the cycle

Modality	Spatial Resolution	Temporal resolution
Echocardiography	2-3 mm	1-5 ms for some modes; typically 20-30ms for 2D
MRI	0.1mm	30-100 ms
Angiography (Fluoroscopy)	0.1mm	1-10 ms
Computed Tomography	0.5mm in x,y; 0.5-0.625mm in z	65-175 ms
SPECT/PET	8-10 mm for SPECT; 3-5 mm for PET	minutes

# Relevant Topics in Computer Vision

# Machine learning in cardiac disease - what physician practices can we mimic?

1. All current **measurements** (cardiac chamber areas, ventricular thickness) are **performed manually**
  1. Severity of a stenosis of the coronary artery
  2. Left ventricular cardiac volumes (and by comparison across the cardiac cycle, ejection fraction)
2. Some disease diagnoses involve **classification of images/videos**



We'll come back to whether these contributions would be seen as valuable

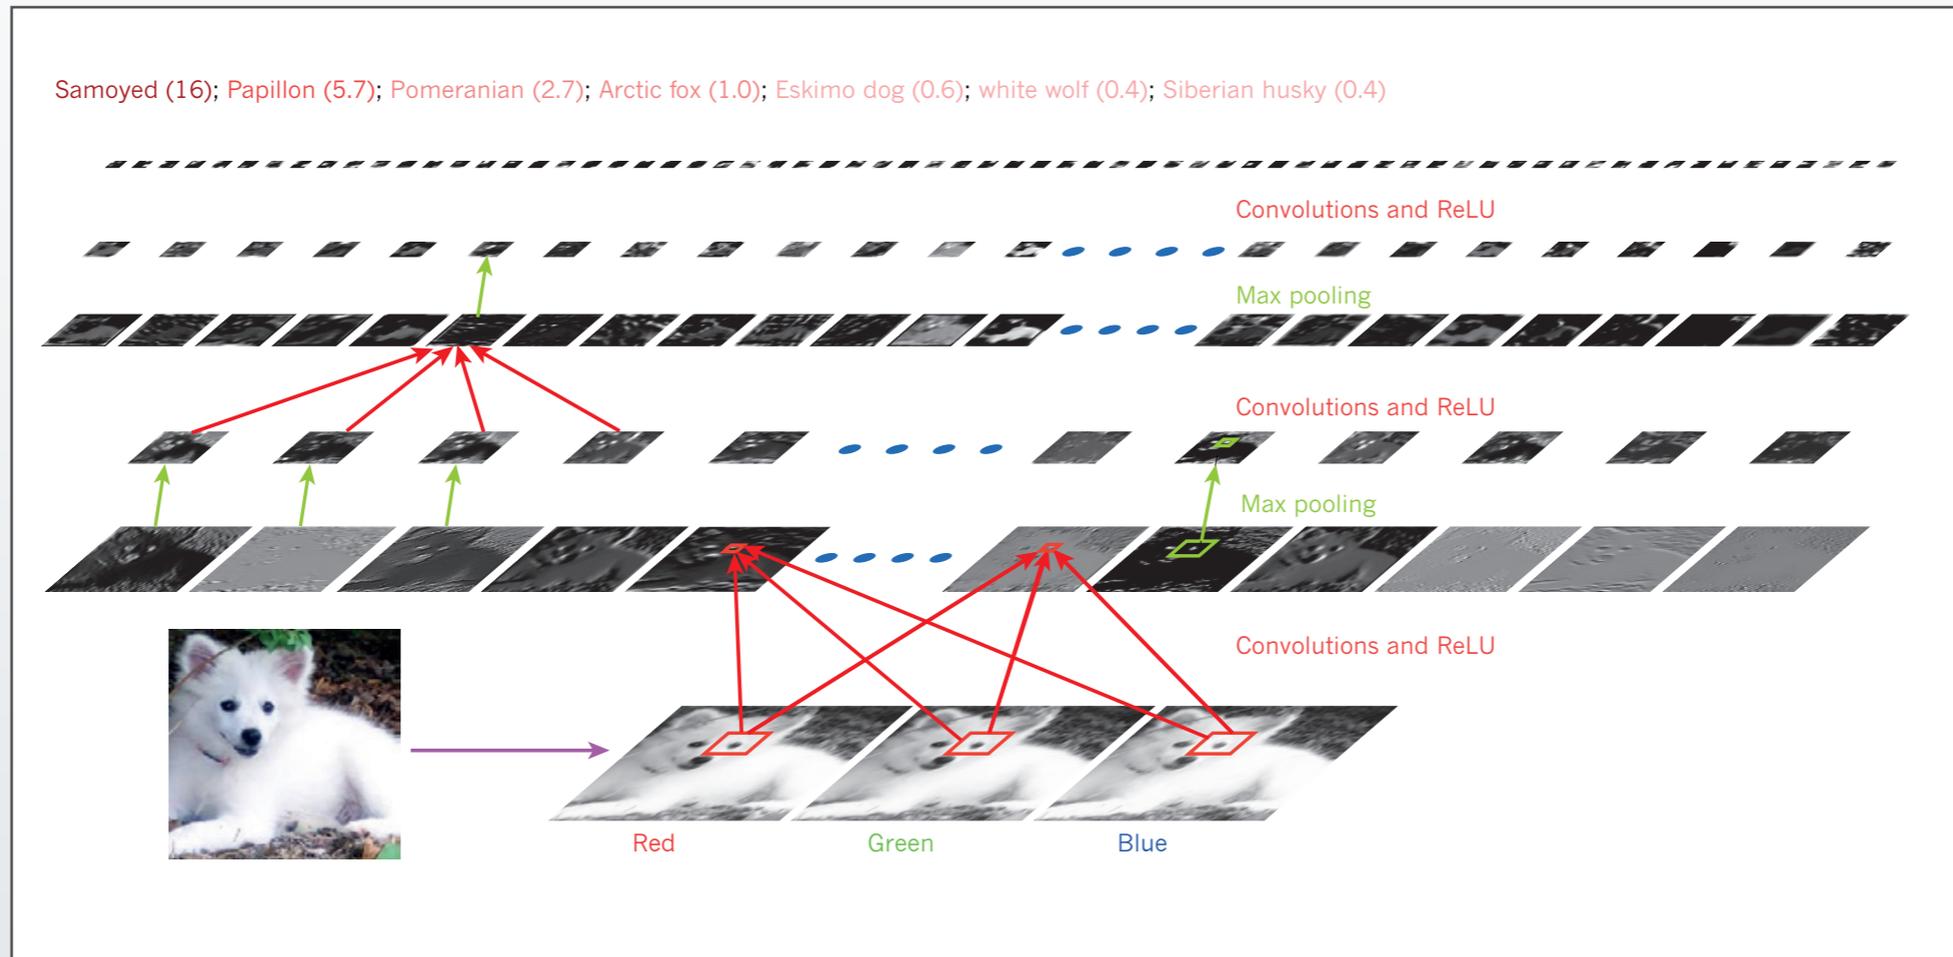
# Many priorities in computer vision are of great interest to cardiac imaging

1. **Image (and video) classification:** assigning a label to an image/video
2. **Semantic segmentation:** associating each pixel in an image with a class label
3. **Image registration:** mapping different sets of images onto one coordinate system

# Image classification: an obvious task to mimic

1. Many simple **disease recognition tasks** exist in medicine - and can be carried out by an experienced radiologist in 2 minutes or less
  1. e.g. lung cancer or not
  2. pneumonia or not
  3. breast cancer or not
  4. fluid around the heart or not
2. Many of the first successes in medical image classification have involved situations with **very large data sets, already labeled in the context of routine clinical care**
  1. Chest x-rays
  2. Mammograms
3. **Barriers to data export and sharing** have limited the size of many other data sets

# Image classification: convolutional neural networks have rekindled an interest in automated medical image classification



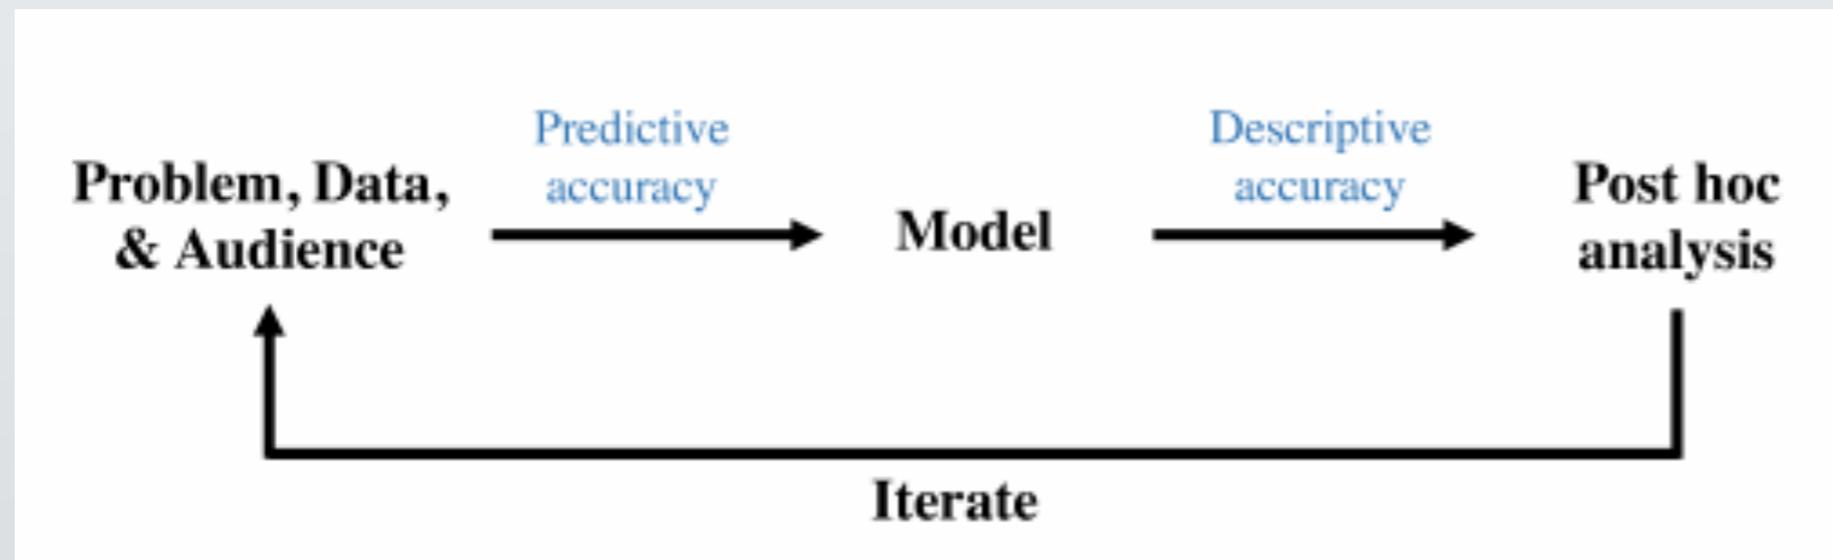
1. Representation learning
2. No need for hand-engineering of features
3. Transfer learning: important in training data poor scenarios

# Image classification: will anyone use it?

1. If a radiologist takes 2 minutes to read a study, how much benefit is there to automate the process
2. **Liability** is an enormous reason why we don't task-shift image interpretation to less-skilled personnel - radiologists are among the most sued physicians
3. So it is unlikely we will see the benefits of other disciplines where a machine will be **permitted to independently read a study** BUT:
  1. If there are 1000 X-rays to (over)read or if the hospital is in off-hours (overnight, weekend), a machine can pre-read and decide **what's most urgent to look at**
  2. An independent read should **catch some missed diagnoses**
4. The calculus may change in **resource poor settings**

# Image classification: explaining the diagnosis

1. All imaging-based medical decisions have typically required a **corresponding human confirmation of a visual finding**
2. In some cases, the need is unambiguous: **you can't take a biopsy of a tumor you can't localize**; nor can you submit a report that doesn't localize the abnormality
3. Increasingly **patients and providers share in decisions**: requiring both to be convinced of the validity of the conclusion
4. CNNs raise some concerns as to whether a simple explanation can always be given

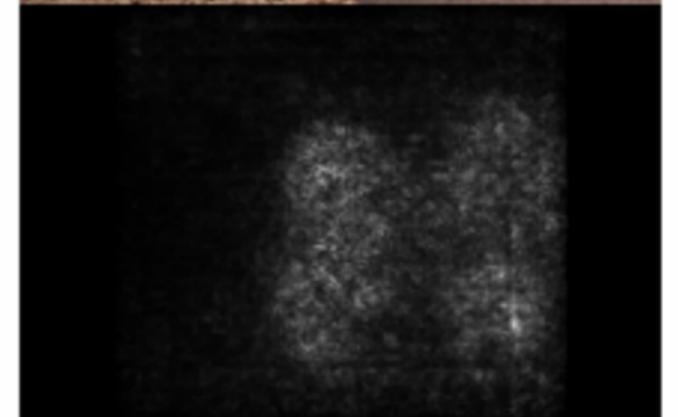


# Image classification: explaining the diagnosis

1. Different strategies
  1. Find input images that **maximally activate a given class score** and compare them according to some interpretable property
  2. Visualize how the network **responds to a specific input image**
2. Saliency map (Simonyan, Vedaldi, Zisserman, 2014):
  1. **Class model visualization**: generate an image that maximizes the (regularized) class score
  2. **Image-specific class-specific saliency map**: plots the derivative of the score function for a given class with respect to each pixel



**ostrich**



# VISUALIZING DEEP NEURAL NETWORK DECISIONS: PREDICTION DIFFERENCE ANALYSIS

**Algorithm 1** Evaluating the prediction difference using conditional and multivariate sampling

**Input:** classifier with outputs  $p(c|x)$ , input image  $\mathbf{x}$  of size  $n \times n$ , inner patch size  $k$ , outer patch size  $l > k$ , class of interest  $c$ , probabilistic model over patches of size  $l \times l$ , number of samples  $S$

**Initialization:**  $\text{WE} = \text{zeros}(n \times n)$ ,  $\text{counts} = \text{zeros}(n \times n)$

**for** every patch  $\mathbf{x}_w$  of size  $k \times k$  **in**  $\mathbf{x}$  **do**

$\mathbf{x}' = \text{copy}(\mathbf{x})$

$\text{sum}_w = 0$

    define patch  $\hat{\mathbf{x}}_w$  of size  $l \times l$  that contains  $\mathbf{x}_w$

**for**  $s = 1$  **to**  $S$  **do**

$\mathbf{x}'_w \leftarrow \mathbf{x}_w$  sampled from  $p(\mathbf{x}_w | \hat{\mathbf{x}}_w \setminus \mathbf{x}_w)$

$\text{sum}_w += p(c | \mathbf{x}')$

**end for**

$p(c | \mathbf{x} \setminus \mathbf{x}_w) := \text{sum}_w / S$

$\text{WE}[\text{coordinates of } \mathbf{x}_w] += \log_2(\text{odds}(c | \mathbf{x})) - \log_2(\text{odds}(c | \mathbf{x} \setminus \mathbf{x}_w))$

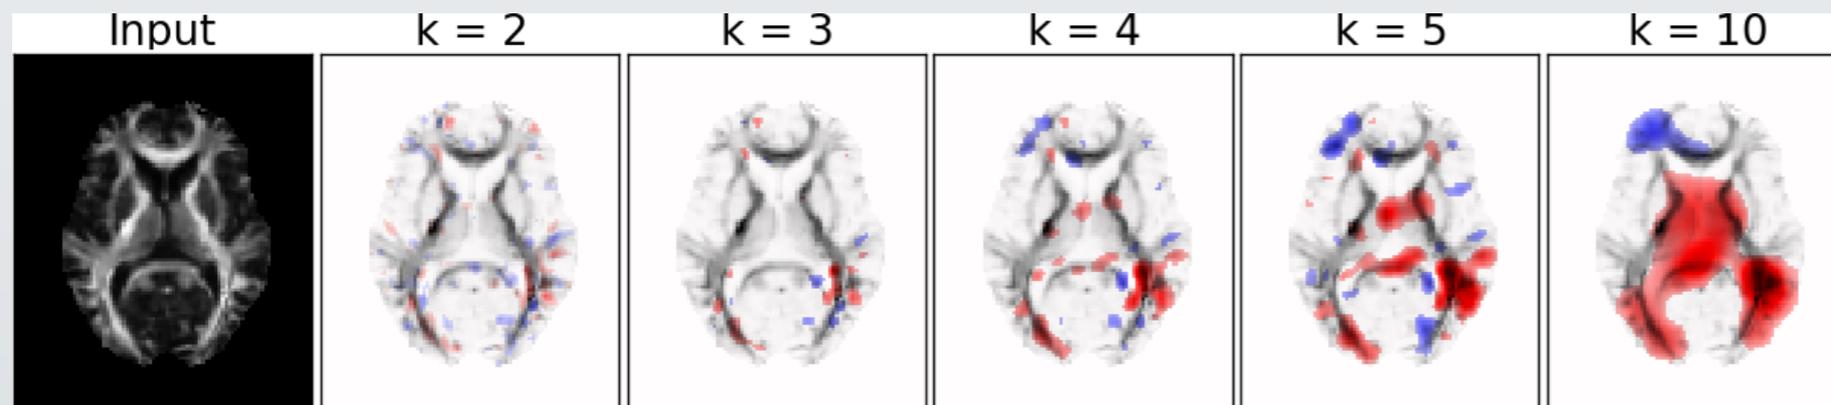
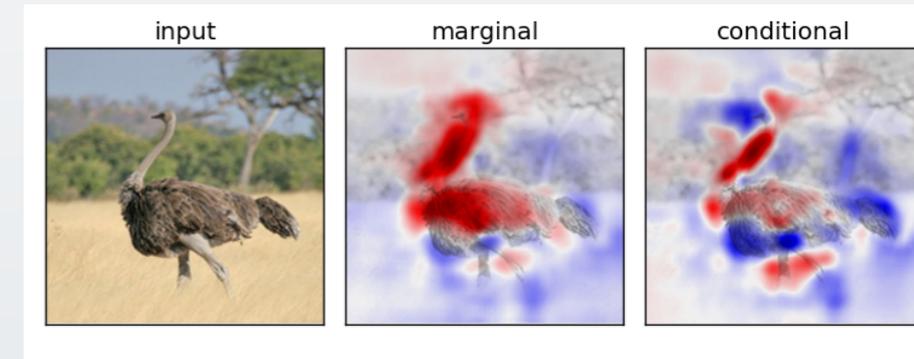
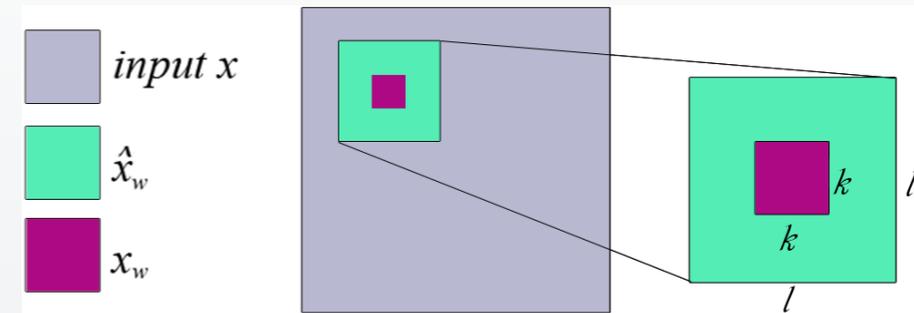
$\text{counts}[\text{coordinates of } \mathbf{x}_w] += 1$

**end for**

**Output:**  $\text{WE} / \text{counts}$

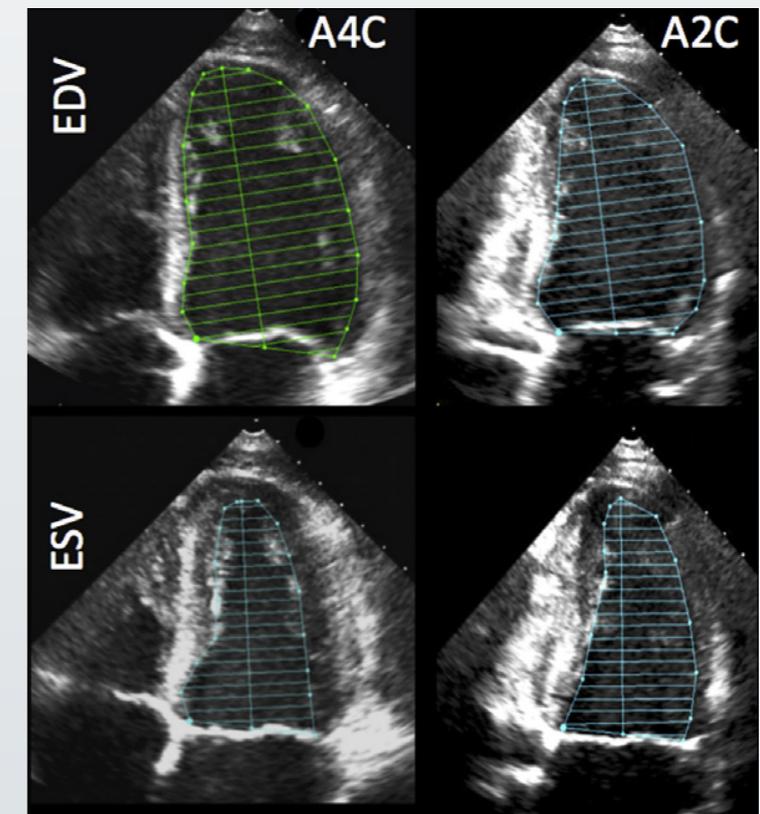
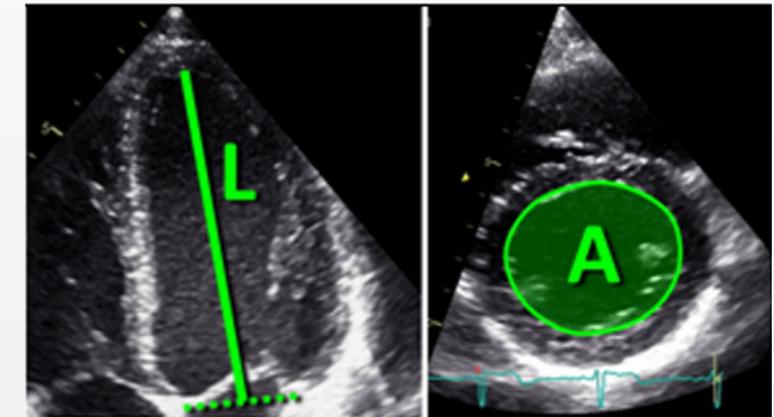
▷ evaluate classifier

▷ point-wise division



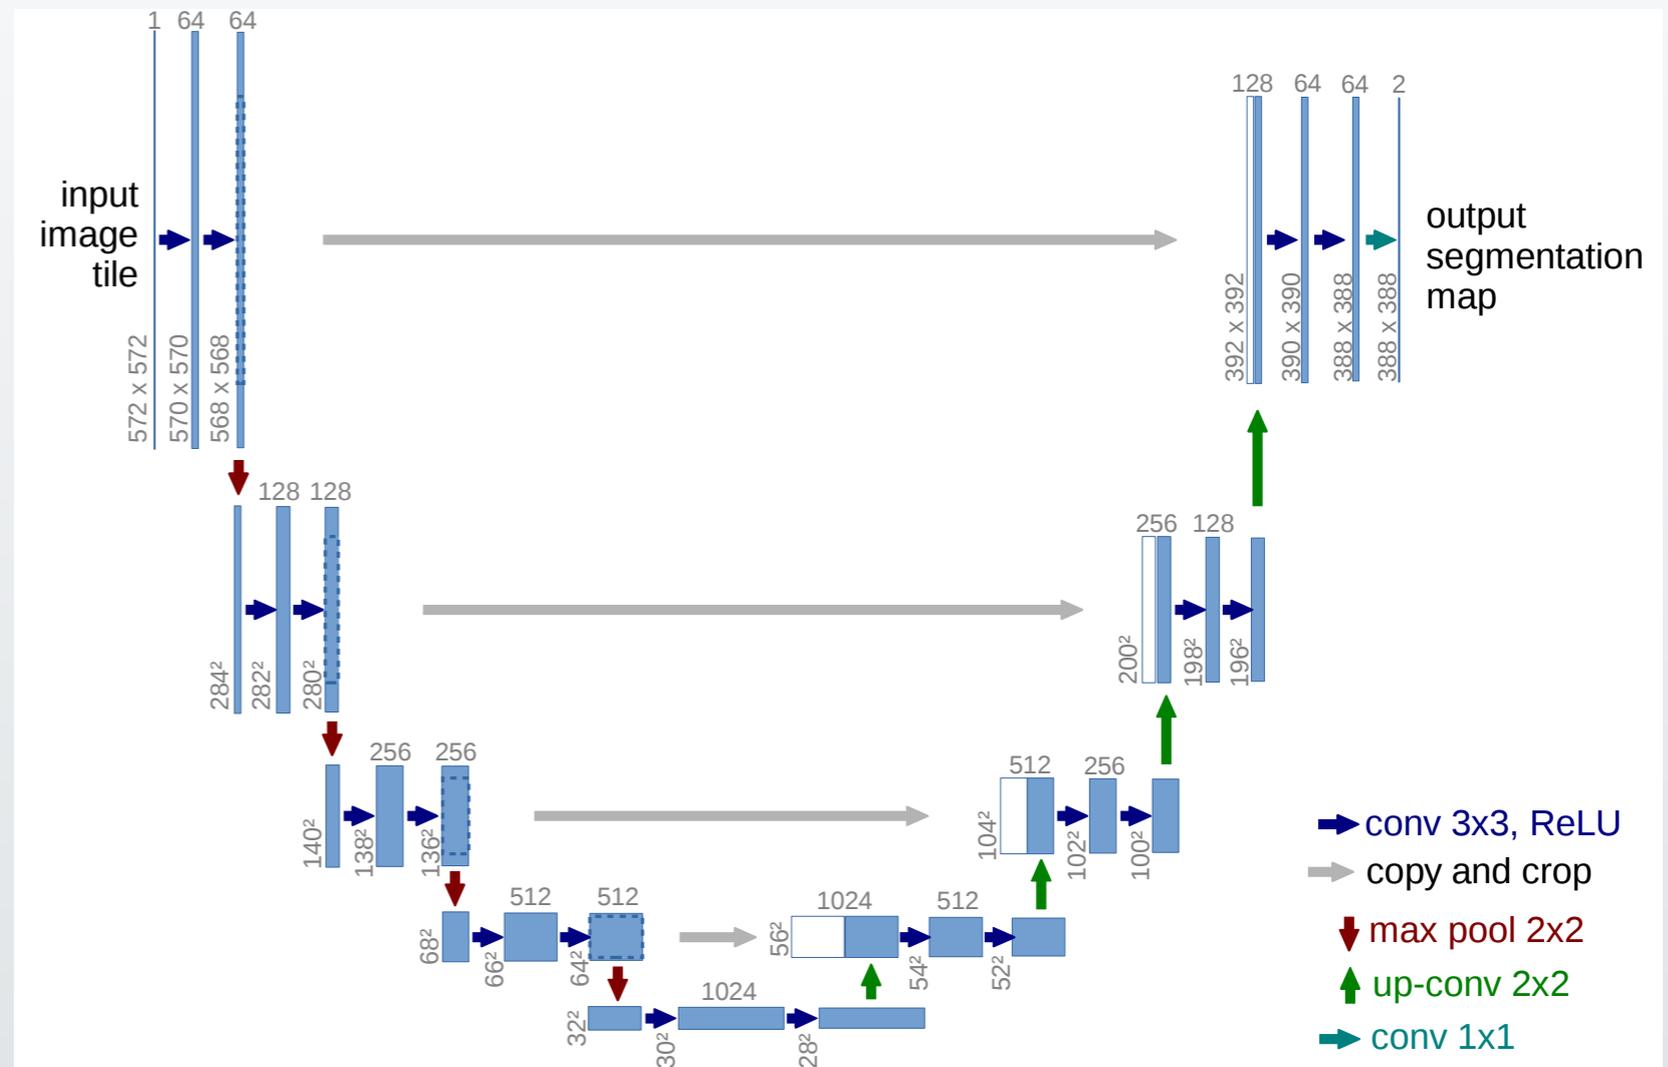
# Semantic segmentation: localizing structures

1. **Semantic segmentation**: assigning labels to individual pixels
2. **Quantification of cardiac function** involves estimating the heart's volume at its upper and lower limits
  1. This requires **delineating the boundaries of the heart**: a segmentation task
3. A diagnosis (classification) may also require **limiting the field of view** to a given structure
4. Many radiology reports involve **providing measurements** (lengths, areas) for basic structures

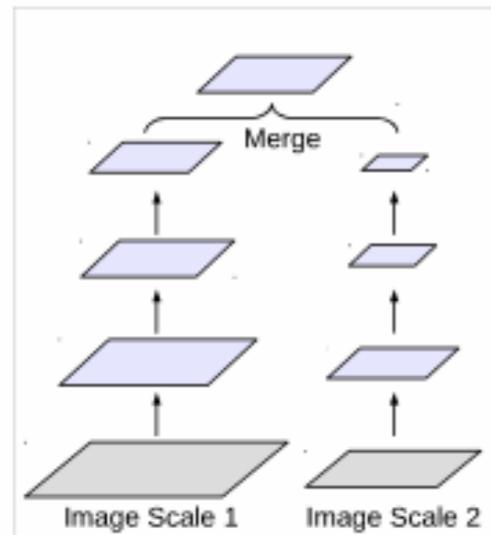


# The U-Net Architecture for Semantic Segmentation

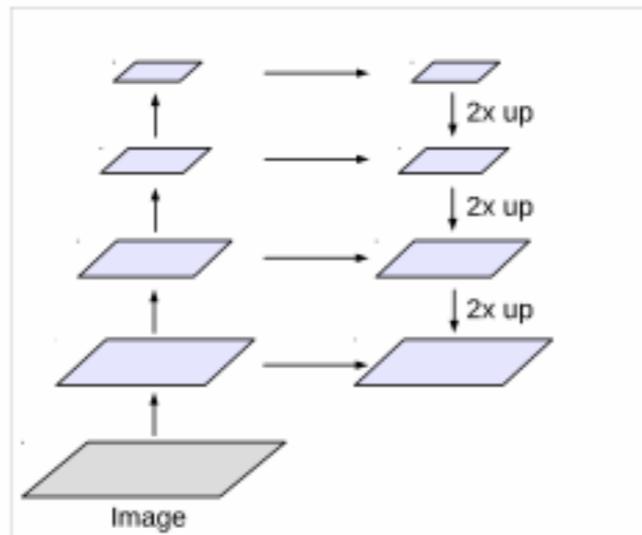
1. **The contraction/ downsampling layer** provides a representation of the context of the image
2. **The expanding/ upsampling layer** maps contextual features to the appropriate localization
3. **Skip connections** concatenate images from the downsampling layer to the upsampling one



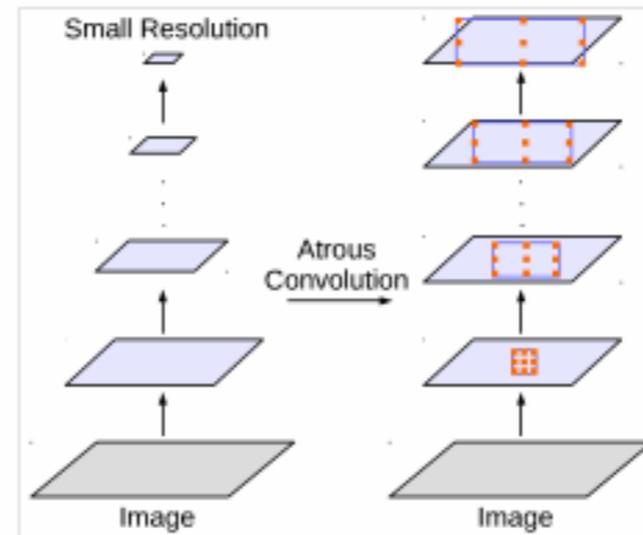
# Various architectures help incorporate global features and contextual interactions



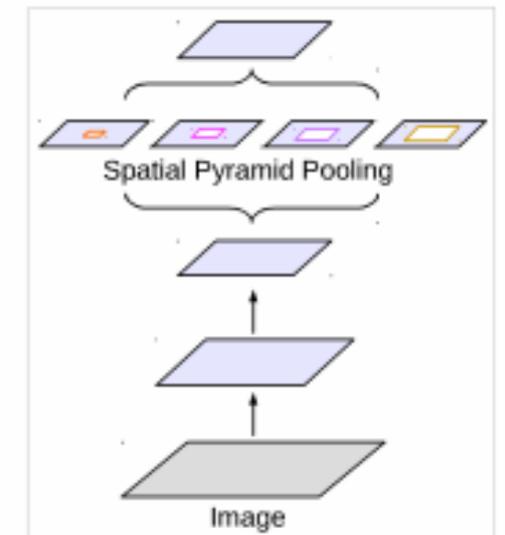
(a) Image Pyramid



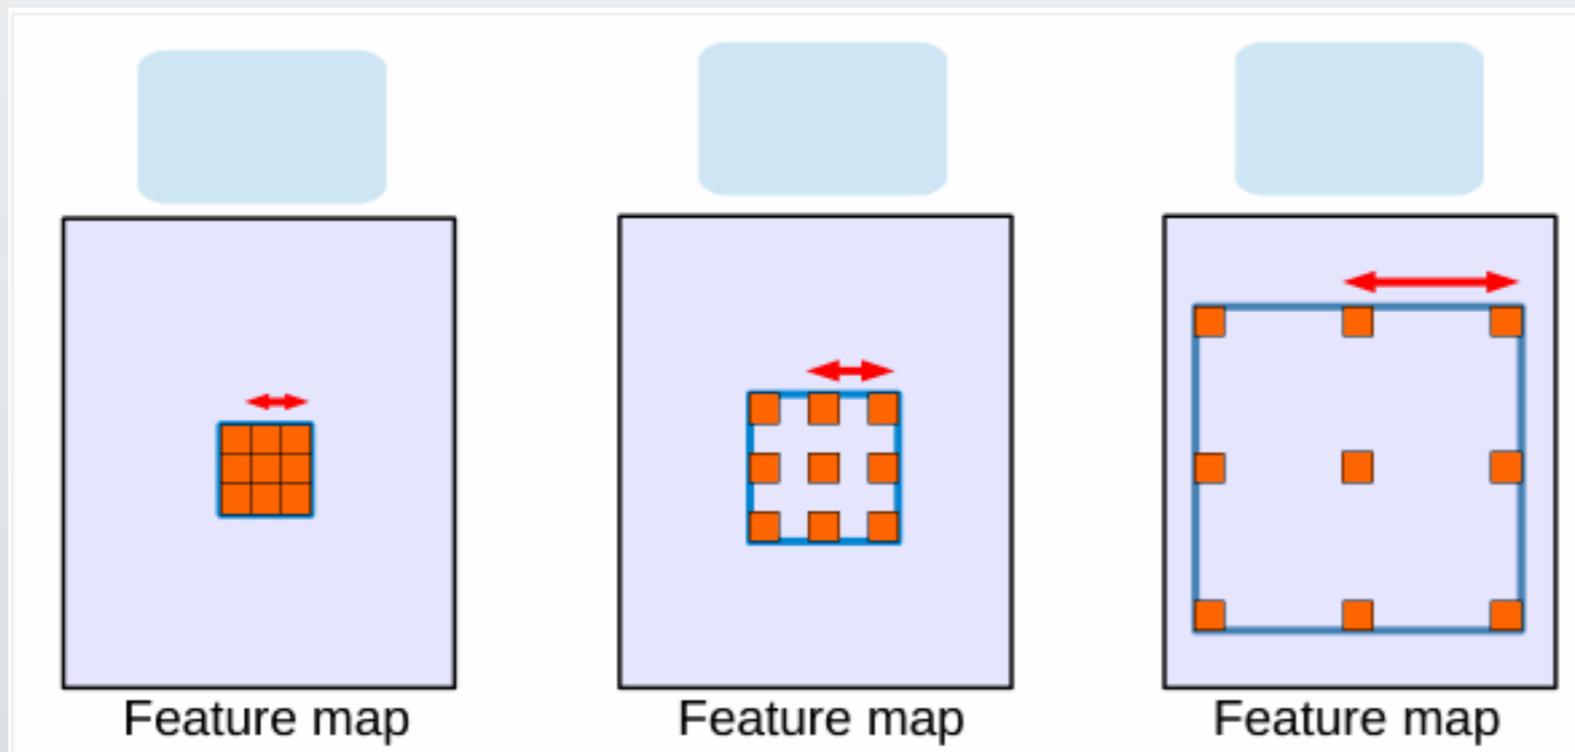
(b) Encoder-Decoder



(c) Deeper w. Atrous Convolution



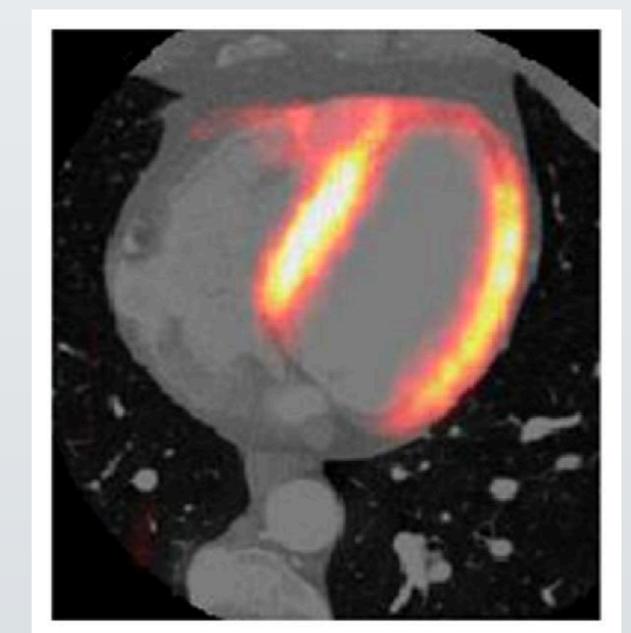
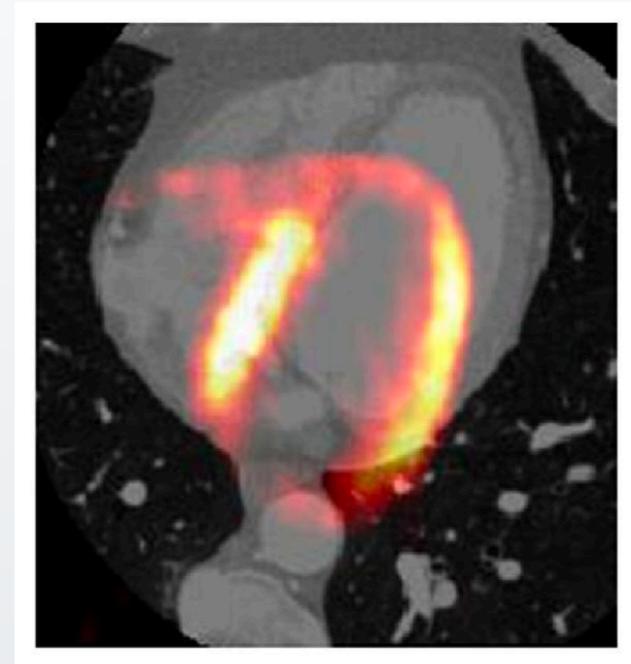
(d) Spatial Pyramid Pooling



Chen et al, arXiv, 2017

# Image registration: aligning different images

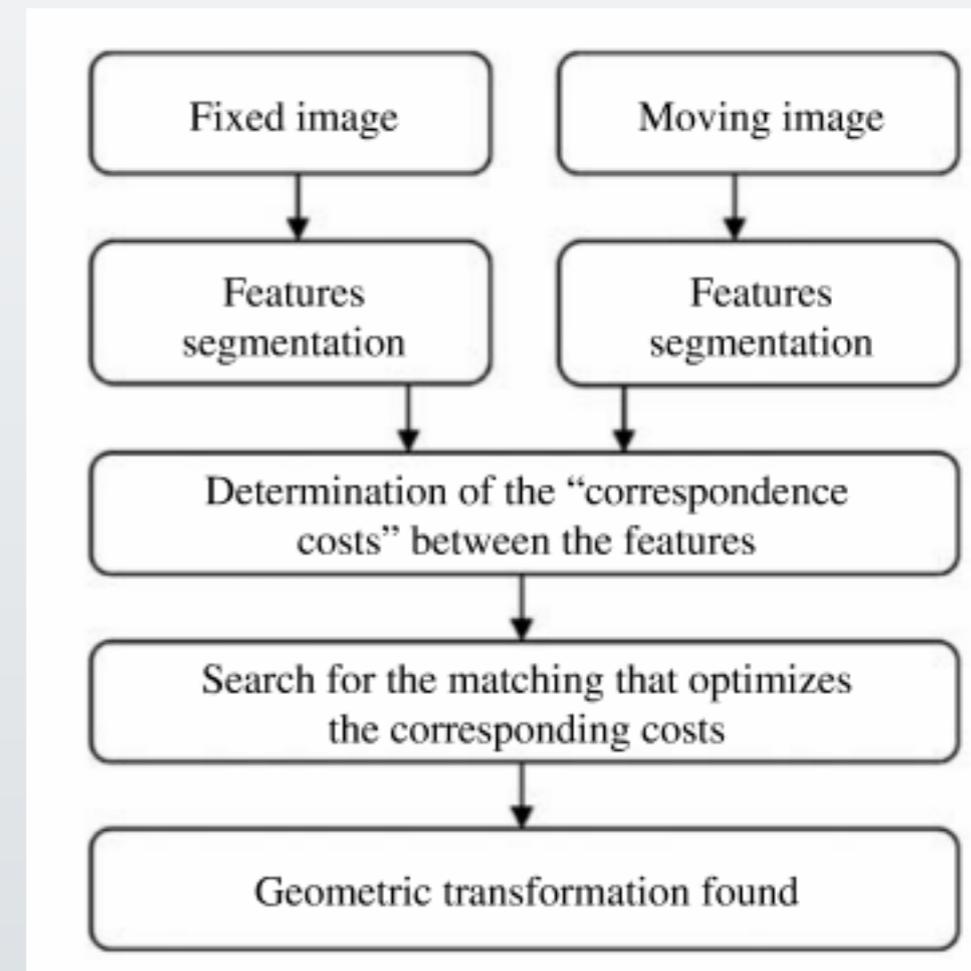
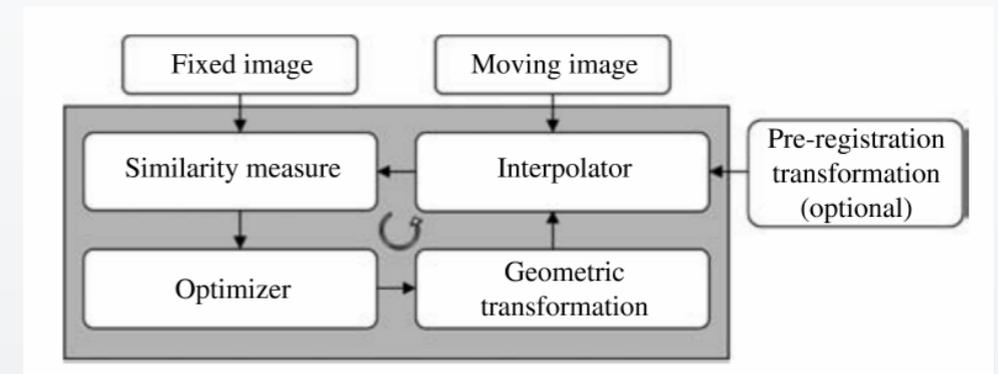
1. There is sometimes a need to **merge images from different modalities** or from **different time points** in the same study
2. In cardiac imaging this is relevant for merging a study with poor spatial resolution with one which is superior but may lack functional information: PET + CT
3. The low temporal resolution of PET also requires **averaging across many cardiac cycles** (ECG-based gating) and there may be movement of the thorax (breathing) during this time



# Image registration methods: pre- and post-CNNs

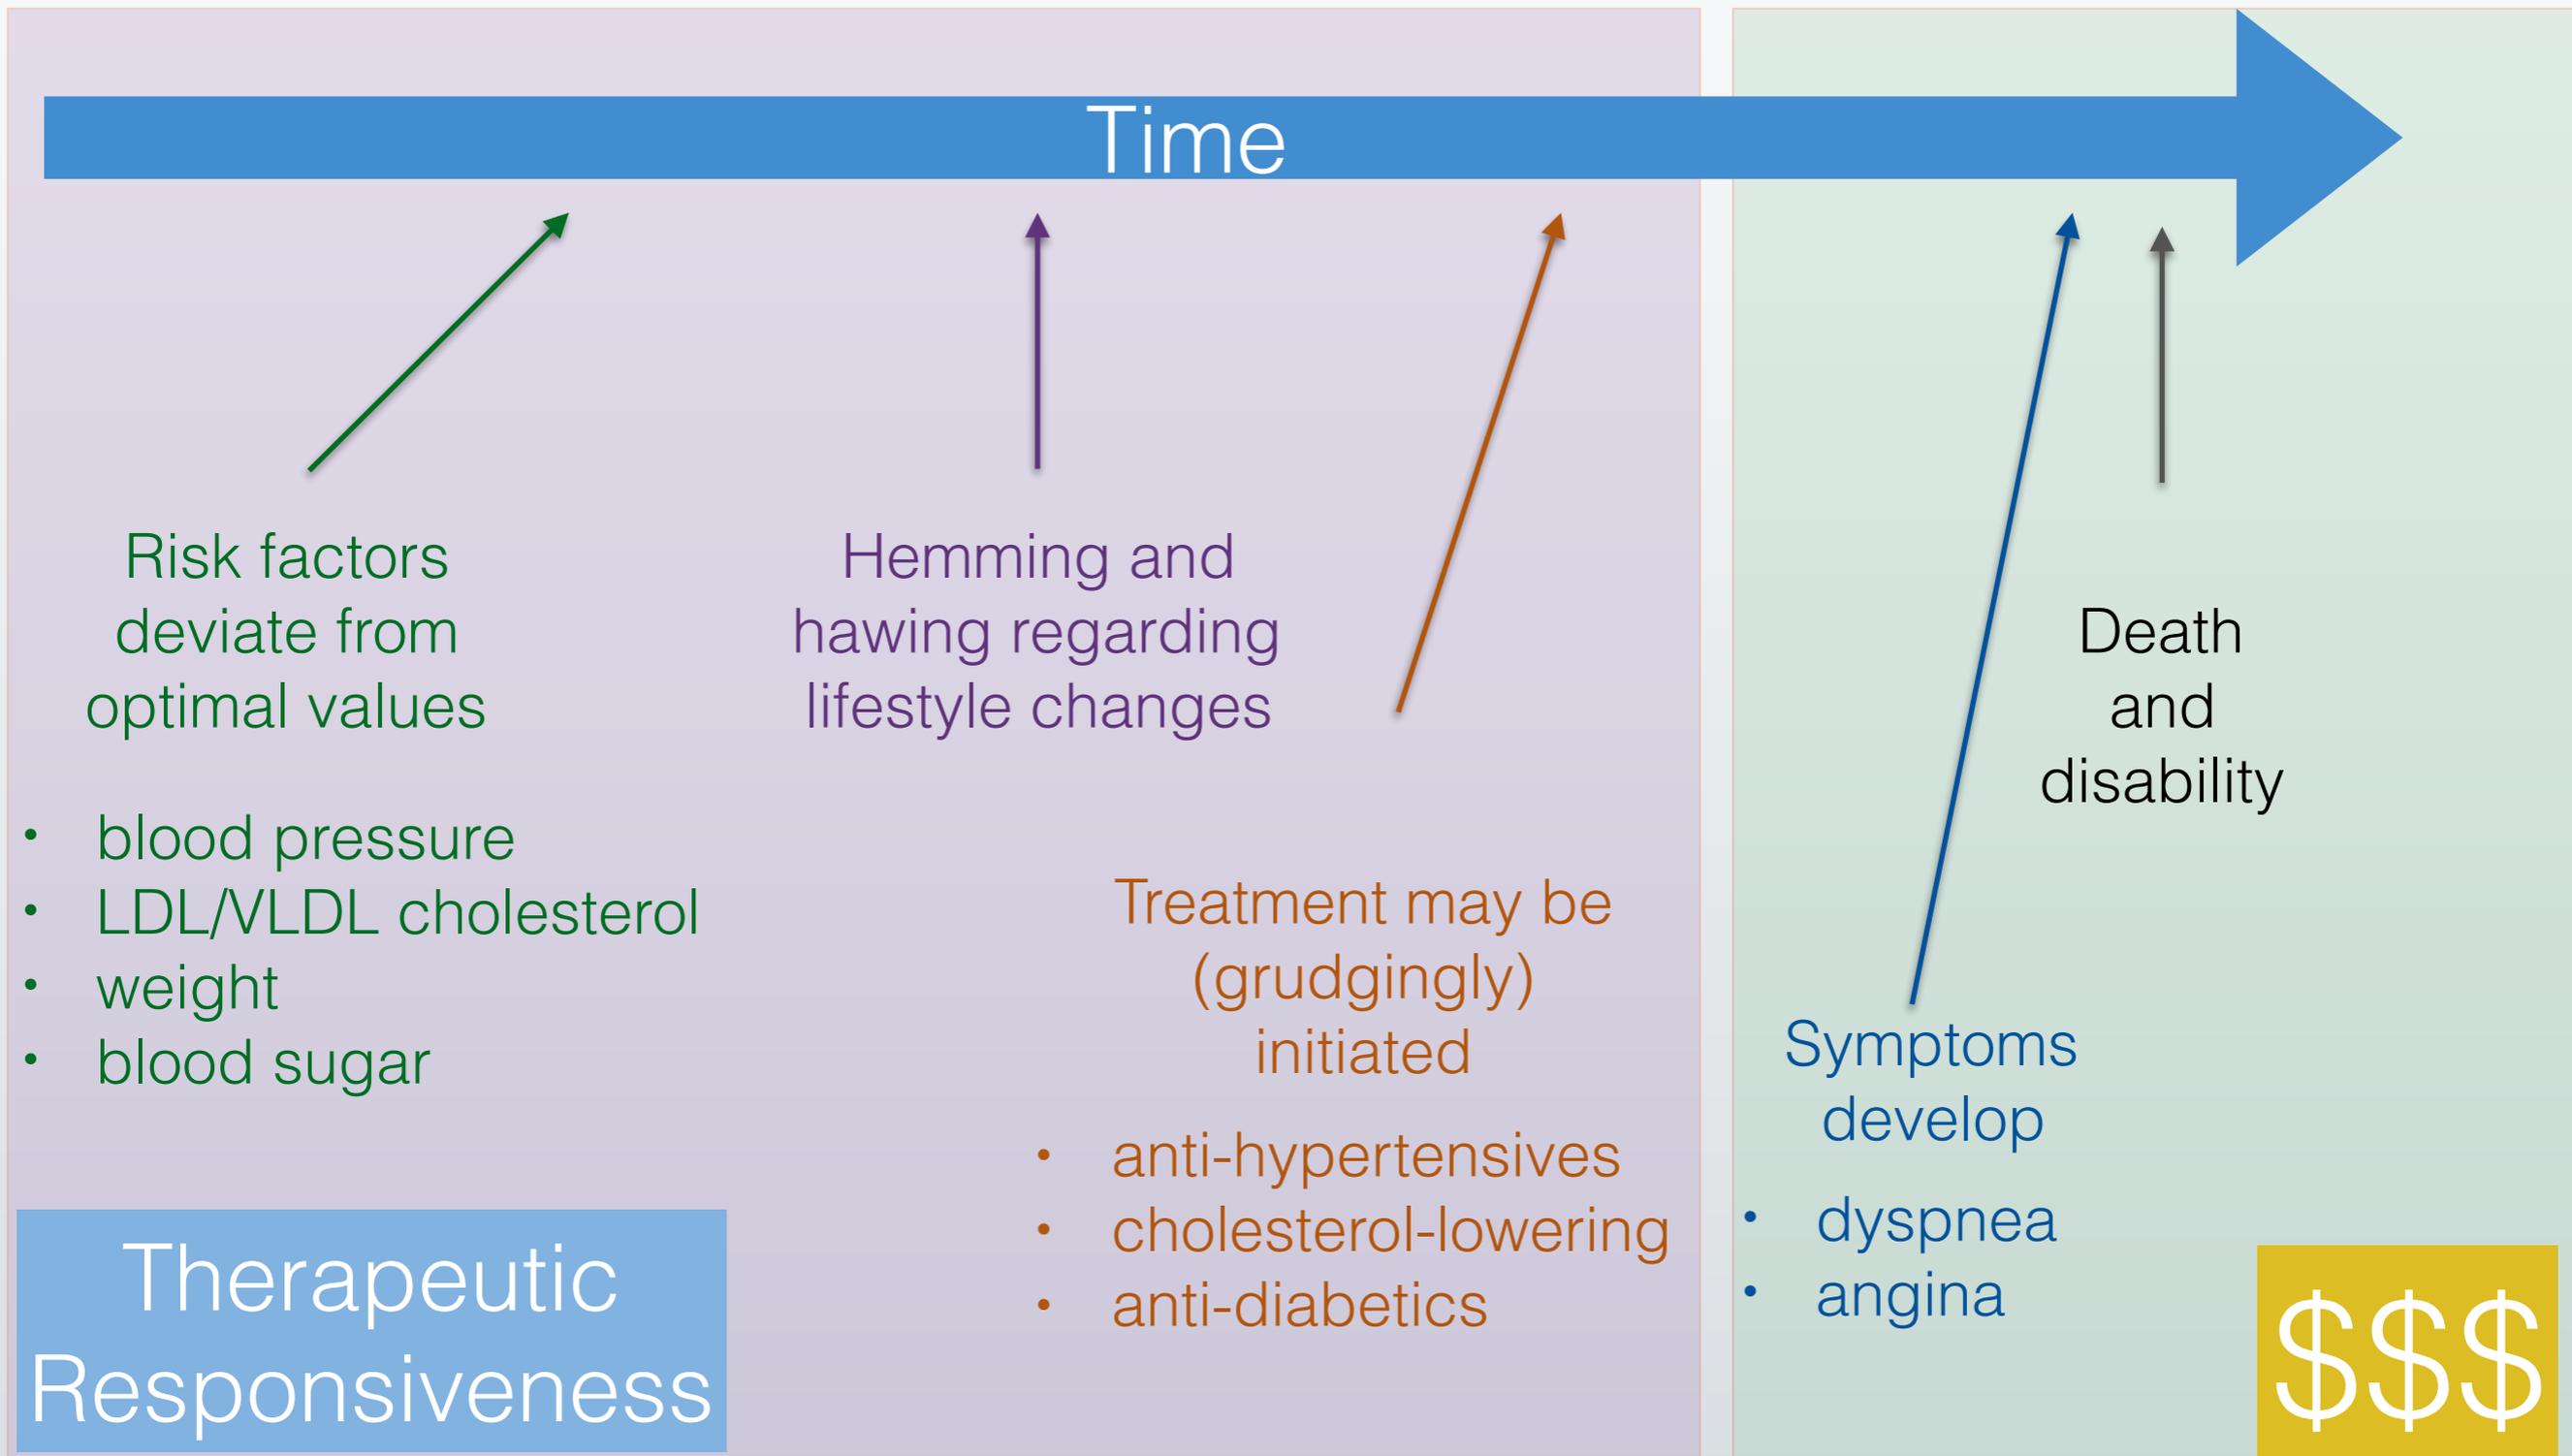
## 1. Classification

1. Intensity domain vs. frequency domain
  2. Raw intensities vs. feature-based
  3. Global (whole image) vs. local (region of interest)
  4. Type of transformation used to relate one image to the other
  5. Monomodal vs. multimodal
2. Additional choice of **similarity metric** as well as algorithms to search parameter space for geometric transformation
  3. **Conditional variational autoencoder** have been explored to learn geometric transformations between pairs of images (Krebs ... Mansi, arXiv:1812.07460, 2018)



# **A Fully Automated Echocardiogram Interpretation Pipeline**

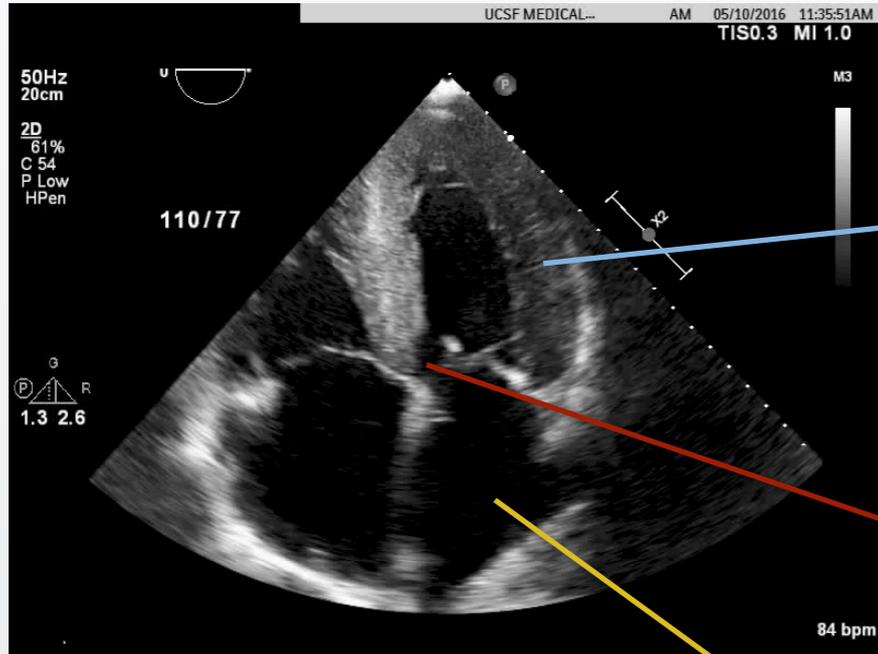
# The Failure of our Current Approach to Cardiovascular Disease



# What sort of solution are we looking for?

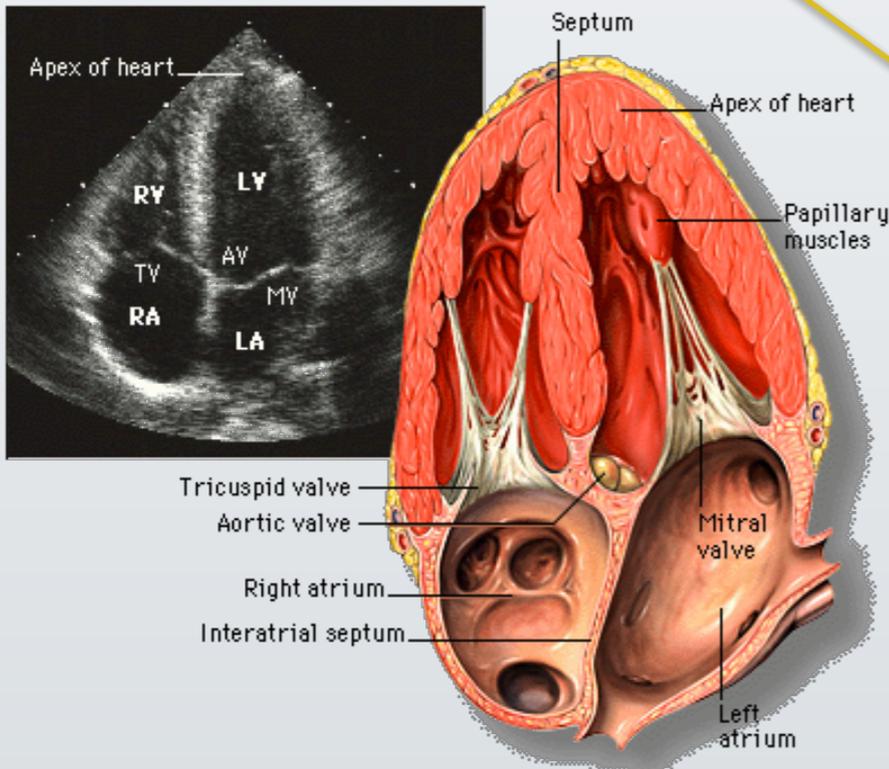
1. **Low-cost quantitative metrics** that are indicative of disease progression and reflect the onset of these tissue-level changes
2. Should be **specific** to the disease process:
  1. expressive: capture complex underlying processes (molecular, cellular, imaging ...)
  2. multidimensional: can't readily be "gamed"
3. Should be **ameliorated** with therapy (c.f. genetic risk)

# Simple cardiac ultrasound views provide a quantitative metric of early disease progression



Left ventricular mass increases with disease

Left ventricular function diminishes with disease



Left atrial volume increases with disease

# A role for automated interpretation at the “low risk - high reward” portion of the spectrum

Non-skilled acquisition  
(primary care)

Low cost handheld ultrasound

Automated interpretation

Early in disease course

Decision support regarding initiation  
or intensification of therapy

Low liability



Skilled sonographer

Costly full ultrasound system

Expert cardiologist interpretation

Late in disease course

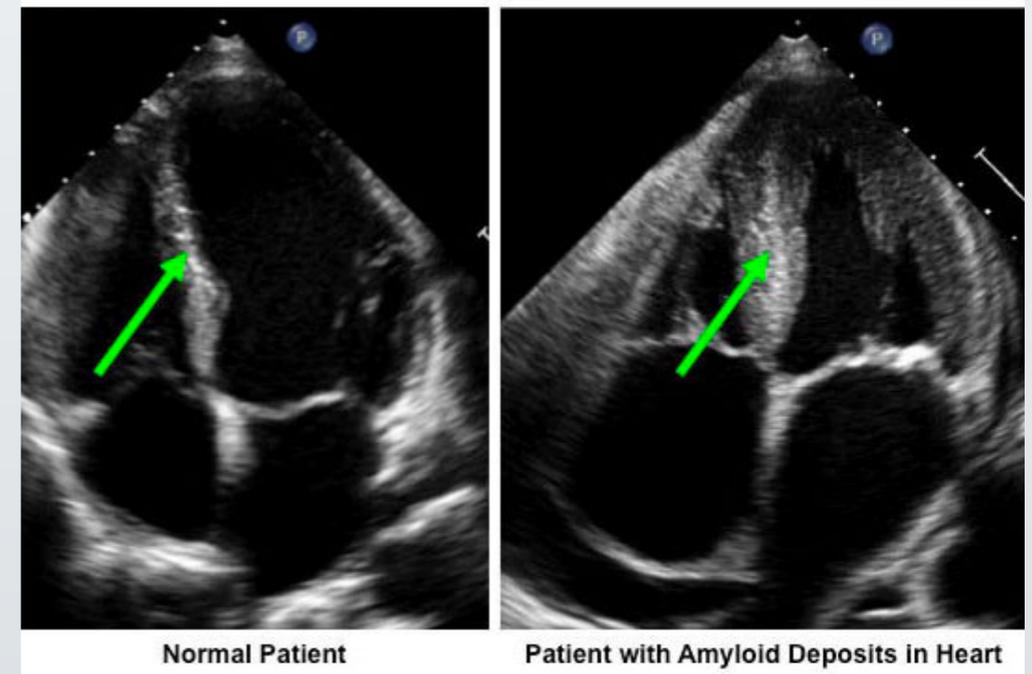
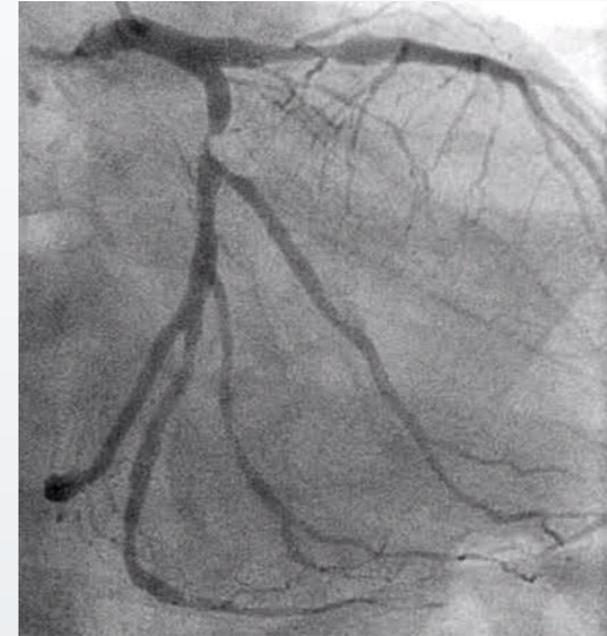
Difficult decisions regarding  
surgery

High liability



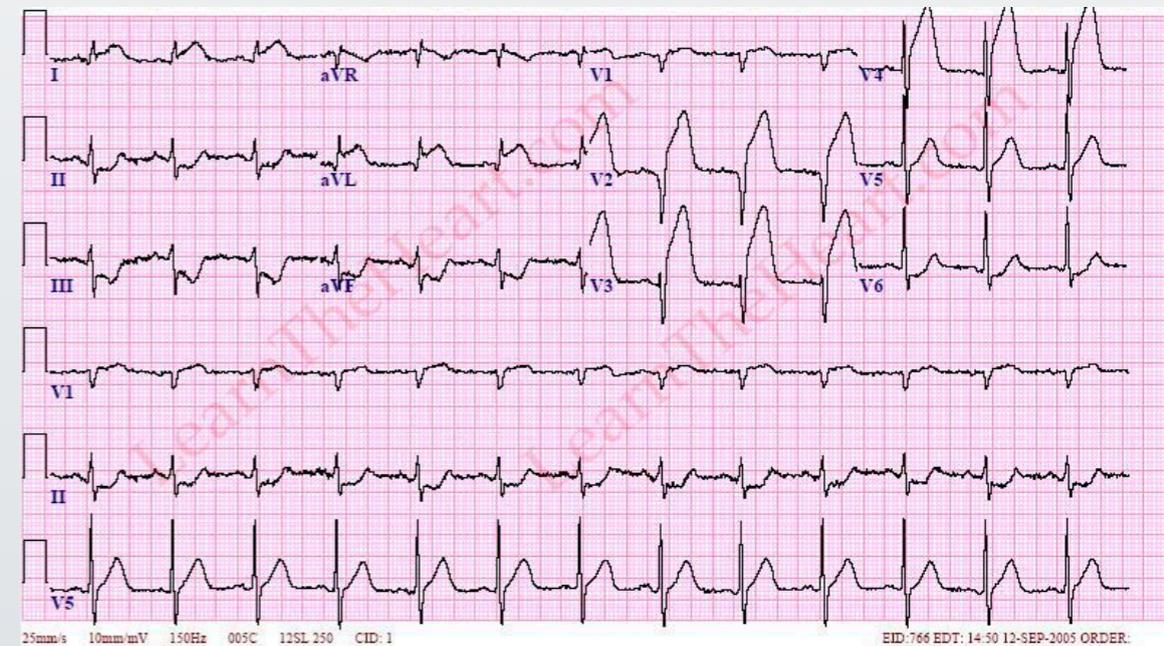
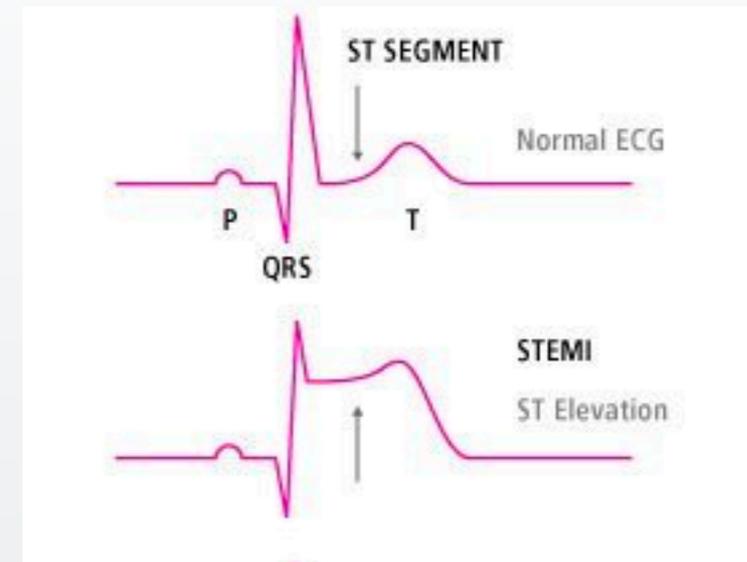
# Machine learning in cardiac disease - what should we be focusing on?

1. Enabling much **greater volumes of data** (tracking, clinical trial) by:
  1. **Reducing costs of acquisition**
  2. **Augmenting interpretation of simply acquired data** - i.e. diagnosing abnormalities of relaxation without Tissue Doppler
  3. Automating interpretation to **reduce costs**
2. **Surveillance within a hospital system:** patient identification for therapies
3. **Triage:** automated interpretation (lessons from ECGs in the ambulance/emergency room)
4. Can we go beyond what humans can see?
  1. Quantitative tracking of **intermediate states of disease** and assessing treatment response
  2. Recognizing meaningful **subclasses of disease** that differ in prognosis and treatment



# A role for rapid triage: the electrocardiogram in myocardial infarction

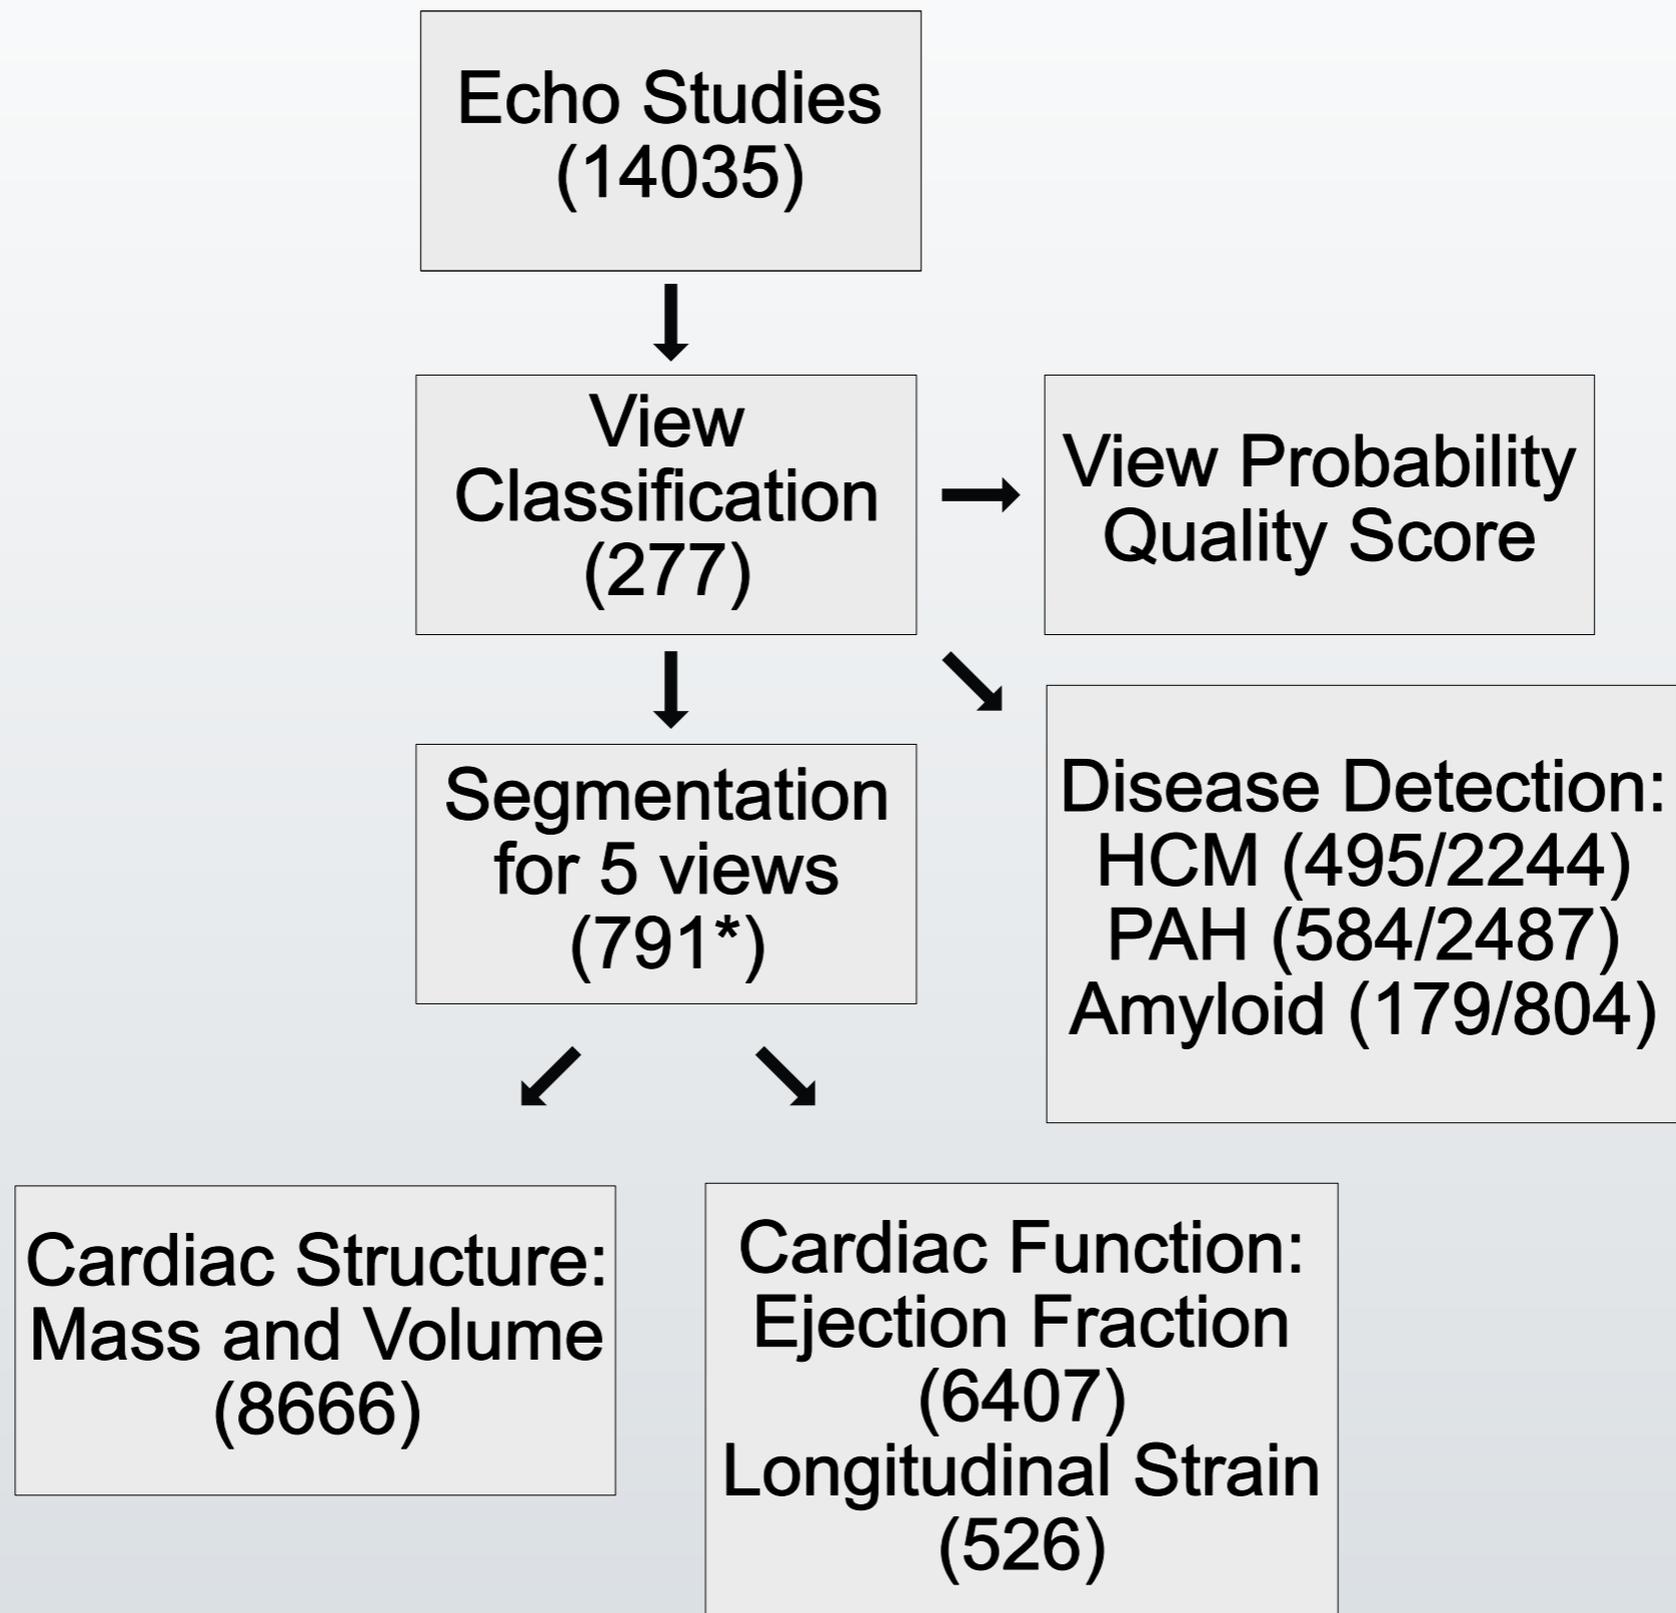
1. The ST-elevation myocardial infarction pattern in the ECG arises from complete obstruction of blood flow to portions of the heart
2. In the early 2000's it was recognized that any delay in angioplasty and stenting would result in irreversible damage to the heart
3. The previous approach - with a cardiologist reviewing the ECG before any action was taken was replaced with a rapid triage system by ambulance personnel or ED physicians
4. The cardiac catheterization team would be “activated” by non-cardiologists and needed to arrive to the hospital within 30 minutes
5. Some subsets of these activations were false positives



# What is in an Echo Study?

1. Typically a collection of up to **70 videos of the heart taken over multiple cardiac cycles and focusing on different viewpoints** (requiring ~45 min by a skilled sonographer)
2. Heart can be visualized from **>10 different views** - though not truly discrete classes as sonographer can zoom and angle the probe to focus on structures of interest. These are typically unlabeled.
3. Still images are typically included to enable manual measurements
4. UCSF performs 12,000-15,000 echo studies per year; BWH performs 30,000-35,000 studies
5. There are **>7,000,000 echos** performed annually in the Medicare population alone
6. There are likely 100,000,000's of archived echos

# An Automated (low-cost!) Approach to Echo Interpretation





## EchoCV Team



Berkeley  
UNIVERSITY OF CALIFORNIA

### Developers

Jeffrey Zhang  
Rahul Deo

### Wise Computer Vision

#### Gurus

Ruzena Bajcsy  
Alyosha Efros

### Herceptin Data

#### Acquisition

Mandar Aras  
Eugene Fan  
Kirsten Fleischmann  
Michelle Melisko  
ChaRandle Jordan  
Atif Qasim

### Project Design

Jeffrey Zhang  
Rahul Deo  
Geoff Tison  
Sanjiv Shah

### Image Labeling

Rahul Deo  
Sravani Gajjala  
Francesca Delling

### Strain Computation

Lauren Beusslink-Nelson  
Sanjiv Shah  
Atif Qasim  
Mats Lassen

### Computer Vision

#### Consultants

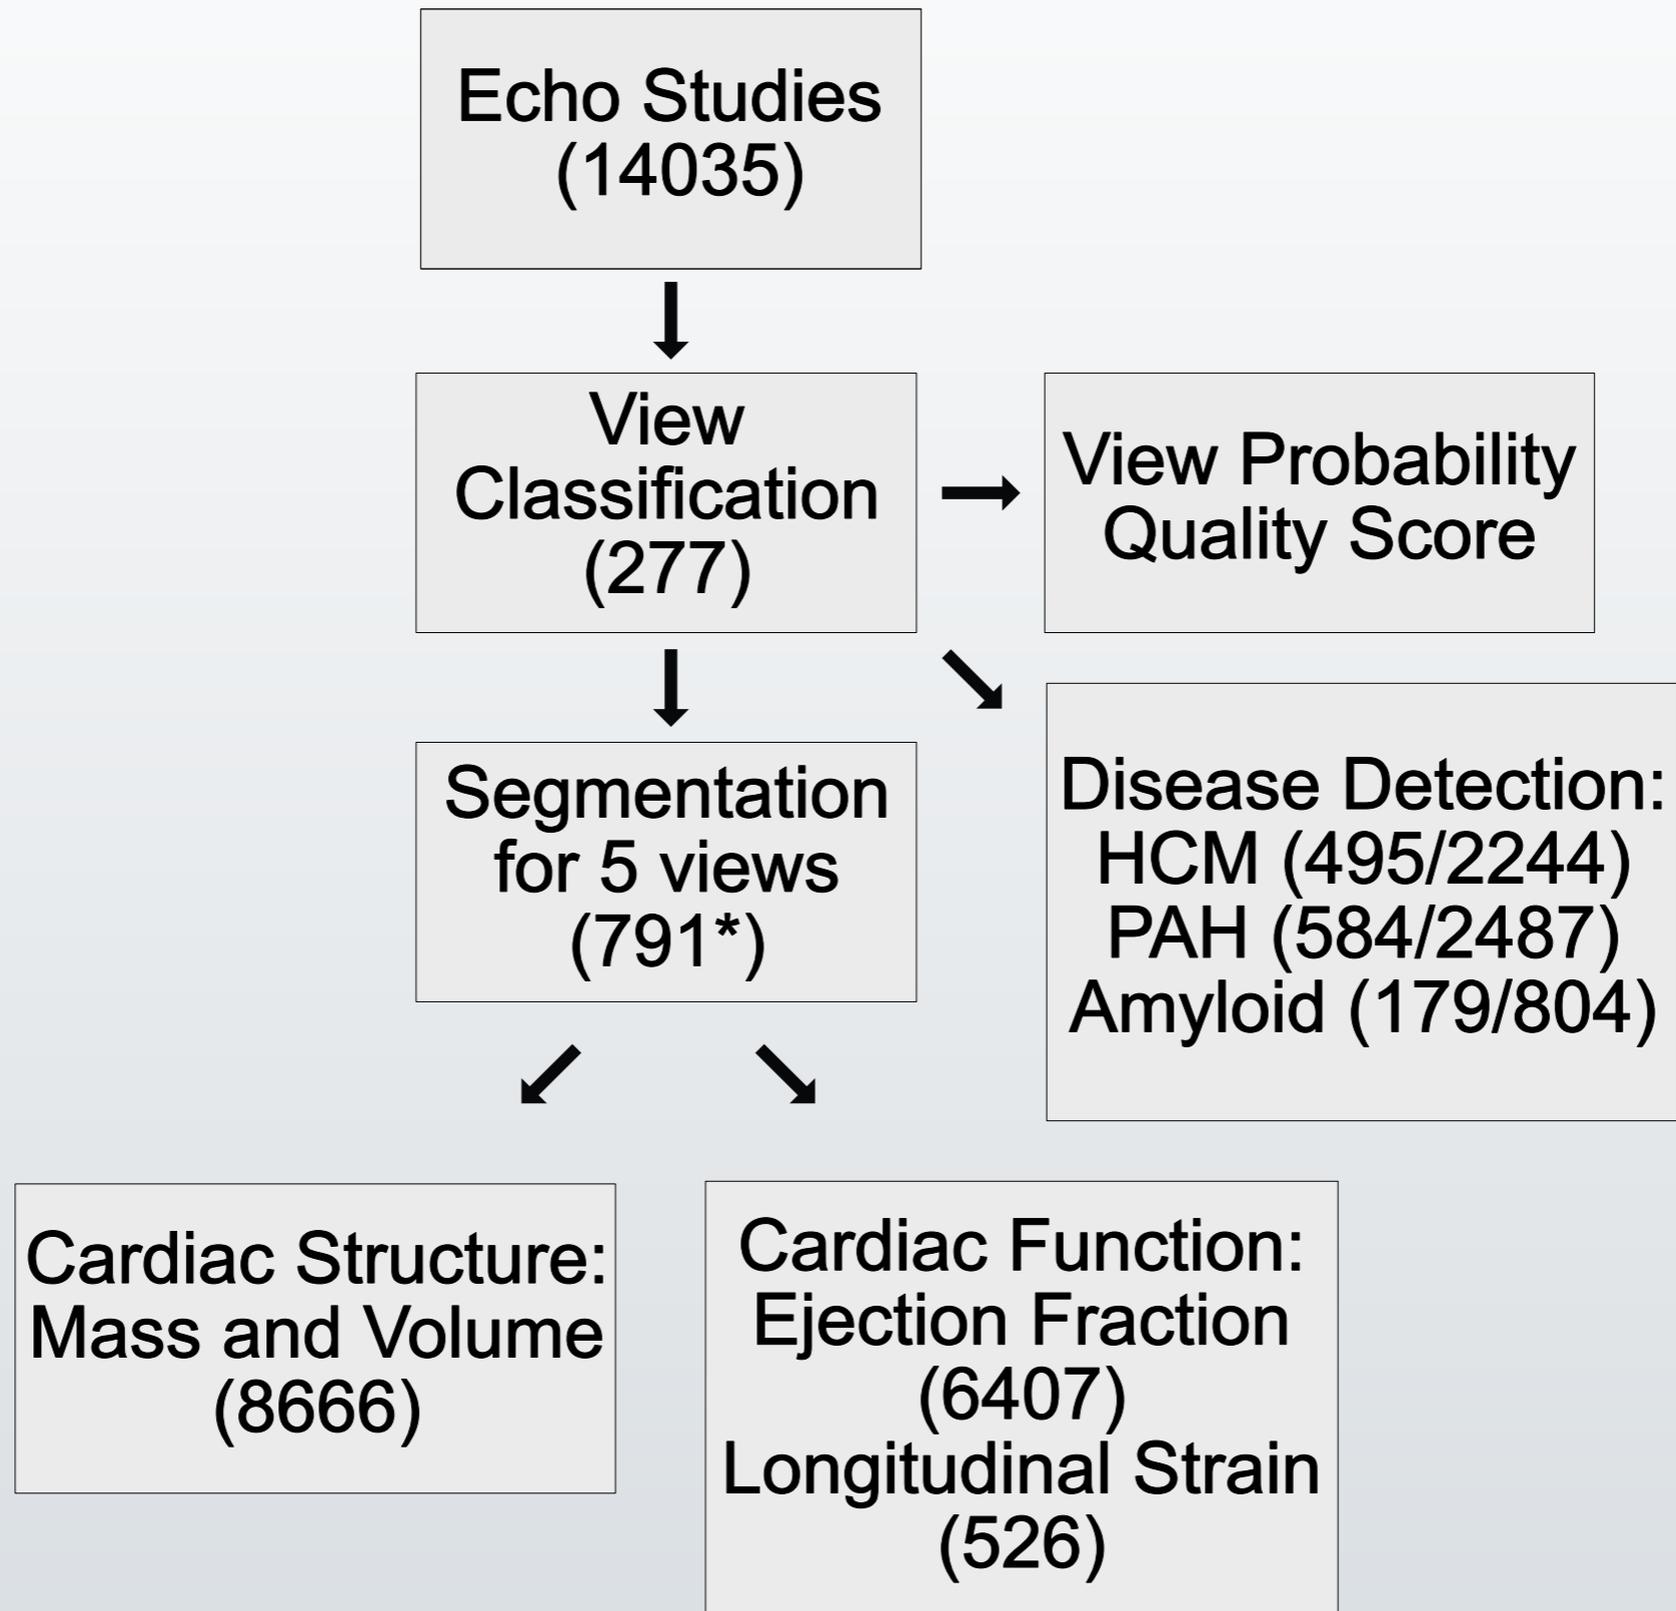
Laura Hallock  
Pulkit Agrawal

### Image Segmentation

Rahul Deo  
Sravani Gajjala  
Geoff Tison



# An Automated (low-cost!) Approach to Echo Interpretation



# View Classification - Someone Already Got To It Before Us



ELSEVIER

Contents lists available at [ScienceDirect](#)

## Information Fusion

journal homepage: [www.elsevier.com/locate/infus](http://www.elsevier.com/locate/infus)

## A fused deep learning architecture for viewpoint classification of echocardiography

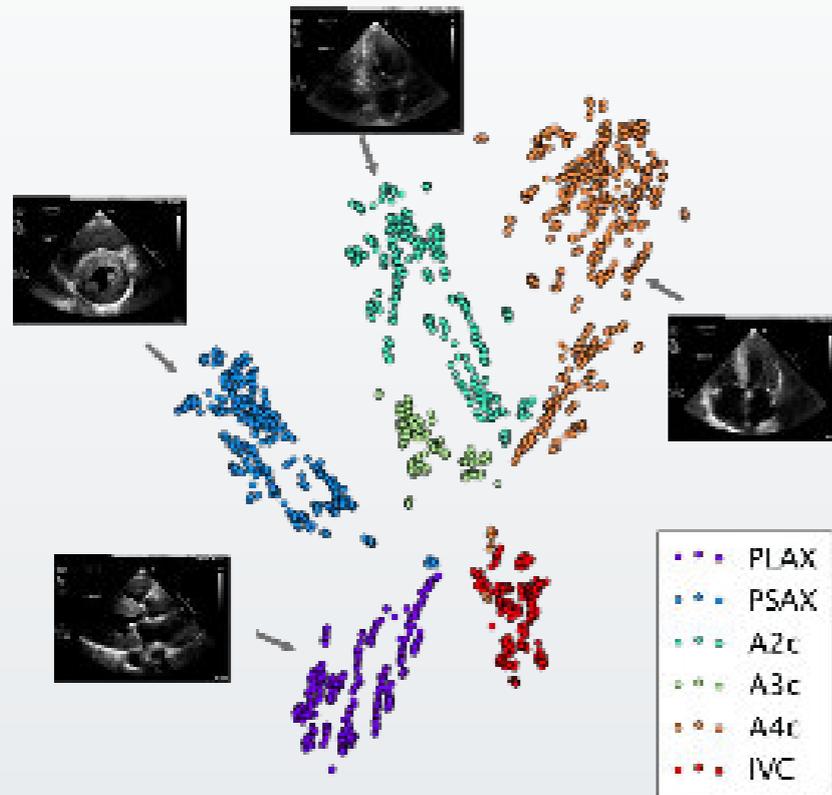
Xiaohong Gao<sup>a,\*</sup>, Wei Li<sup>a,b</sup>, Martin Loomes<sup>a</sup>, Lianyi Wang<sup>c</sup>

<sup>a</sup> Department of Computer Science, Middlesex University, London, NW4 4BT, UK

<sup>b</sup> Institute of Biomedical Engineering, Fuzhou University, Fuzhou, China

<sup>c</sup> Cardiovascular Centre, Tsinghua University Hospital, Beijing, China

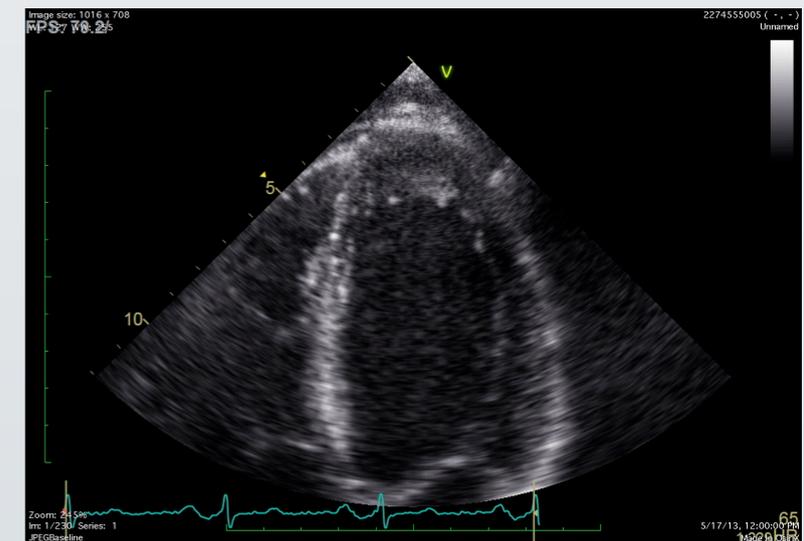
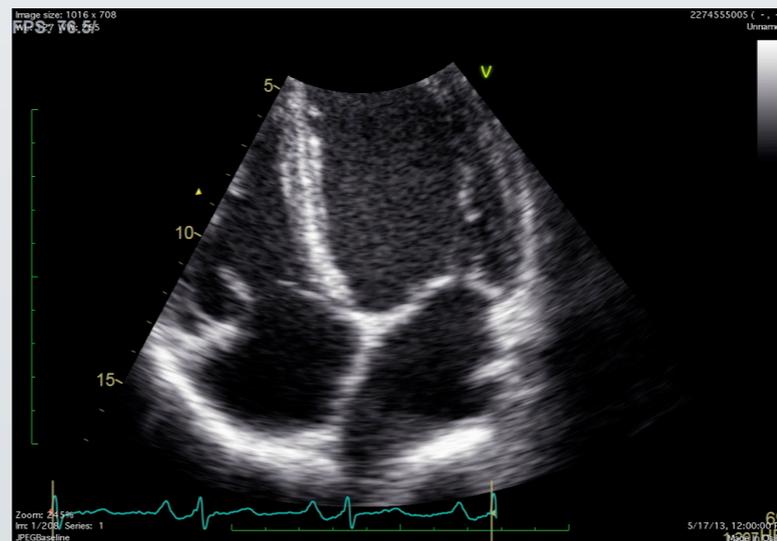
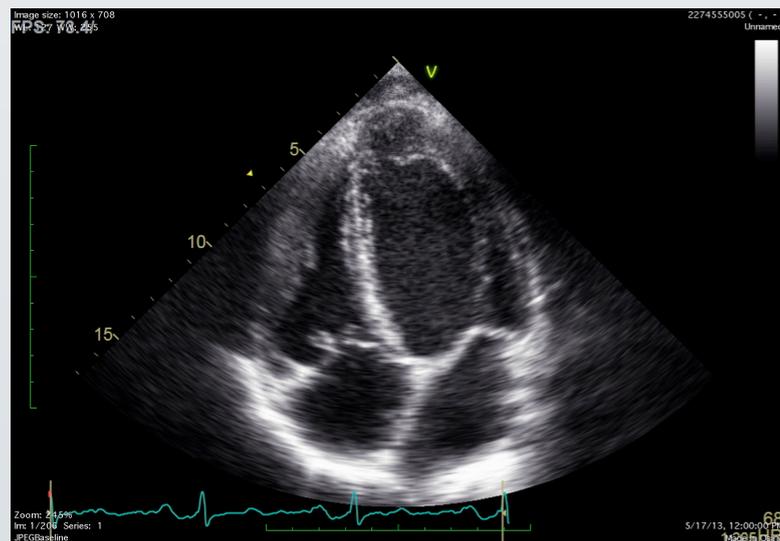
# View Classification - Our Take



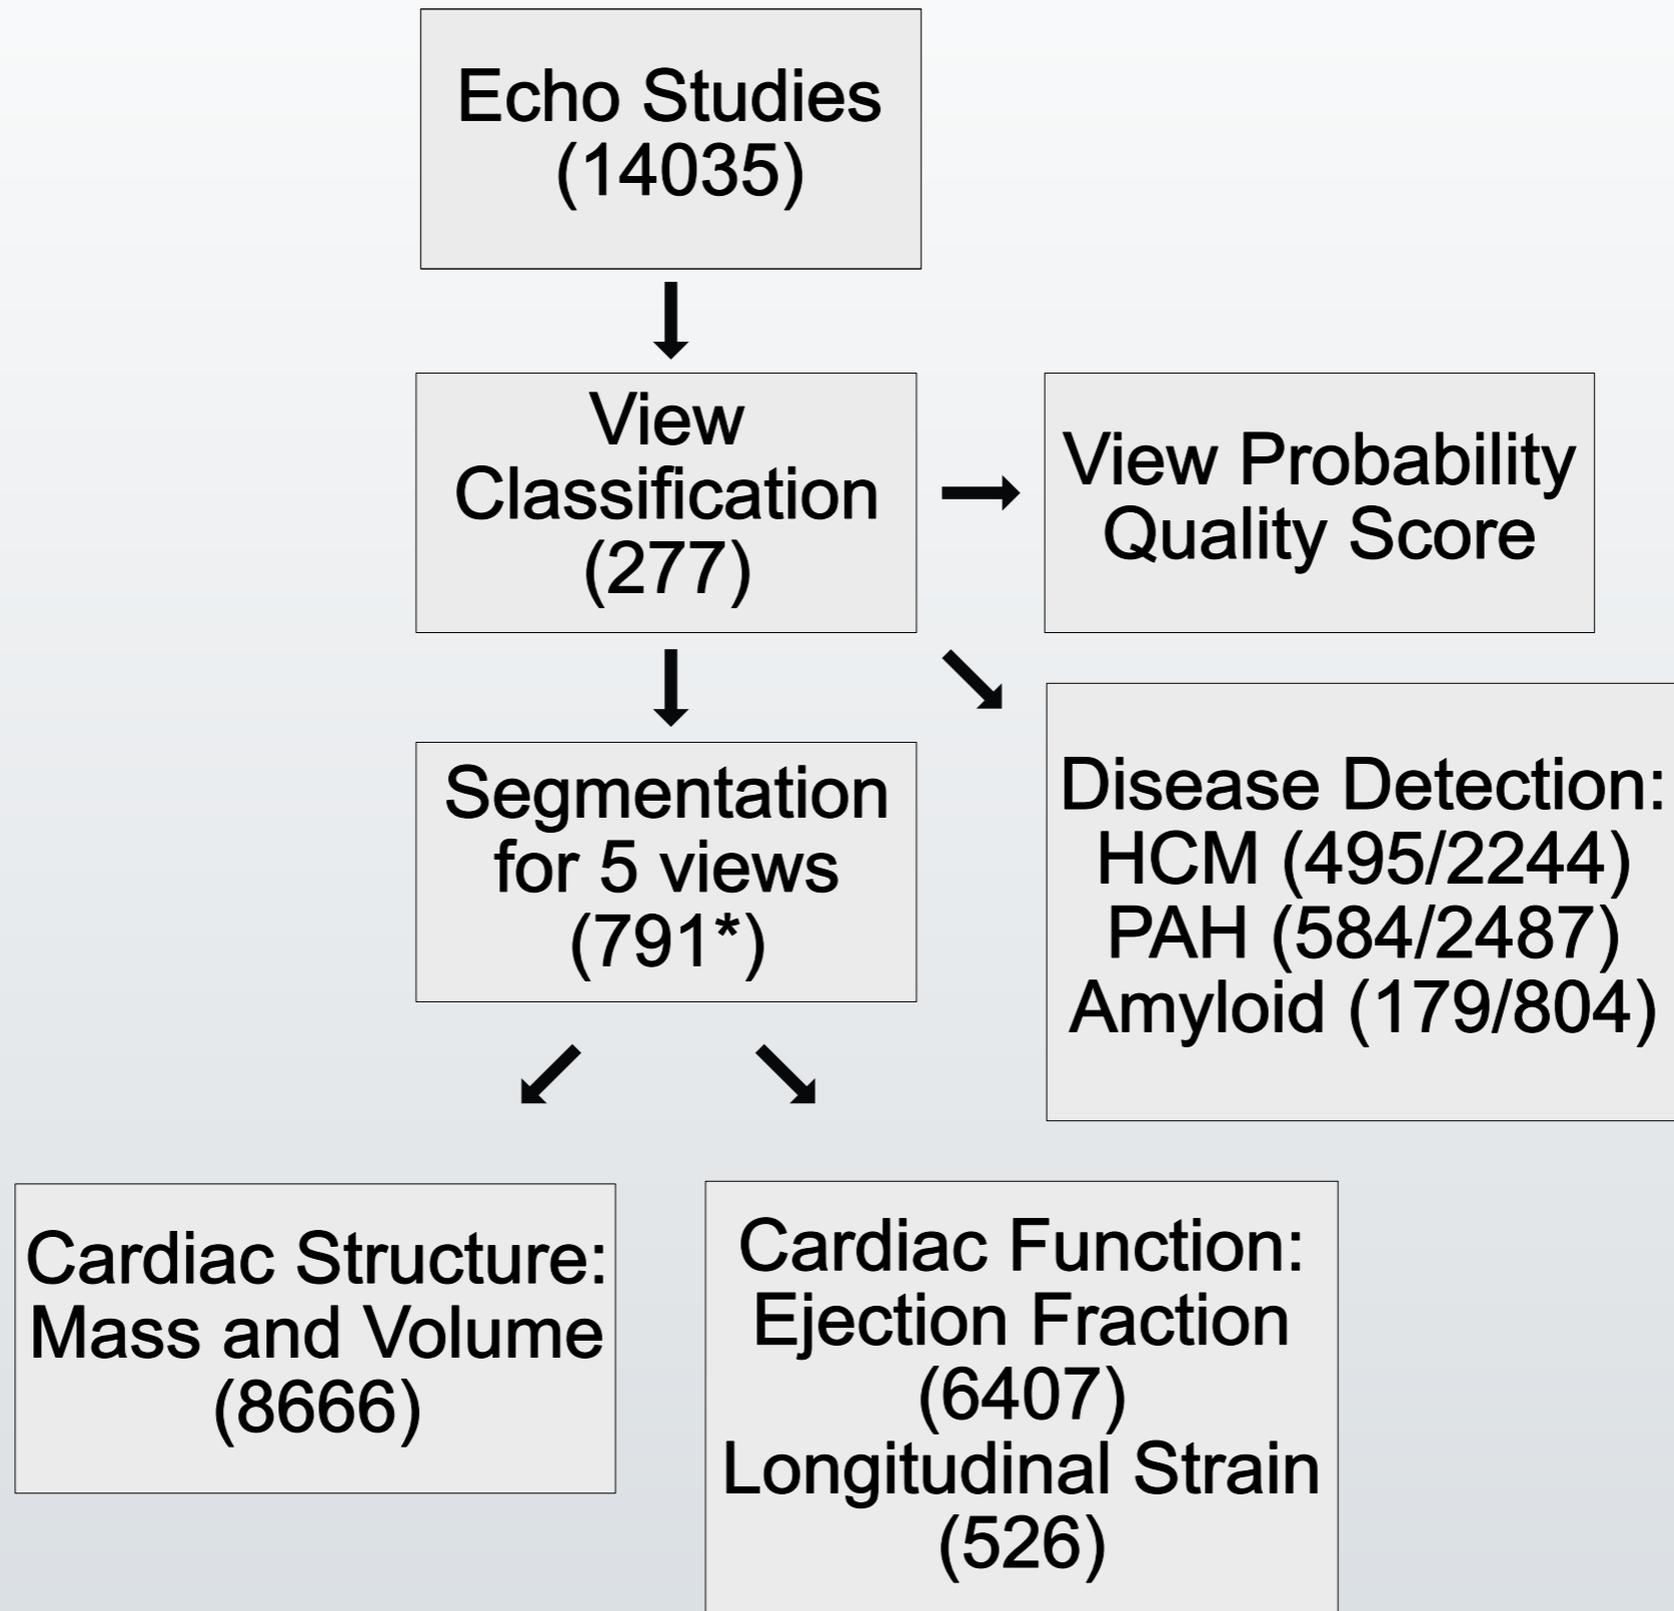
Prediction of echo views for individual images

Ground Truth	PLAX.remote	PLAX	PLAX.zoom of LA	PLAX.centered on LA	RV inflow	PSAX.apex	PSAX.PapMuscle	PSAX.MV	PSAX.AoV	PSAX.AoV zoom	A2c.no occlusions	A2c.occluded LA	A2c.occluded LV	A3c.no occlusions	A3c.occluded LA	A3c.occluded LV	A4c.no occlusions	A4c.occluded LA	A4c.occluded LV	A5c	Subcostal	Suprasternal	Other		
PLAX.remote	364	45	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
PLAX	13	1155	10	82	3	0	1	0	27	0	0	9	0	10	1	0	0	0	0	0	0	0	0	0	0
PLAX.zoom of LA	0	0	81	75	0	0	0	4	0	1	0	0	0	0	0	0	6	0	0	0	0	4	0	0	0
PLAX.centered on LA	0	21	52	107	2	0	0	0	7	2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
RV inflow	0	18	4	18	266	2	15	2	8	0	7	0	0	0	0	0	1	0	5	0	0	0	0	0	0
PSAX.apex	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
PSAX.PapMuscle	0	17	0	0	41	26	936	129	13	2	1	10	0	0	0	0	0	19	21	0	2	0	1	0	0
PSAX.MV	0	1	0	5	7	18	0	195	29	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
PSAX.AoV	0	17	3	14	0	0	0	4	766	20	0	0	0	0	0	0	0	1	0	1	0	15	2	0	0
PSAX.AoV zoom	0	0	29	1	0	0	1	8	39	290	0	0	0	0	0	0	0	4	0	0	0	3	13	0	0
A2c.no occlusions	0	1	0	0	0	4	0	0	9	0	1070	90	36	80	8	0	38	17	0	23	2	0	0	0	0
A2c.occluded LA	0	0	0	0	0	0	0	0	0	0	67	362	2	0	9	9	44	0	1	0	0	0	0	0	0
A2c.occluded LV	0	0	0	0	0	0	0	2	0	0	64	0	10	0	1	0	10	0	0	0	0	0	0	0	0
A3c.no occlusions	0	0	0	0	0	0	0	0	0	0	20	5	0	435	80	30	3	0	0	0	0	0	0	0	0
A3c.occluded LA	0	0	0	0	0	0	0	0	0	0	1	3	12	21	161	0	0	1	0	0	0	0	0	0	0
A3c.occluded LV	0	0	0	0	0	0	0	0	0	0	9	11	0	0	0	0	0	0	0	0	0	0	0	0	0
A4c.no occlusions	0	0	0	0	0	0	2	0	0	0	90	1	0	0	0	0	1530	45	104	93	0	0	0	0	0
A4c.occluded LA	0	0	0	0	0	0	3	0	0	22	32	0	0	0	0	0	28	543	20	12	0	0	0	0	0
A4c.occluded LV	0	0	0	0	0	0	6	0	0	0	0	0	1	0	0	0	3	4	47	20	0	0	0	0	0
A5c	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	14	0	9	185	0	0	0	0	0
Subcostal	3	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	9	1	9	15	930	1	24	0	0
Suprasternal	0	0	0	8	0	0	4	0	0	0	0	0	0	0	0	0	0	0	0	0	36	146	16	0	0
Other	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	3255	0

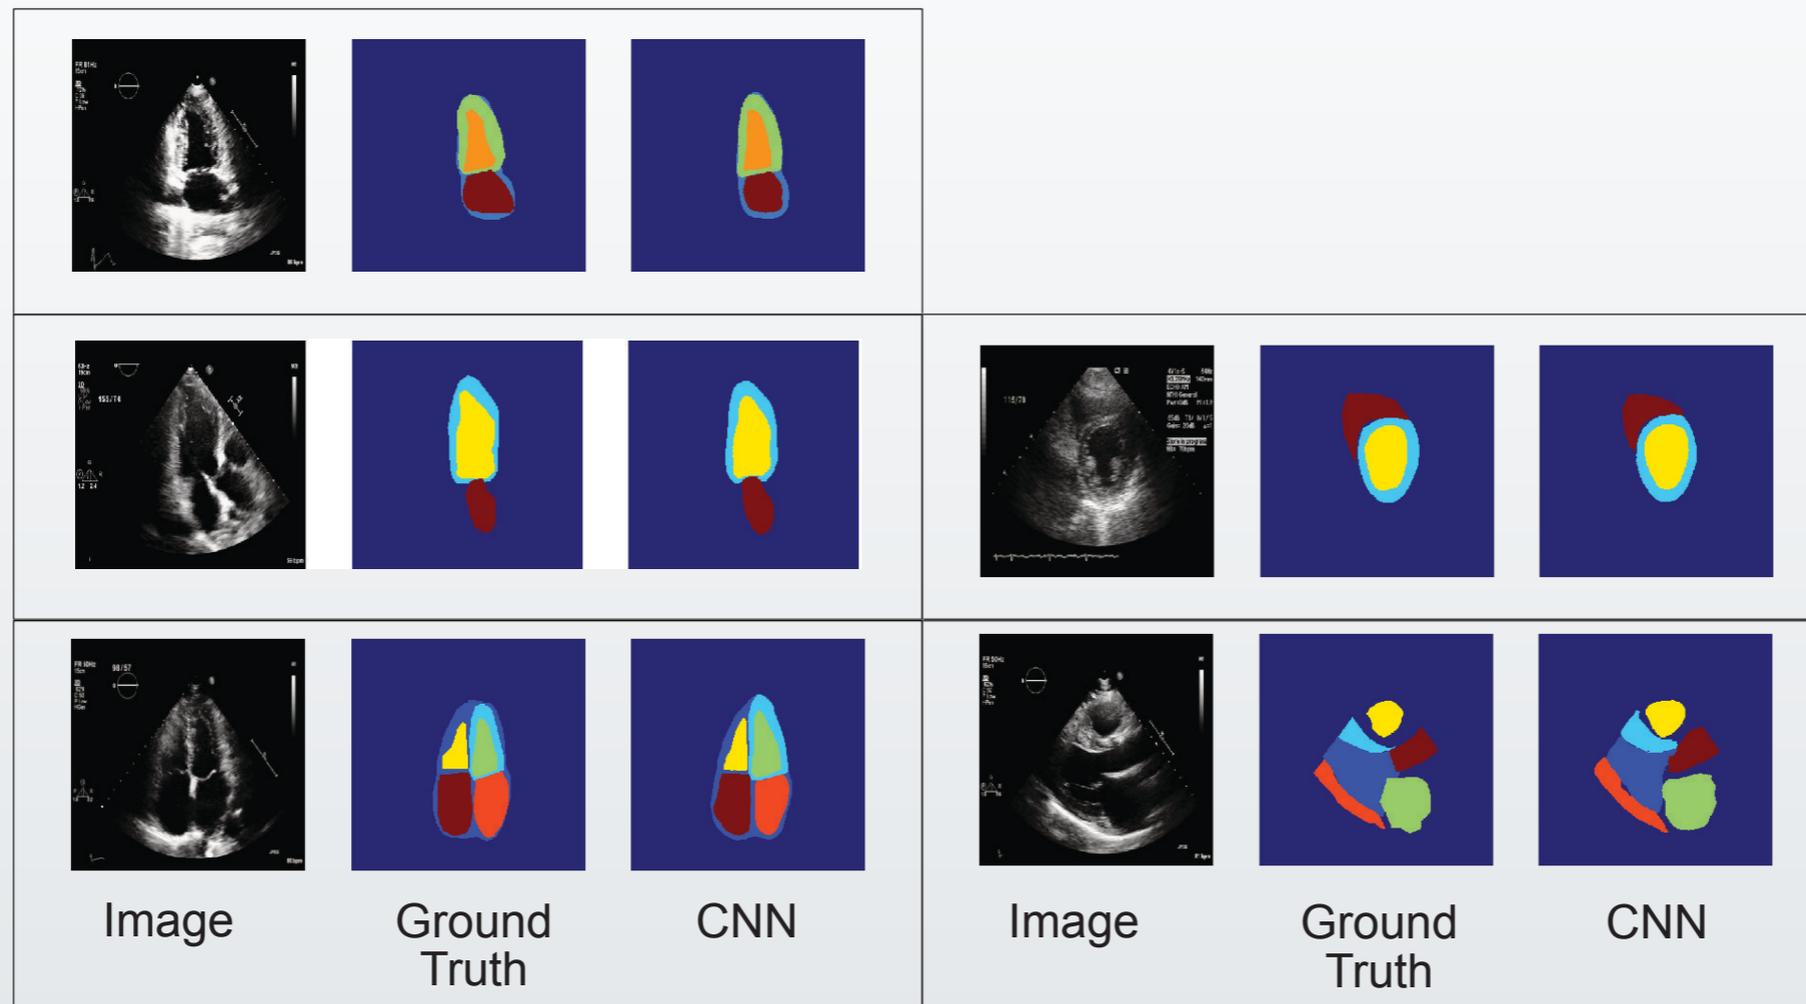
Predictions



# An Automated (low-cost!) Approach to Echo Interpretation



# Segmentation using Convolutional Neural Networks



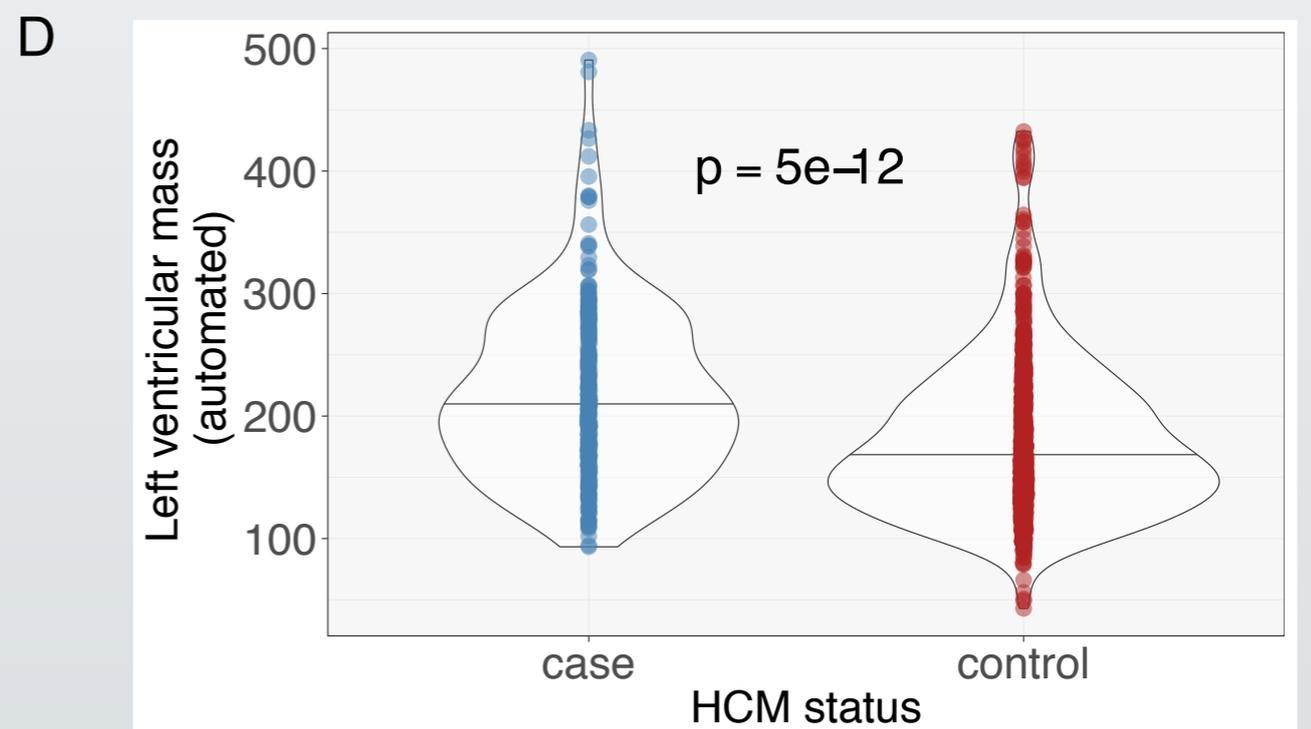
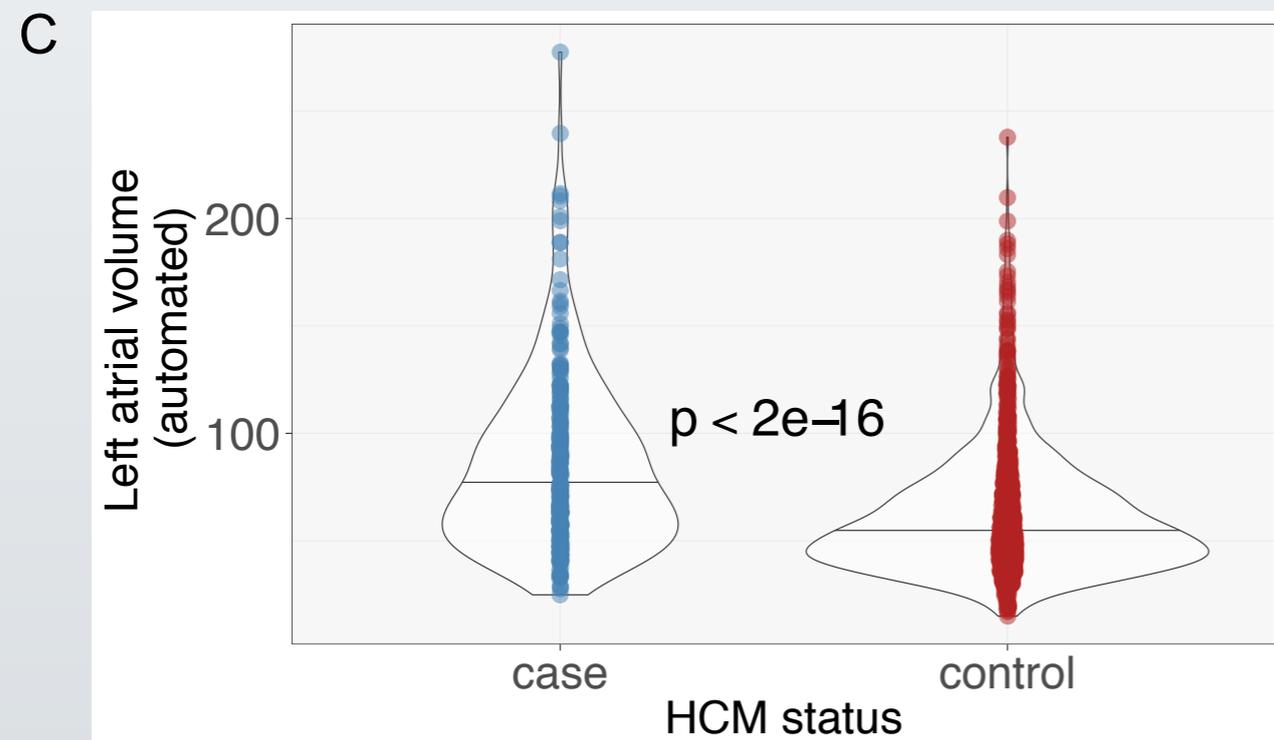
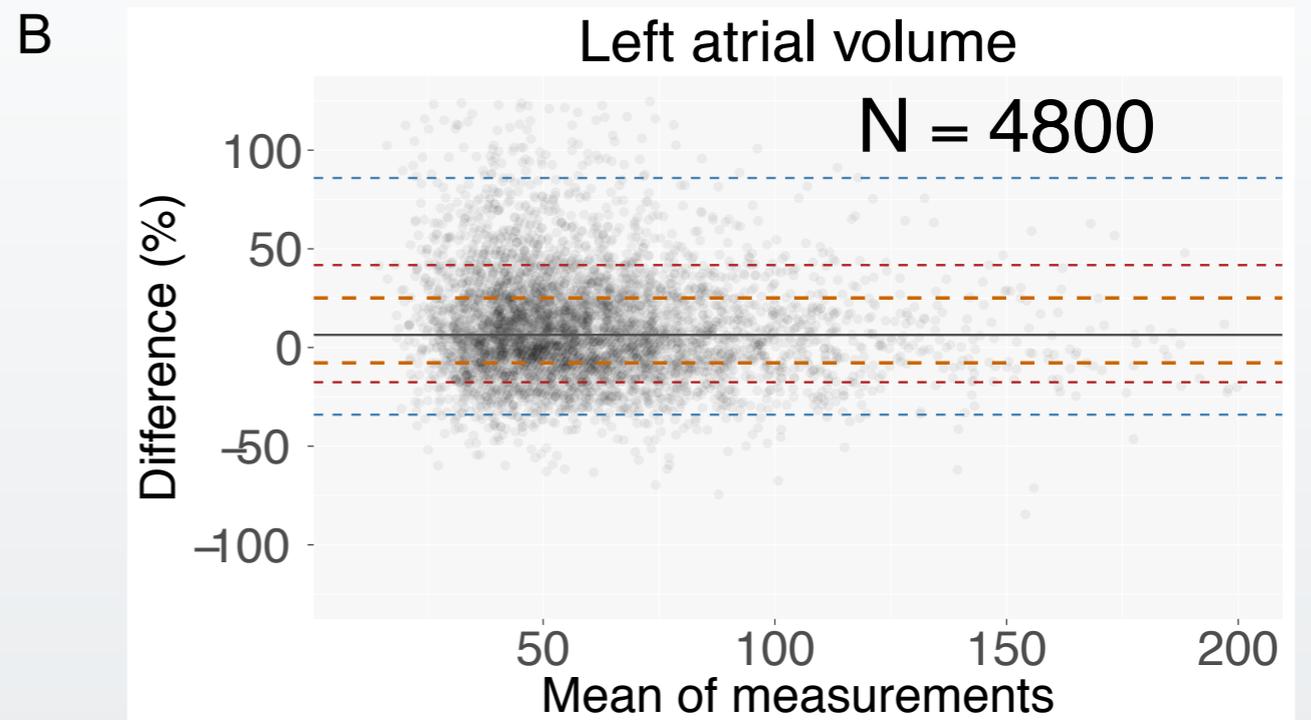
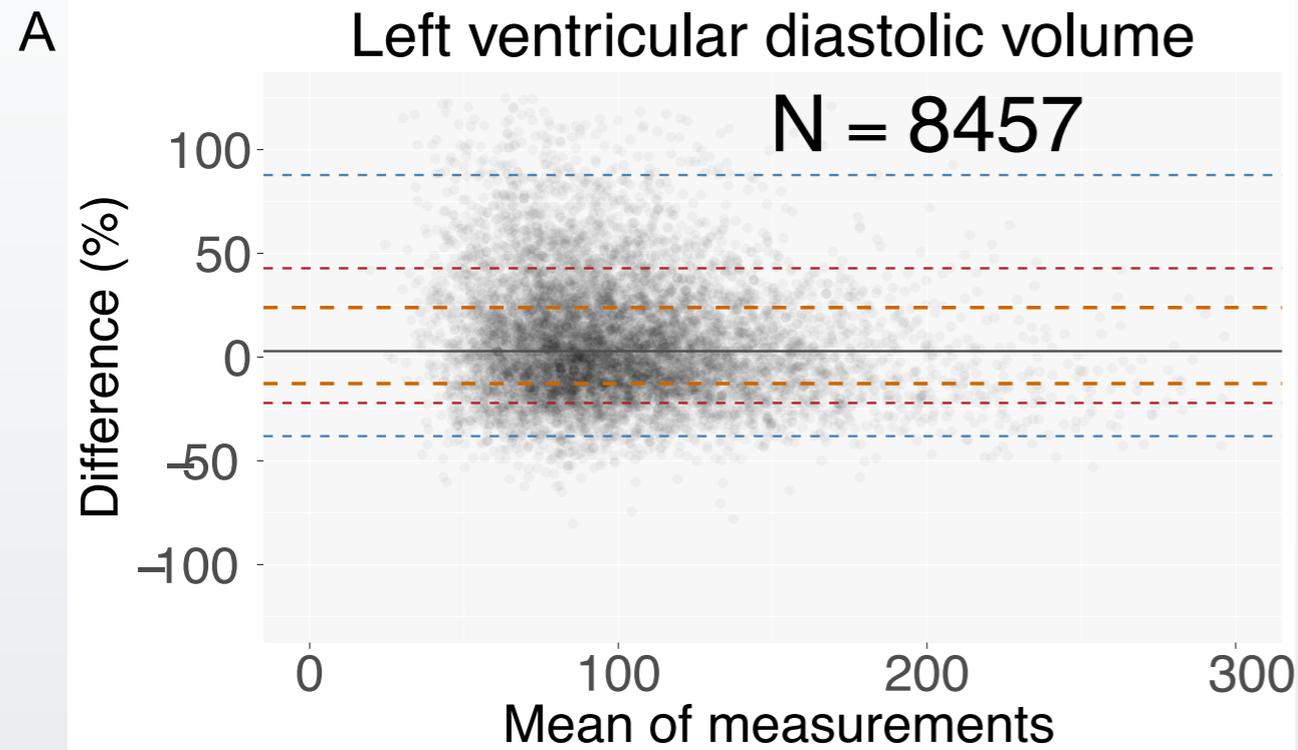
For all views, only 100-200 manually traced images were used for training  
We segment every frame of every video

# Deriving “Real World” Measurements: Comparisons to Thousands of Studies from the UCSF Clinical Echo Laboratory

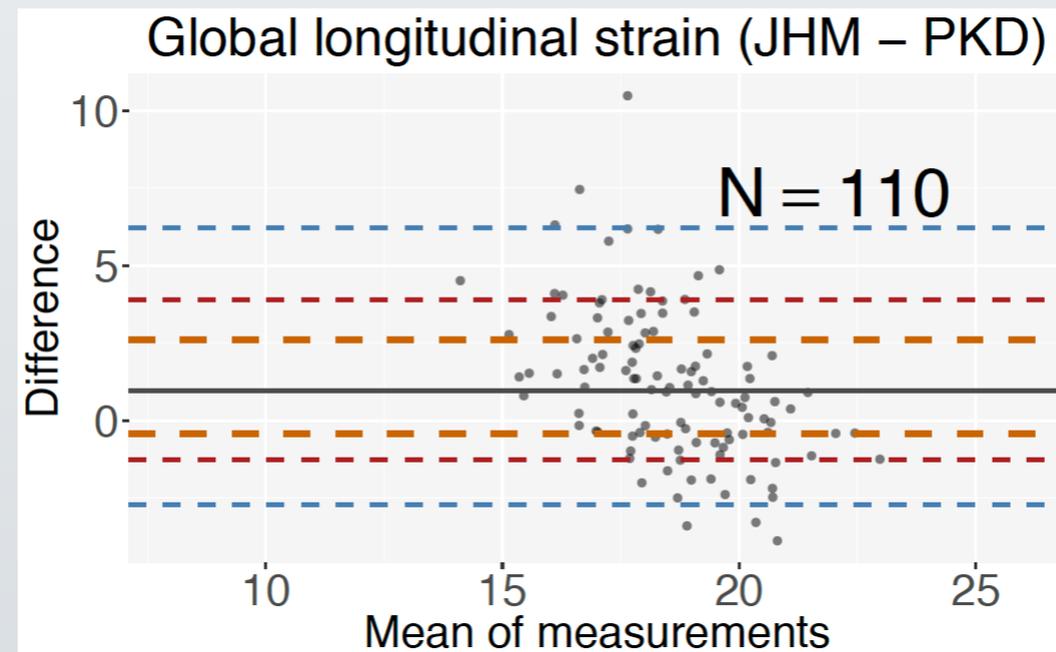
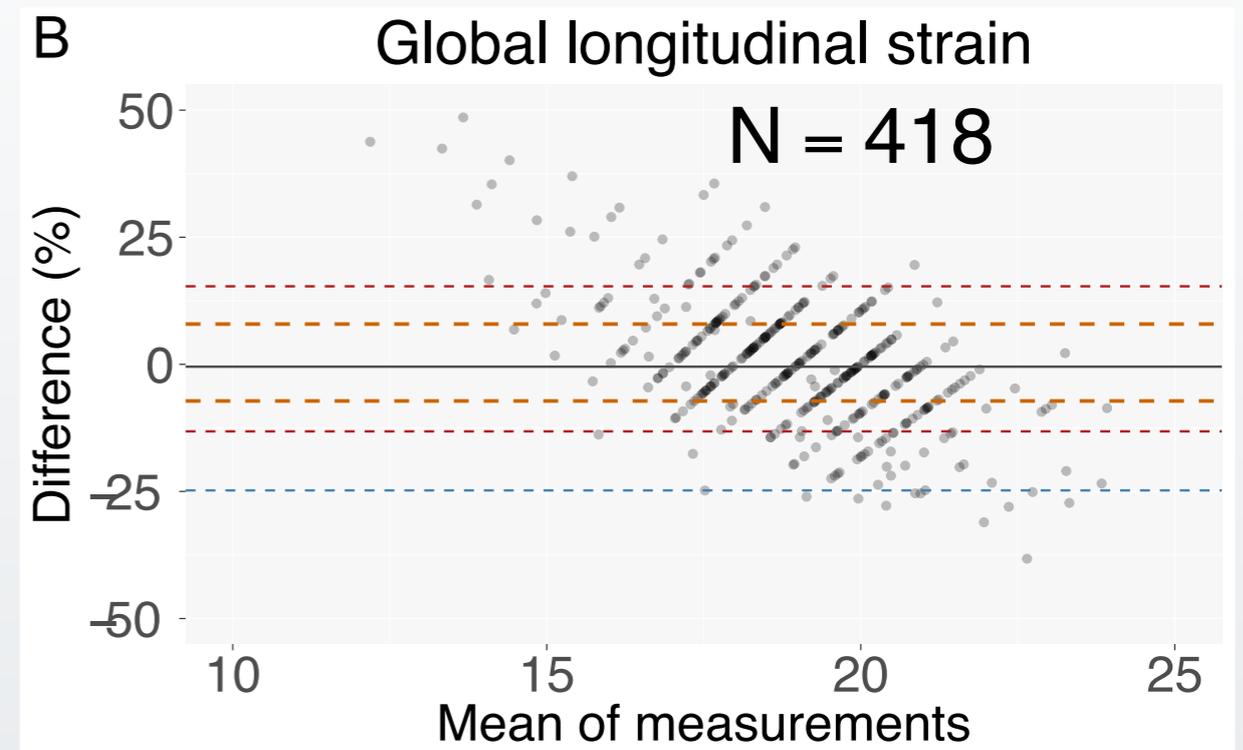
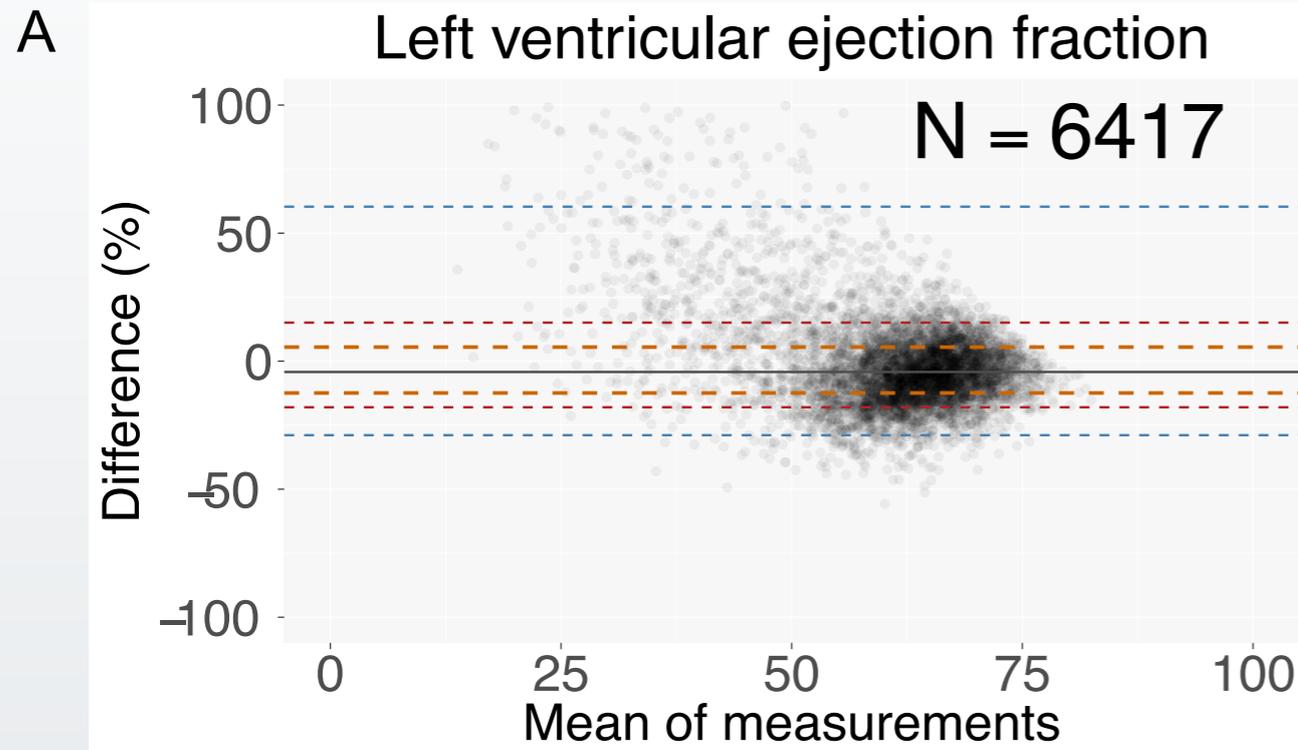
Metric	Number of Echo Studies Used for Comparison	Median Value (IQR)	Absolute Deviation - % of Manual (Automated vs. Manual Measurement)		
			50	75	95
Left atrial volume	4800	52.6 (40.0-71.0)	16.1	29.3	66.2
Left ventricular diastolic volume	8457	92.1 (71.8-119.1)	17.2	30.5	68.0
Left ventricular systolic volume	8427	33.2 (24.1-46.8)	26	47	108
Left ventricular mass	5952	148.0 (117.3-159.9)	15.1	27.6	61
Left ventricular ejection fraction	6407	64.8 (58.3-59.41)	9.7	17.2	39.9
Global longitudinal strain	418	19.0 (17.0-21.0)	7.5	13.6	30.8
Global longitudinal strain (Johns Hopkins PKD study)	110	18.0 (16.0-20.0)	9.0	17.1	39.4

We can make all of the common measurements for B-mode echo.

# Deriving “Real World” Measurements



# Assessing Cardiac Function



# Are clinicians really a gold standard?

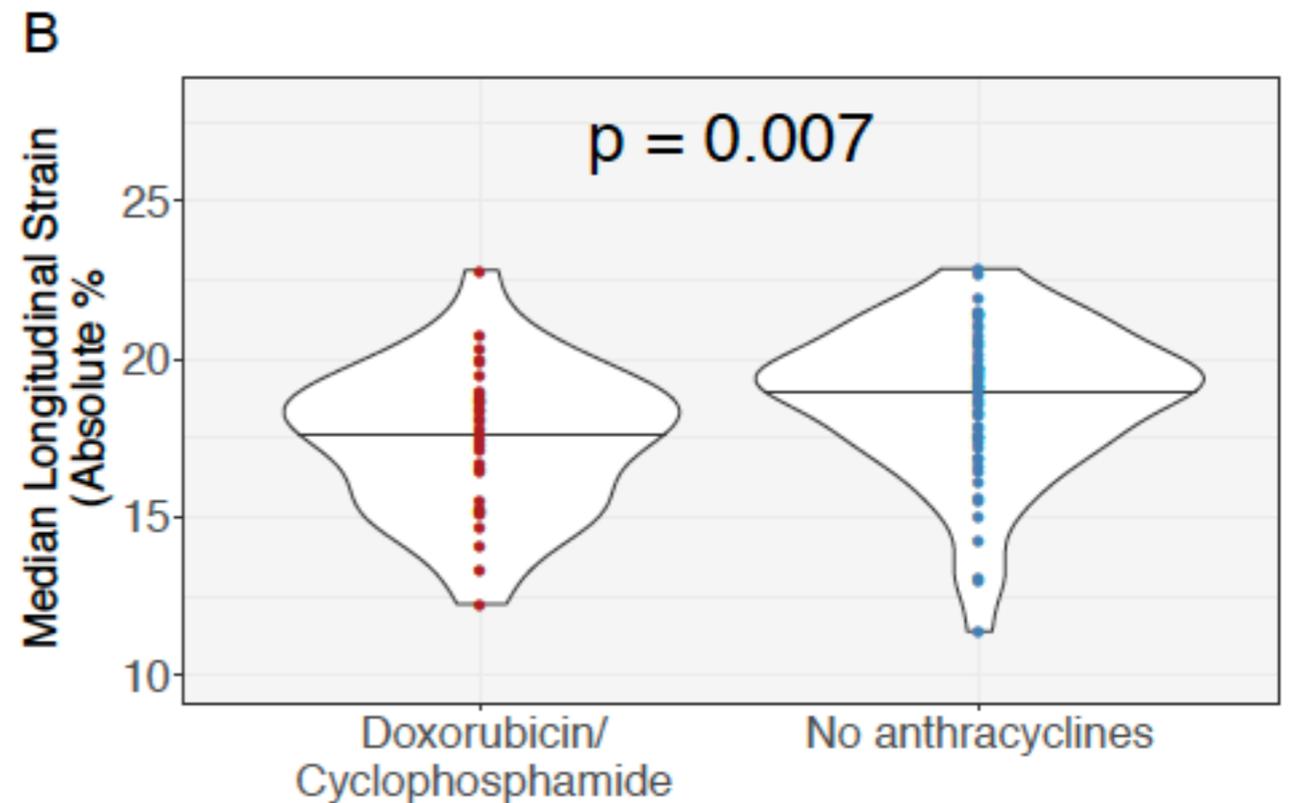
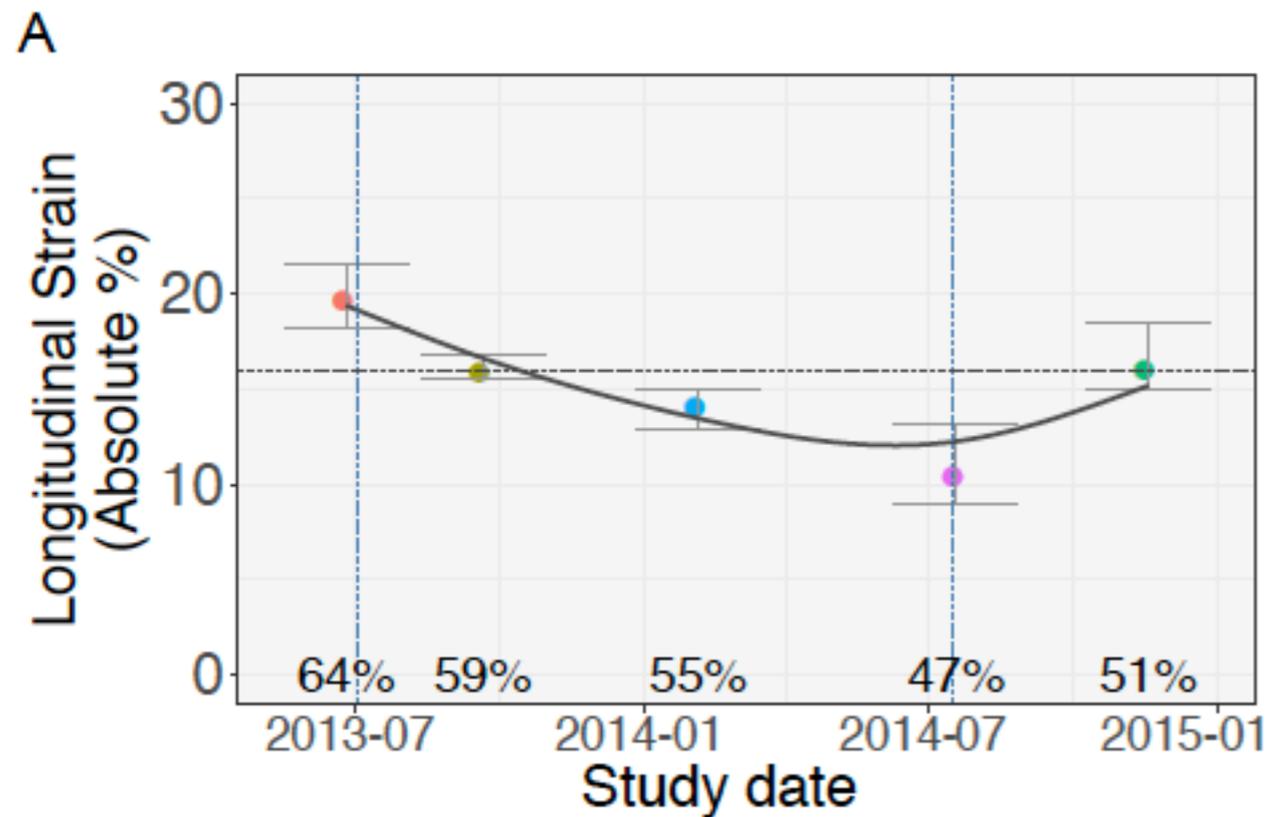
1. A typical echocardiogram reader will manually trace the heart in 4-6 “representative frames” and that will be the gold standard
2. There is wide variability in measurement values from physician to physician - up to 8-9% for the ejection fraction, which has a normal value of 60%
3. How do we show an improvement?
  1. Compare to an average of multiple readers
  2. Compare to a gold-standard imaging system (e.g. MRI)
  3. Demonstrate utility in outcomes

# Internal Measures of Consistency

Comparison	N	Correlation – Manual vs. Manual (p-value)	Correlation – Automated vs. Automated (p-value)
Left atrial volume vs. left ventricular mass	4012	0.54 (<2e-16)	0.56 (<2e-16)
Left ventricular mass vs. left ventricular diastolic volume	5874	0.62 (<2e-16)	0.61 (<2e-16)
Left ventricular mass vs. left ventricular systolic volume	5856	0.58 (<2e-16)	0.55 (<2e-16)
Left atrial volume vs. left ventricular diastolic volume	4748	0.48 (<2e-16)	0.56 (<2e-16)
Left atrial volume vs. left ventricular systolic volume	4738	0.49 (<2e-16)	0.46 (<2e-16)
Left atrial volume vs. left ejection fraction	4720	-0.22 (<2e-16)	-0.23 (<2e-16)
Left ventricular mass vs. global longitudinal strain	243	-0.16 (0.01)	-0.27 (<2e-16)
Left ventricular mass vs. left ejection fraction	5123	-0.28 (<2e-16)	-0.28 (<2e-16)
Left ventricular diastolic volume vs. global longitudinal strain	326	-0.15 (0.006)	-0.17 (0.002)
Left ventricular systolic volume vs. global longitudinal strain	326	-0.29 (<2e-16)	-0.27 (<2e-16)
Left ventricular ejection fraction vs. global longitudinal strain	251	0.37 (<2e-16)	0.32 (<2e-16)

# Longitudinal Strain in Longitudinal Studies

## Tracking Patients on Herceptin Chemotherapy

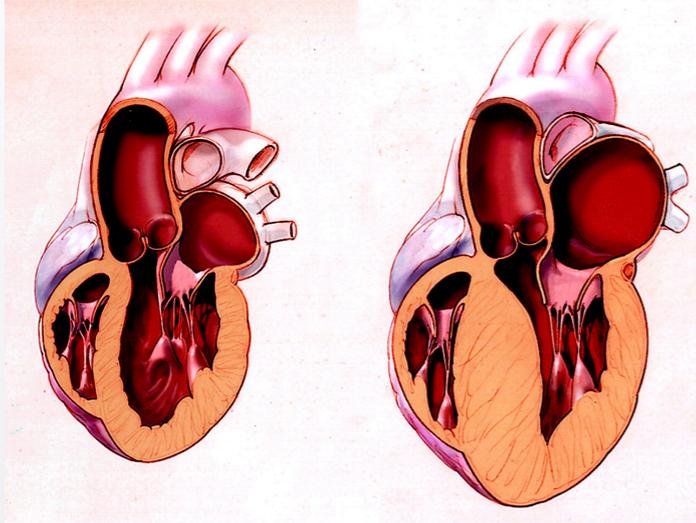


A vision of low cost serial monitoring of patients at risk of cardiac dysfunction: hypertension, obesity, diabetes

# Automated Disease Detection - What's the Point?

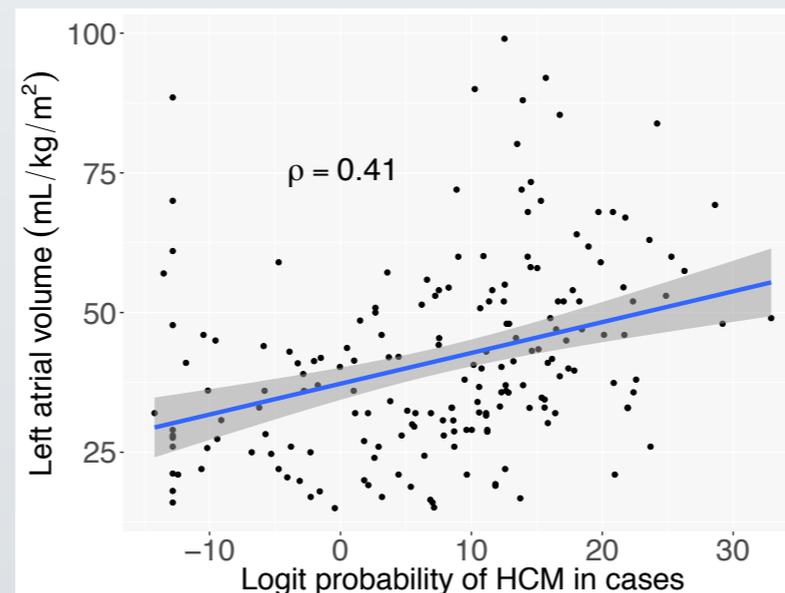
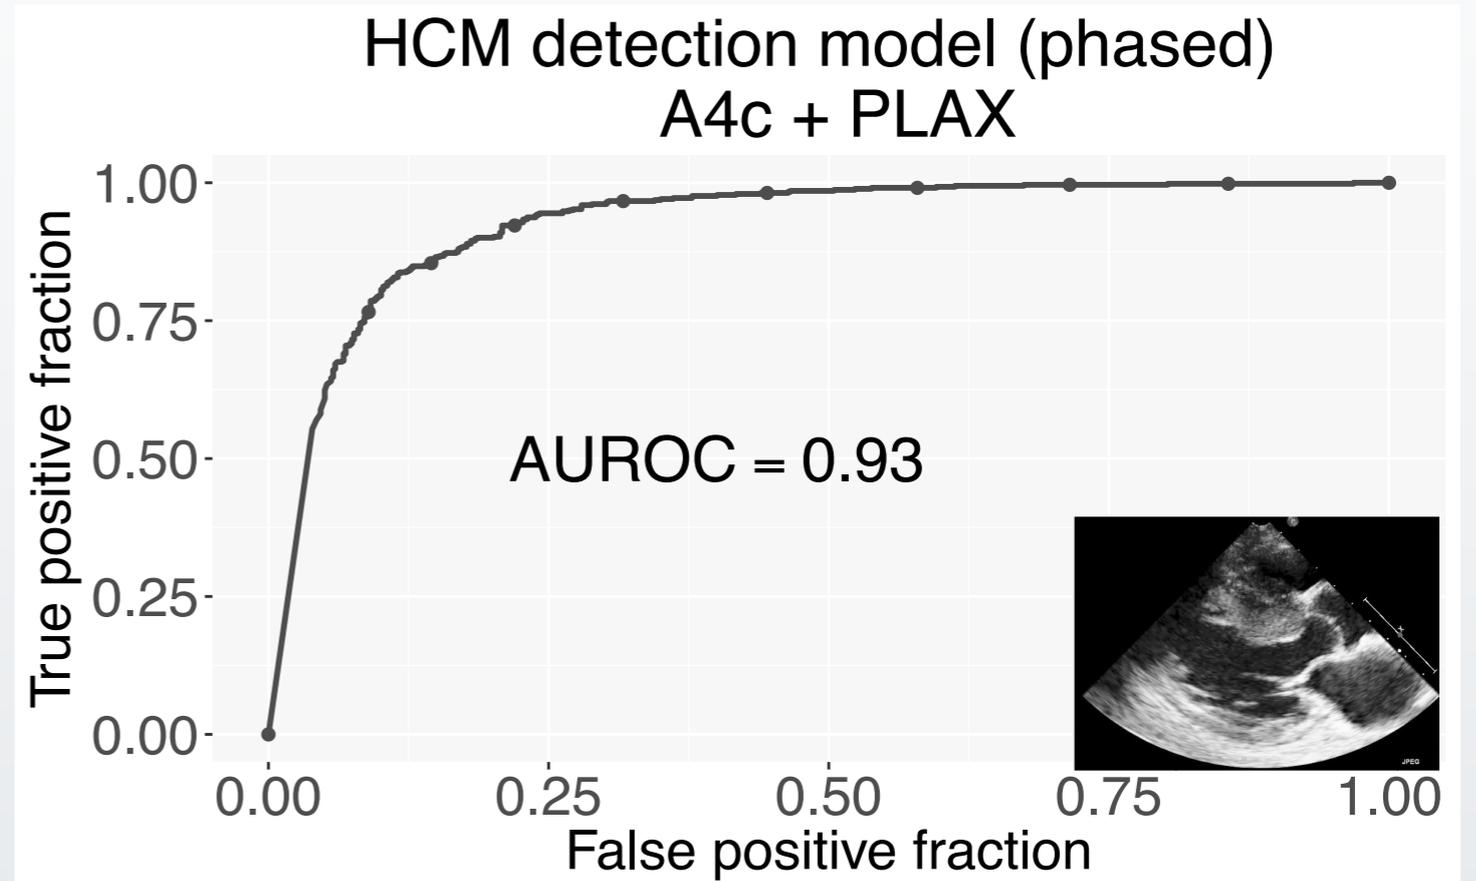
1. Several rare diseases would benefit from referral to a **cardiologist or specialty center**
2. These diagnoses tend to be **missed at centers that see them infrequently**
3. We hypothesized that we could implement “disease detection modules” based on these same simple views
4. This would again be themed as **“decision support”** - not definitive diagnoses

# A Model for Hypertrophic Cardiomyopathy

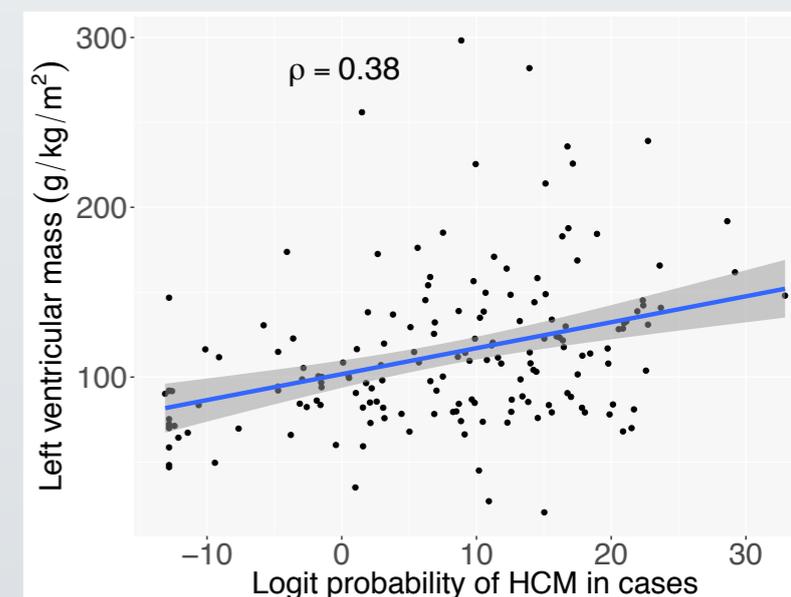


1:500 individuals

- Leading cause of **sudden death** in young athletes
- Inherited in families
- Can result in unstable heart rhythms, heart failure, and stroke
- Management involves behavioral changes, medication, and preventive implantation of a defibrillator

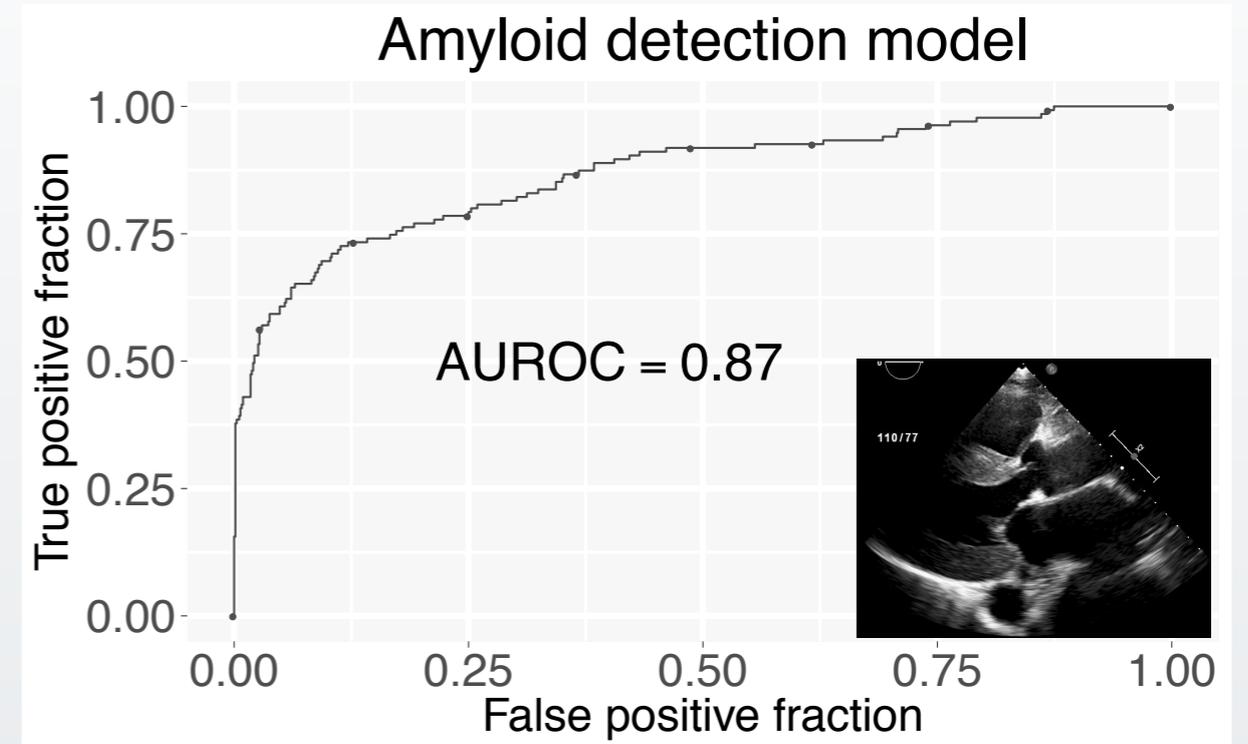


Left atrial volumes

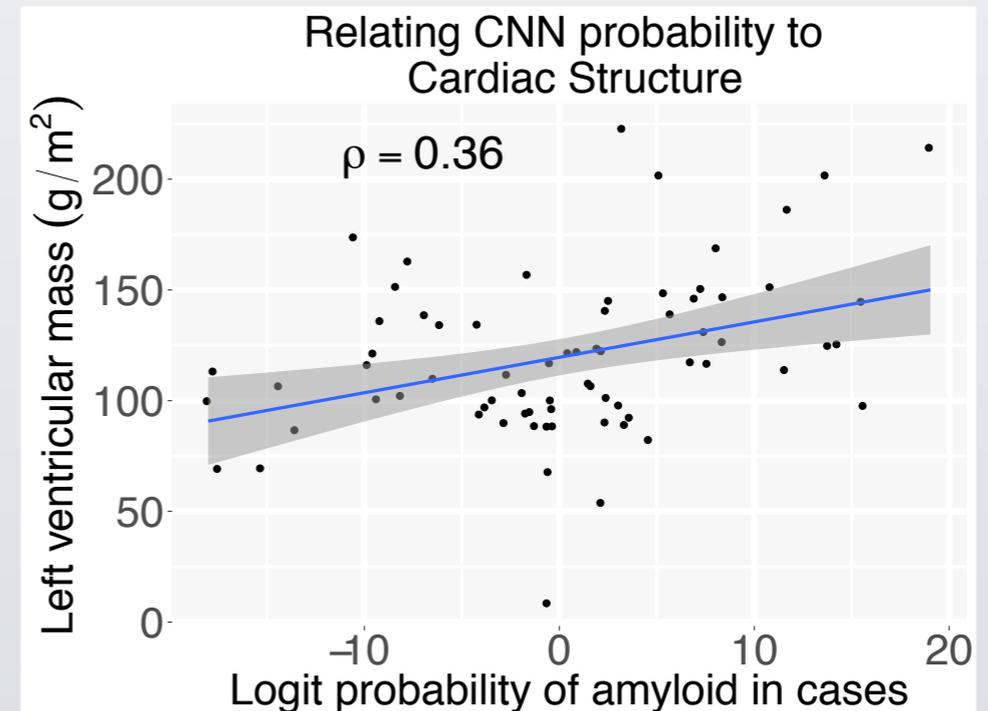


Left ventricular mass

# A Model for Cardiac Amyloidosis



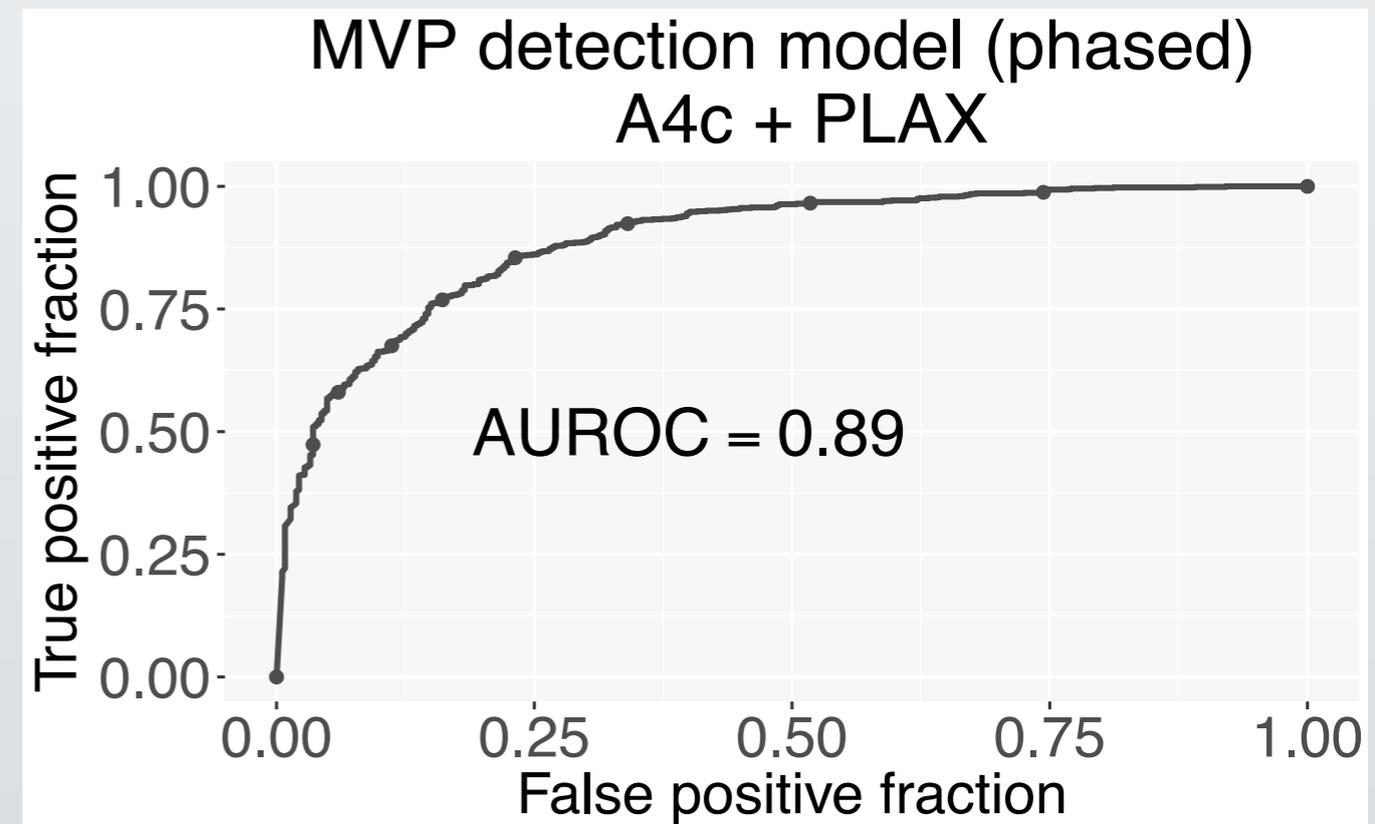
- “Senile amyloidosis” is a common cause of heart failure in the elderly ... but often missed
- Can be inherited in families
- Can result in unstable heart rhythms, heart failure, and stroke
- Management involves medication, and preventive implantation of a pacemaker/defibrillator



# A Model for Mitral Valve Prolapse



- A disease characterized by abnormal myxomatous thickening of the valve leaflets
- Seen in 1% of the population
- Can progress to severe valve disease and is sometimes seen with arrhythmia and sudden death



# What Next - Clinical Deployment!!!

1. UCSF has filed provisional patent for our system
2. Code and weights are freely available for academic/nonprofit use on Bitbucket: <https://bitbucket.org/rahuldeo/echocv>
3. We don't have FDA approval as a diagnostic - but proof-of-concept studies are needed to test the value of integrating automation into the clinical workflow
4. **Enabling National and Global Clinical Deployment**
  1. Brandon Fornwalt: Geisinger Health System
  2. Patrick Gladding: The University of Auckland, NZ

# Musings about the future

# Some predictions for the future of cardiac imaging: following the path of ECG interpretation

1. Routine measurements will be made in an automated way - with a visual check of segmentation quality
2. Some automated diagnoses may happen at point-of-care: assessment of heart function, dangerous accumulation of fluid around the heart
3. Until image acquisition is facilitated, the benefits of automated interpretation will be muted

# Where there is greater uncertainty ...

1. Ideally, we should be using automated interpretation to **elevate medicine beyond the current practice** - but that requires much larger data sets and imaging more often (i.e. time course) than what is currently performed (and reimbursed)
2. **Pharmaceutical companies have motivation to perform high frequency serial imaging** to assess whether there are any benefits to medications in a shortened Phase II trial - accurate scalable quantification will be needed
3. **Surveillance of daily studies** may be useful to enable identification of individuals who may be eligible for clinical trials or newly approved therapies (e.g. cardiac amyloidosis)

# Subclassification, risk models, ... and the challenge of demonstrating utility

1. There is no question most disease classifications are crude ... and **finer distinctions can be made between disease states**
2. There is also no question that survival models are crude, and **better predictive models should be possible** with imaging data and emerging algorithms
3. Unfortunately, physicians are only interested in classifications or risk models **that will change practice** ... and require evidence to justify this
4. So until we have more data, we are left with the status quo and a bunch of research manuscripts

**What about the biology?**

# What is missing in medicine?



1. **Low-cost quantitative metrics** that are indicative of disease progression and reflect the onset of these tissue-level changes
2. Should be **specific** to the disease process:
  1. expressive: capture complex underlying biological processes
  2. multidimensional: can't readily be "gamed"
3. Should be **ameliorated** with therapy (c.f. genetic risk)

# Immediate challenges

1. **Inaccessible biology:** the tissue of interest in CHD is not accessible ... how then to build biological assays
2. **Expensive:** most detailed biological measurements are costly (c.f. imaging, proteomics, DNA sequencing)
3. **Problematic to train:**
  - **sample size:** models that quantify complex biological processes will need to be high-dimensional ... but these will require very large sample size to train and validate
  - **time:** CHD develops slowly over time but a new biological assay requires prospective enrollment



# Expanding phenotypic space

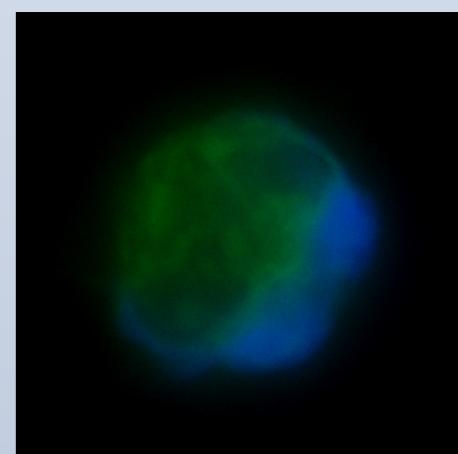
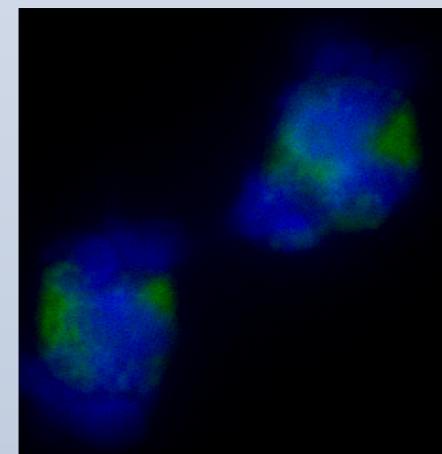
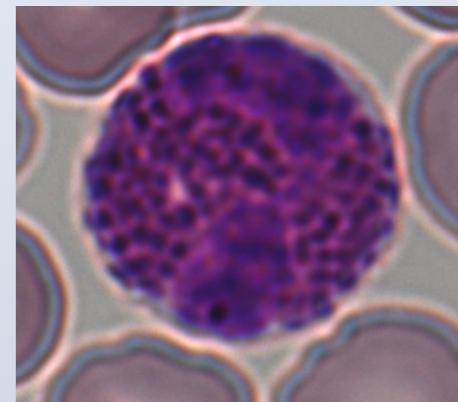
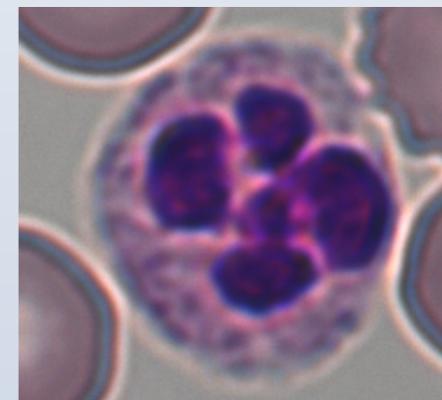
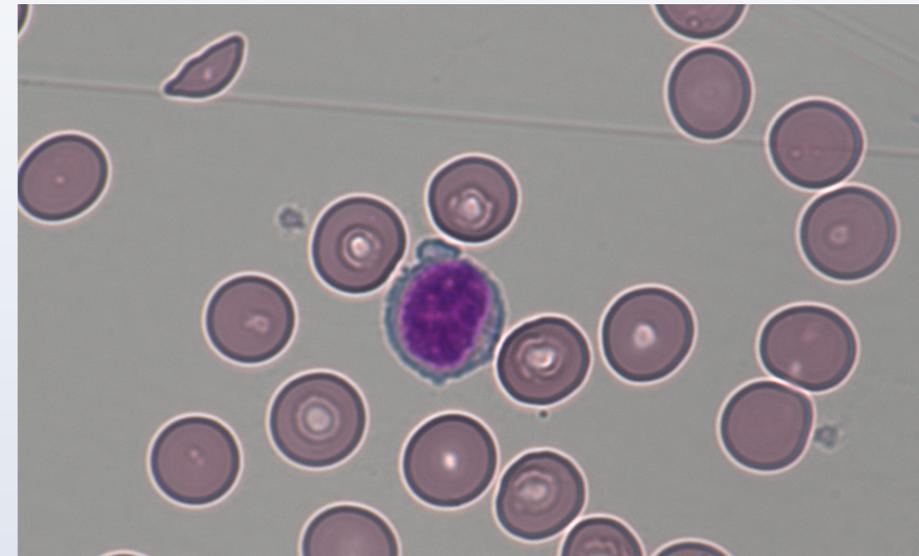
1. Current clinical data sets lack the scale and expressivity needed to reflect underlying biological processes
2. Discoveries from UK Biobank, Partners Biobank, Vanderbilt, Geisinger, etc., are all limited by the underlying low dimensionality of phenotypic information - and that will not be solved by sample size (more of the same)
3. But these studies were exorbitant and have taken decades to accrue the current sample size... how do we improve on this?
4. We need a data type that has the dimensionality to capture biological heterogeneity and complexity and yet can still be collected in a very scalable manner (cf. representation learning needs)
5. It became very clear we need to stay clear of sequencing technologies and costly medical imaging

# A focus on individual circulating blood cells

1. **Causally implicated in CHD pathogenesis**
  1. Involvement of neutrophils, monocytes, and lymphocytes in disease pathogenesis (plaque pathology; plaque pathology); genetic models in mice
  2. CANTOS trial
  3. Accelerated atherosclerosis in autoimmune disorders
  4. Association of clonal hematopoiesis and early myocardial infarction
2. **Accessible:** accessible in a blood draw
3. **Precedence for utility:** Existing predictive models exist for CAD using WBC/RBC characteristics
4. **Express many of the proteins implicated by genetic analysis in atherosclerosis:** e.g. LDLR, LPL, FADS1/2
5. **Reflect many pathways found in diverse cell types:** autophagy, phagocytosis, free radical dissipation

# Cell morphology rather than genomics

1. Takes advantage of **computer vision advances** in characterizing subtle distinctions between cell types and states at low cost
2. Can analyze **tens of thousands of individuals cells** per participant
3. **Fluorescent dyes** permit characterization of organelles (mitochondria, ER, Golgi), cytoskeleton (actin), nucleic acid (DNA, RNA)
4. Can be connected to **gene expression** to clarify underlying functional abnormalities





# Recruitment workflow – 1000-1200 patients per month (12,000-15,000 per year)



Cardiology clinic



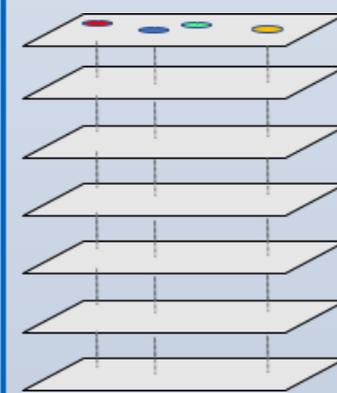
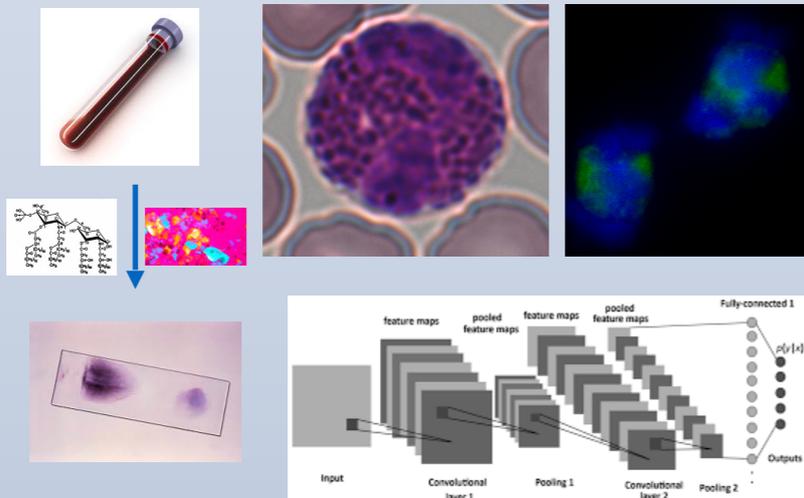
General medicine



Primary care

**Primary assays:** low cost, reproducible, expressive, rapid, responsive to therapy, interpretable

**Secondary assays:** costly, less robust, limited expressivity, non-responsive, non-biological



Somatic sequencing (CHIP)

GWAS (GRS)

Novel Devices

ECG

PET

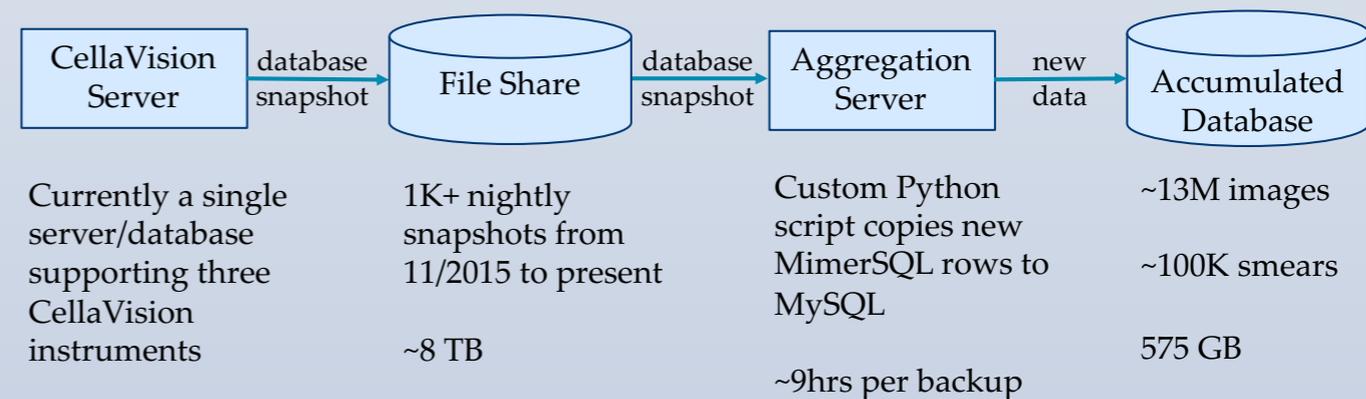
Single cell RNA-Seq

WGS

# Scaling up to incorporate longitudinal data: tapping into the pathology lab

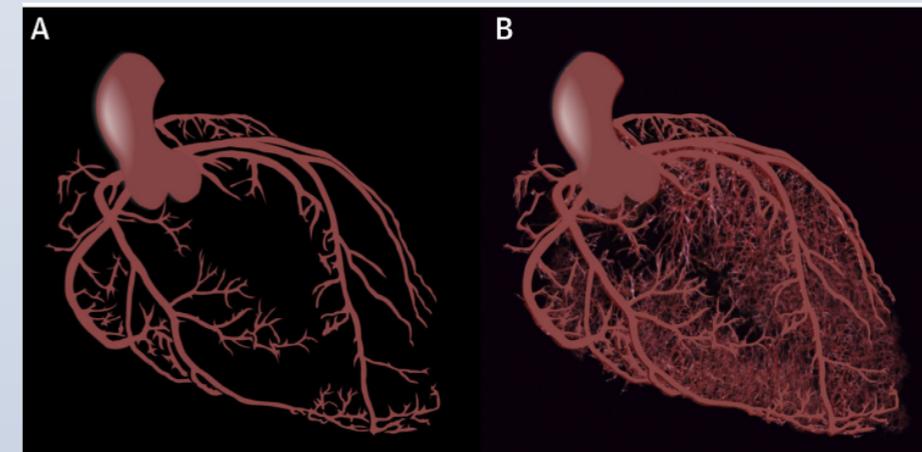
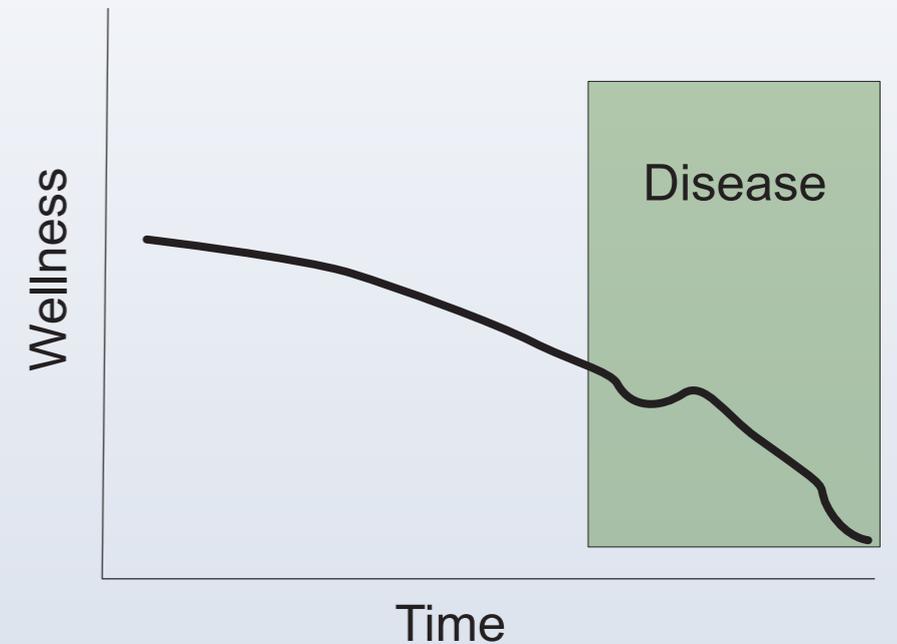
1. Data has been collected since **2015**
2. **>3.5 million data records with 1800 added a day**
3. We connect cellular data to **full medical data** via API
4. Digitized slides from **100,000 patients - 13 million images**
5. Prospectively doing this for all acute MI patients

<b>Sysmex</b>  3.5M FCS files from the XE-5000  740K raw files from the XN-9000 + 1800/day	<b>CellaVision</b>  13M images 18% Lymphocytes (in active DB)  100K smears + 135/day
<b>Abbot Sapphire</b>  170K+ FCS files  160K+ WBC TYP files	<b>Siemens Advia</b>  65K+ FCS files  7.7M raw files



# Summary of our approach

1. A permissive recruitment scheme to enable rapid accrual of **tens of thousands of patients** per year all with **expressive phenotyping and full medical records**
2. Use of cell morphology/cell counter data to massively **expand phenotypic space** at low cost using perturbations and diverse readouts
3. Overlapping of **multiple phenotypic scales** in different cohorts to convert costly, tissue-localized phenotypes (e.g. PET, CHIP sequencing) into lower cost (TTE, cell imaging) models
4. **API-based cohort identification** to allow rapid identification of patients of interest
5. **Automated curation of the medical record** into a vehicle for machine learning and causal inference



The macro- and microcirculation

Taqueti and DiCarli, 2018

# Acknowledgements

