Machine Learning

Uwe Schmidt
MPI-CBG & CSBD

#DLBC18

https://xkcd.com/1838/

Machine Learning

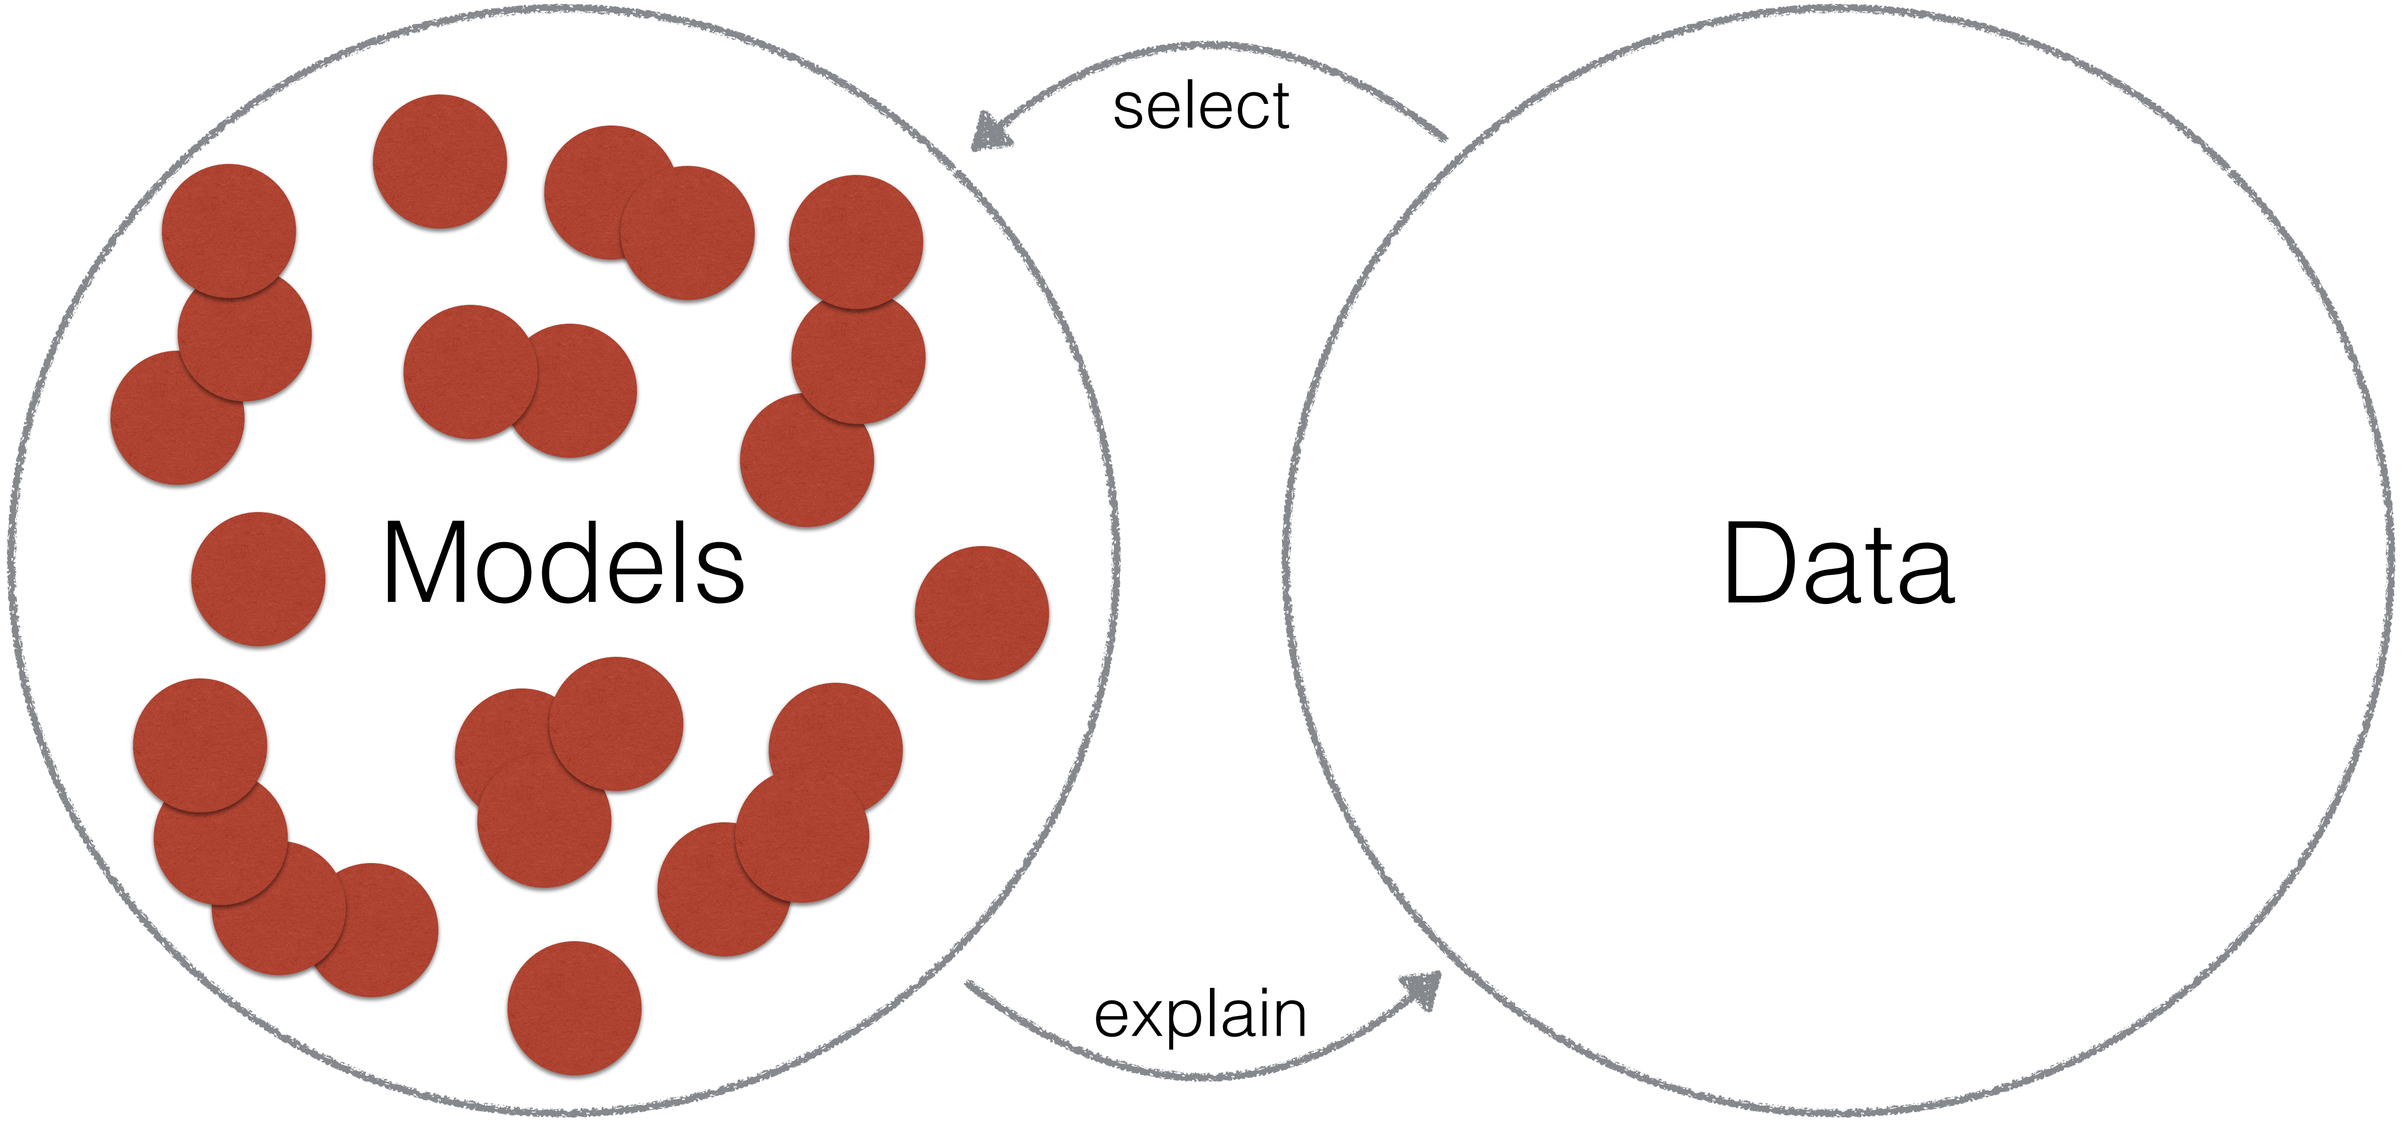

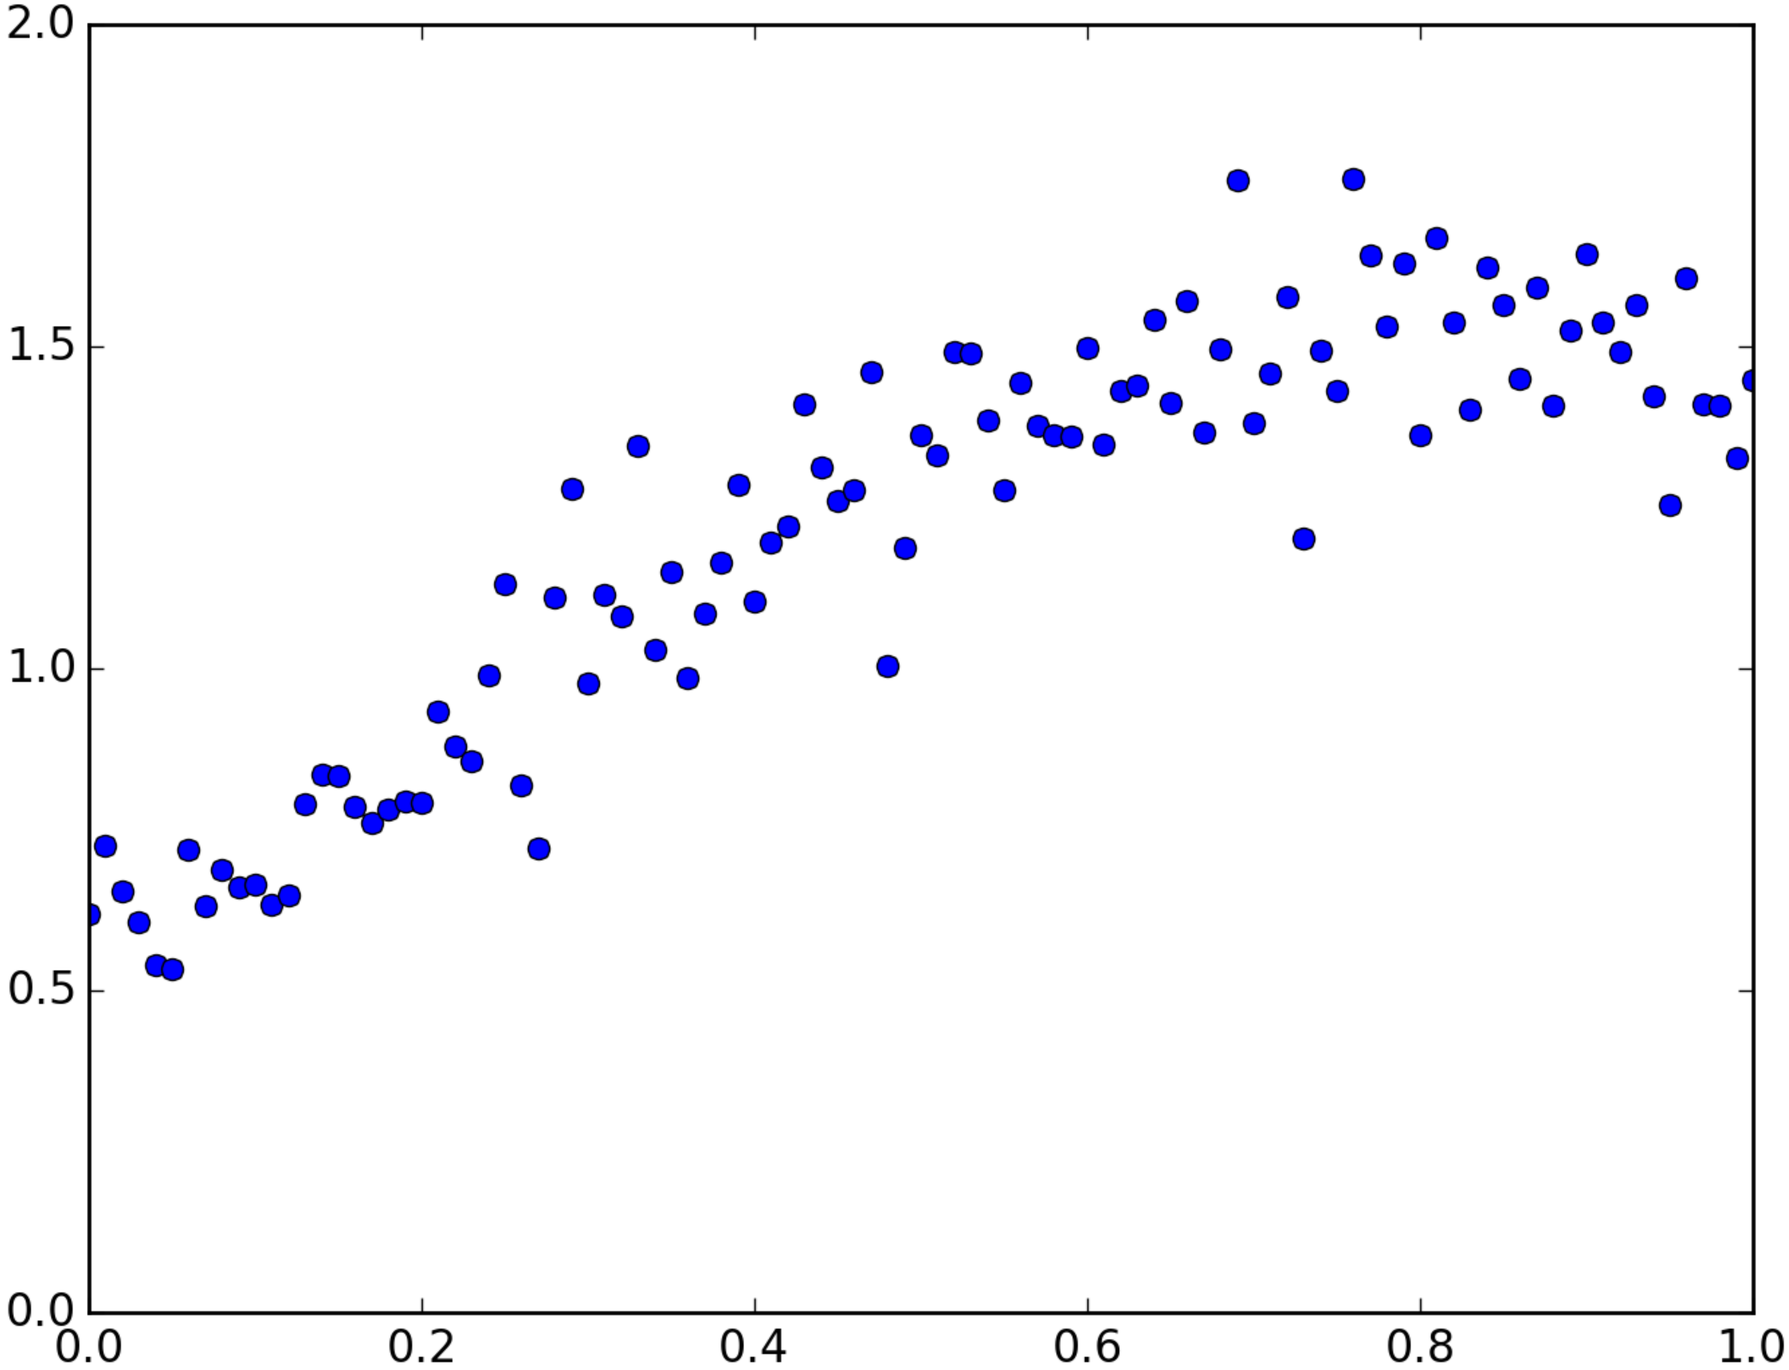
2.0
1.5
1.0
0.5
0.0
0.0
0.2
0.4
0.6
0.8
1.0

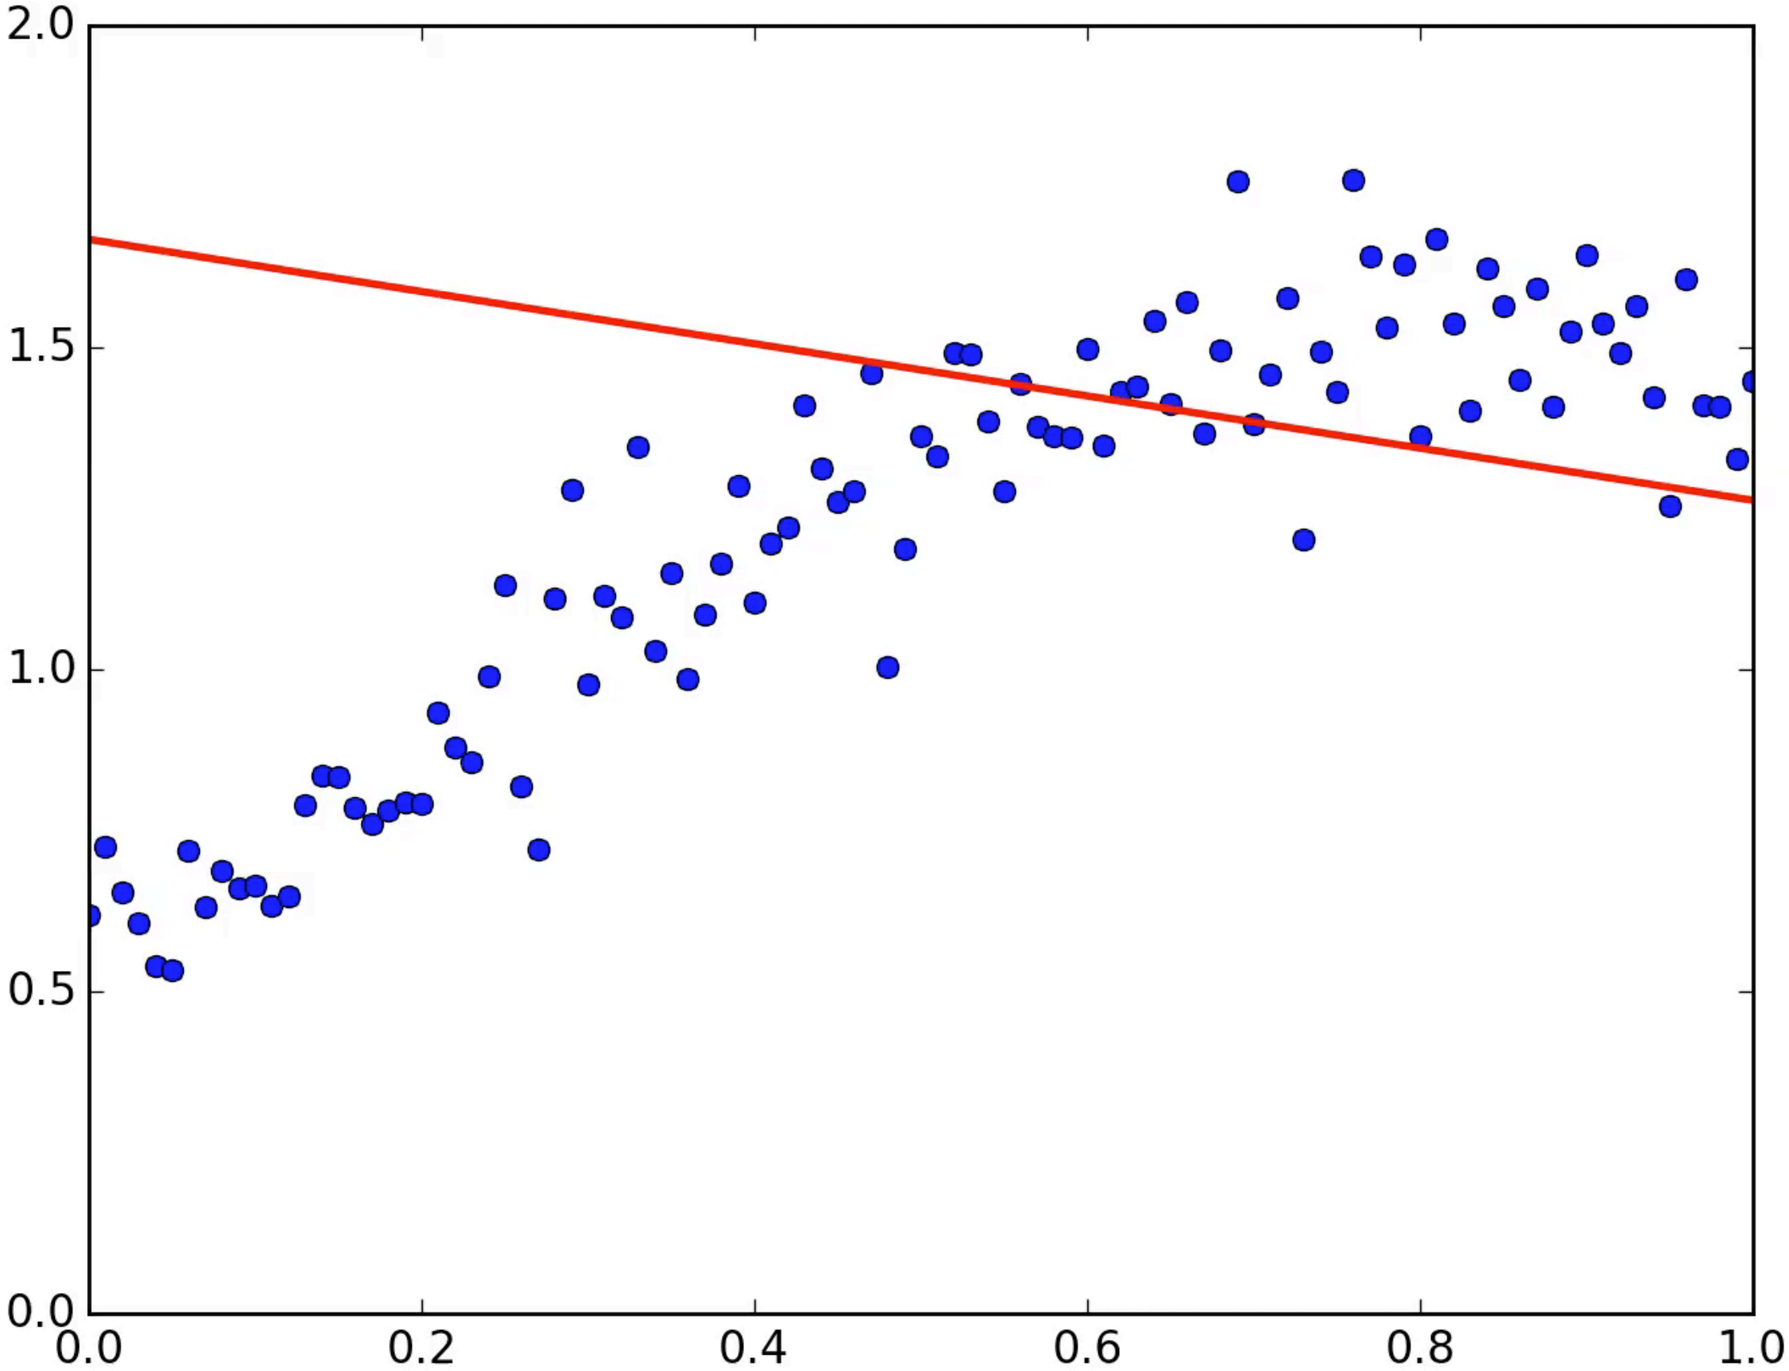
2.0
1.5
1.0
0.5
0.0
0.0
0.2
0.4
0.6
0.8
1.0

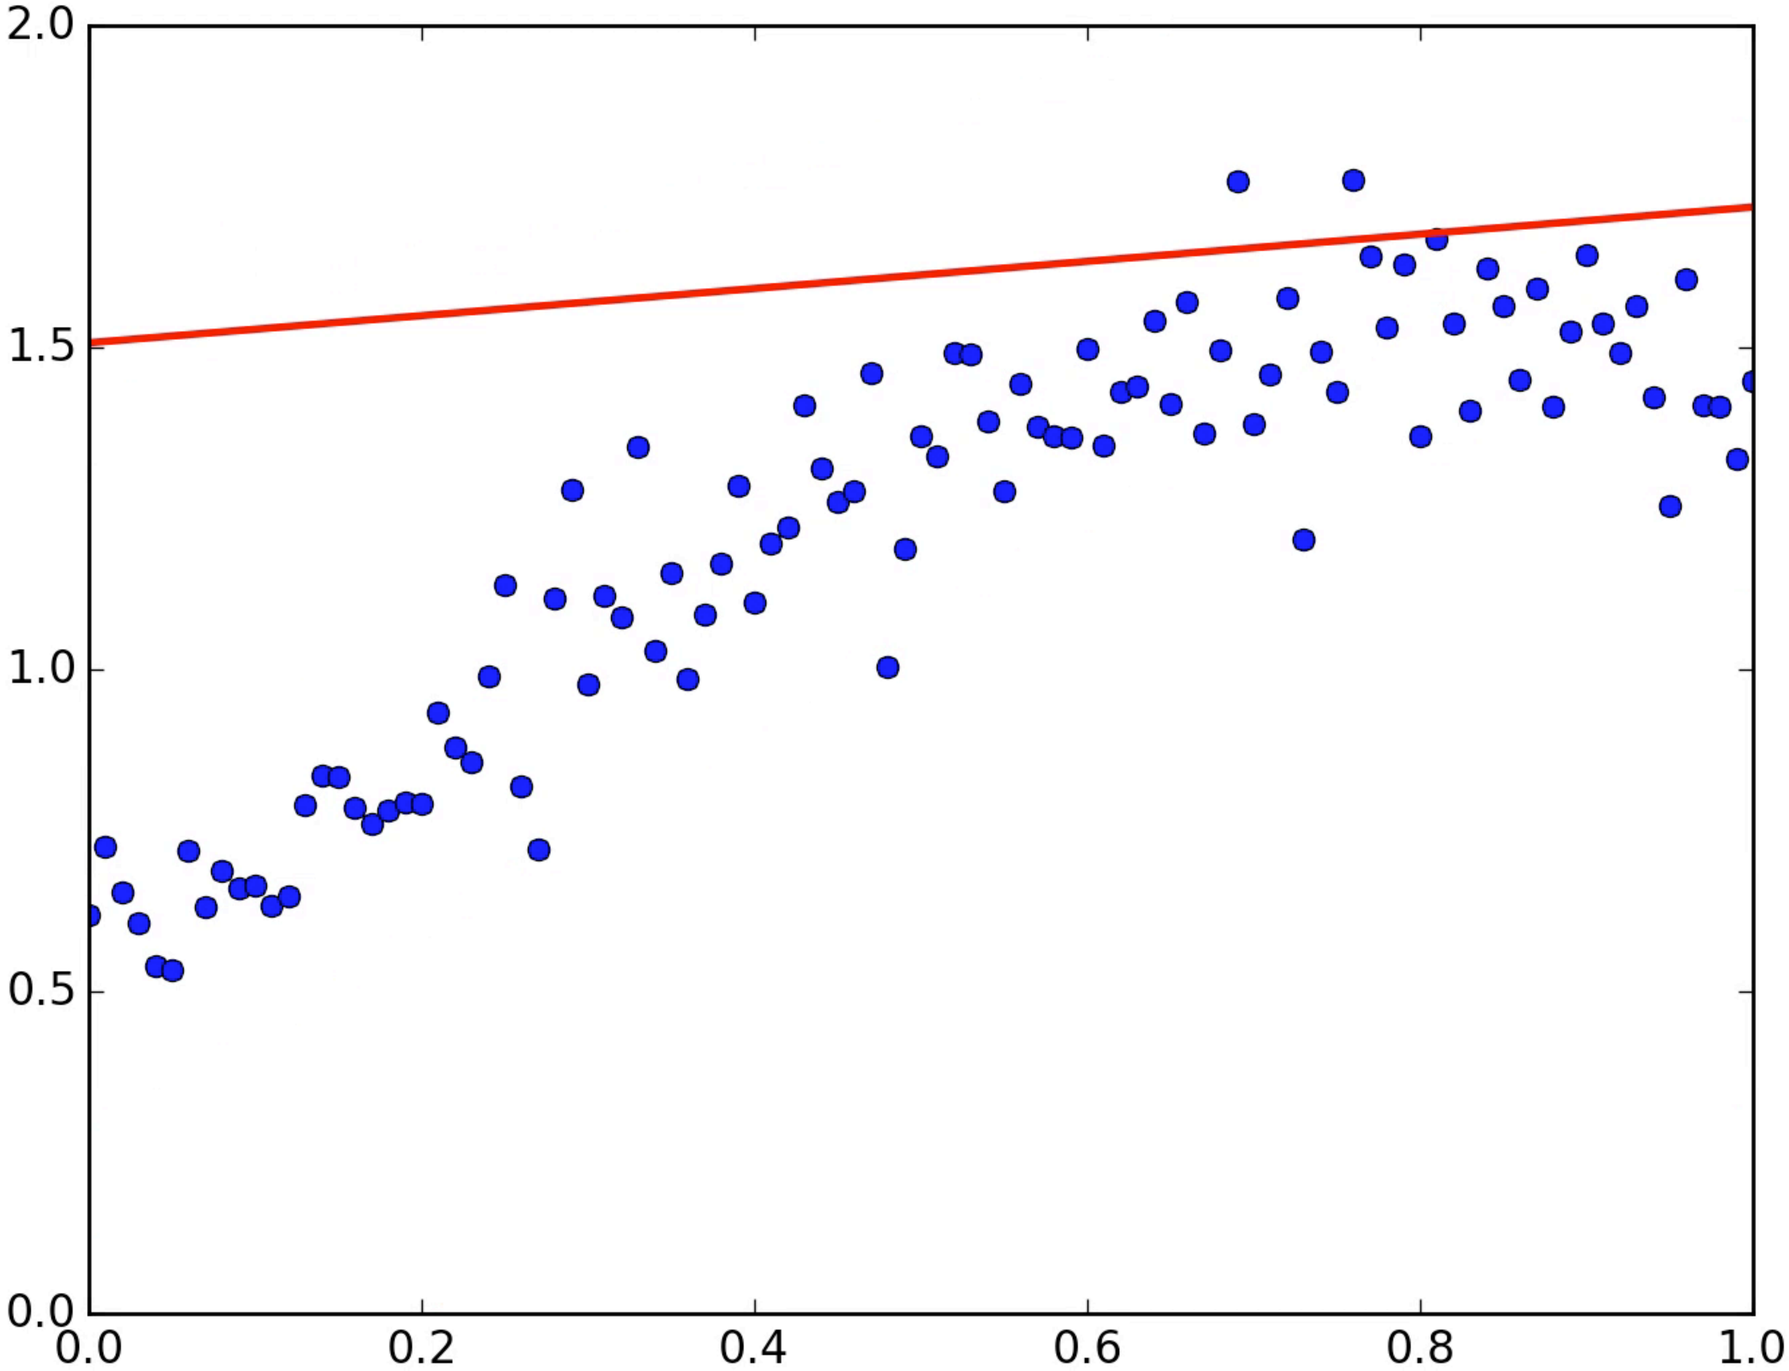
2.0
1.5
1.0
0.5
0.0
0.0
0.2
0.4
0.6
0.8
1.0

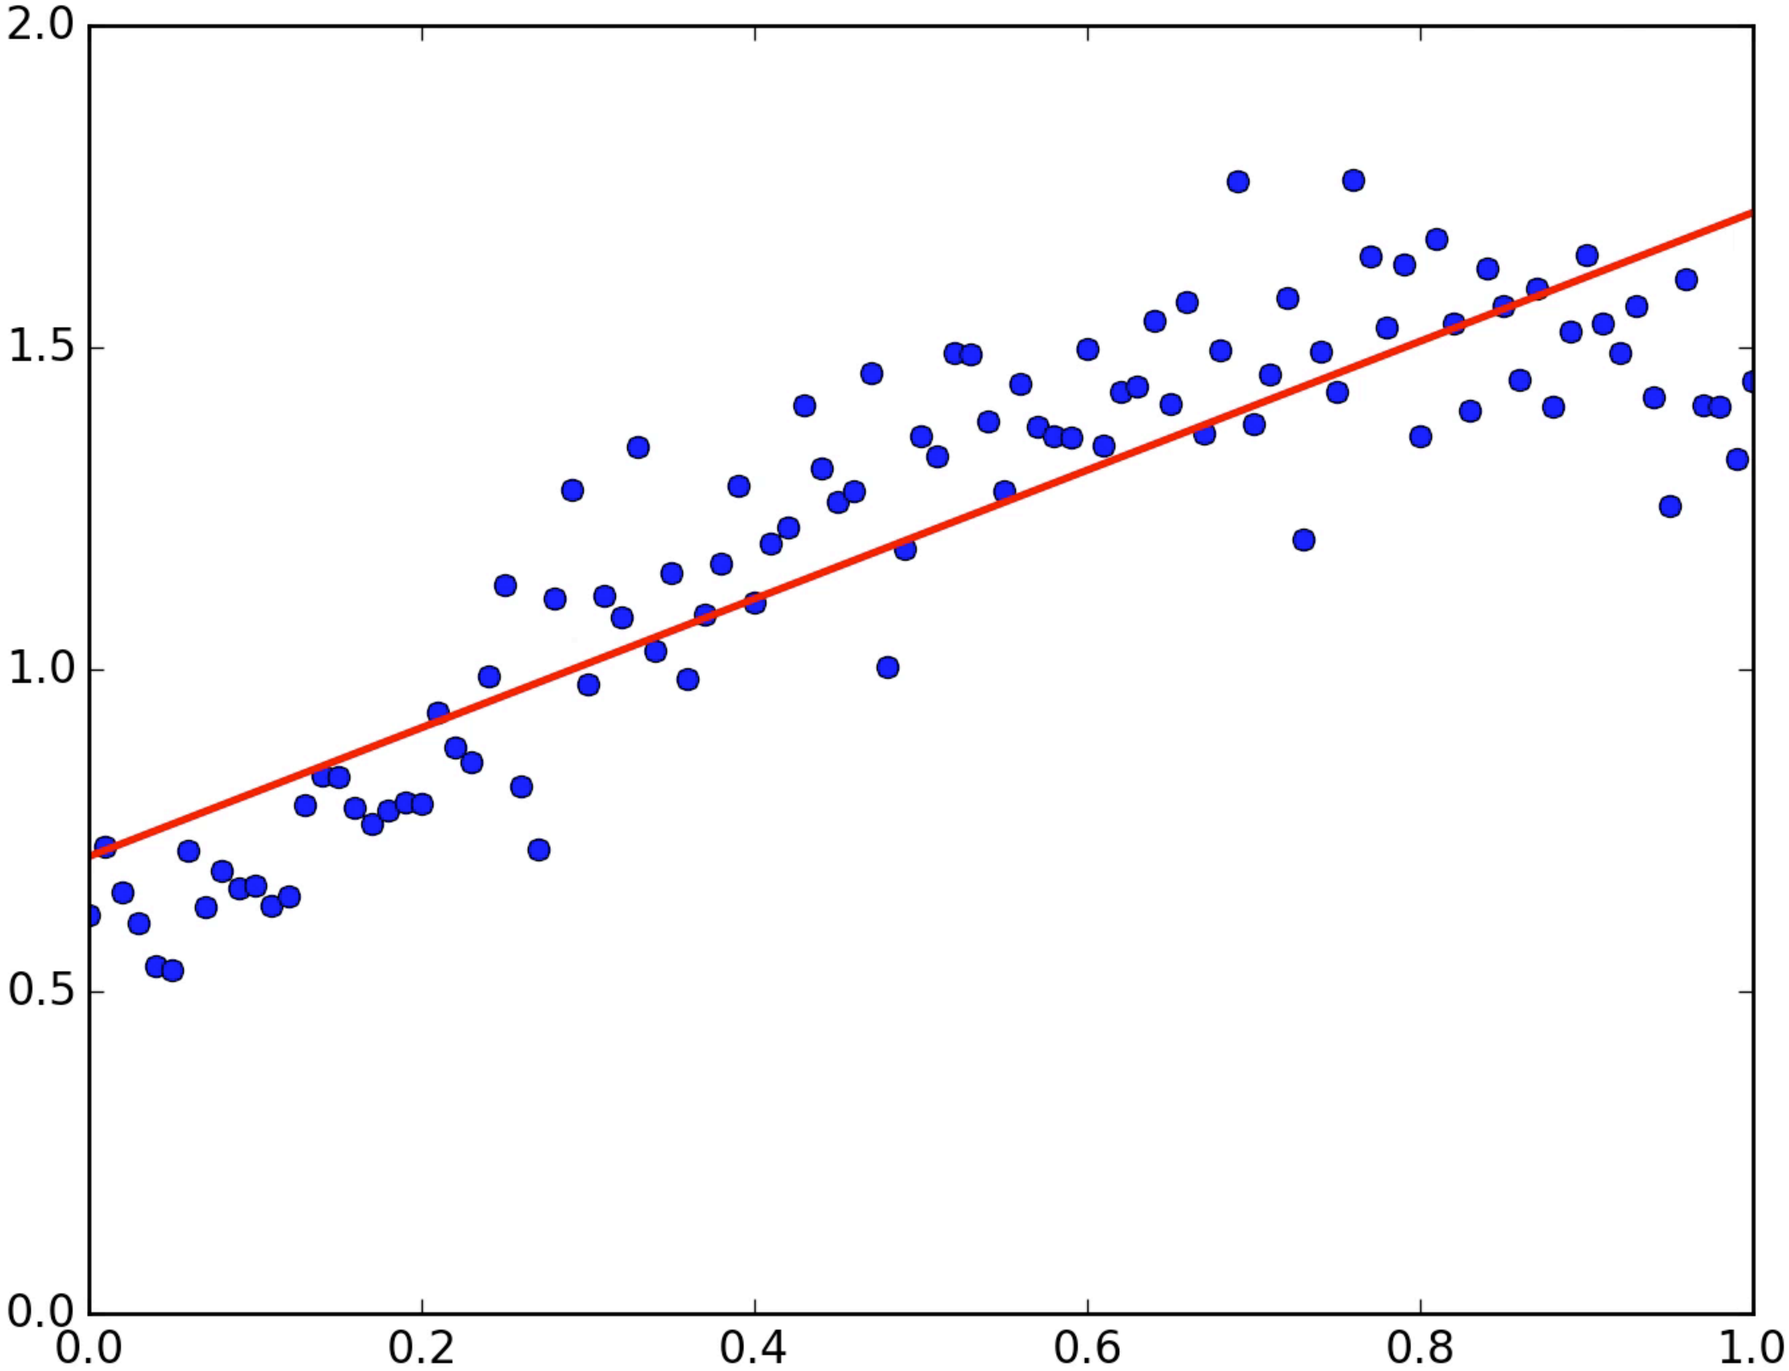
2.0
1.5
1.0
0.5
0.0
0.0
0.2
0.4
0.6
0.8
1.0

Regression Demo

Types of Learning

- Unsupervised learning
- Supervised learning
- Reinforcement learning

Supervised Learning

- Model $y = f(x; w)$
- Loss function $L\big(y^{\mathrm{GT}}, f(x; w)\big)$
- Data $\mathcal{D} = \big\{(x_1, y_1), \ldots, (x_n, y_n)\big\}$
- Cost function $R(w) = \frac{1}{N} \sum_i L\big(y_i, f(x_i; w)\big)$
- Learning $\hat{w} \approx \arg\min_w R(w)$

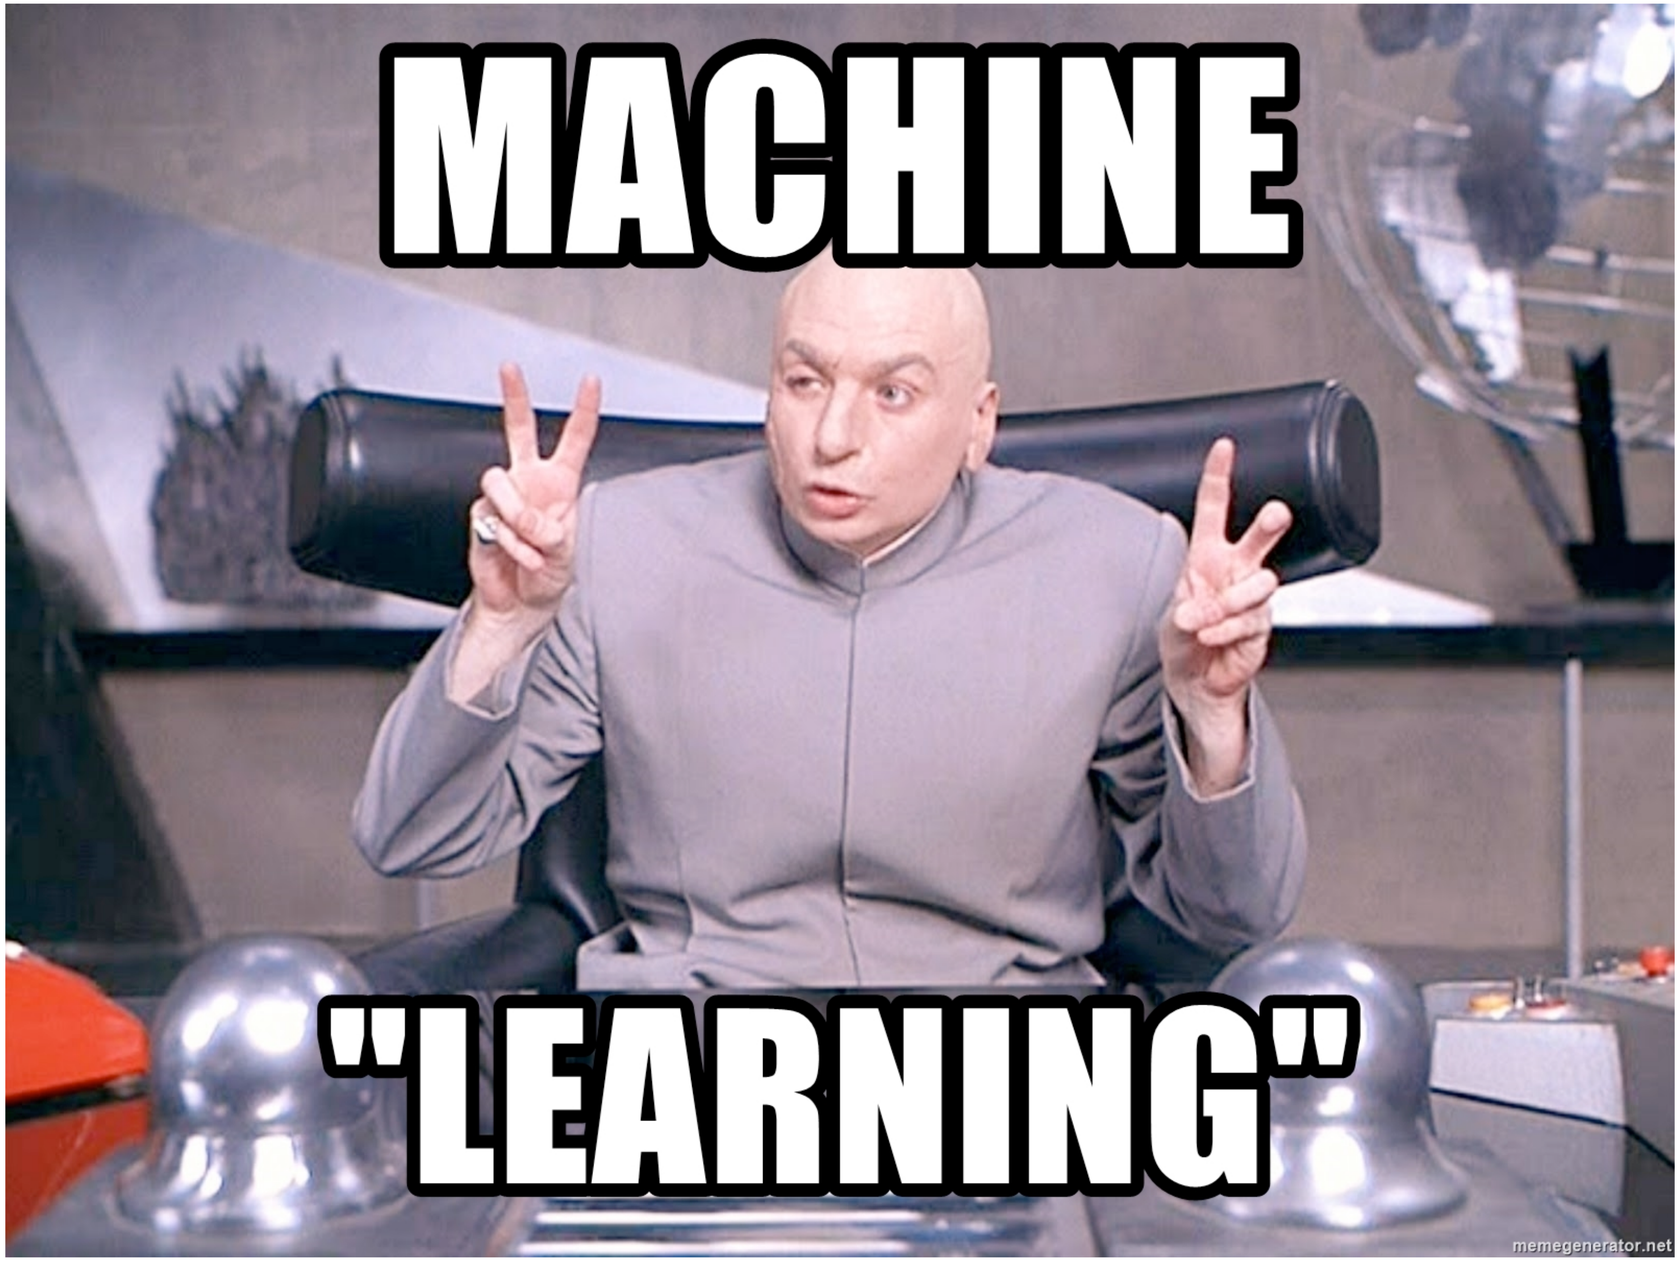
MACHINE
"LEARNING"
memegenerator.net

Supervised Learning

- Regression
  - continuous output, e.g. restoration
- Classification
  - discrete output, e.g. segmentation

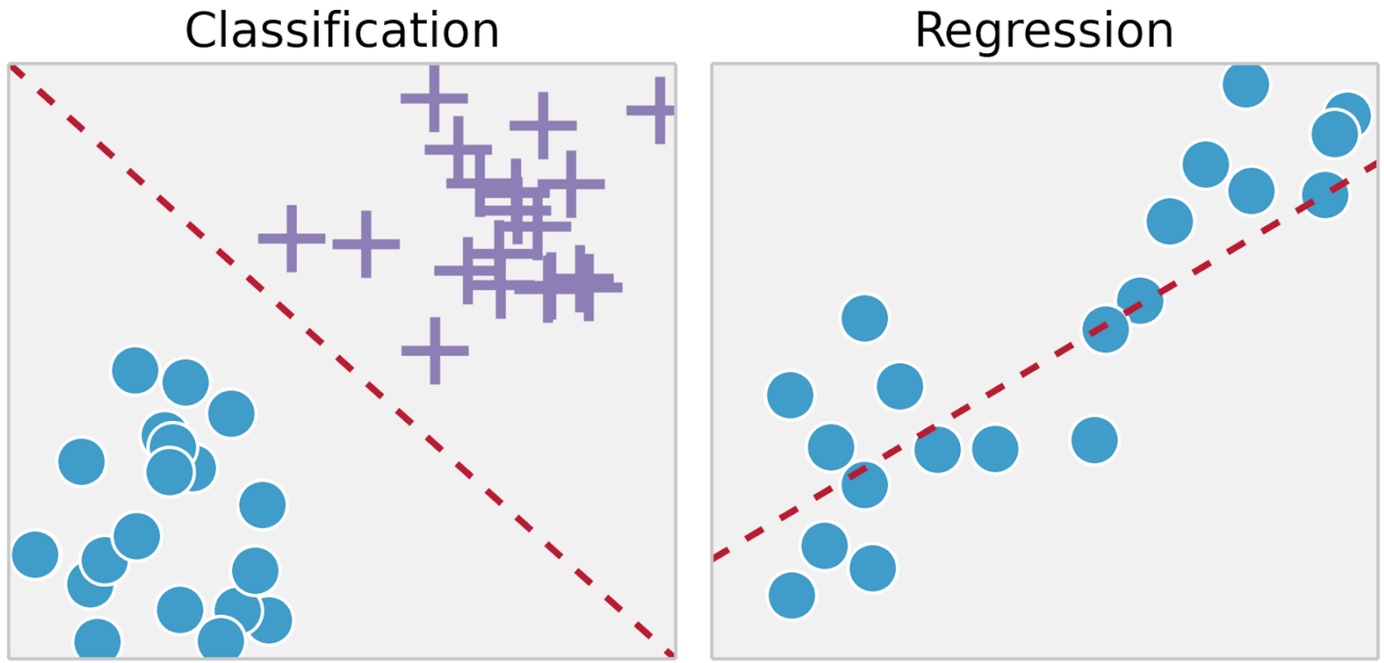

Why use (supervised) ML?

- Learned model can be used to predict output for novel inputs
- Learn predictive models that are difficult, time-consuming, or even impossible to do create manually
- Model is different for each user or changes frequently

Reasons to not use ML

- Simple model is good enough
  - easy to tune parameters
  - interpretable
- Complex learned model often hard to interpret
- Hard to collect necessary training data
- Training can be slow and annoying

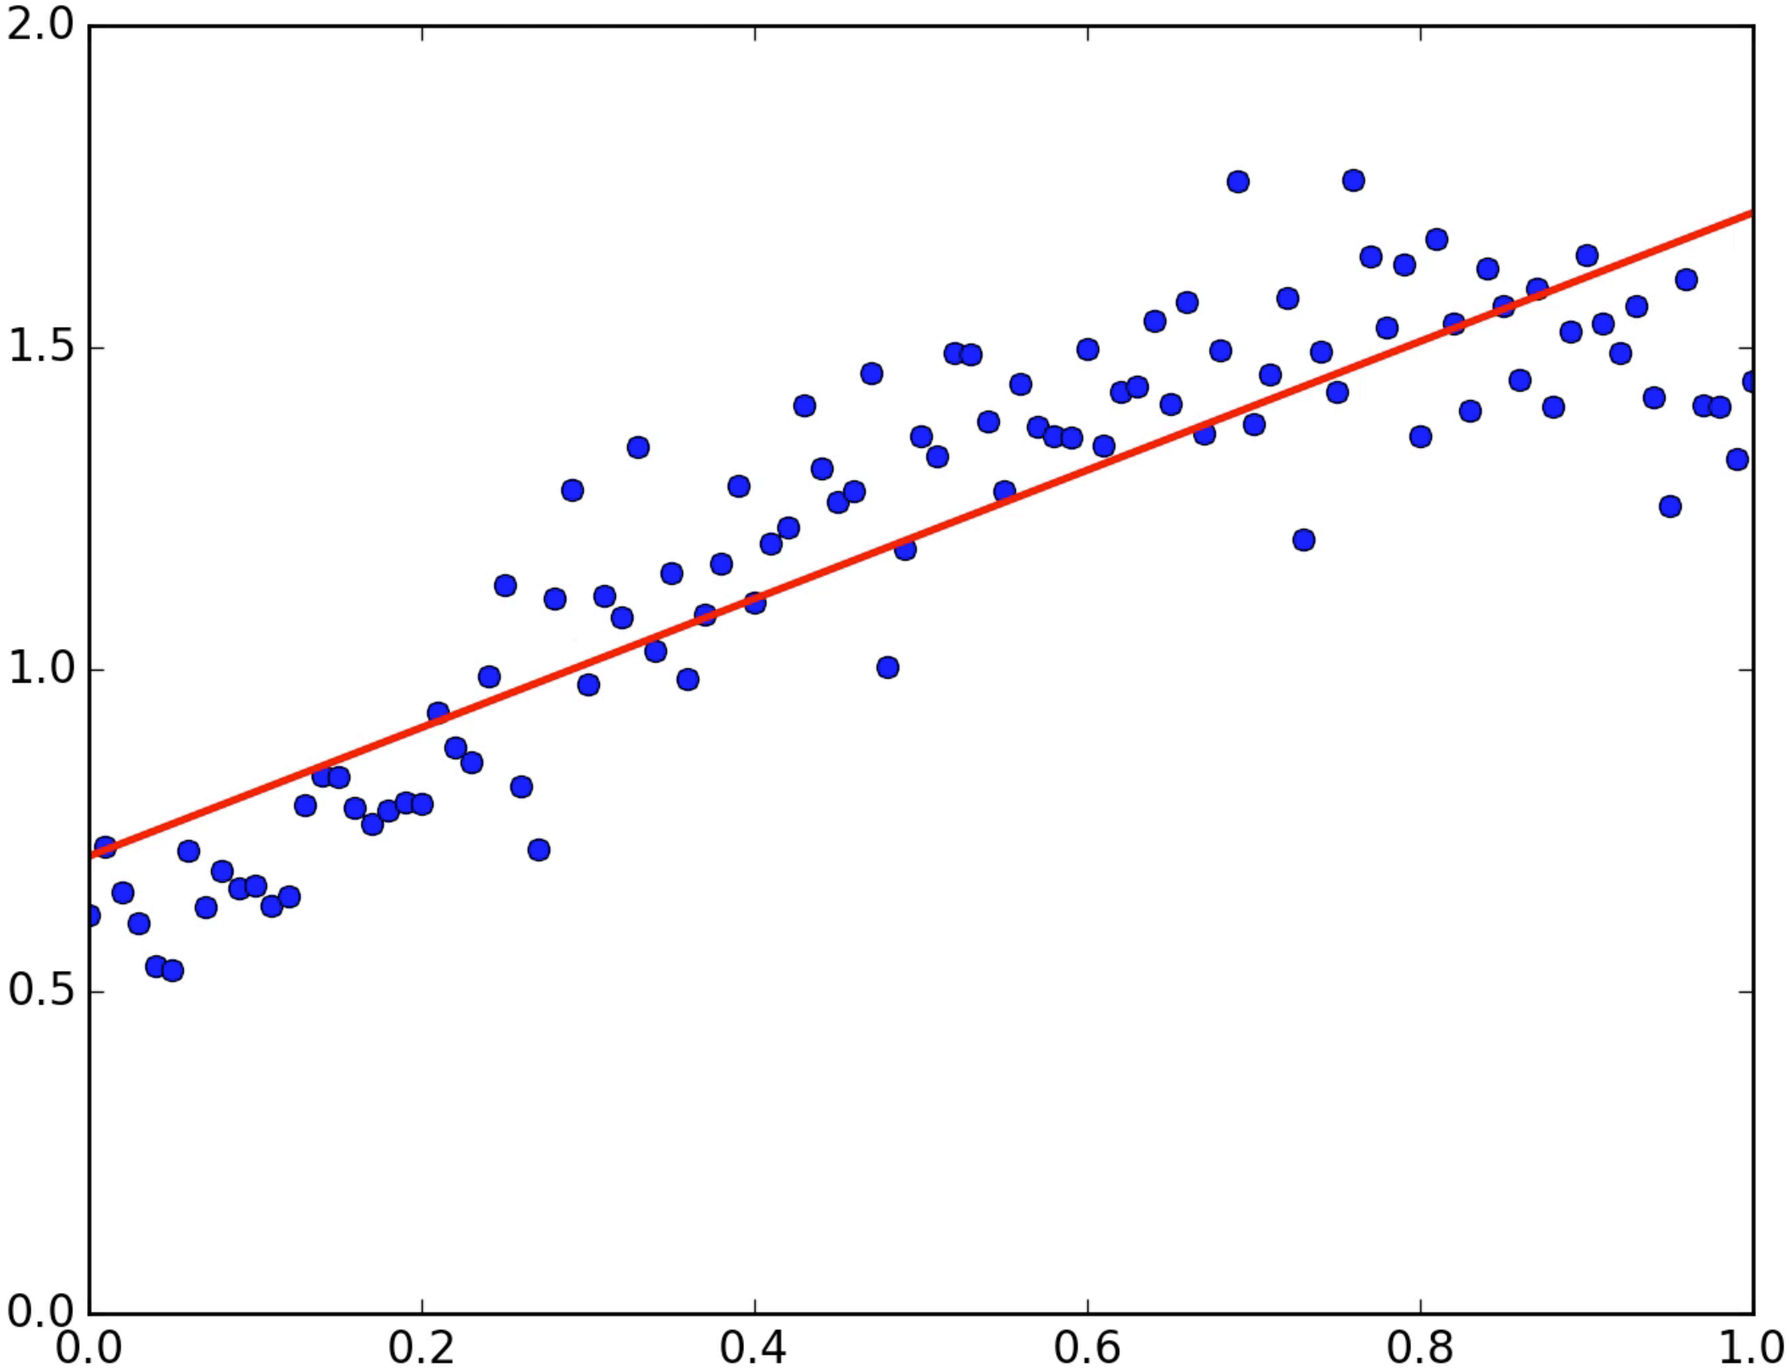
2.0
1.5
1.0
0.5
0.0
0.0
0.2
0.4
0.6
0.8
1.0

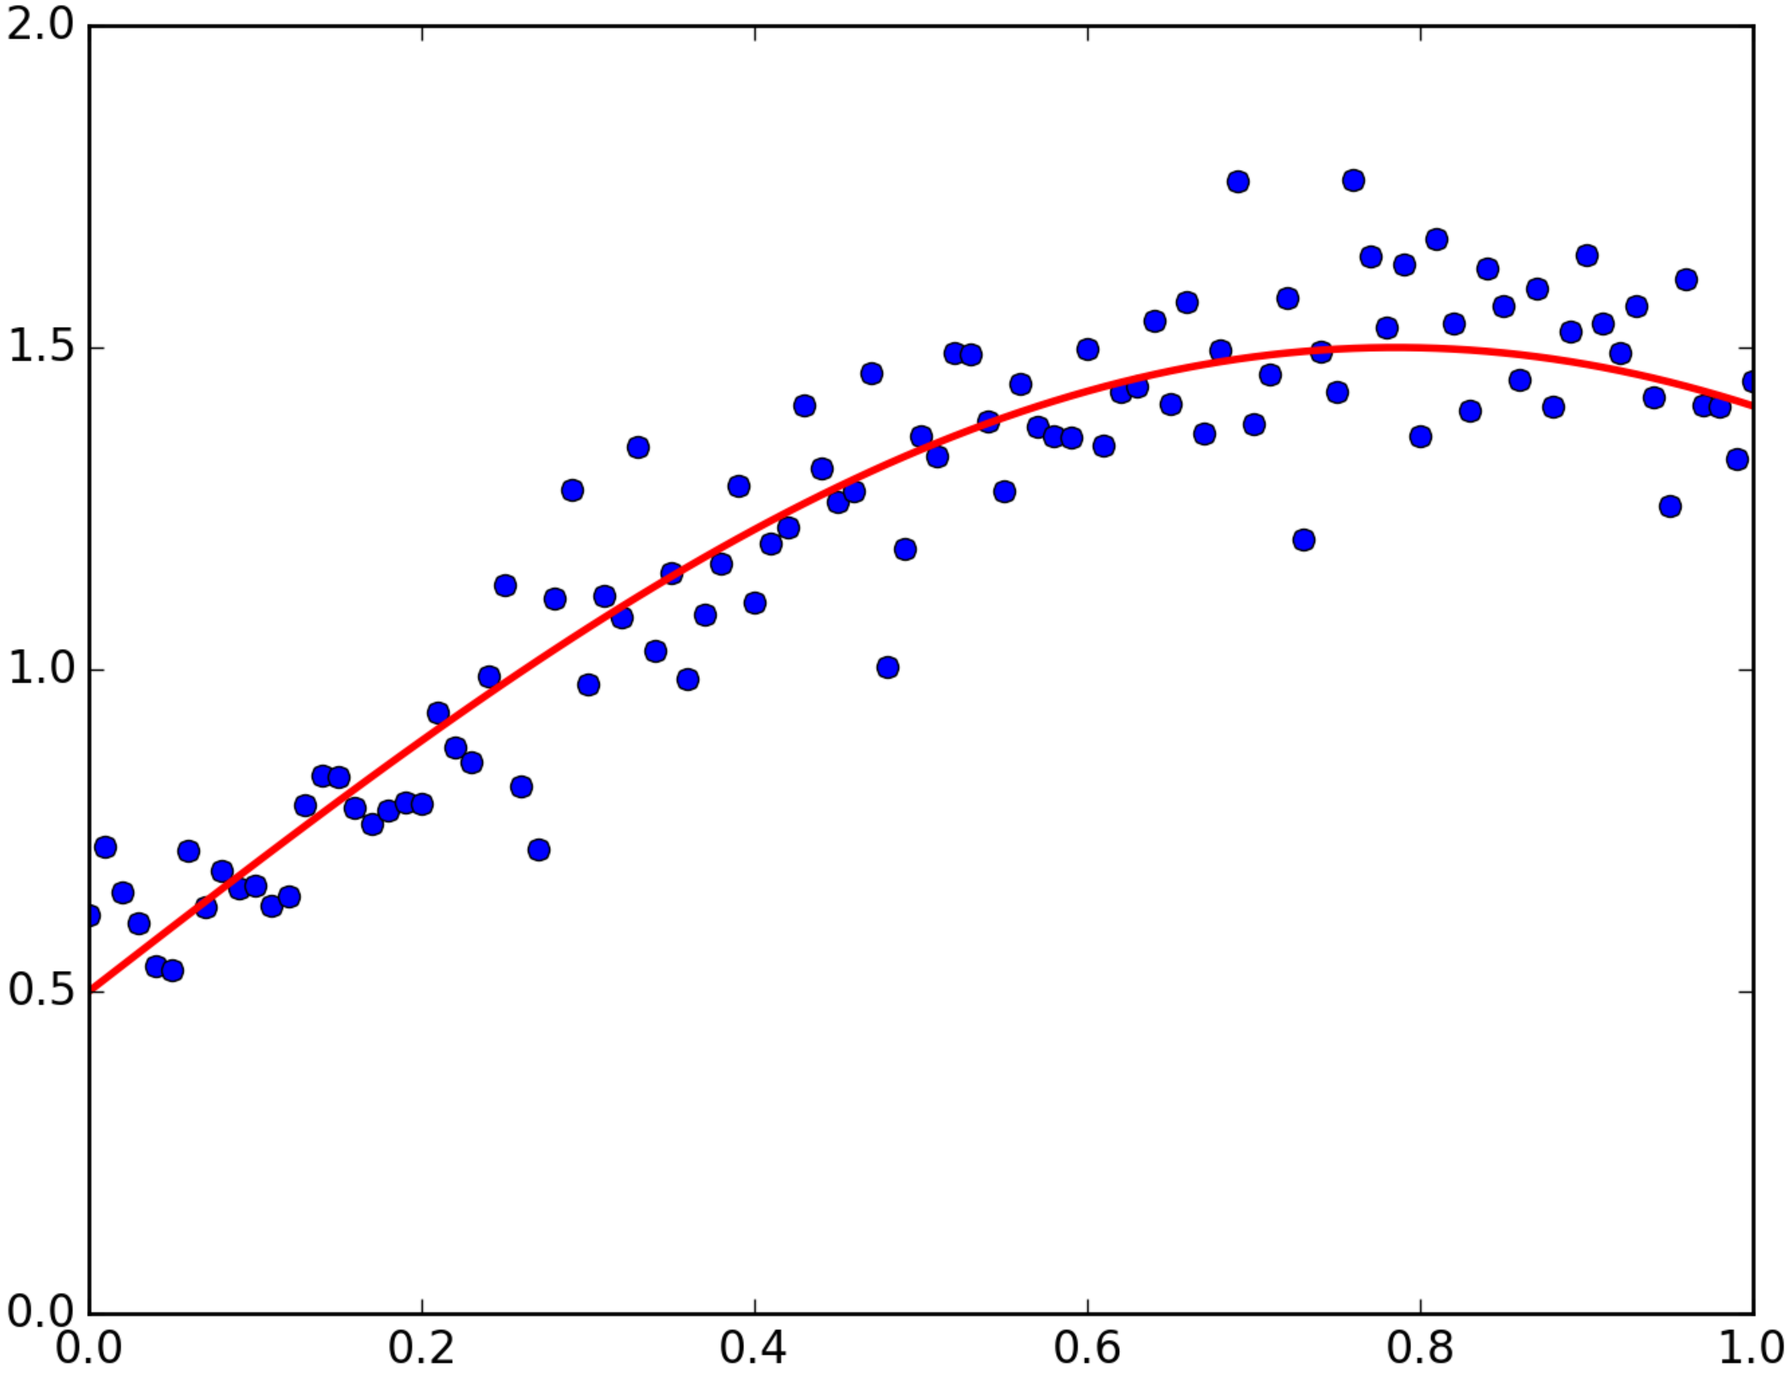
2.0
1.5
1.0
0.5
0.0
0.0
0.2
0.4
0.6
0.8
1.0

How to select a model (class)?

Acquisition

microscope

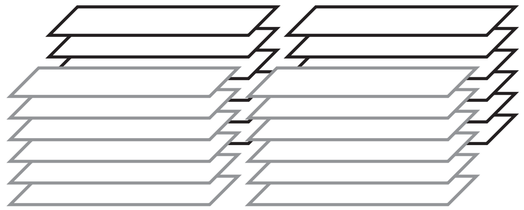

Restoration

fusion, denoising, deconvolution, etc.

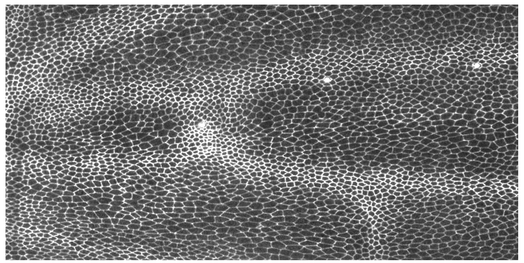

Segmentation

e.g. cells, membranes

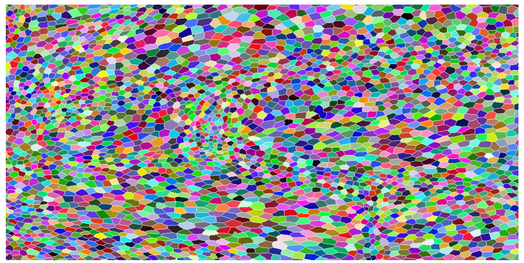

Tracking

e.g. cell lineages

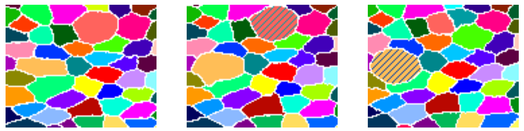

time

Images: Corinna Blasse

Use models to explain observations

Acquisition

microscope

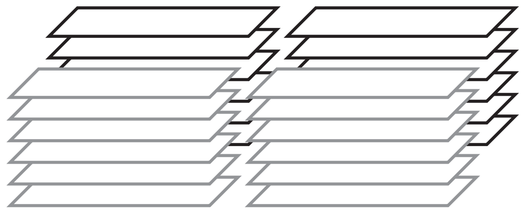

Restoration

fusion, denoising, deconvolution, etc.

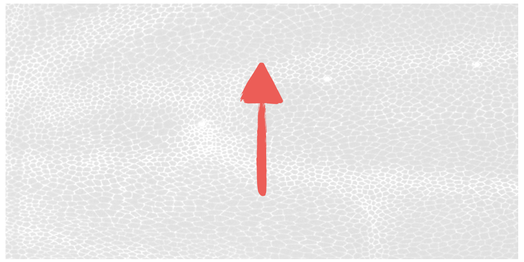

Segmentation

e.g. cells, membranes

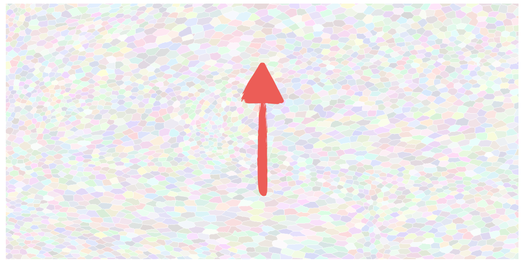

Tracking

e.g. cell lineages

Temporal Model

dynamics, processes (e.g. cell divisions)

Shape & Appearance

e.g. circles, intensities

Image Formation

noise, PSF, image prior, view geometry, etc.

Simulation: generate observations from models

Images: Corinna Blasse

How to select a model (class)?

- Model class known (e.g., from physics)
  - can be mathematically inconvenient to use
  - unfortunately, reality often more complicated
- Generic model class (that can approximate (m)any tasks/functions)
  - e.g. Support Vector Machines, Decision trees (random forests)
  - 🔥 Neural networks (**deep learning**)

Decision Tree

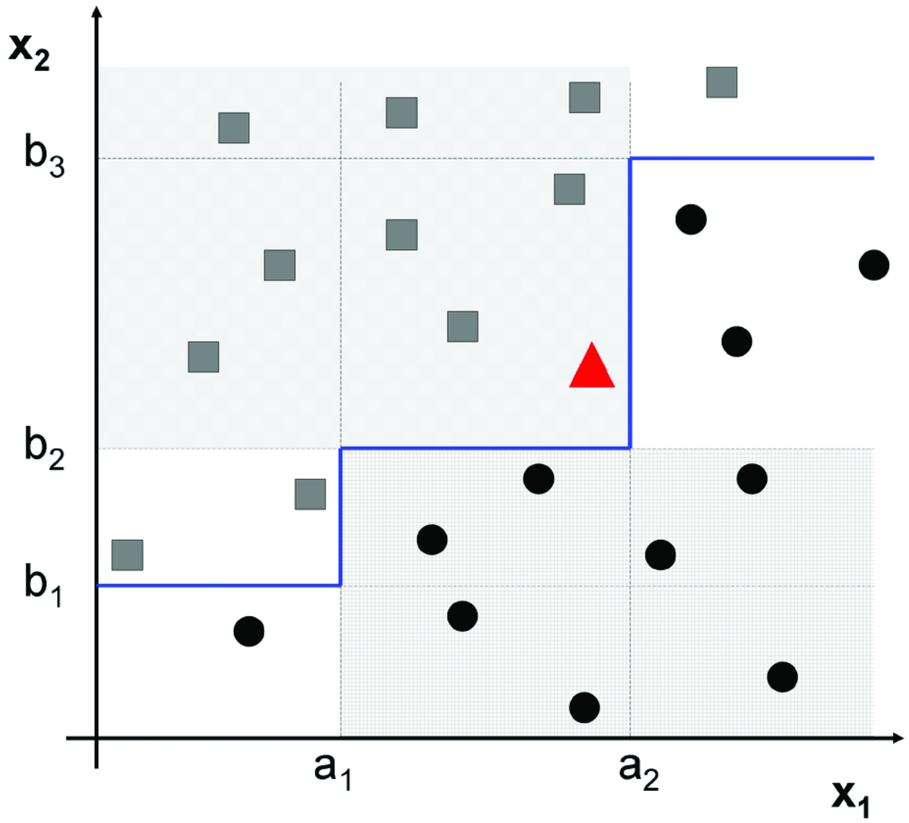

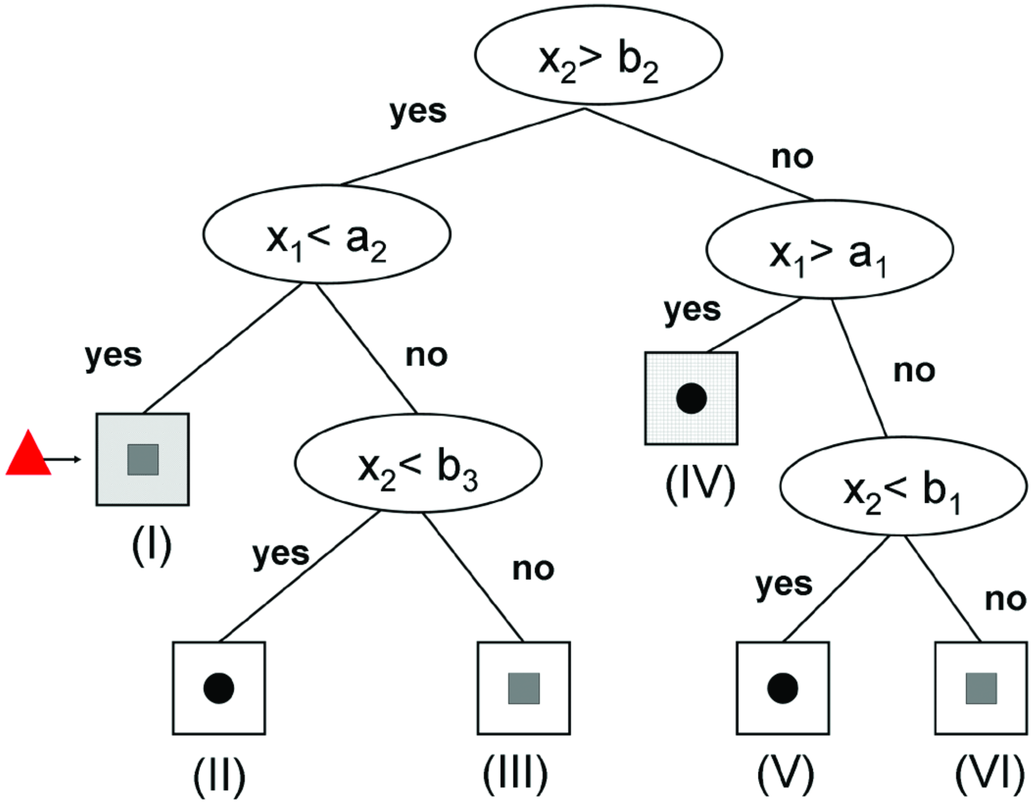

Neural Network

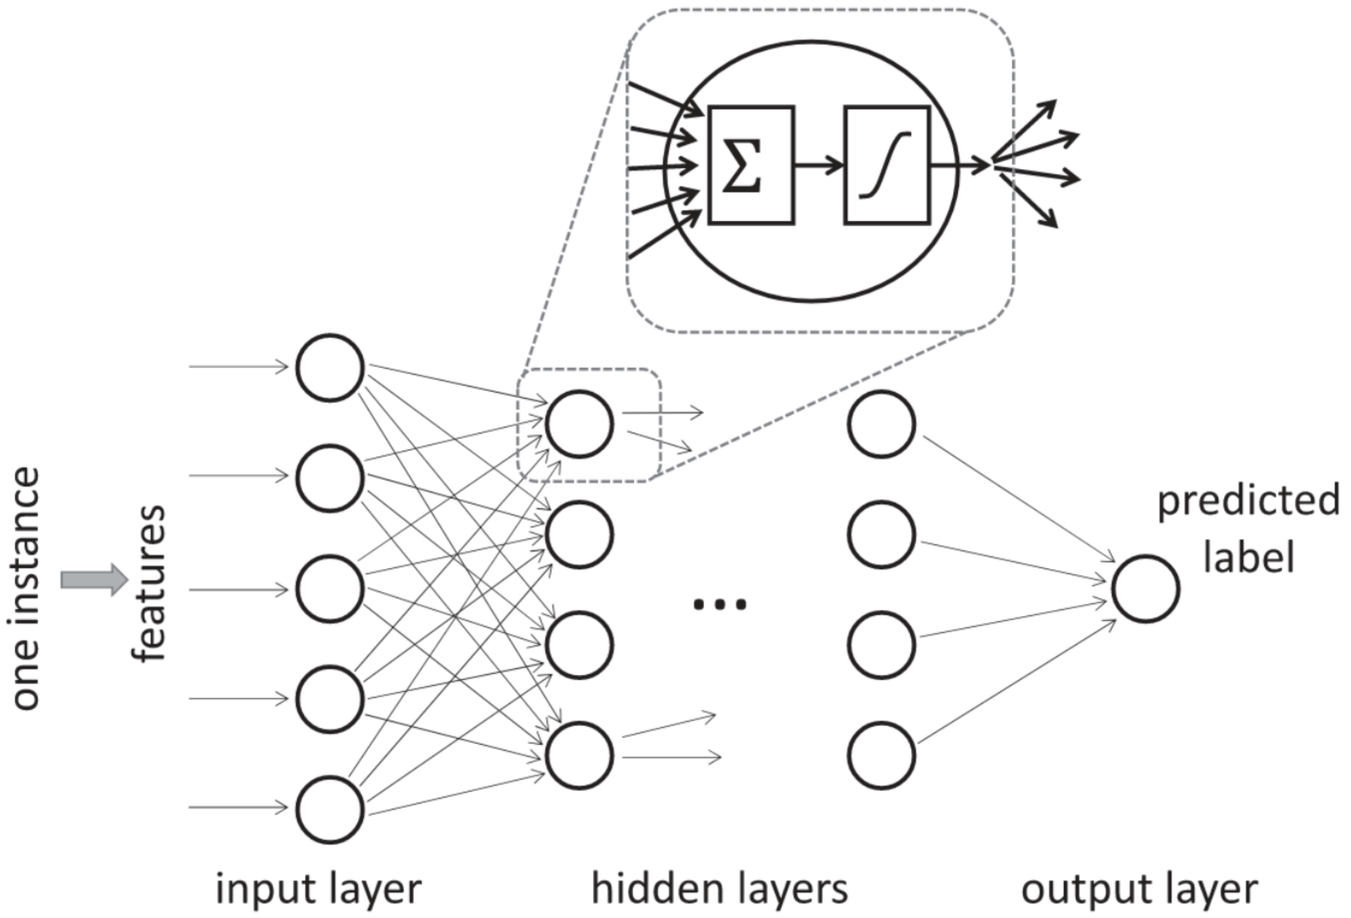

scikit-learn.org

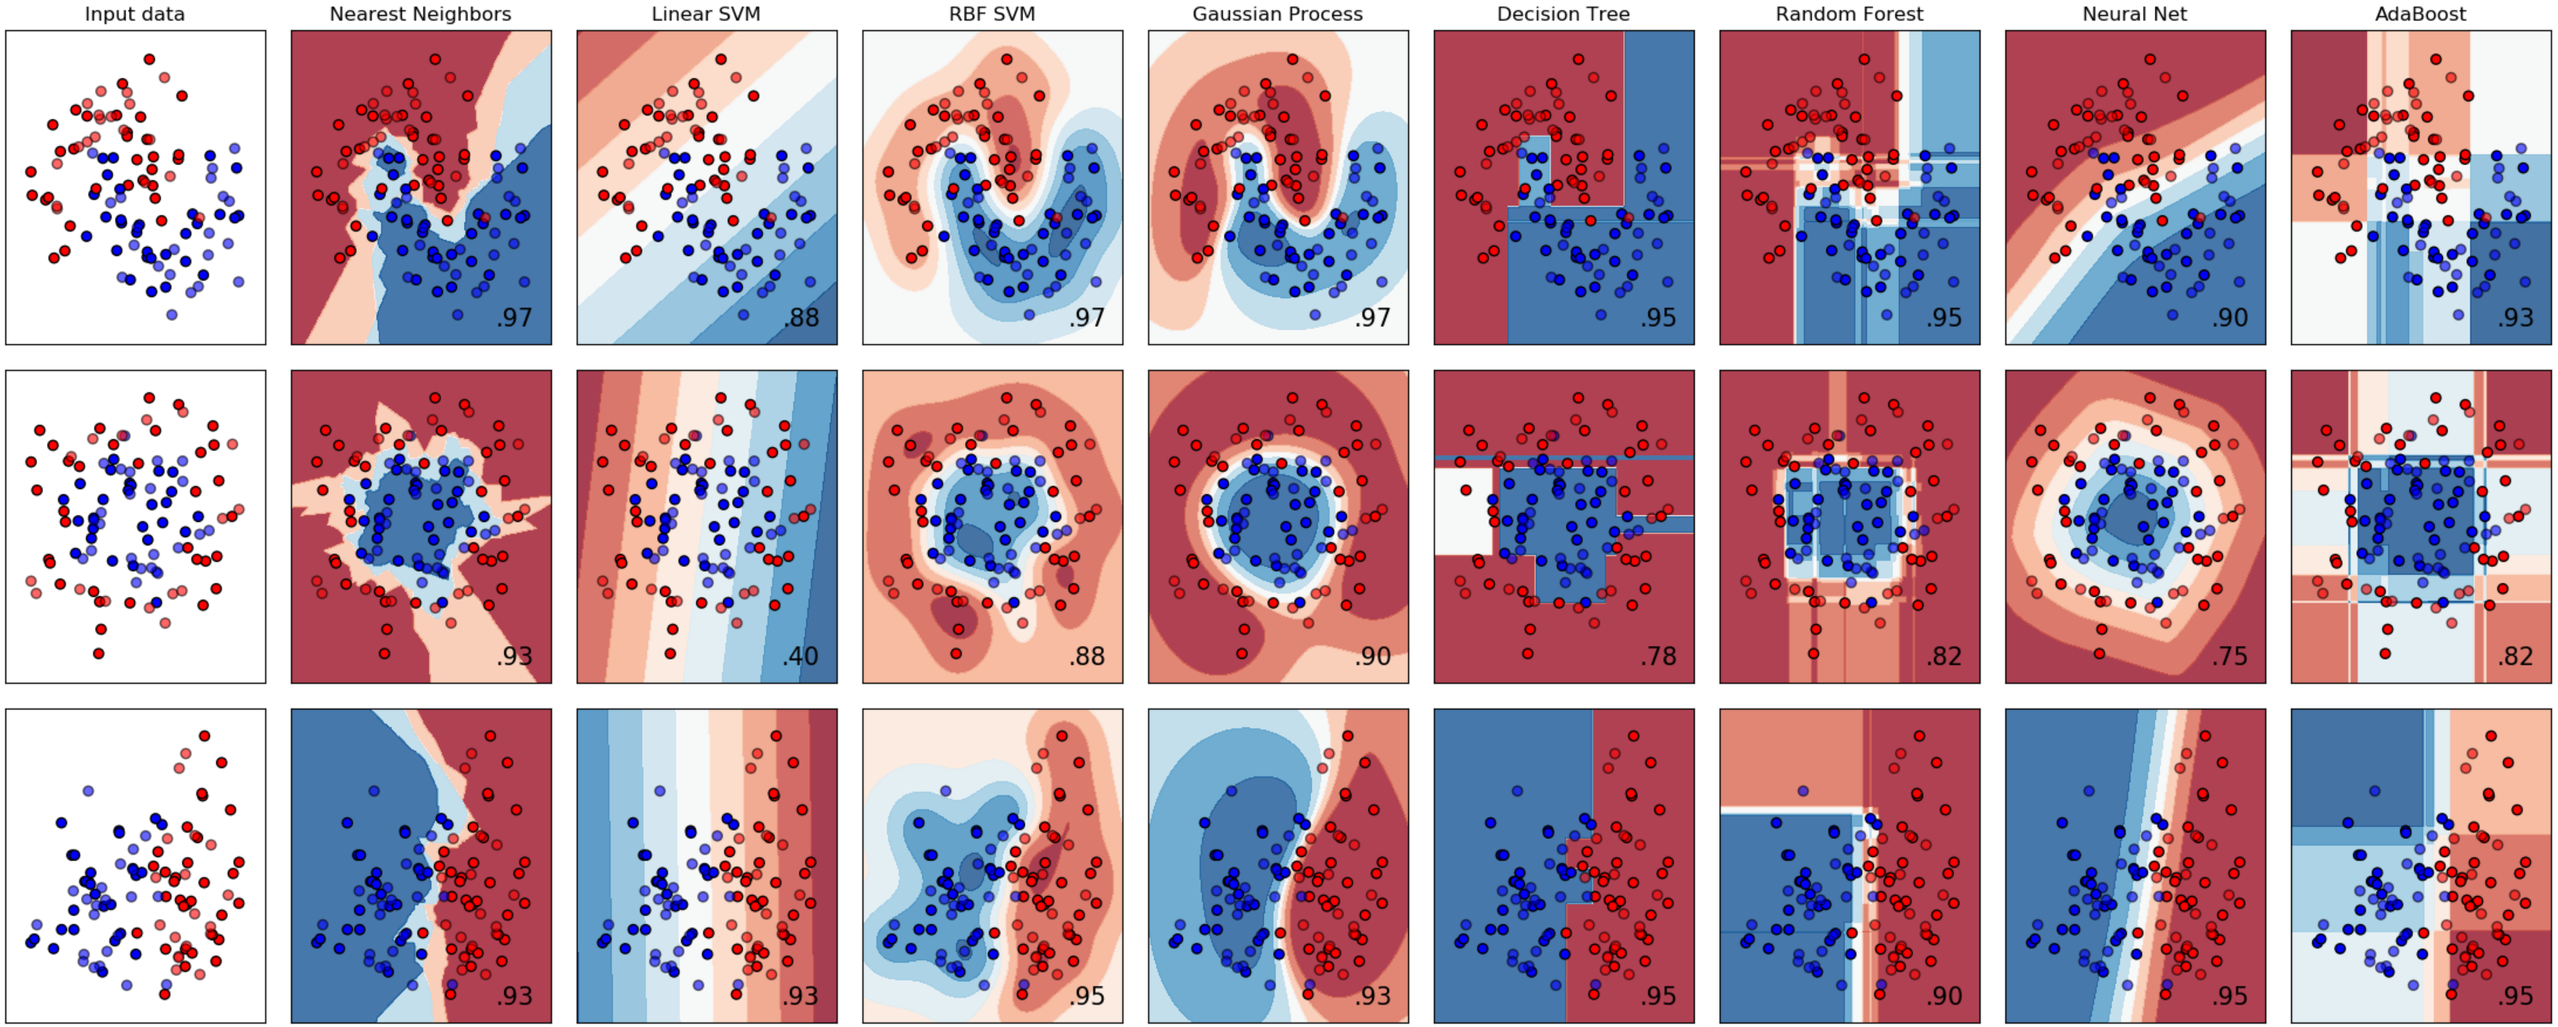

http://scikit-learn.org/stable/auto_examples/classification/plot_classifier_comparison.html

playground.tensorflow.org

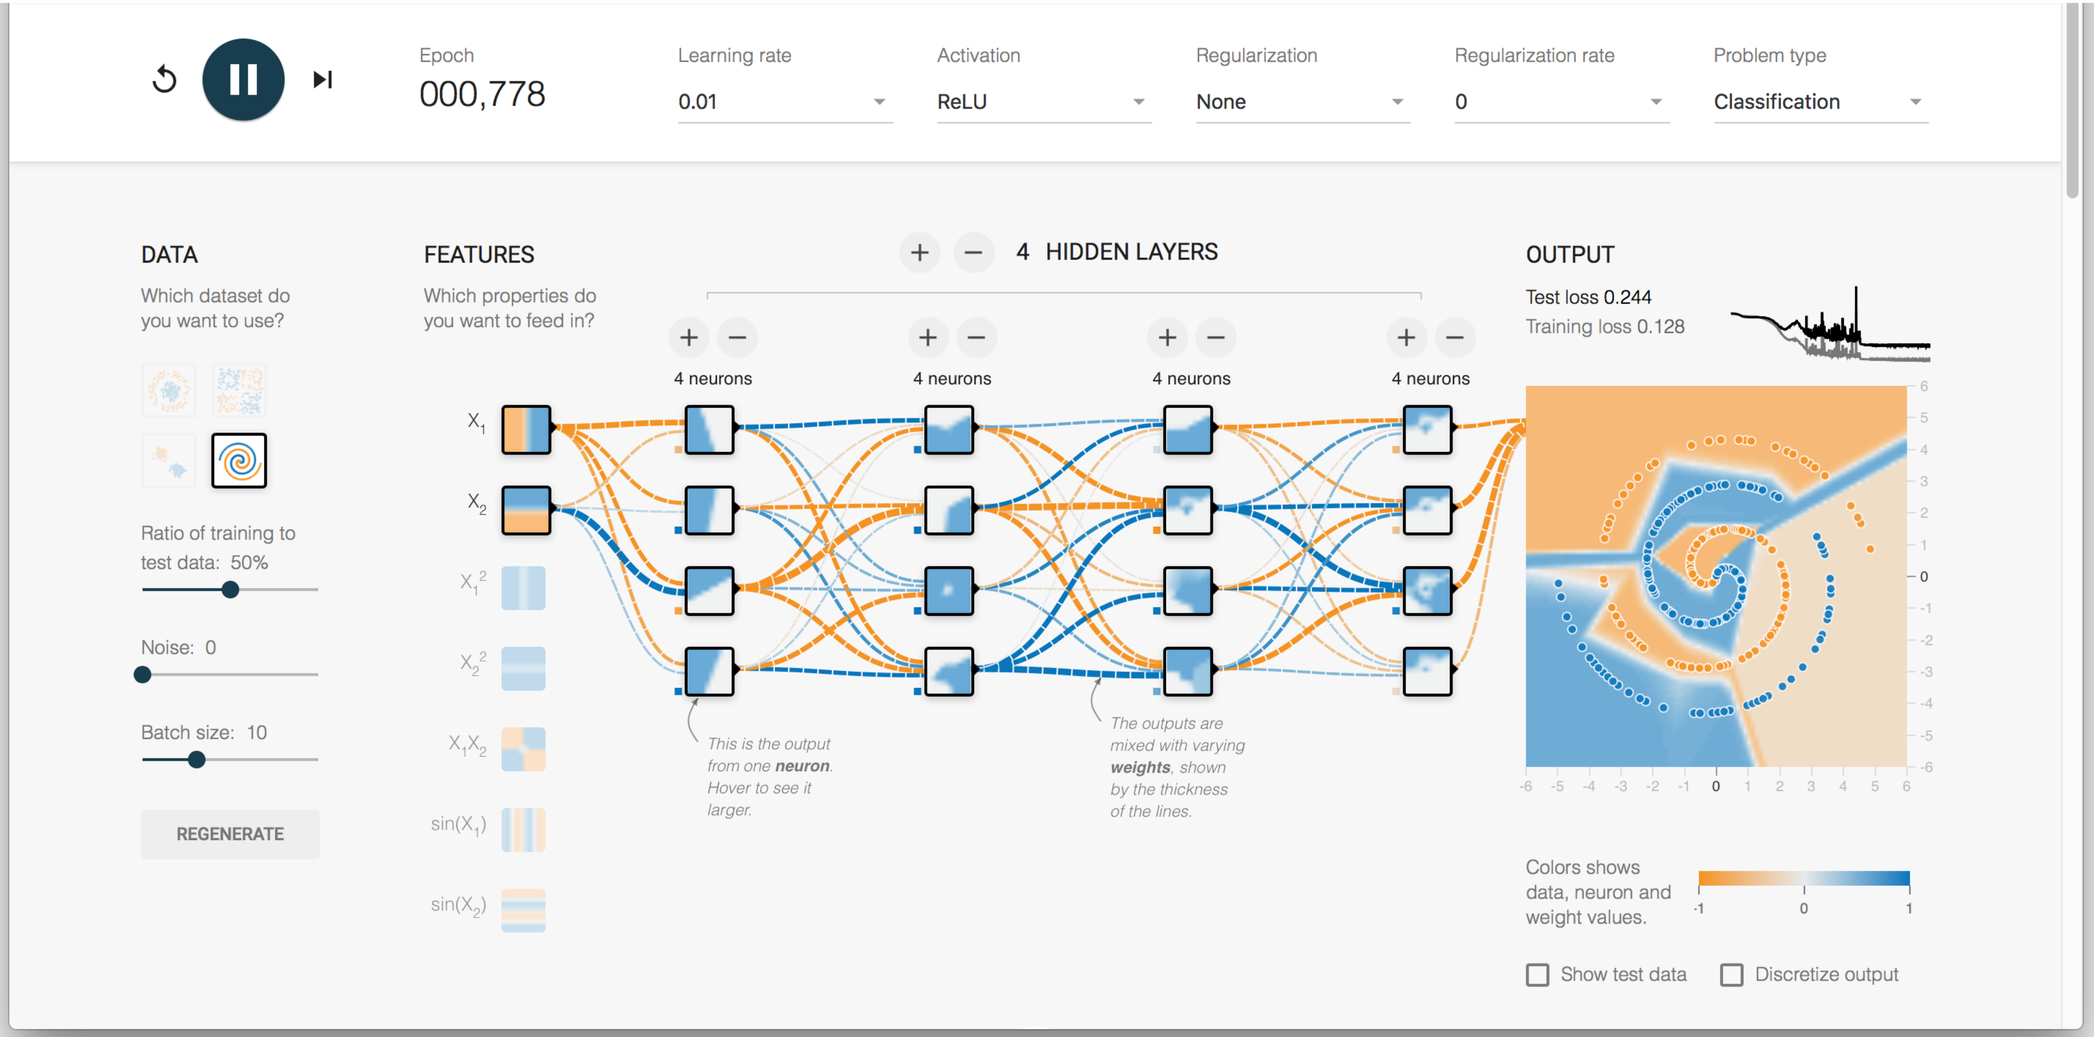

Deep Learning

Merge

Upsampling

Activation

Pooling

Convolution

Synced Follow
AI Technology & Industry Review - www.syncedreview.com || www.jiqizhixin.com || Subscribe: http://goo.gl/Q4cP3B
Dec 13, 2017 · 4 min read

LeCun vs Rahimi: Has Machine Learning Become Alchemy?

The medieval art of alchemy was once believed capable of creating gold and even human immortality. The trial-and-error method was however gradually abandoned after pioneers like Issac Newton introduced the science of physics and chemistry in the 1700s. But now, some machine learning researchers are wondering aloud whether today's artificial intelligence research has become a new sort of alchemy.

The debate started with Google's Ali Rahimi, winner of the the Test-of-Time

International Conference on Machine Learning (ICML) 2018

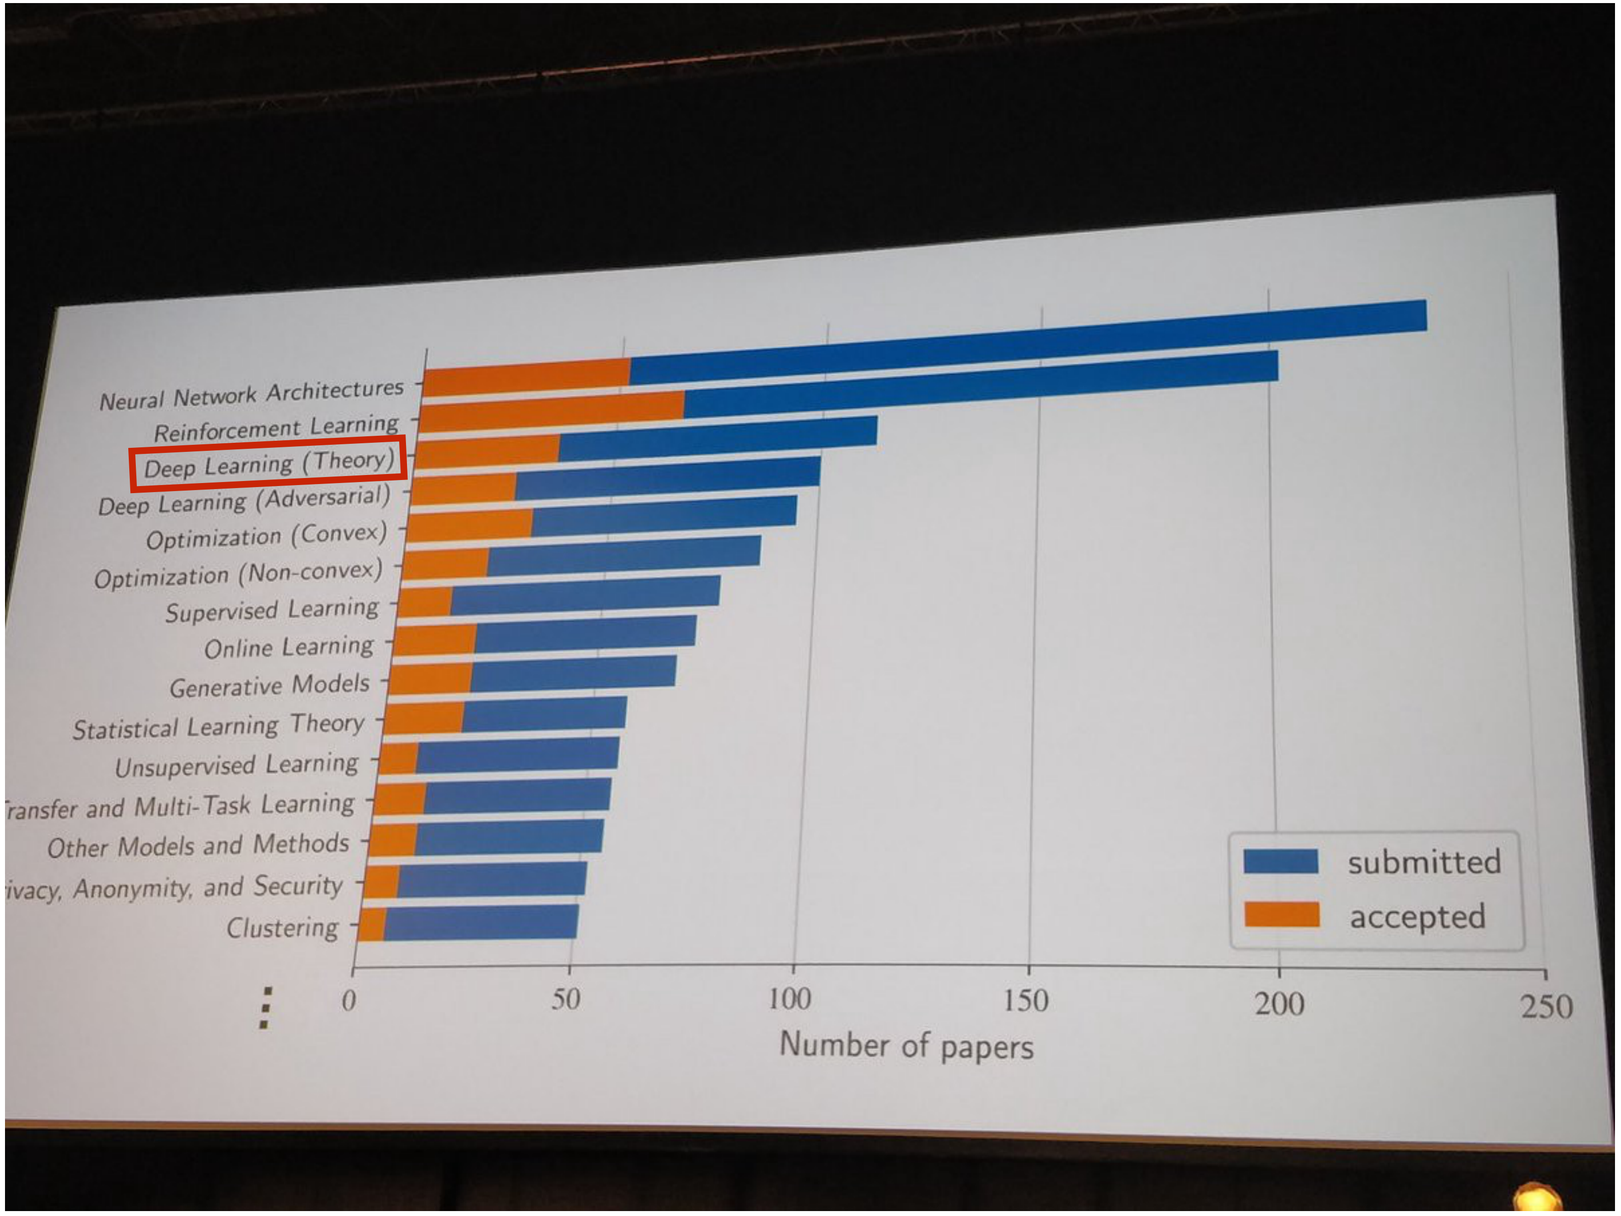

ML/DL Status Quo

- Libraries and frameworks quite mature, require little programming
  - e.g. *Scikit-Learn*, *Tensorflow*, *Keras*, *PyTorch*, etc.
- Problem formulation and data acquisition remain challenging
  - choosing a good model (class) for your problem
    - model debugging requires experience
  - acquiring good data in sufficient quantities
- Early days: theory of how and why deep learning works so well