CREATIS Medical Imaging Research Laboratory
www.creatis.insa-lyon.fr
UNIVERSITÉ DE LYON

Machine learning in cardiology: why do we need realistic synthetic databases?

N. Duchateau

CardioFunXion winter school - Lyon, FR - 19/02/2018

PHILIPS
upf.

Let’s dream = easy data collection and sharing

Simulation is complex, can’t we do everything from data?

- Yes

So why do we need synthetic data?

- mmmh… let’s ask Matlab:

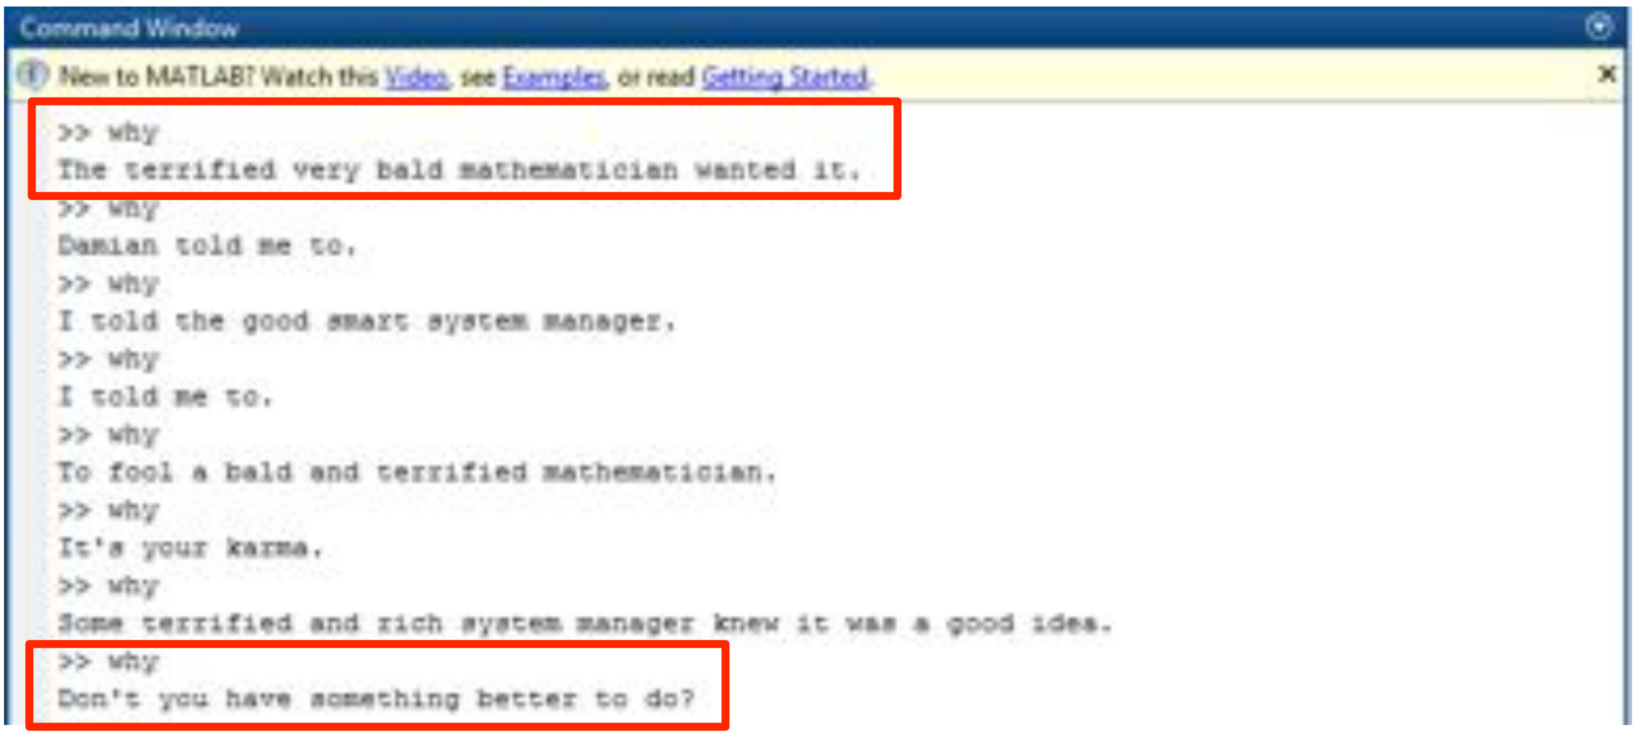

Machine learning = learning from samples

Copeland M https://blogs.nvidia.com/ 2016

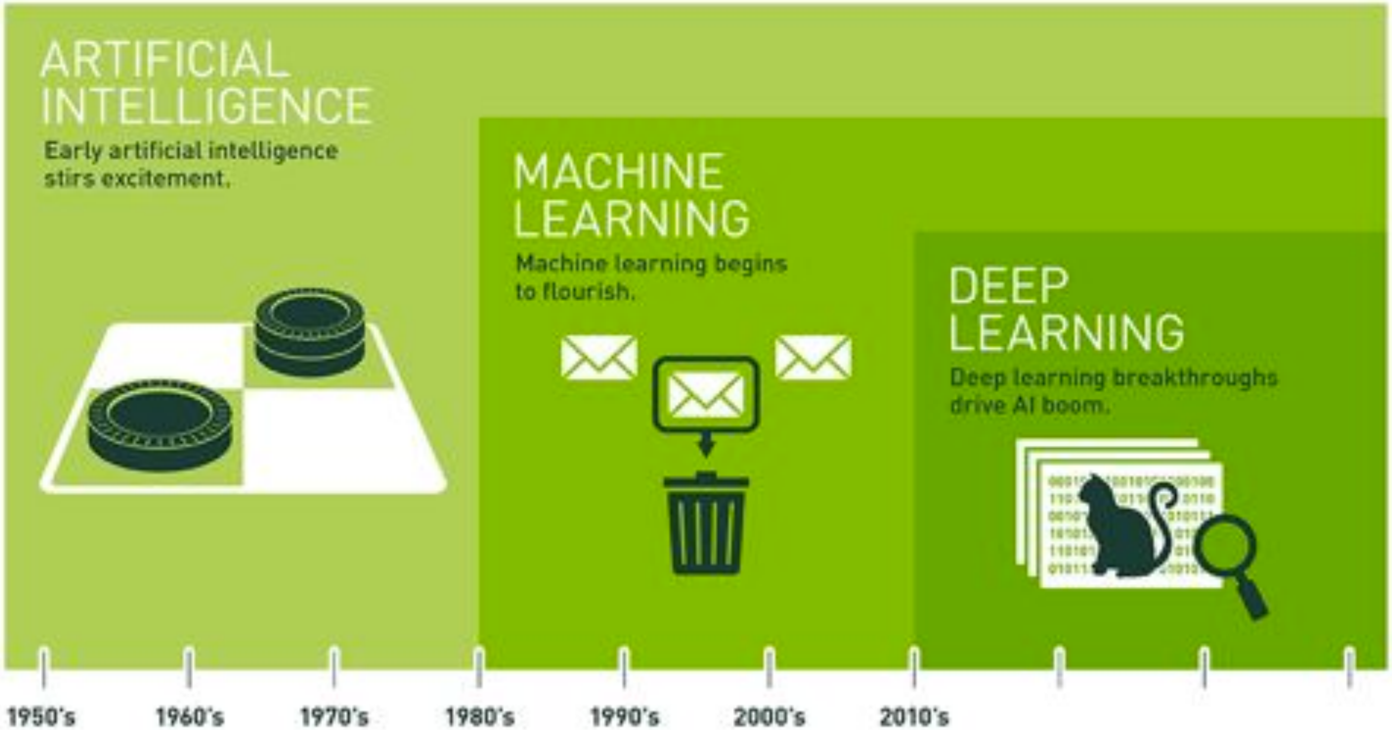

Computers reason as humans

Computers learn without being explicitly programmed

Arthur Samuel (1959)

Networks that adapt themselves to data

Machine learning = learning from samples

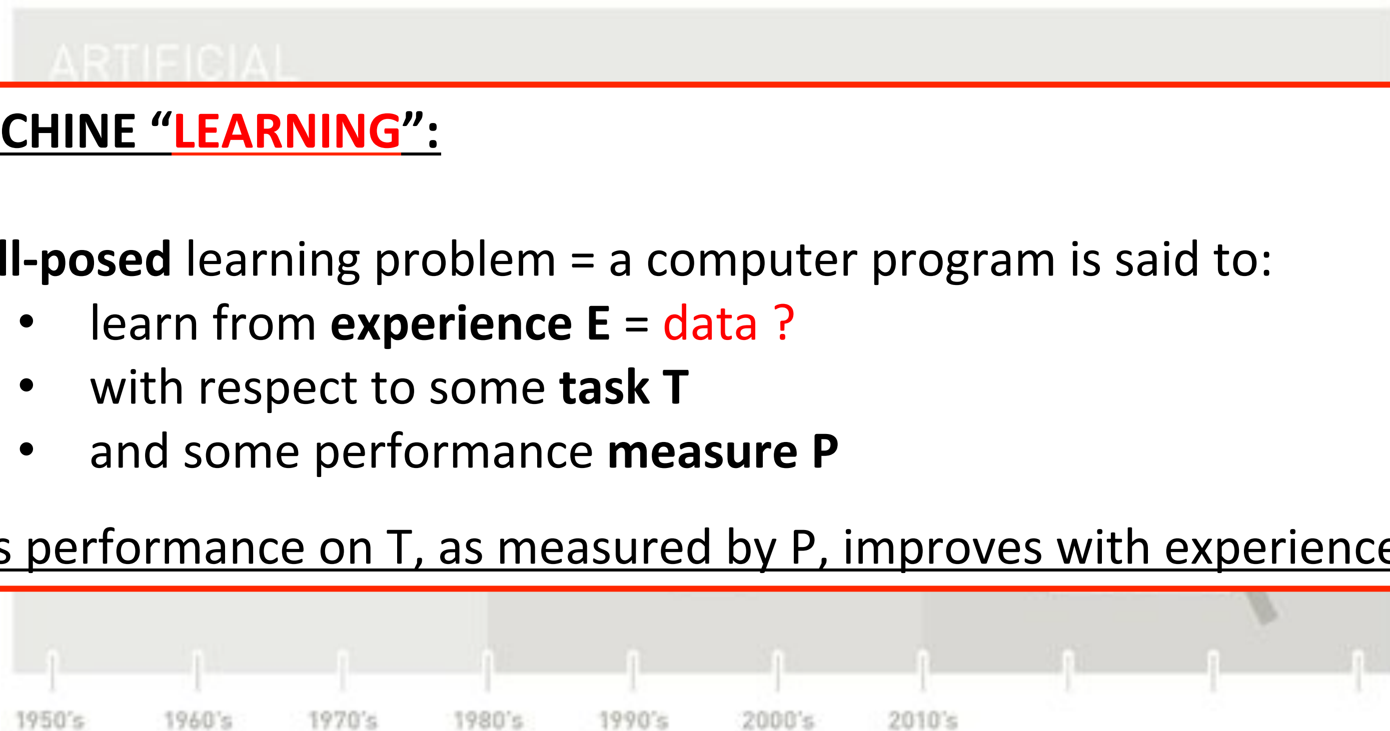

MACHINE "LEARNING":

Well-posed learning problem = a computer program is said to:

- learn from **experience E** = data ?
- with respect to some **task T**
- and some performance **measure P**

if its performance on T, as measured by P, improves with experience E.

Computers reason as humans

~~**Computers** learn without being explicitly programmed~~

Arthur Samuel (1959)

Networks that adapt themselves to data

Machine learning = learning from samples

How many samples? = complexity of the question

<u>Kernel regression (exact):</u>

5 samples

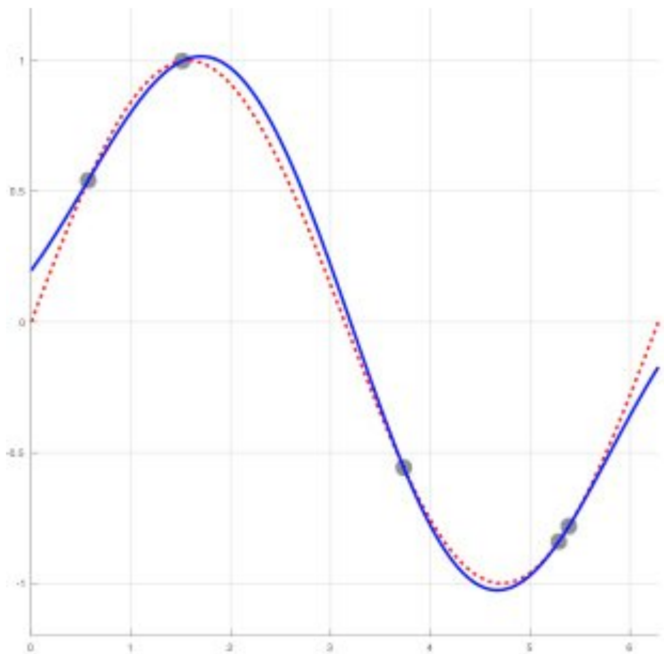

15 samples

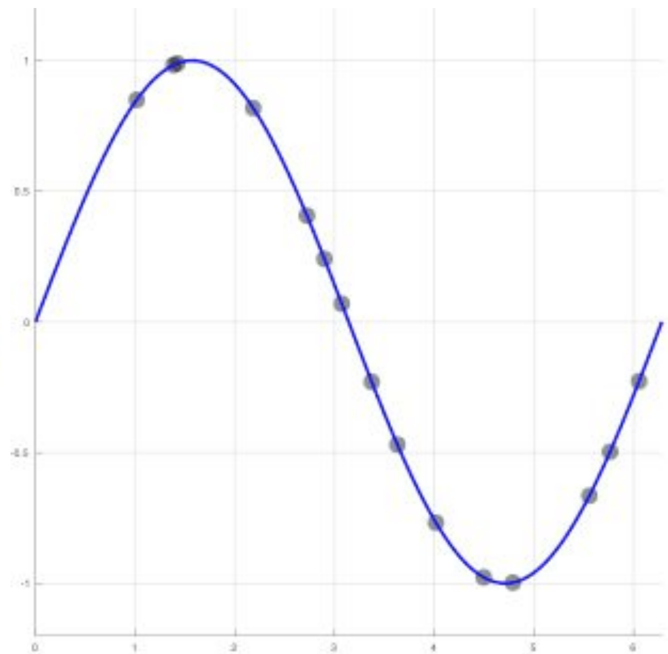

100 samples

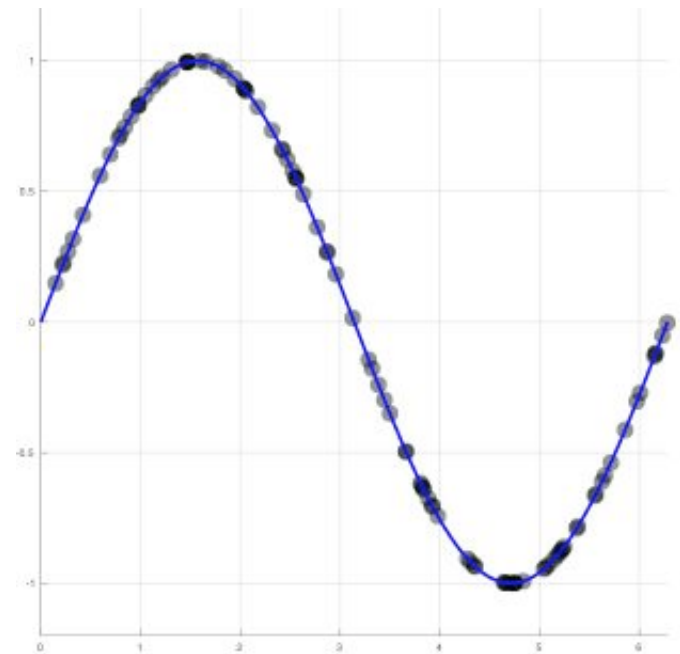

Machine learning = learning from samples

How many samples? **= complexity of the question**

IMAGENET

Russakovsky O et al. *Int J Comput Vis* 2015

1000 object classes…

Machine learning = learning from samples

How many samples? = complexity of the question

IMAGENET

Russakovsky O et al. *Int J Comput Vis* 2015

1000 object classes…
… training = 1.2M !!!
testing = 100k !!!

Machine learning = learning from samples

How many samples? **= complexity of the question**

Only the problem? **Or the model itself?**

Linear model:

2 samples !

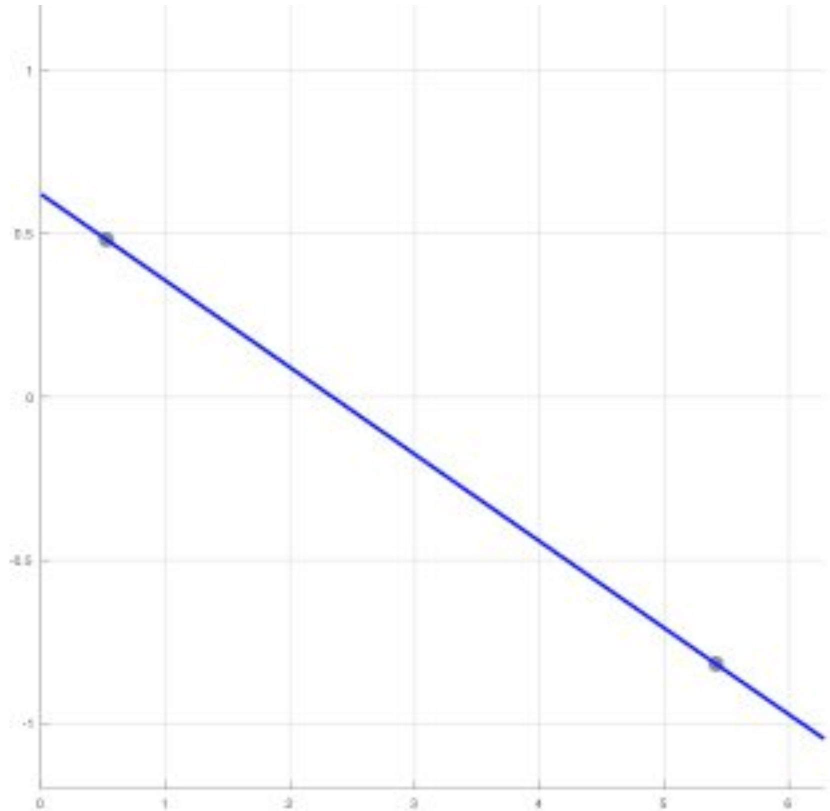

Machine learning = learning from samples

Giro i Nieto X 2017

How many samples? = complexity of the **question**

Only the problem? Or the model itself?

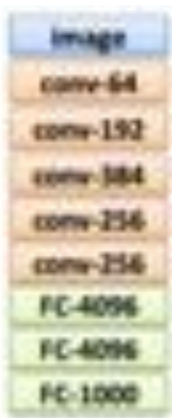

AlexNet
Krizhevsky A et al. *NIPS* 2012

Heusel M 2016

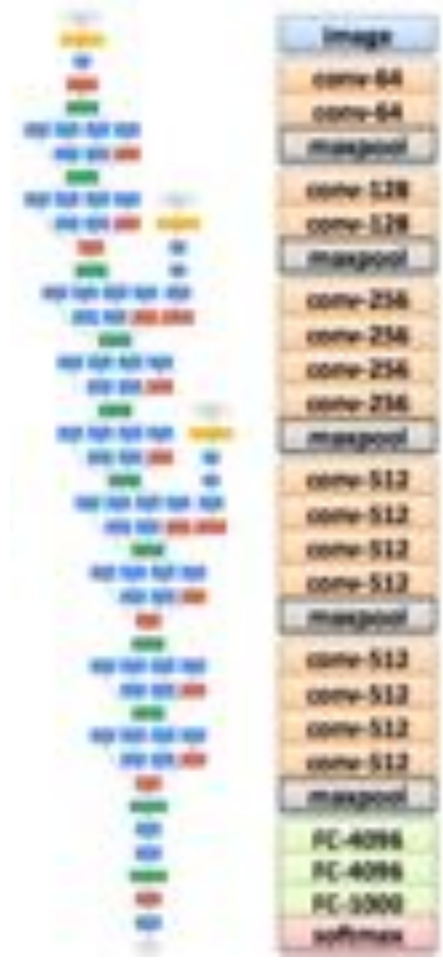

GoogleNet
Szegedy C et al. *CVPR* 2015

Machine learning = learning from samples

How many samples? = complexity of the **question**

Only the problem? Or the model itself?

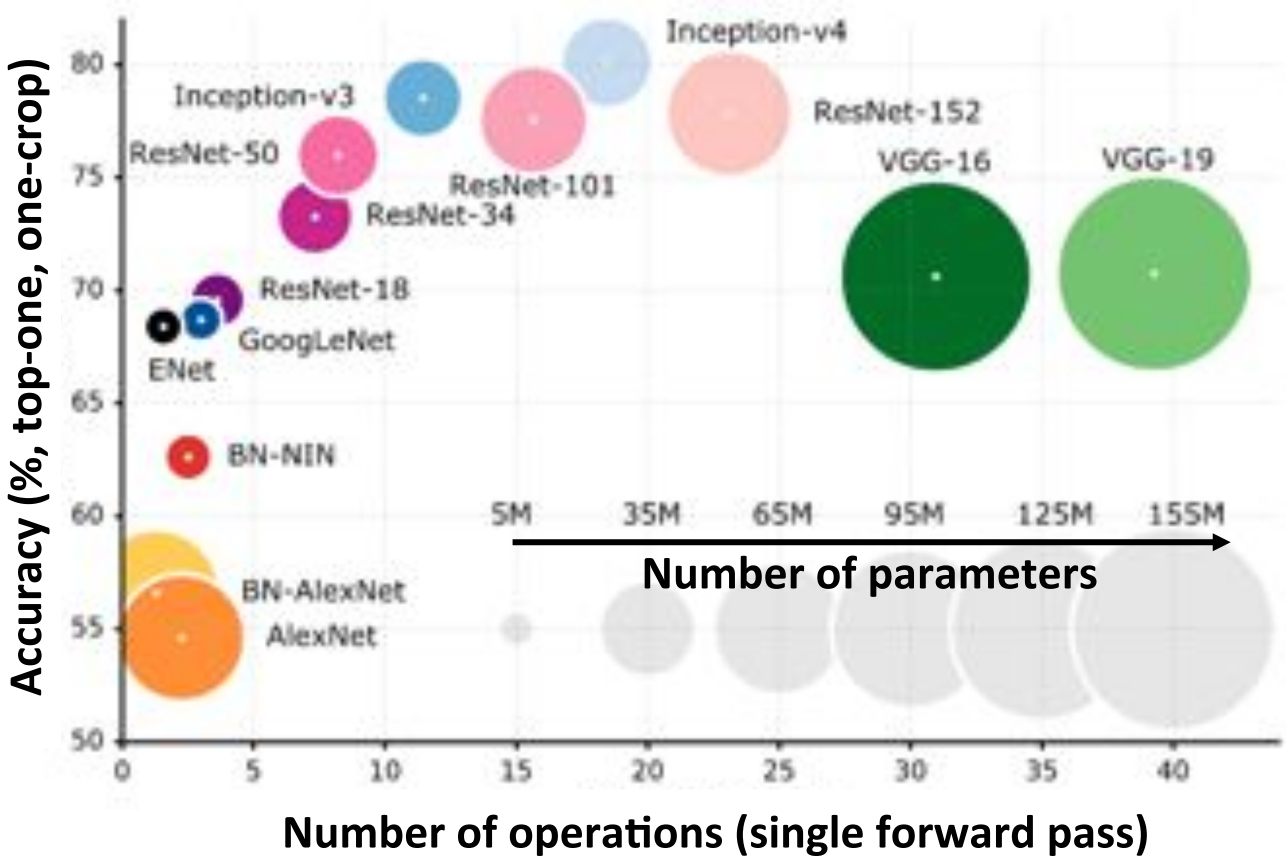

Canziani A et al. *arXiv* 2017

Machine learning = learning from samples

How many samples? = complexity of the **question**

Only the problem? Or the **model** itself?

Or the dimensionality of the data?

- **Input data**
- **Intrinsic dimensionality**

Kernel regression (exact - 2D):

Ground truth Interpolation

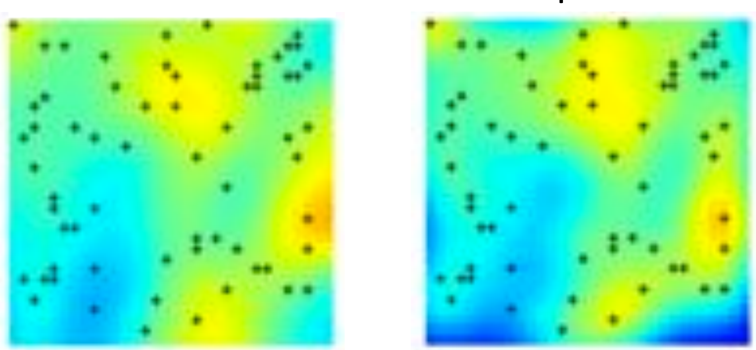

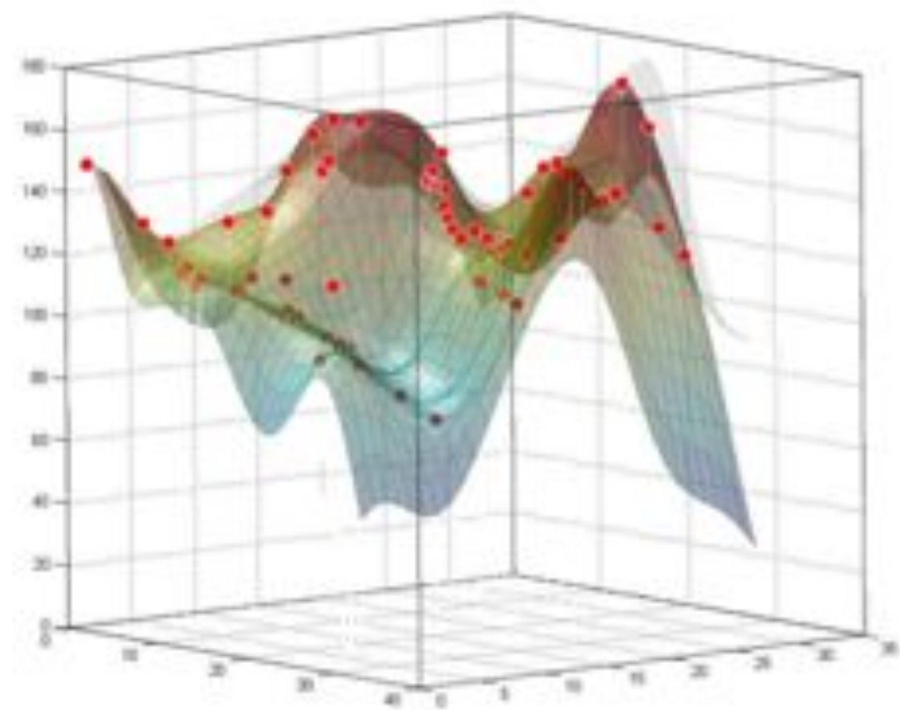

Machine learning = learning from samples

How many samples? = complexity of the **question**

Only the problem? Or the **model** itself?

Or the dimensionality of the data?

Input data

Intrinsic dimensionality

2D images: 256 x 256 pixels (resampled)

IMAGENET

Russakovsky O et al. *Int J Comput Vis* 2015

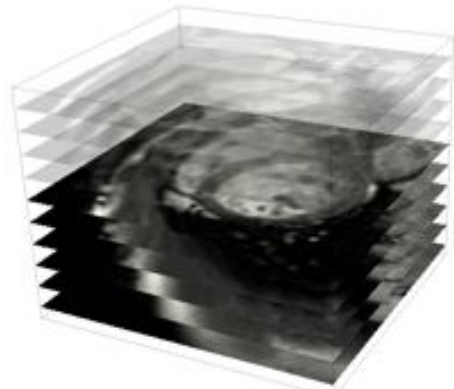

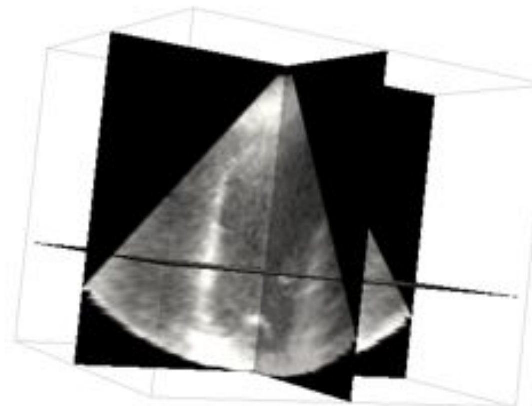

Cine MR ~ 100 x 100 x 14 voxels (3D stack)

3D echo ~ 200 x 200 x 200 voxels

x N_{frames} in the cycle !!!

Machine learning = learning from samples

How many samples? = complexity of the **question**
Only the problem? Or the **model** itself?
Or the **dimensionality** of the data?
Or noise on the data?

<u>Kernel regression (inexact):</u>

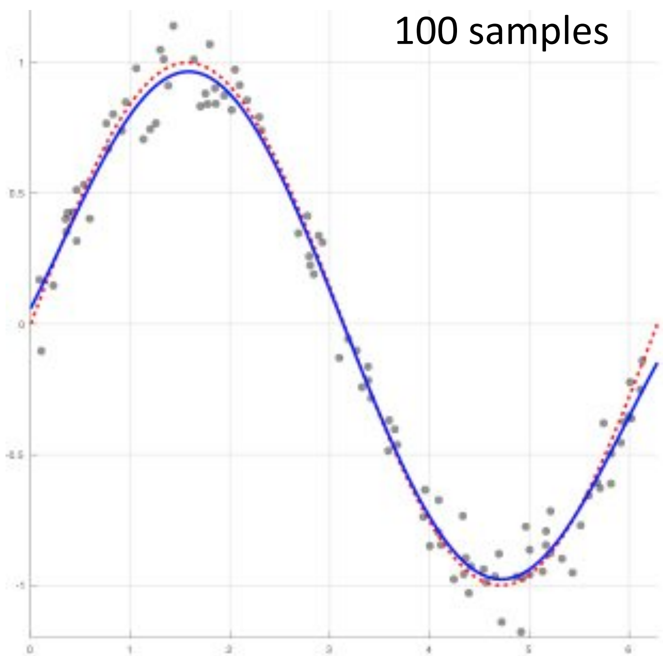

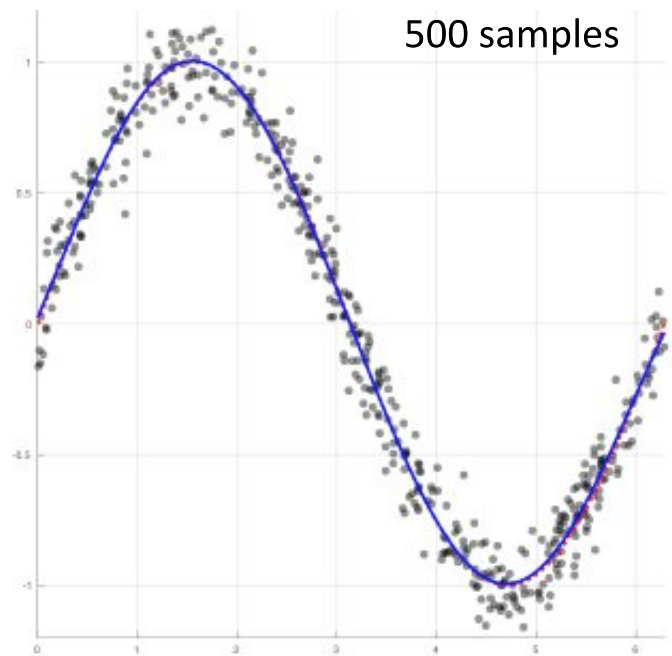

Machine learning in cardiac imaging = what for ?

$N_{samples}$ = ?
- Complexity of the **question**
- **Model** itself
- **Dimensionality** of the data
- **Noise** on the data

Machine learning in cardiac imaging = what for ?

$N_{samples}$ = ?
- Complexity of the **question**
- **Model** itself
- **Dimensionality** of the data
- **Noise** on the data

Do you have labels ?

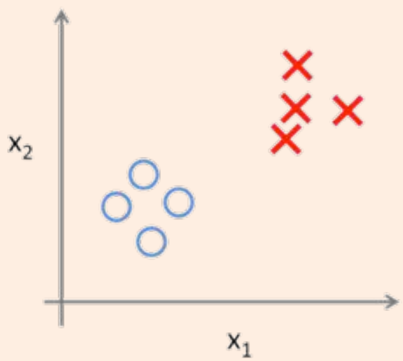

Supervised

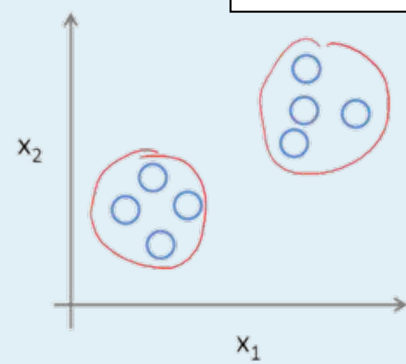

Unsupervised

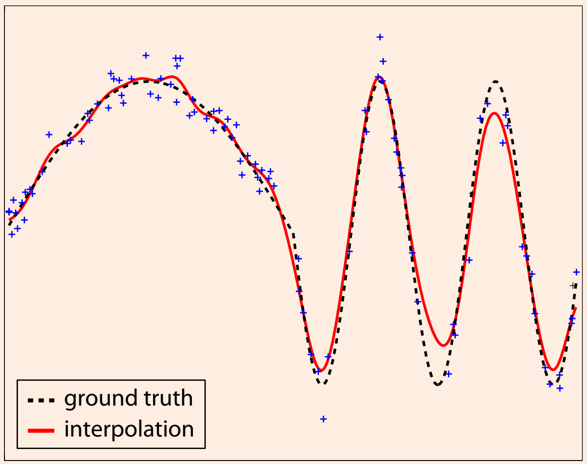

regression

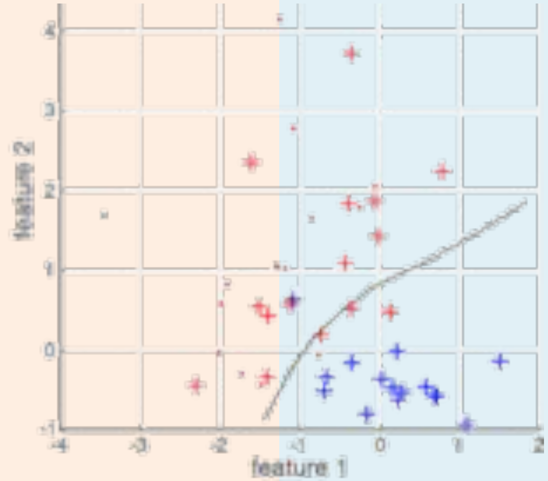

classification / clustering

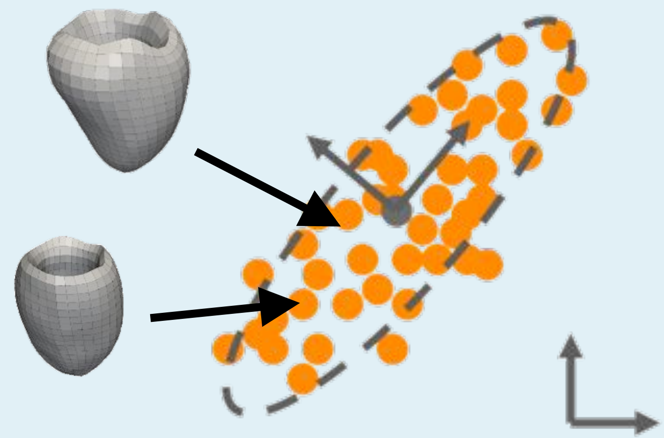

variability analysis

Machine learning in cardiac imaging = what for ?

$N_{samples}$ = ?
- Complexity of the **question**
- **Model** itself
- **Dimensionality** of the data
- **Noise** on the data

Example 1 = Automatic segmentation?

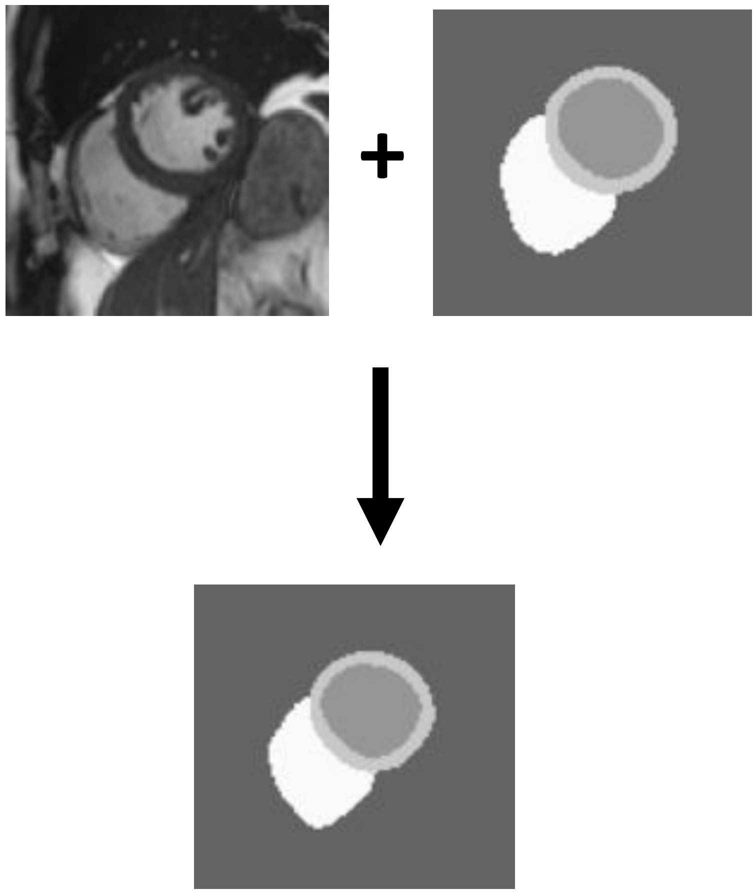

Zhang Q et al. *Ongoing work* 2018

Machine learning in cardiac imaging = what for ?

$N_{samples}$ = ?
- Complexity of the **question**
- **Model** itself
- **Dimensionality** of the data
- **Noise** on the data

Example 1 = Automatic segmentation?

+

Classification
Supervised

Zhang Q et al. *Ongoing work* 2018

Machine learning in cardiac imaging = what for ?

$N_{samples}$ = ?
- Complexity of the **question**
- **Model** itself
- **Dimensionality** of the data
- **Noise** on the data

Example 1 = Automatic segmentation?

Example 2 = Differences between healthy and…?

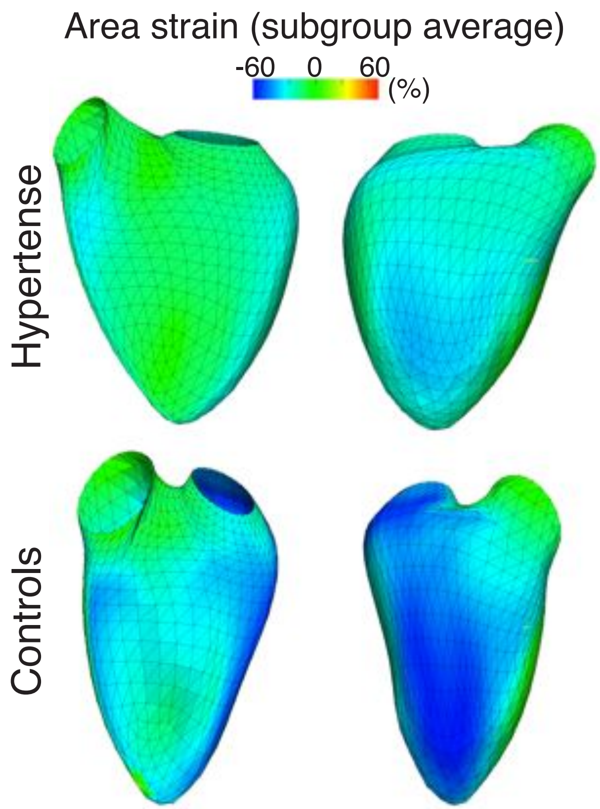

Moceri P et al. *EHJ Cardiovasc Imag* 2017

Machine learning in cardiac imaging = what for ?

$N_{samples}$ = ?
- Complexity of the **question**
- **Model** itself
- **Dimensionality** of the data
- **Noise** on the data

Example 1 = Automatic segmentation?

Example 2 = Differences between healthy and…?

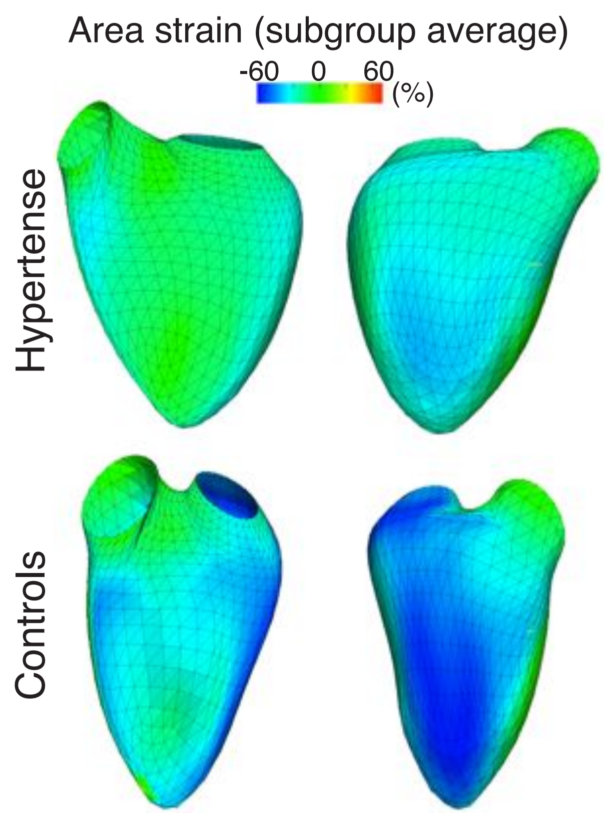

Variability analysis
Supervised

Moceri P et al. *EHJ Cardiovasc Imag* 2017

Machine learning in cardiac imaging = what for ?

$N_{samples}$ = ? - Complexity of the **question**
- **Model** itself
- **Dimensionality** of the data
- **Noise** on the data

Example 1 = Automatic segmentation?

Example 2 = Differences between healthy and…?

Example 3 = Predict infarct location from myocardial deformation?

Deformation (strain from echo)

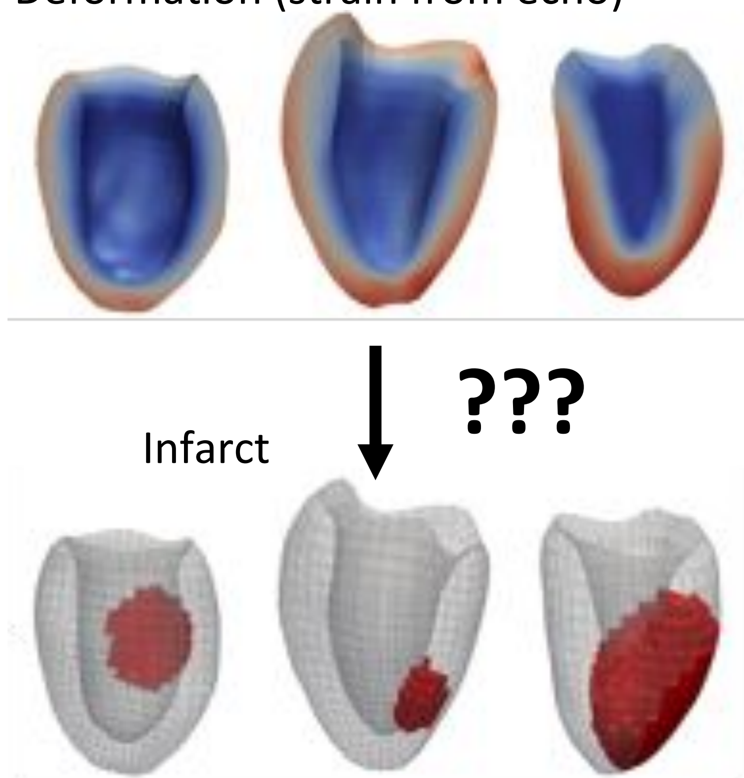

Duchateau N et al. *IEEE TMI* 2016

Machine learning in cardiac imaging = what for ?

$N_{samples}$ = ?
- Complexity of the **question**
- **Model** itself
- **Dimensionality** of the data
- **Noise** on the data

Example 1 = Automatic segmentation?

Example 2 = Differences between healthy and...?

Example 3 = Predict infarct location from myocardial deformation?

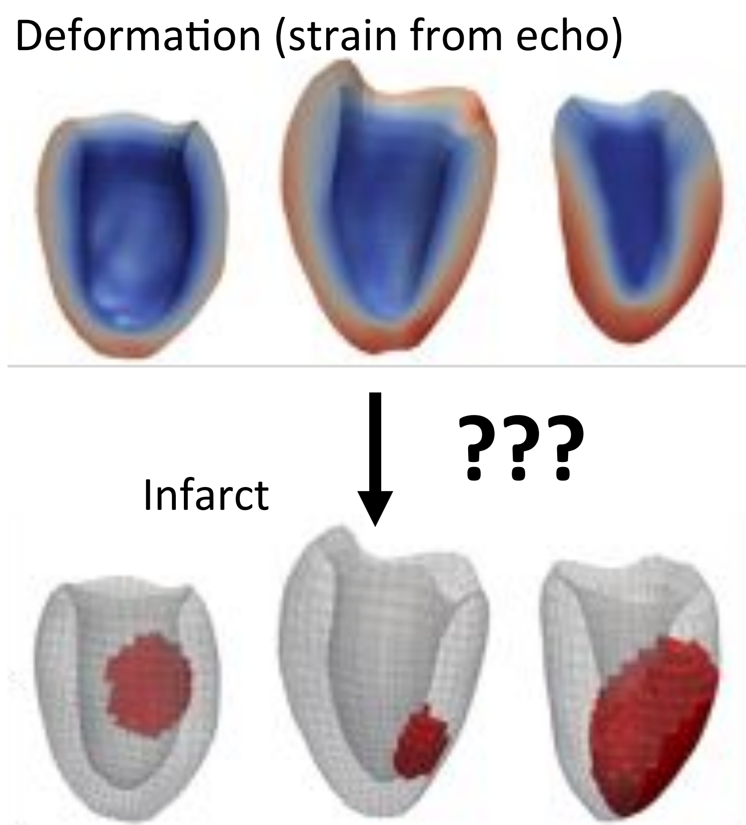

Regression

Supervised

Duchateau N et al. *IEEE TMI* 2016

Machine learning in cardiac imaging = what for ?

$N_{samples}$ = ?
- Complexity of the **question**
- **Model** itself
- **Dimensionality** of the data
- **Noise** on the data

Example 1 = Automatic segmentation?
Example 2 = Differences between healthy and…?
Example 3 = Predict infarct location from myocardial deformation?
Example 4 = Detect outliers in a coherent population?

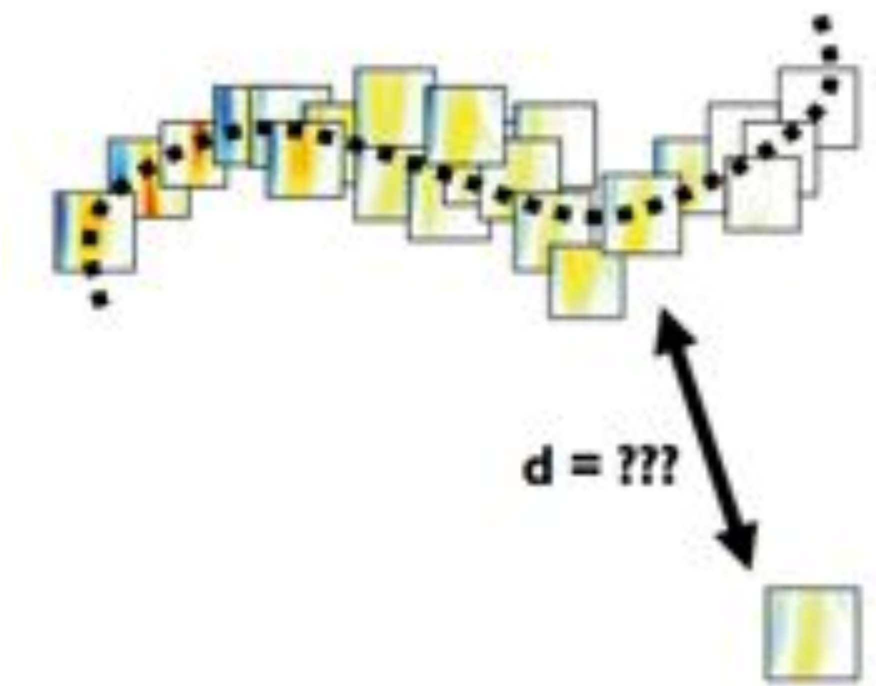

Duchateau N et al. *Med Image Anal* 2012

Machine learning in cardiac imaging = what for ?

$N_{samples}$ = ?
- Complexity of the **question**
- **Model** itself
- **Dimensionality** of the data
- **Noise** on the data

Example 1 = Automatic segmentation?

Example 2 = Differences between healthy and...?

Example 3 = Predict infarct location from myocardial deformation?

Example 4 = Detect outliers in a coherent population?

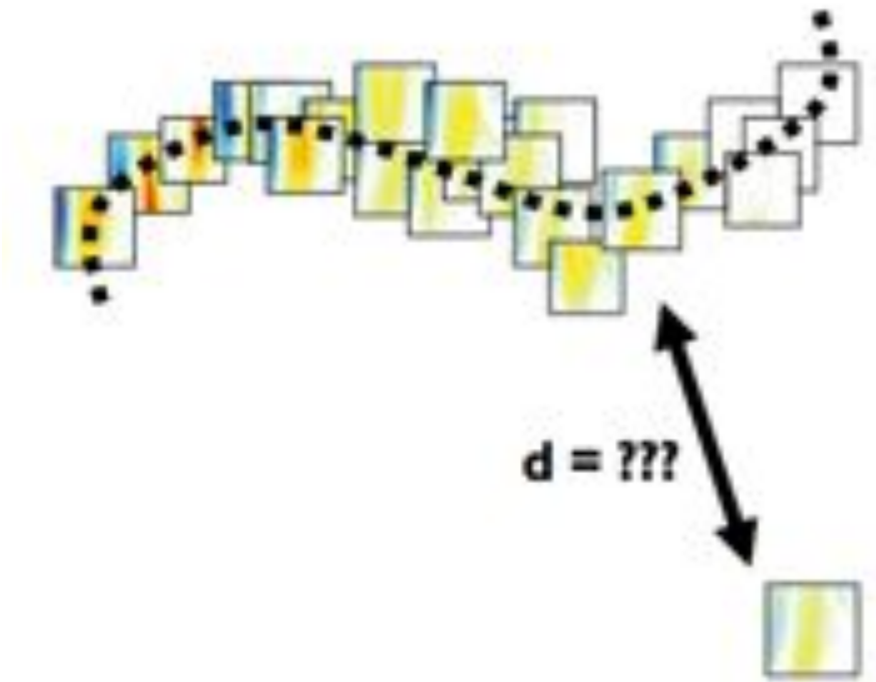

Variability analysis + clustering
Unsupervised

Duchateau N et al. *Med Image Anal* 2012

Machine learning in cardiac imaging = what for ?

$N_{samples}$ = ?
- Complexity of the **question**
- **Model** itself
- **Dimensionality** of the data
- **Noise** on the data

Classification
- Diagnosis
- Segmentation

Regression

Clustering

Variability analysis

Complexity of the question…

Machine learning in cardiac imaging = what for ?

$N_{samples}$ = ?
- Complexity of the **question**
- **Model** itself
- **Dimensionality** of the data
- **Noise** on the data

Classification
- Diagnosis
- Segmentation

Regression

Clustering

Variability analysis

Complexity of the question…

vs. complexity of the disease !

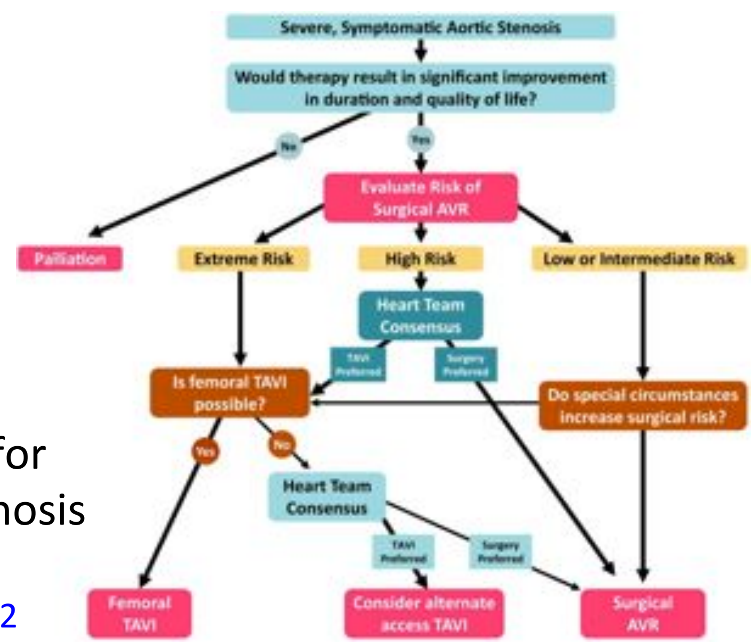

Clinical decision tree for patients with aortic stenosis

Web J et al. *Can J Cardiol* 2012

Machine learning in cardiac imaging = what for ?

$N_{samples}$ = ?
- Complexity of the **question**
- **Model** itself
- **Dimensionality** of the data
- **Noise** on the data

Example: manifold learning

Linear

Principal component analysis

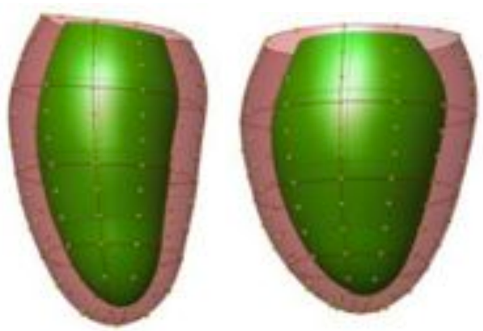

Medrano-Gracia P et al. *JCRM* 2014

Machine learning in cardiac imaging = what for ?

$N_{samples}$ = ?
- Complexity of the **question**
- **Model** itself
- **Dimensionality** of the data
- **Noise** on the data

Example: manifold learning

Linear

Principal component analysis

Medrano-Gracia P et al. *JCRM* 2014

Non-linear

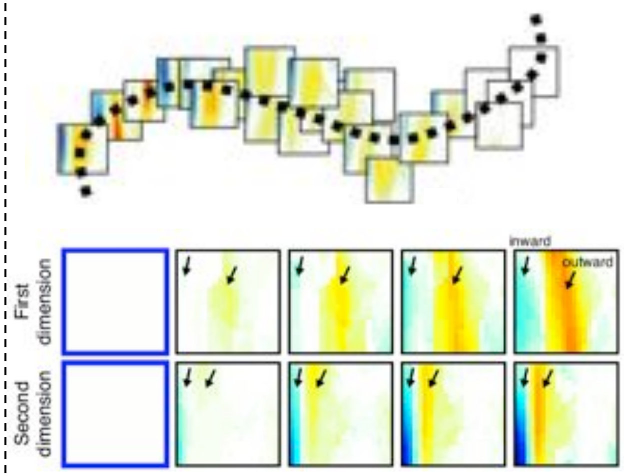

Duchateau N et al. *Med Image Anal* 2012

Deep learning

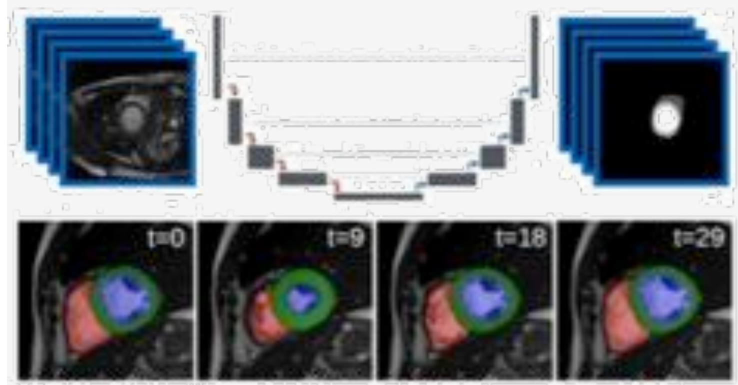

Insensee F et al. *STACOM* 2017

Machine learning in cardiac imaging = what for ?

$N_{samples}$ = ? - Complexity of the **question**
- **Model** itself
- **Dimensionality** of the data
- **Noise** on the data

Images

- Gray level, texture, …

Shapes

- Geometry (meshes, curvature,…), fibers, …

Functional features

- Global: clinical measurements, outcome, …
- Local: **mechanical (motion / deformation)**, electrical, …

Gray level
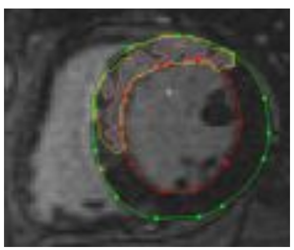
Shape

Fibers

Velocities
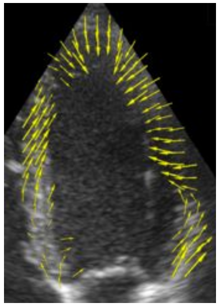
Strain
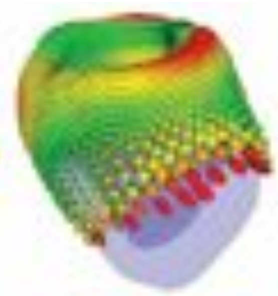
Electrical activation
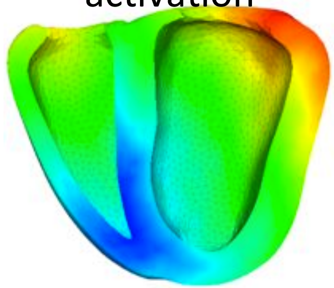

Specificities / constraints ?
- Physiology
- 4D (space + time) … or 5D (longitudinal data)

= **High** dimensionality !!!

Machine learning in cardiac imaging = what for ?

$N_{samples}$ = ? - Complexity of the **question**
- **Model** itself
- **Dimensionality** of the data
- **Noise** on the data

1 sample =
- 1 subject ?
- 1 acquisition / image ?
- 1 measurement ?

Should my sample be exactly on the "population space" ?

➢ Uncertainty in the subject's… **physiological condition** (many factors !!!)
+ measurements

Machine learning in cardiac imaging = what for ?

$N_{samples}$ = ?
- Complexity of the **question**
- **Model** itself
- **Dimensionality** of the data
- **Noise** on the data

1 sample =
- 1 subject ?
- 1 acquisition / image ?
- 1 measurement ?

Should my sample be exactly on the “population space” ?

- ➢ Uncertainty in the subject’s…
- ➢ **Generalization ability of my algorithms !**

Kernel regression (exact)
20 samples

Before simulating: where are we in terms of databases ?

Before simulating: where are we in terms of databases ?

General population

- **Cardiac Atlas Project**
  MESA study (6500+ cases)

  Heckbert S et al. *JACC* 2006
  Medrano-Gracia P et al. *JCRM* 2014

- **European Association of Cardiovascular Imaging**
  EchoNoRMAL study (51000+ cases)
  NORRE study study (1100 cases… 440 with suitable 3D echo)

  Lancellotti P et al. *EHJ-CI* 2013
  Bernard A et al. *EHJ-CI* 2017

- **UK Biobank**
  Planned 100000+ cases
  Currently 5000 cases

  Petersen S et al. *JCMR* 2013

- **etc…**

Before simulating: where are we in terms of databases ?

General population

The reality of specific populations = need for a well-settled question !!!

Physiologist's point-of-view ?

- Coherent chosen population
- Understanding physiology
- Identified clinical question

- "Small" population
- Too "controlled" population?

+ multi-center data

Cohort's point-of-view ?

- Much more heterogeneous samples
- Classification (trial)
- … or discovery / question to find?

- Large population
- Real-life population?

Generalization ability of algorithms?

PhD of Sanchez-Martinez S @UPF Barcelona

Simulation approaches for learning: realistic enough ?

Simulation approaches for learning: realistic enough ?

Level #1 = gathering practice

MACHINE "LEARNING":

Well-posed learning problem = a computer program is said to:

- learn from **experience E** = data ?
- with respect to some **task T**
- and some performance **measure P**

if its performance on T, as measured by P, improves with experience E.

Level #1 = gathering practice

MACHINE "LEARNING":

Well-posed learning problem = a computer program is said to:

- learn from **experience E** = data ?
- with respect to some **task T**
- and some performance **measure P**

if its performance on T, as measured by P, improves with experience E.

Level #1 = gathering practice

Shakil O et al. *J Cardiothorac Vasc Anesth* 2012

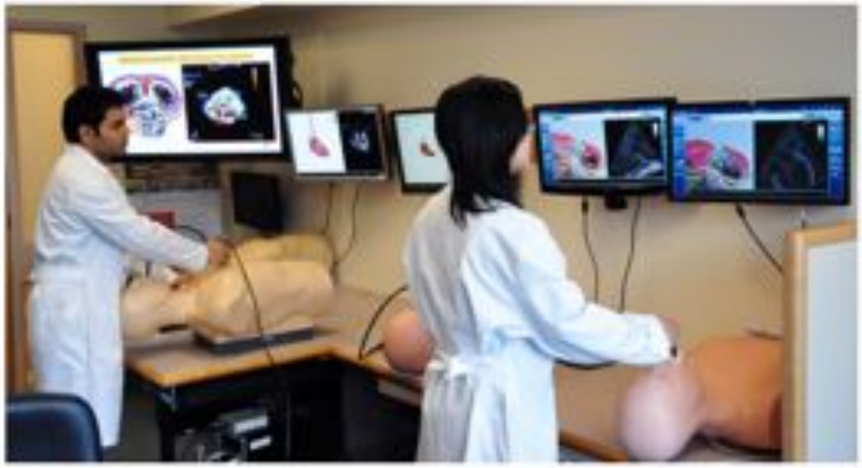

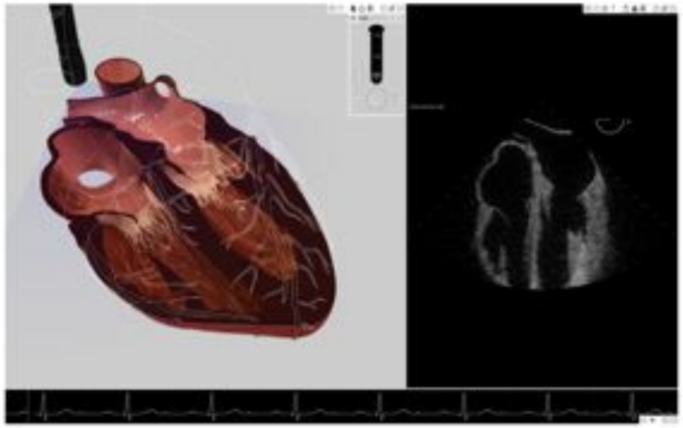

Heartworks (UK)

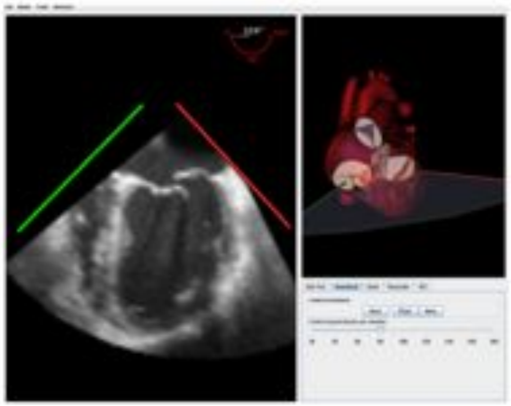

EchoCom (DE)

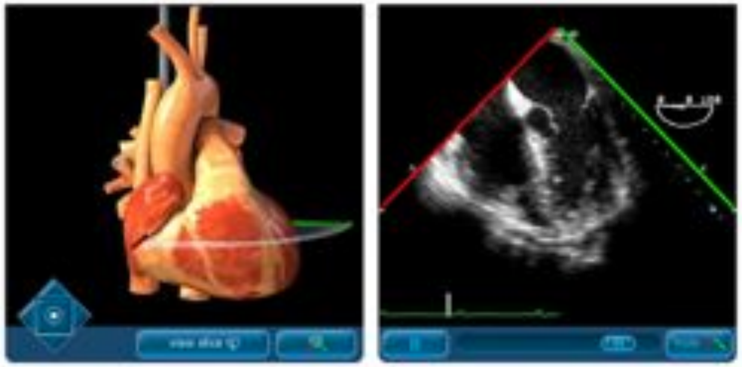

http://pie.med.utoronto.ca/TEE/

Level #1 = gathering practice

Epicardio / Schiller

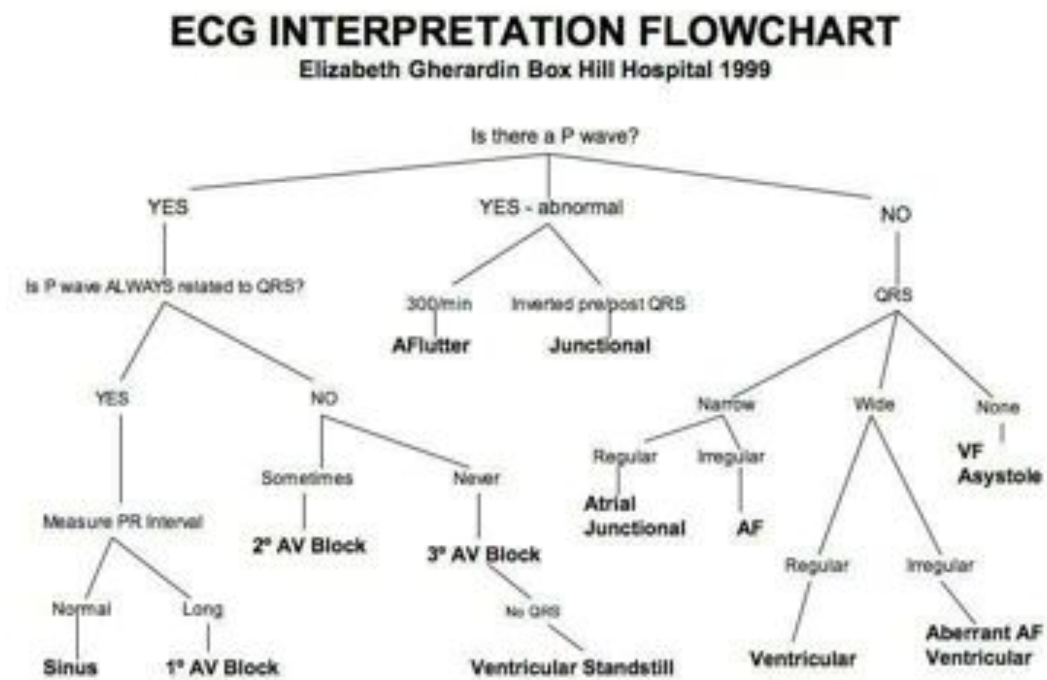

Level #2 = « realistic »… for partial validation

Controlled environment

- Population variability conditioned by few variables
- Good quality data

Level #2 = « realistic »… for partial validation

Controlled environment

- Population variability conditioned by few variables
- Good quality data

➢ **Shared purpose with animal models…**

Level #2 = « realistic »… for partial validation

Controlled environment

- Population variability conditioned by few variables
- Good quality data

➢ **Shared purpose with animal models…**

… but with more data + in-silico

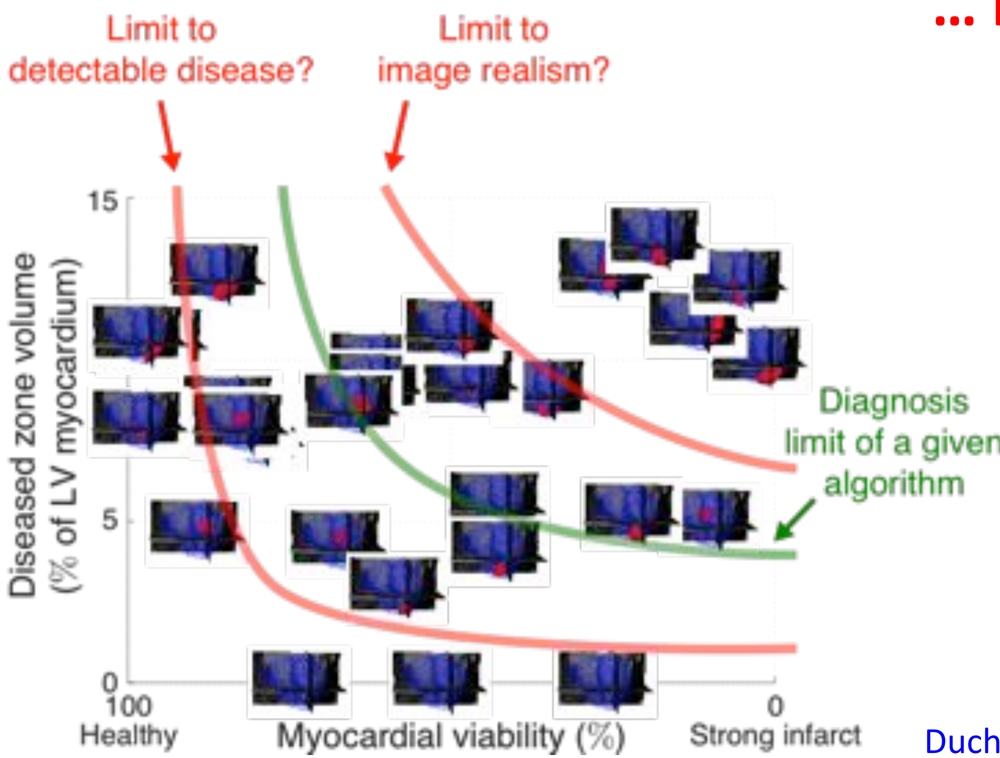

Duchateau N et al. *IEEE TMI* 2017

Level #3 = realistic enough to be mixed with real data

… common strategy for neural networks

= data augmentation

Level #3 = realistic enough to be mixed with real data

… common strategy for neural networks
= data augmentation

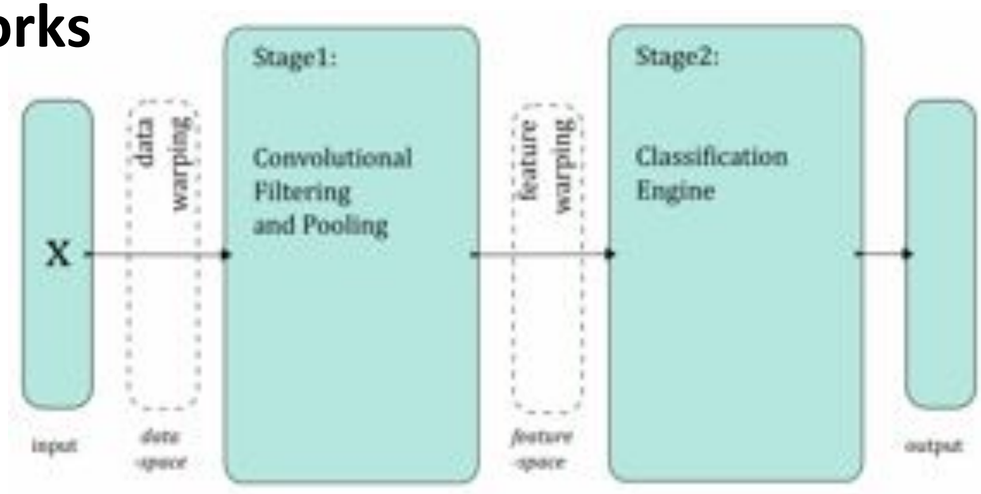

Wong S et al. *IEEE DICTA* 2016

Ratner A et al. *https://hazyresearch.github.io/snorkel/blog/tanda.html*

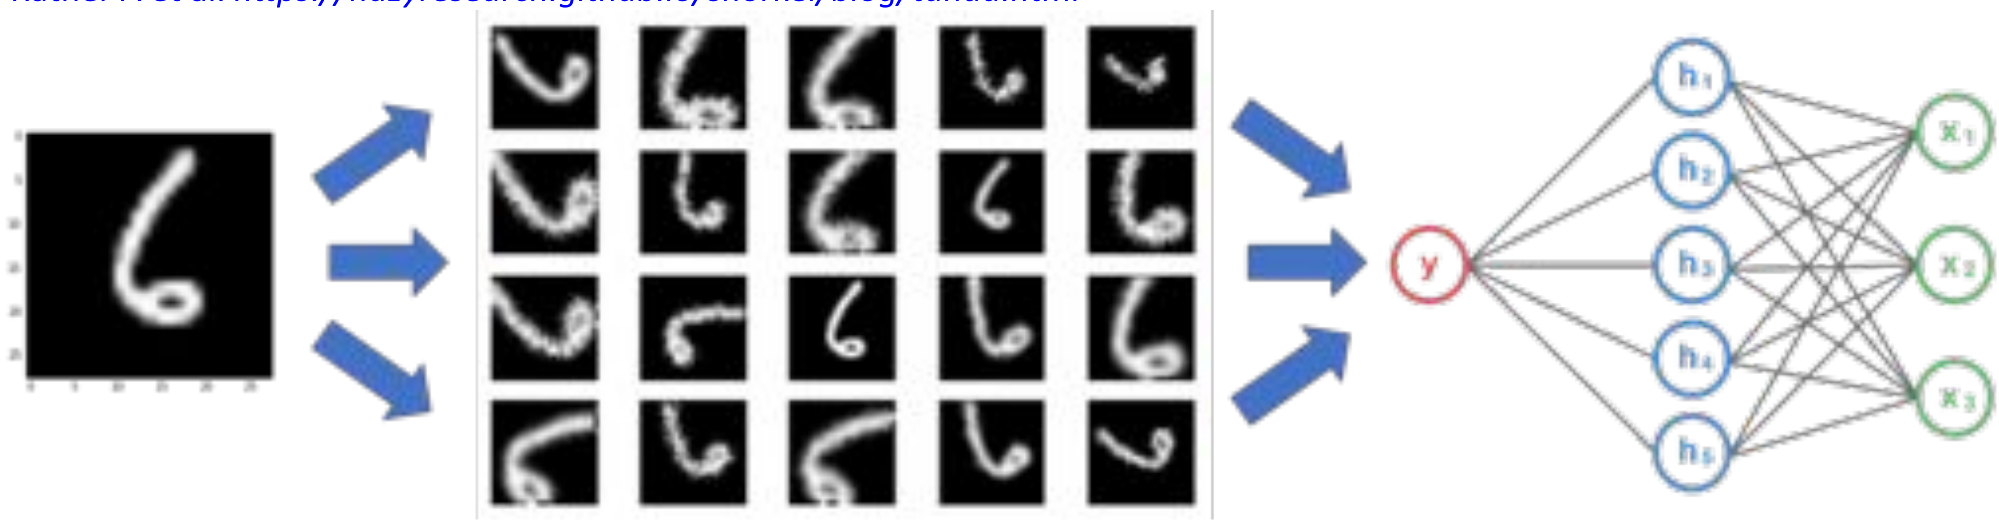

Level #3 = realistic enough to be mixed with real data

Ratner A et al. *https://hazyresearch.github.io/snorkel/blog/tanda.html*

			Without	With	
CIFAR-100	1*	DenseNet	17.18	-	Random shifts, flips
	2*	Wide ResNets	20.50	-	Random shifts, flips
	3*	ResNet (1001-Layer)	22.71	33.47	Random shifts, flips
	4*	FractalNet	23.30	35.34	Random shifts, flips
	5	Fast and Accurate Deep Network Learning by Exponential Linear Units	-	24.28	
	6	Spatially-Sparse Convolutional Neural Networks	24.3	-	Affine transformations
	7*	ResNet with Stochastic Depth (1202-Layer)	24.58	37.80	Random shifts, flips
	8	Fractional Max-Pooling	26.39	-	Randomized mix of translations, rotations, reflections, stretching, and shearing operations, and random RGB color shifts
	9*	ResNet (110-Layer)	27.22	44.74	Random shifts, flips
	10	Scalable Bayesian Optimization Using Deep Neural Networks	27.4	-	Hue, saturation, scalings, horizontal flips

Level #3 = realistic enough to be mixed with real data

Ratner A et al. *https://hazyresearch.github.io/snorkel/blog/tanda.html*

- Rotations
- Scaling / Zooms
- Brightness
- Color Shifts
- Etc…

- Synonymy
- Positional Swaps
- Etc…

Level #3 = realistic enough to be mixed with real data

Ratner A et al. *https://hazyresearch.github.io/snorkel/blog/tanda.html*

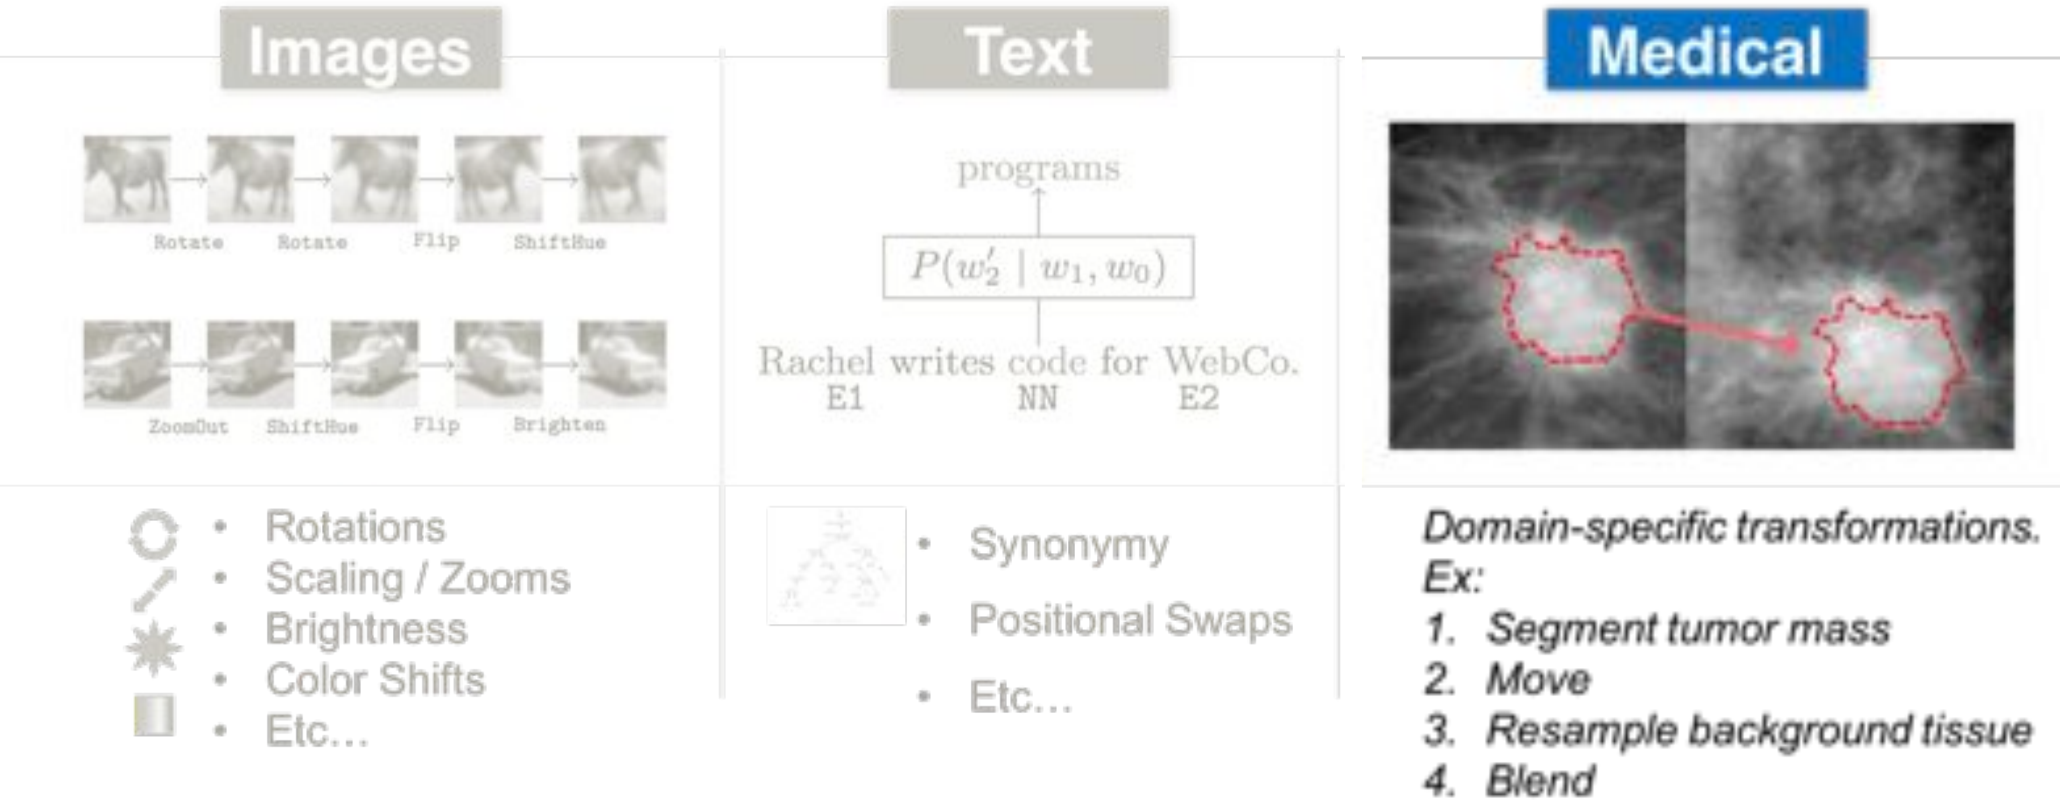

+ our current assest = realistic physiological models !

- complexity of disease representations…

Level #3 = realistic enough to be mixed with real data

Richter S et al. *ECCV* 2016

Synthetic images…
from video games !!!

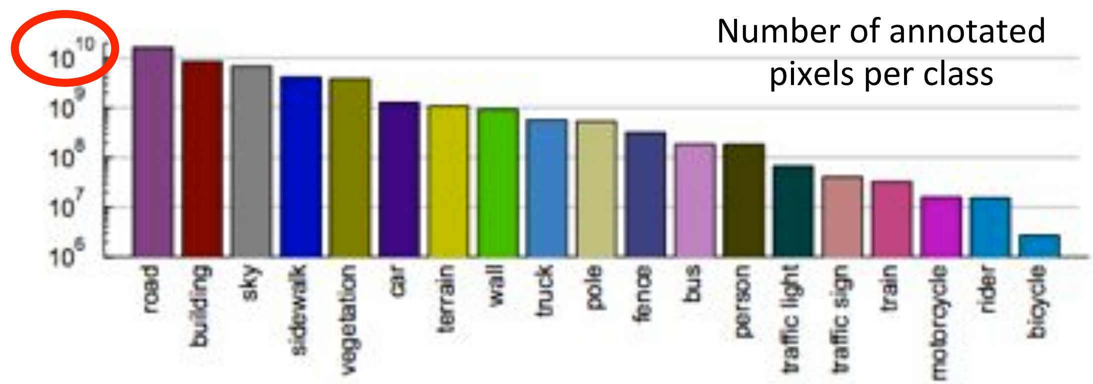

Number of annotated pixels per class

Level #3 = realistic enough to be mixed with real data

Heimann T et al. *Med Image Anal* 2014

In-silico
- 10000 images !

In-vivo
- 68 sequences
- 22 patients

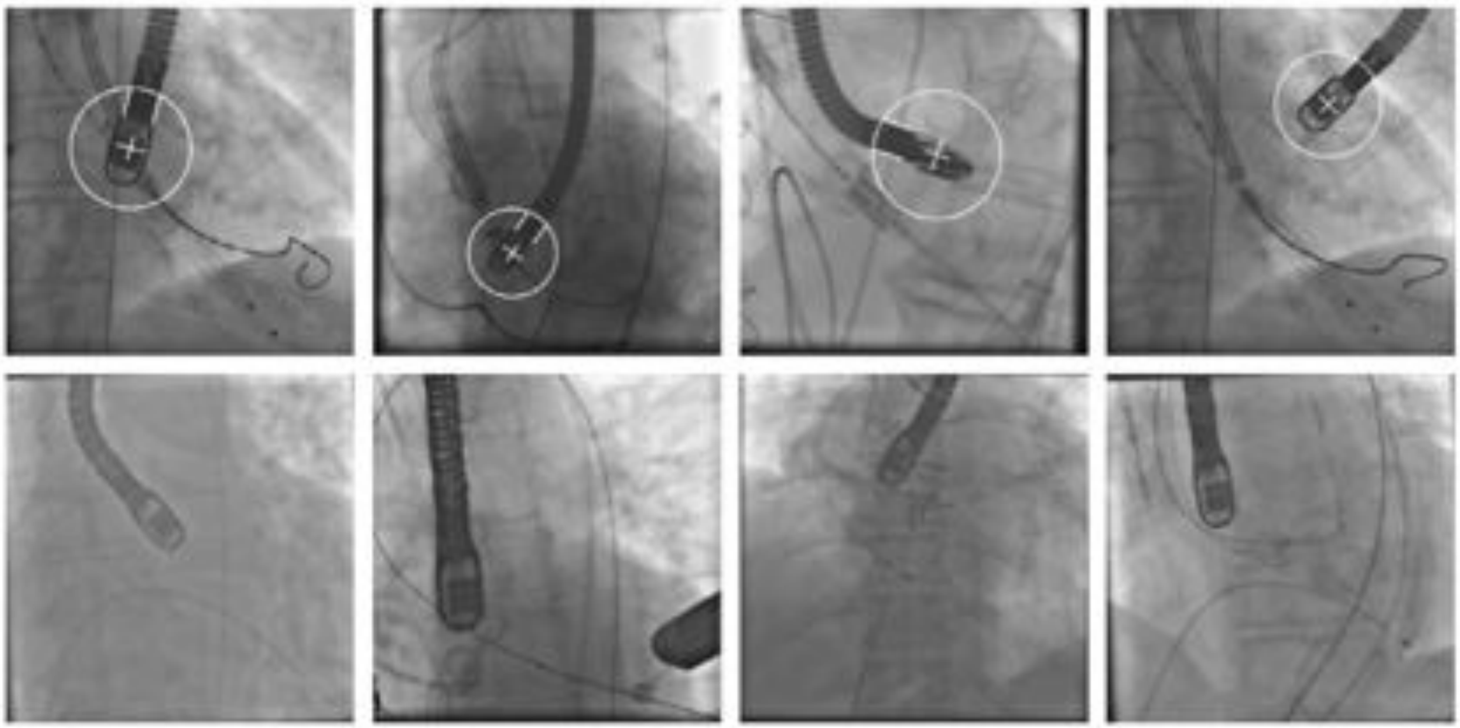

Key issue = how far are we from real data ?

Ways to quantify this = vs. your task **(features extraction, diagnosis, …)**

Key issue = how far are we from real data ?

Ways to quantify this = vs. your task (features extraction, diagnosis, ...)

A whole field of research = transfer learning

Large database

Small database
(but yours)

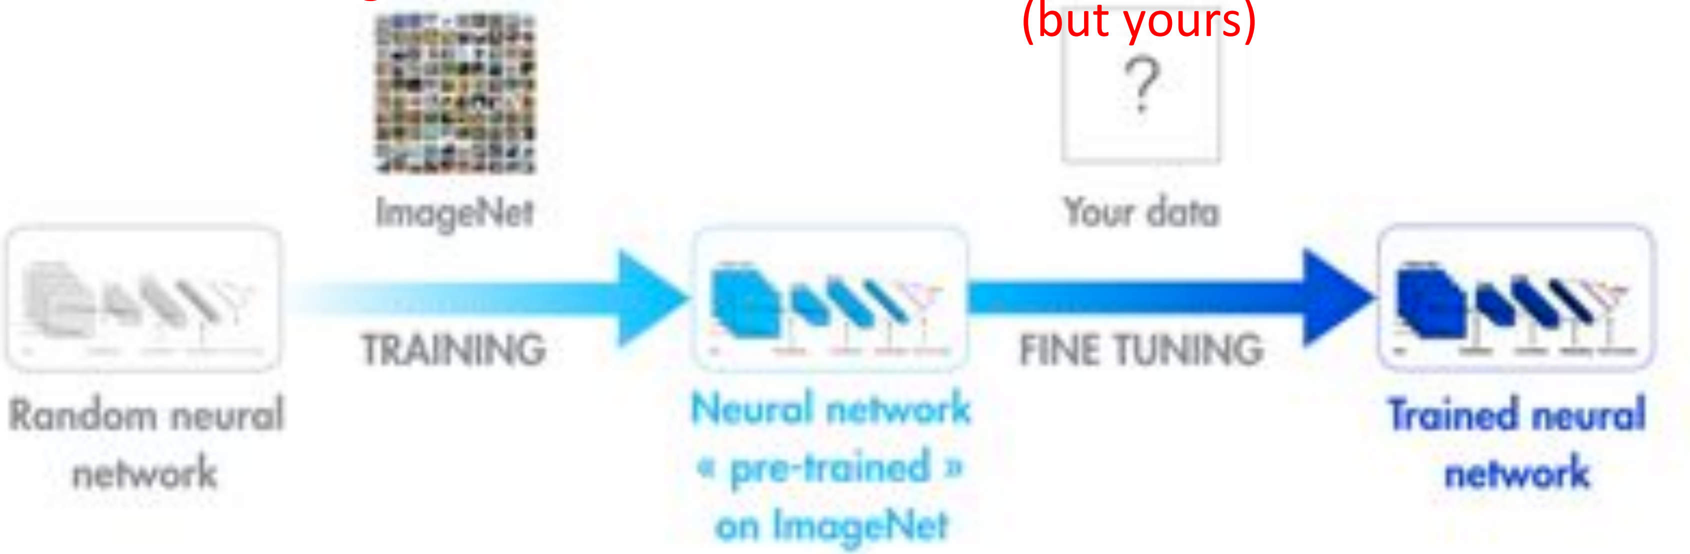

https://medium.com/owkin/transfer-learning-and-the-rise-of-collaborative-artificial-intelligence-41f9e2950657

Key issue = how far are we from real data ?

Ways to quantify this = vs. your task (features extraction, diagnosis, …)

A whole field of research = transfer learning

What's the actual challenge ?

Your data lie in a specific space

➢ Bring close enough simulation and real data spaces

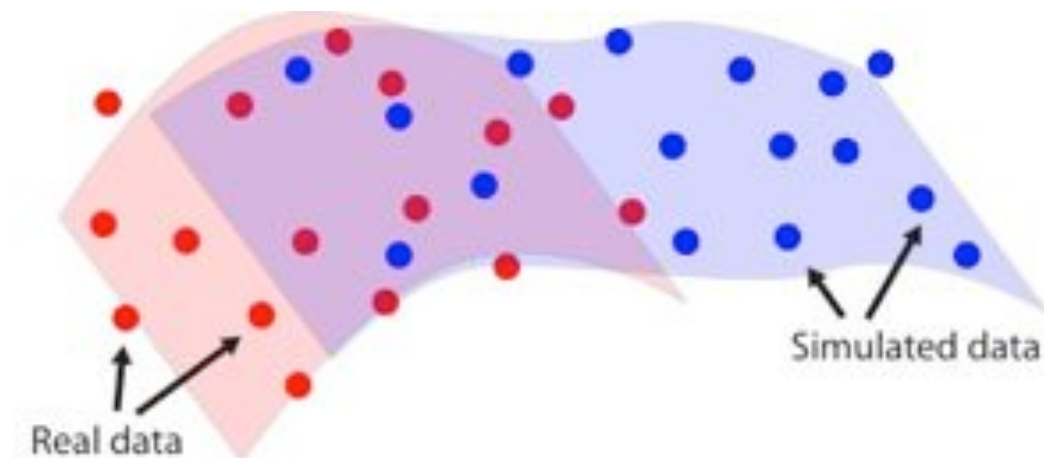

Examples = cardiac synthetic databases for machine learning

Examples = cardiac synthetic databases for machine learning

Cabrera-Lozoya R et al. *IEEE T Biomed Eng* 2016

Electrophysiology signals

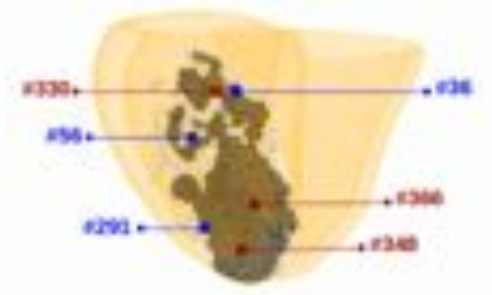

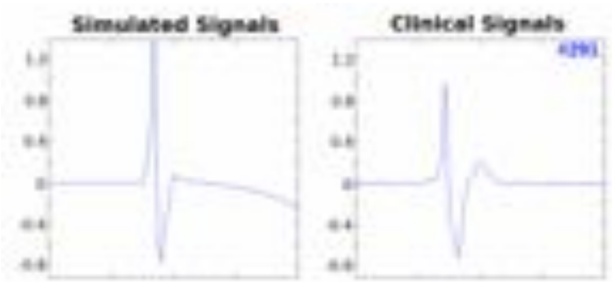

Examples = cardiac synthetic databases for machine learning

Cabrera-Lozoya R et al. *IEEE T Biomed Eng* 2016

Electrophysiology signals

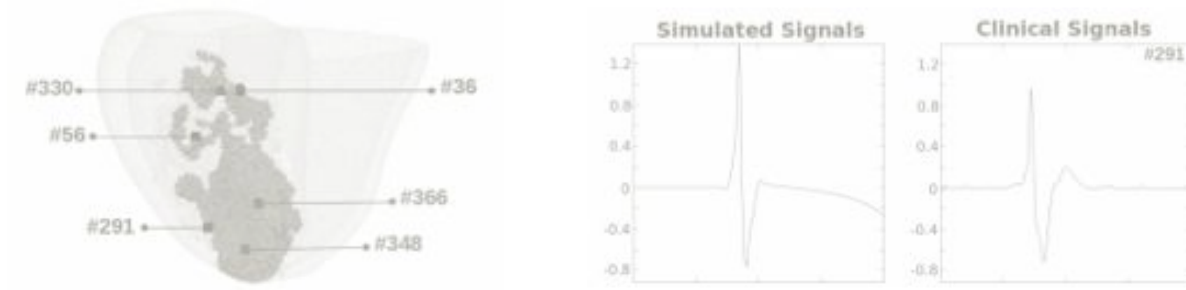

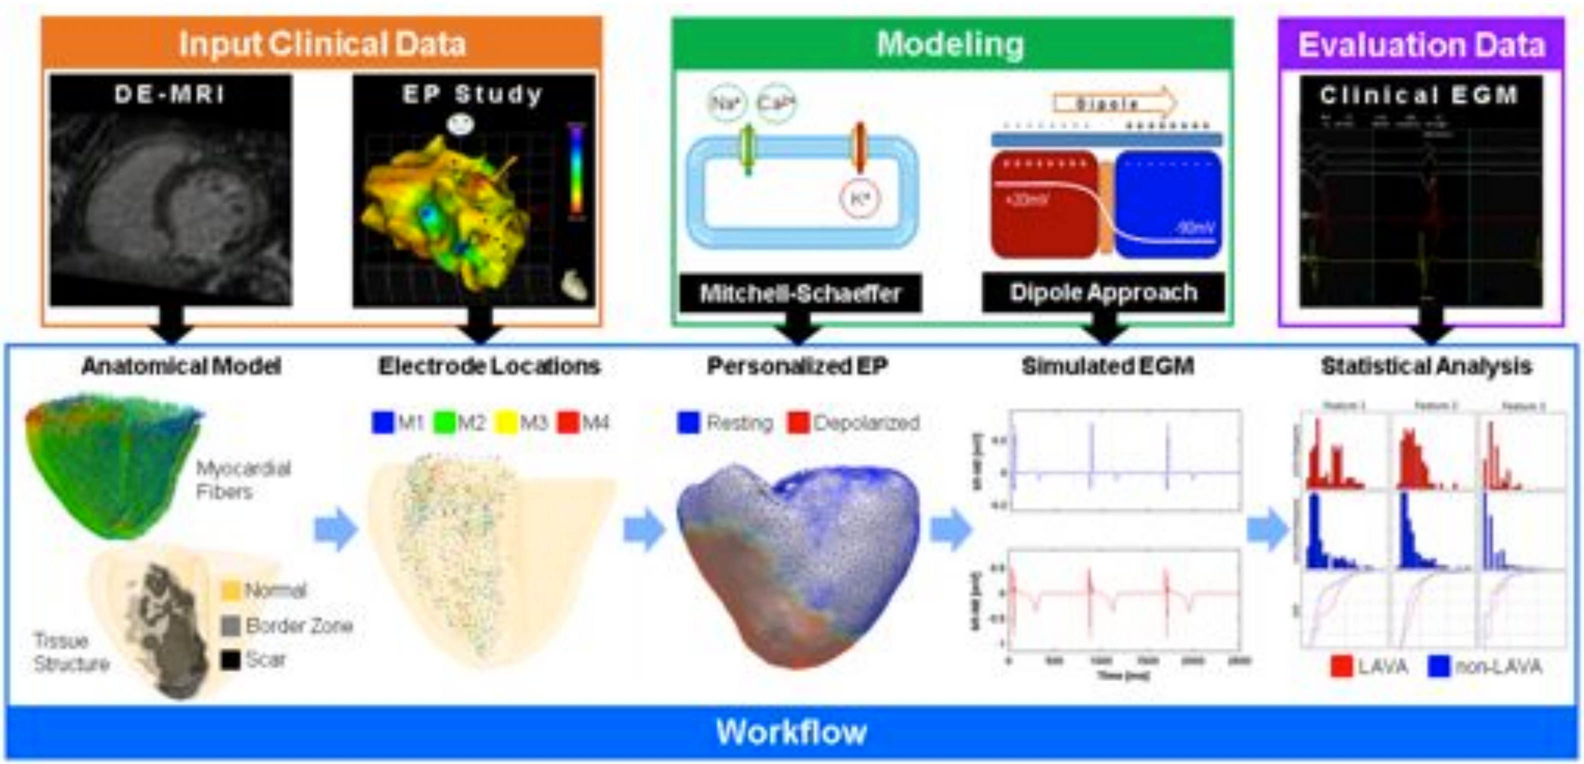

Examples = cardiac synthetic databases for machine learning

Duchateau N et al. *IEEE T Med Imag* 2016
Rumindo K et al. *FIMH* 2017

Mesh sequences
N=500 & 200

Deformation (strain from echo)

???

Infarct

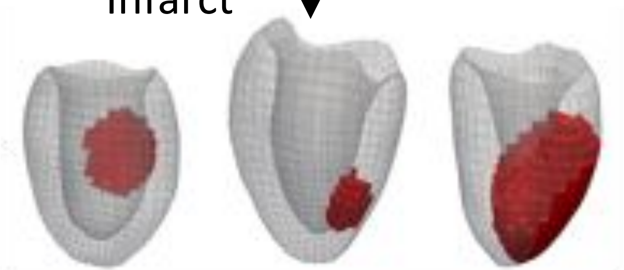

Examples = cardiac synthetic databases for machine learning

Duchateau N et al. *IEEE T Med Imag* 2016
Rumindo K et al. *FIMH* 2017

Mesh sequences
N=500 & 200

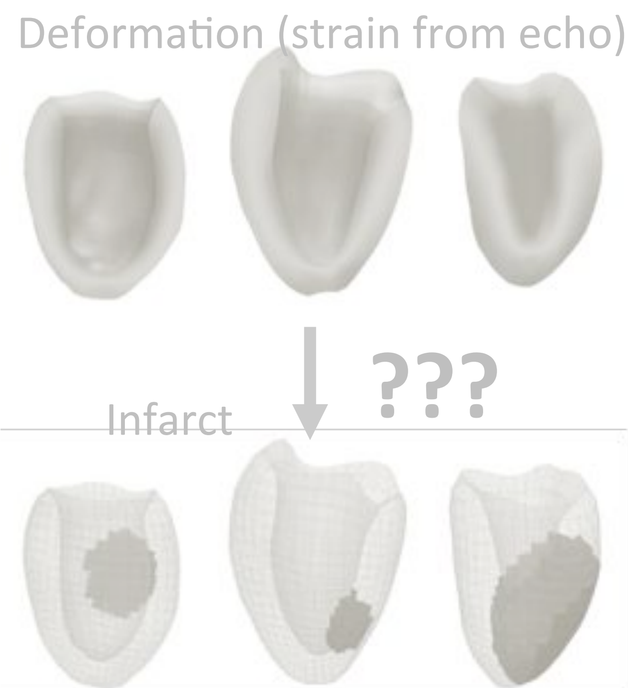

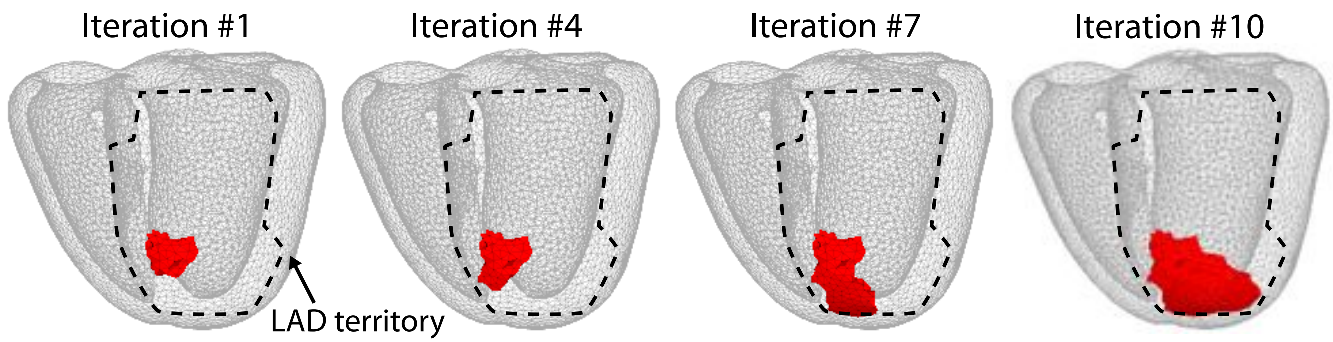

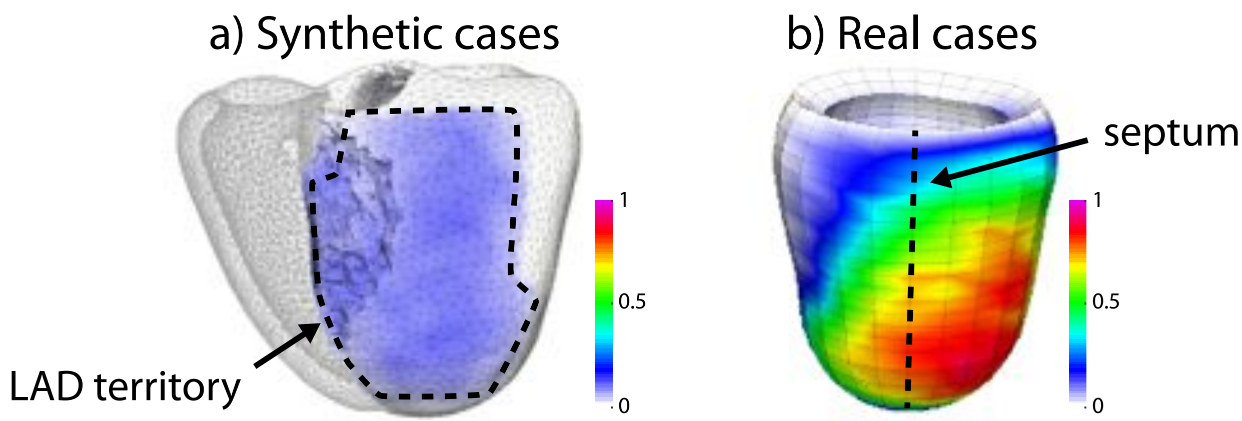

Examples = cardiac synthetic databases for machine learning

Duchateau N et al. *IEEE T Med Imag* 2016
Rumindo K et al. *FIMH* 2017

Mesh sequences
N=500 & 200

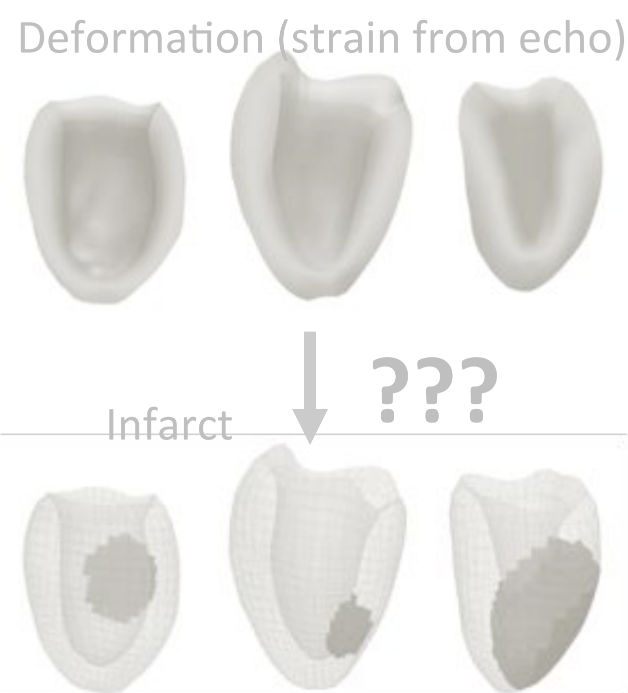

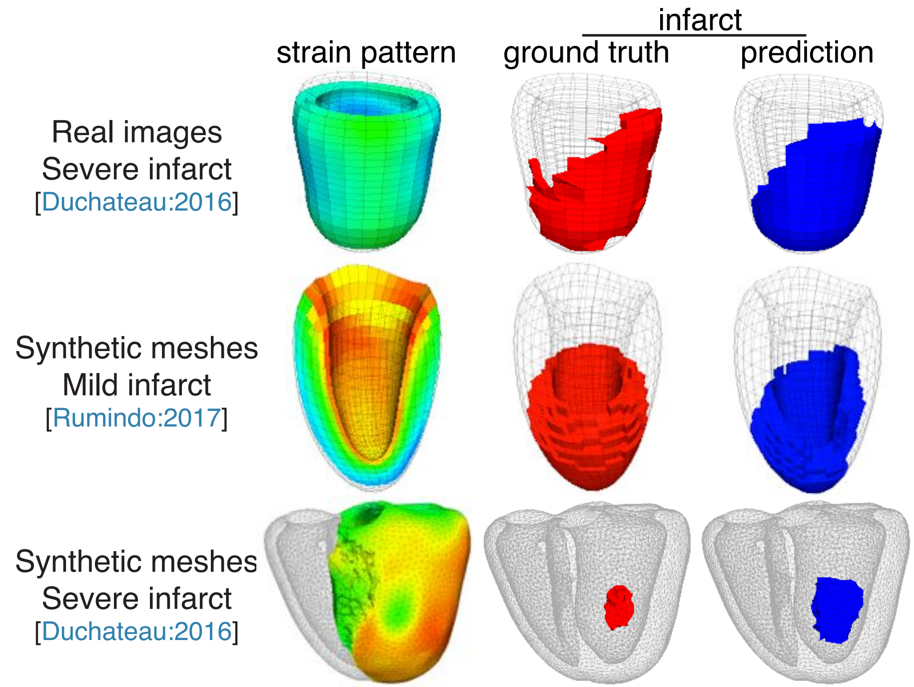

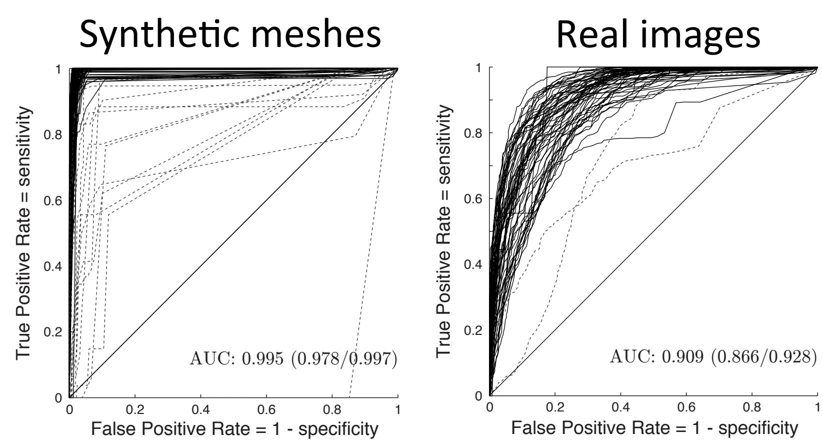

Examples = cardiac synthetic databases for machine learning

Prakosa A et al. *IEEE T Biomed Eng* 2014

Image sequences
N=144-180 / case

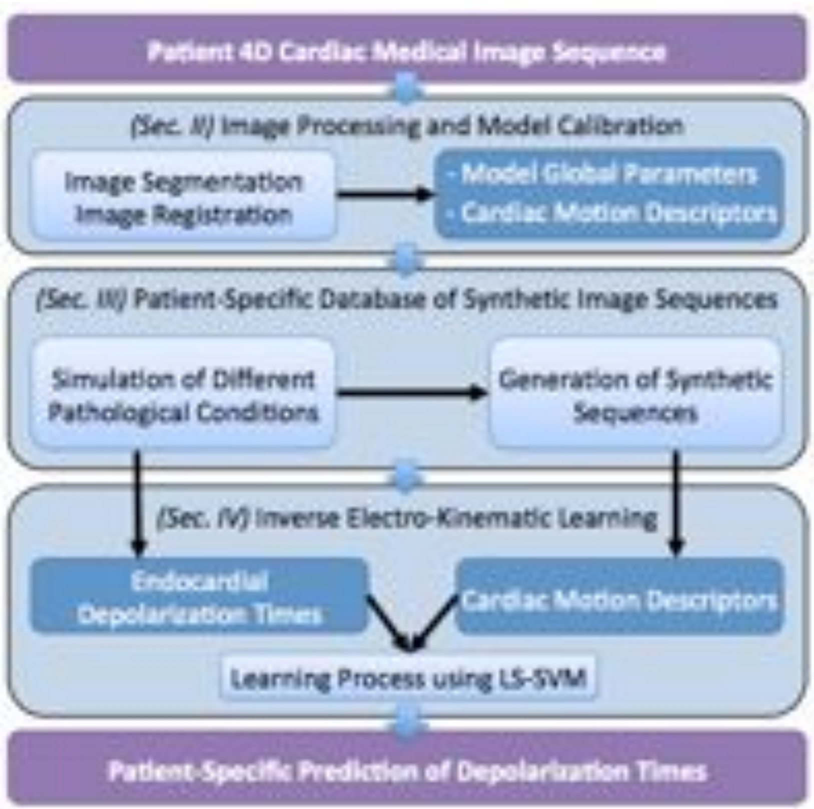

Examples = cardiac synthetic databases for machine learning

Prakosa A et al. *IEEE T Biomed Eng* 2014

Image sequences
N=144-180 / case

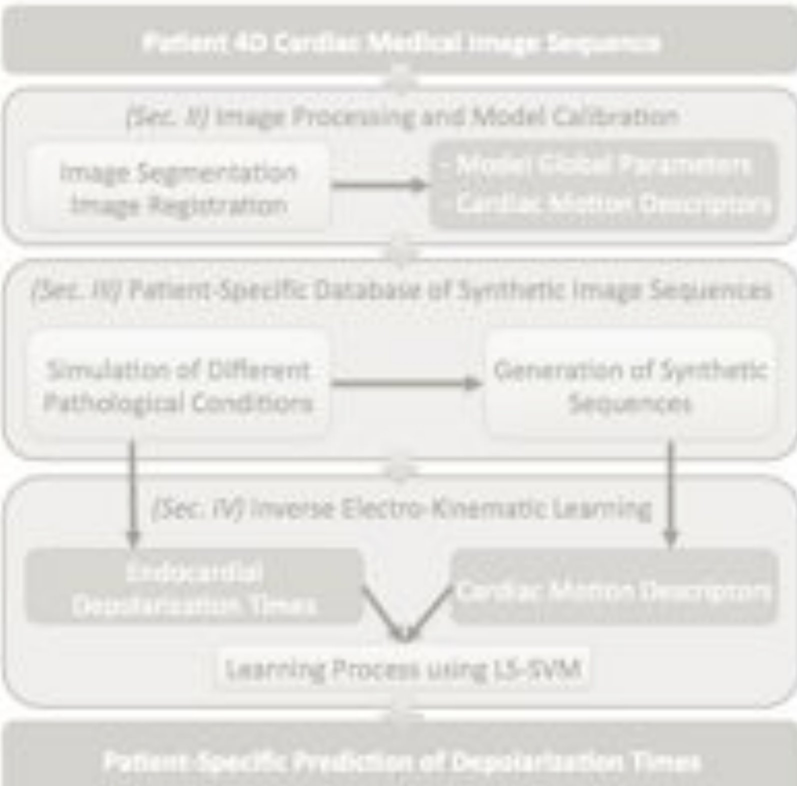

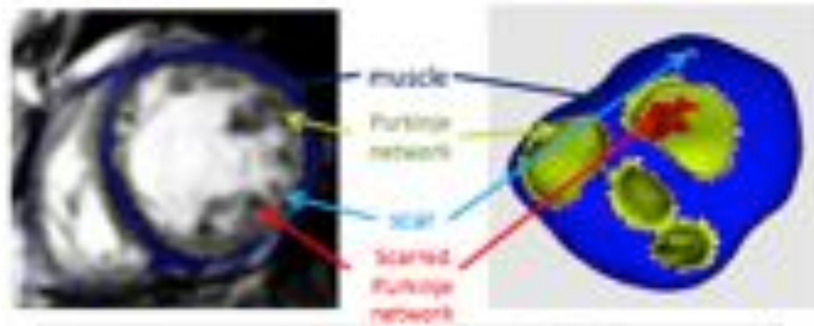

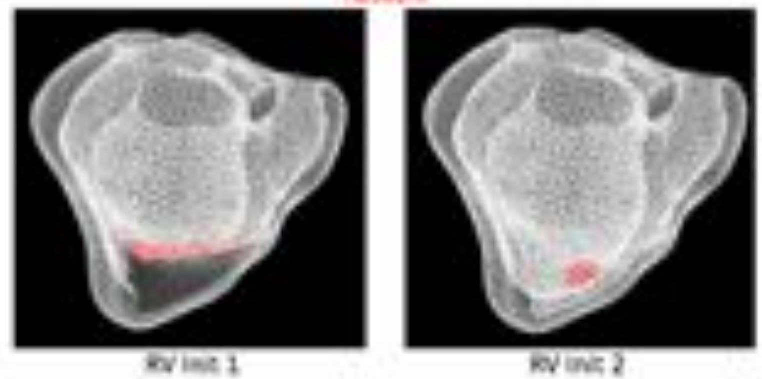

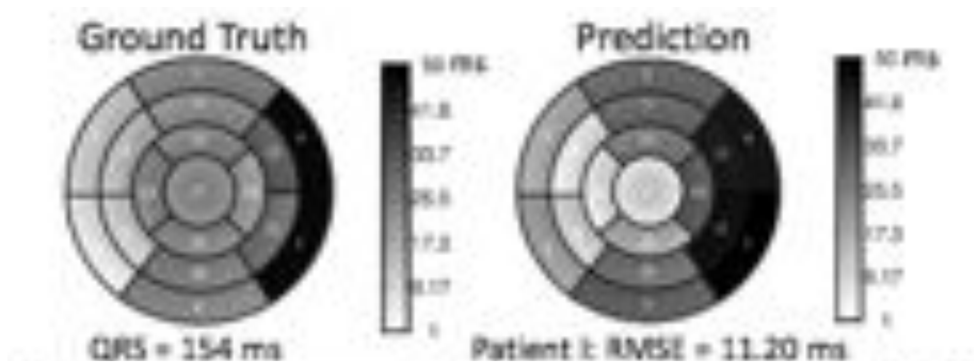

Examples = cardiac synthetic databases for machine learning

Duchateau N et al. *IEEE T Med Imag* 2017

Image sequences
N=465

Cine MRI (LAD infarct)
[Duchateau:2017]

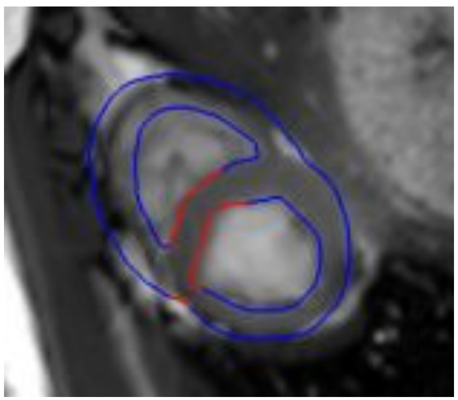

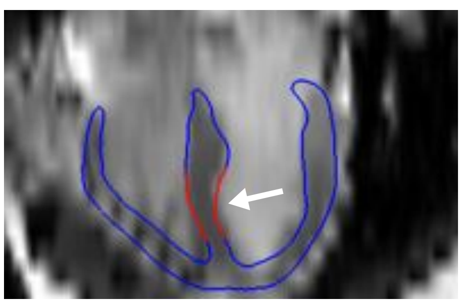

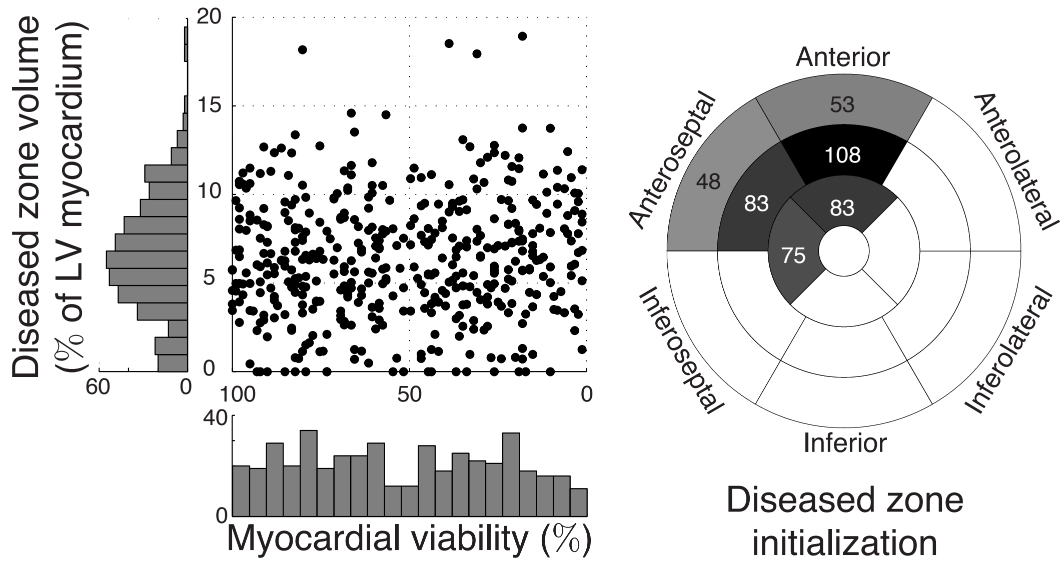

Examples = cardiac synthetic databases for machine learning

Duchateau N et al. *IEEE T Med Imag* 2017

Image sequences
N=465

Cine MRI (LAD infarct)
[Duchateau:2017]

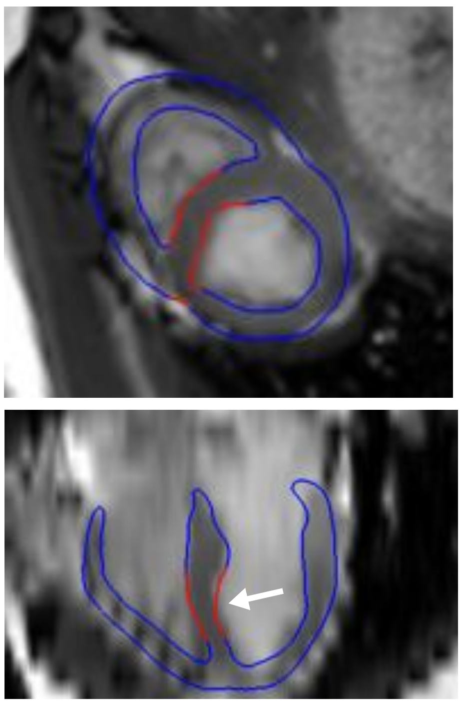

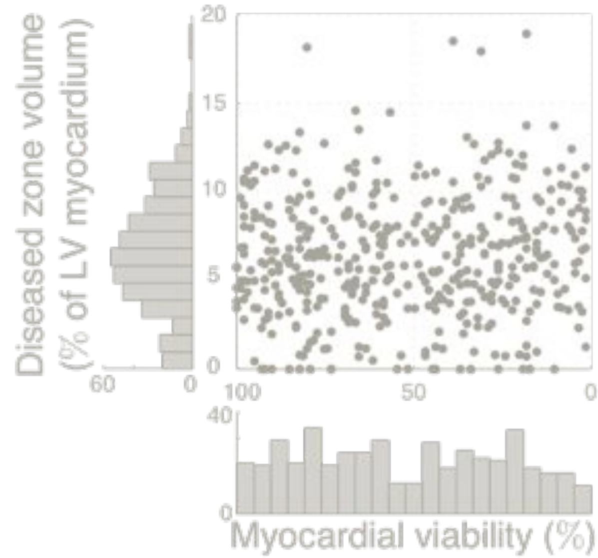

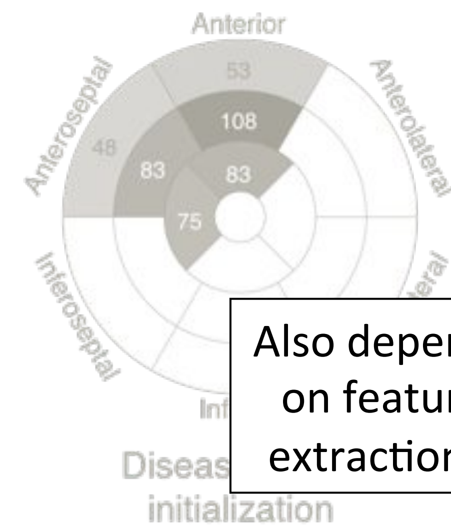

Also depends on feature extraction...

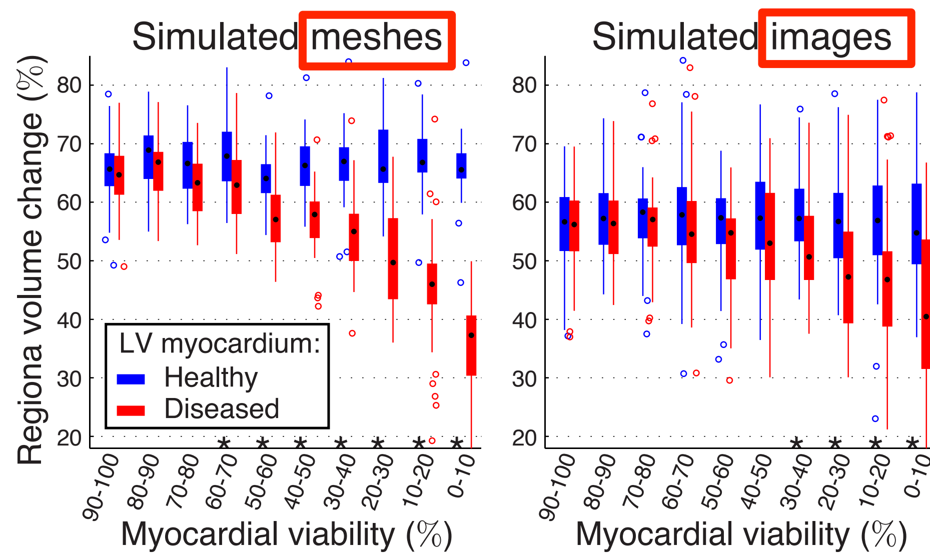

Conclusions

There is space for improvement…

- **Personalisation** to the existing samples
- Simulating **pathologies** = shape & function changes
- More **real-life images** = variety of textures, no warping, adding artifacts
- **Database** design = which cases to simulate

Conclusions

There is space for improvement…

- **Personalisation** to the existing samples
- Simulating **pathologies** = shape & function changes
- More **real-life images** = variety of textures, no warping, adding artifacts
- **Database** design = which cases to simulate

… but this hugely depends on …

- Complexity of the **question**
- **Model** itself
- **Dimensionality** of the data
- **Noise** on the data

Conclusions

There is space for improvement…

- **Personalisation** to the existing samples
- Simulating **pathologies** = shape & function changes
- More **real-life images** = variety of textures, no warping, adding artifacts
- **Database** design = which cases to simulate

… but this hugely depends on …

- Complexity of the **question**
- **Model** itself
- **Dimensionality** of the data
- **Noise** on the data

… and the next advances in databases gathering + sharing

Thanks… questions ?

Take-home = bridging the gap between simulations and real data

PRIMES UNIVERSITE DE LYON

Programme Avenir Lyon Saint-Etienne PALSE

Université Claude Bernard Lyon 1