Machine Learning

CMSC 422- Project Discussion

Soheil Feizi

University of Maryland
Department of Computer Science

Supervised Learning

Optimization Landscape of Deep Learning

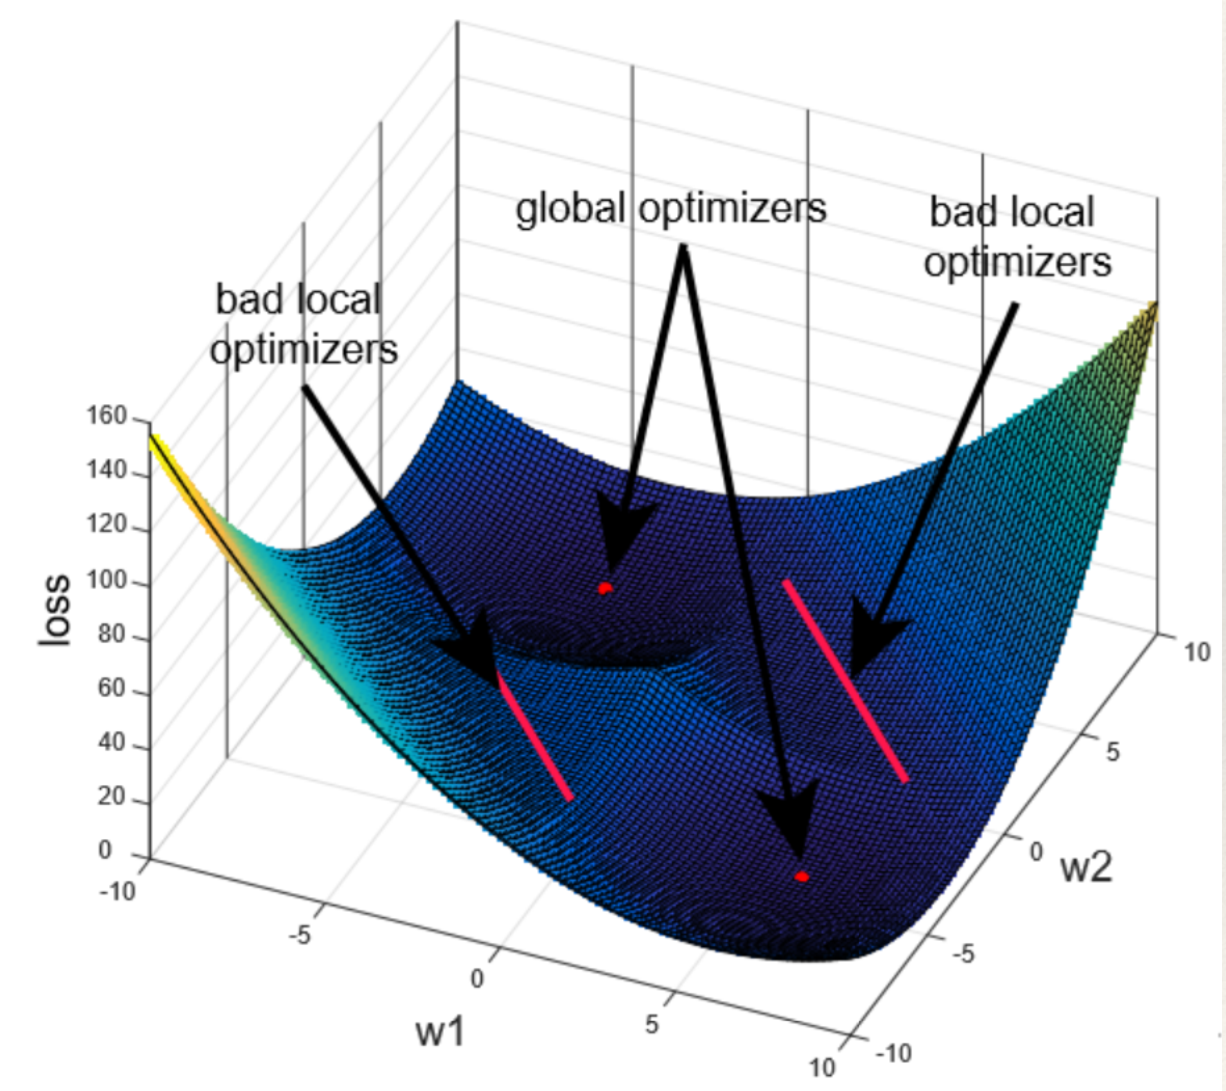

Supervised Learning

Generalization in Deep Learning (# required samples for training)

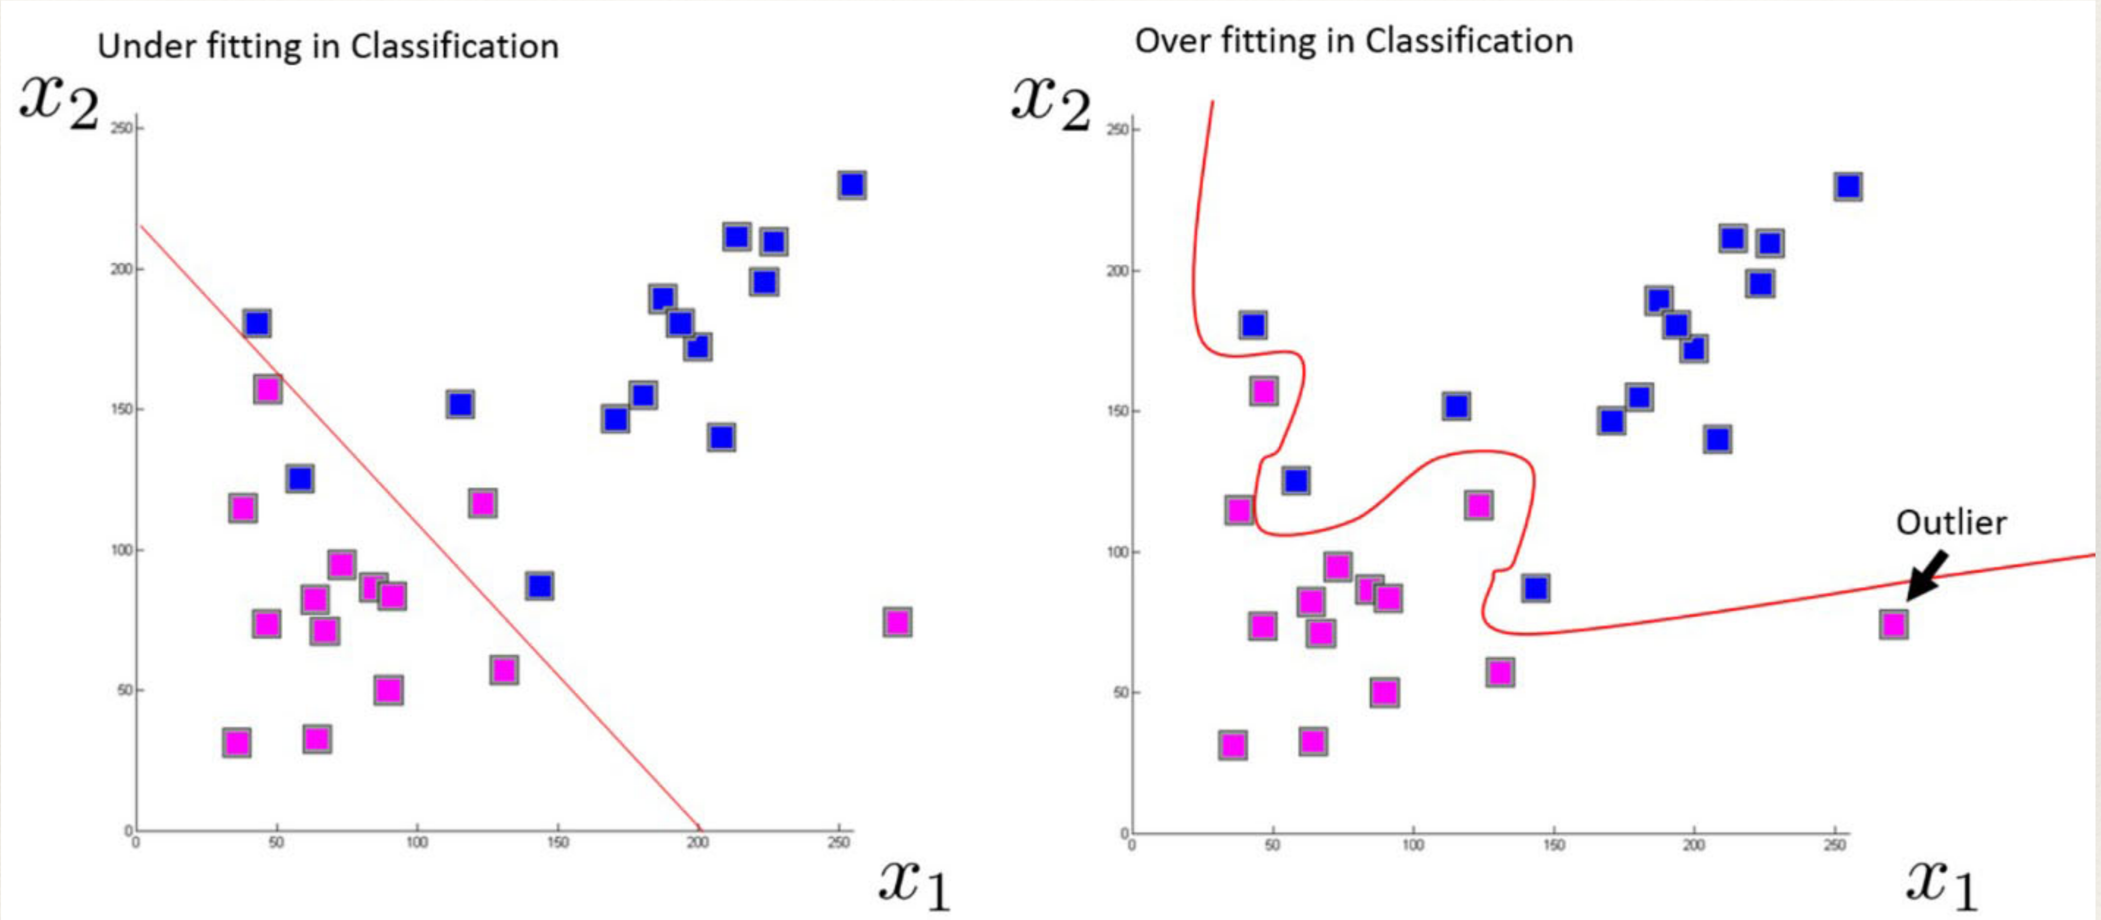

Supervised Learning

Effect of depth in deep learning

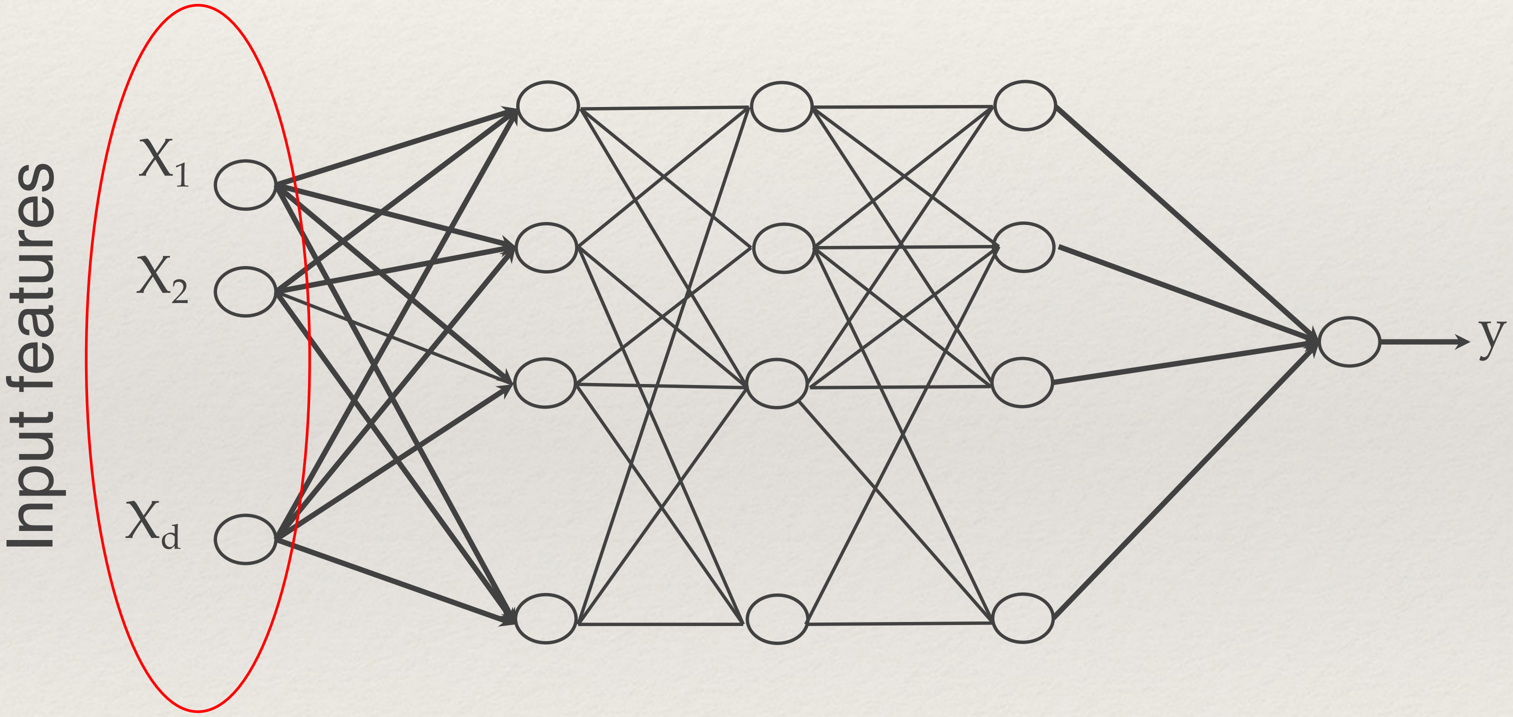

Supervised Learning

Adversarial Examples

data

adversarial noise

noisy data

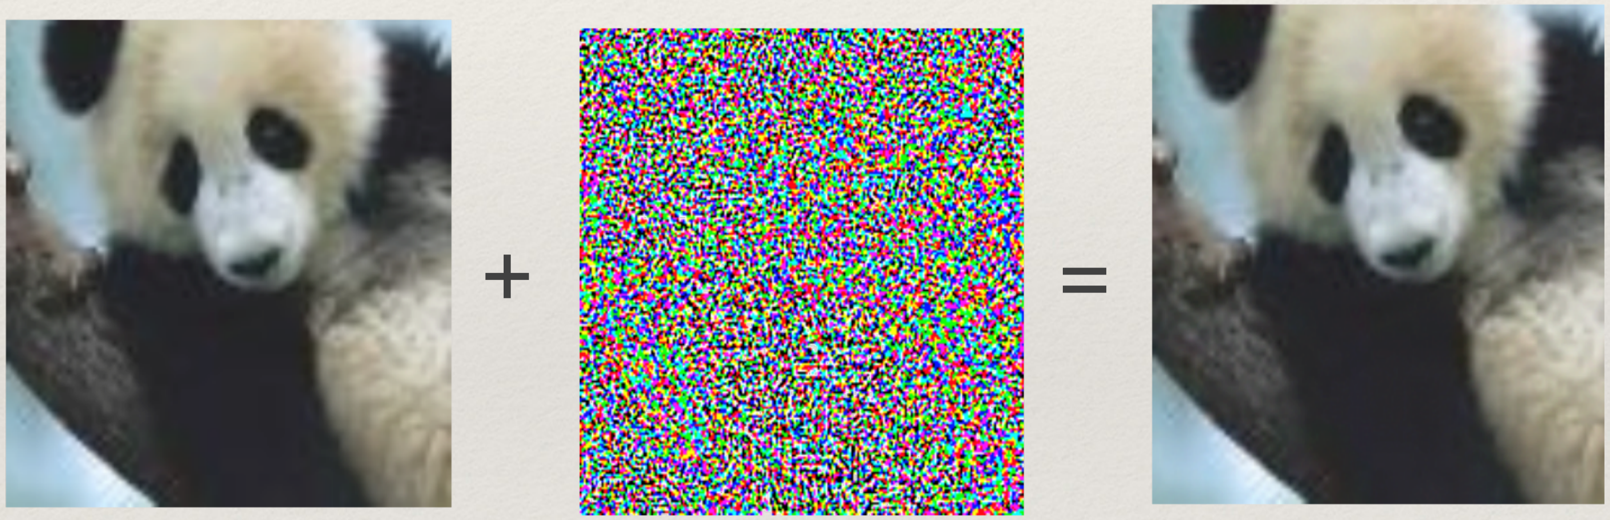

"panda"
57.7% confidence

"gibbon"
99.3 % confidence

Supervised Learning

Interpretability (features and samples)

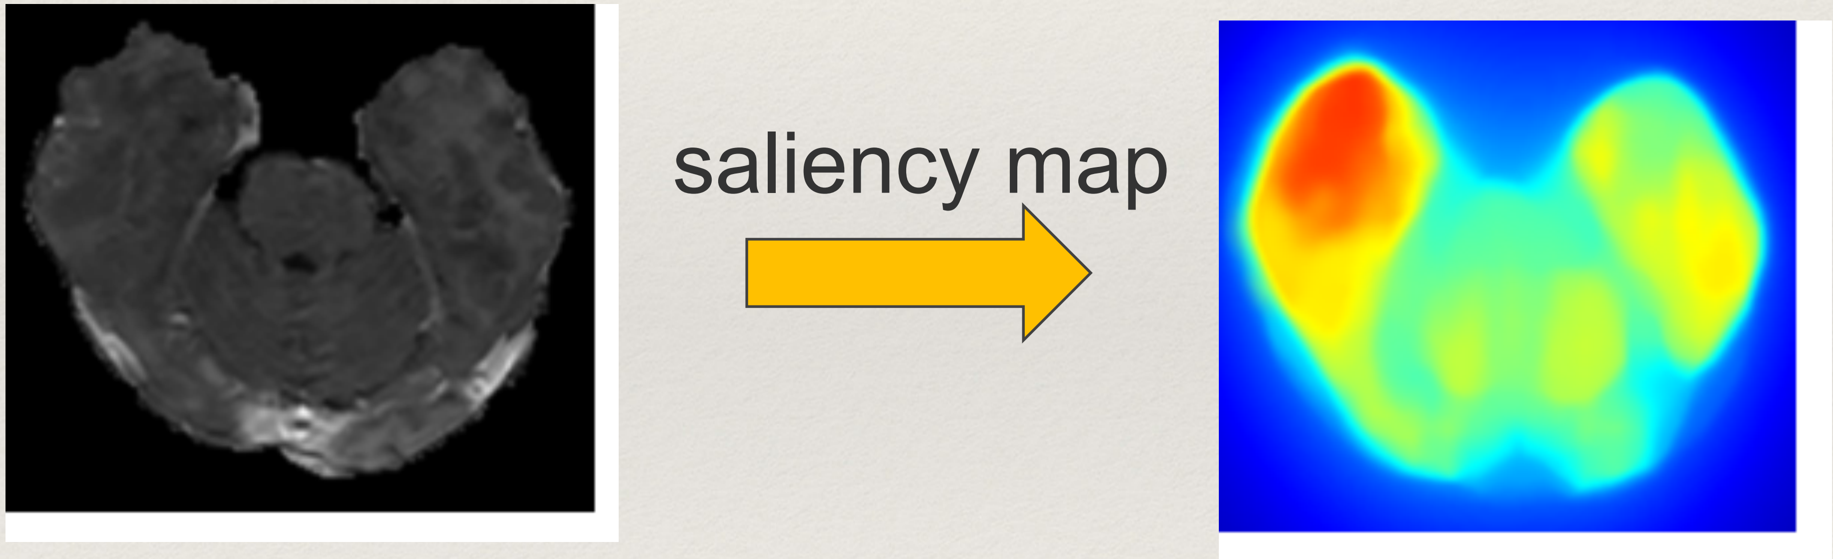

saliency map

low-grade glioma

Supervised Learning

Bayesian Deep Learning

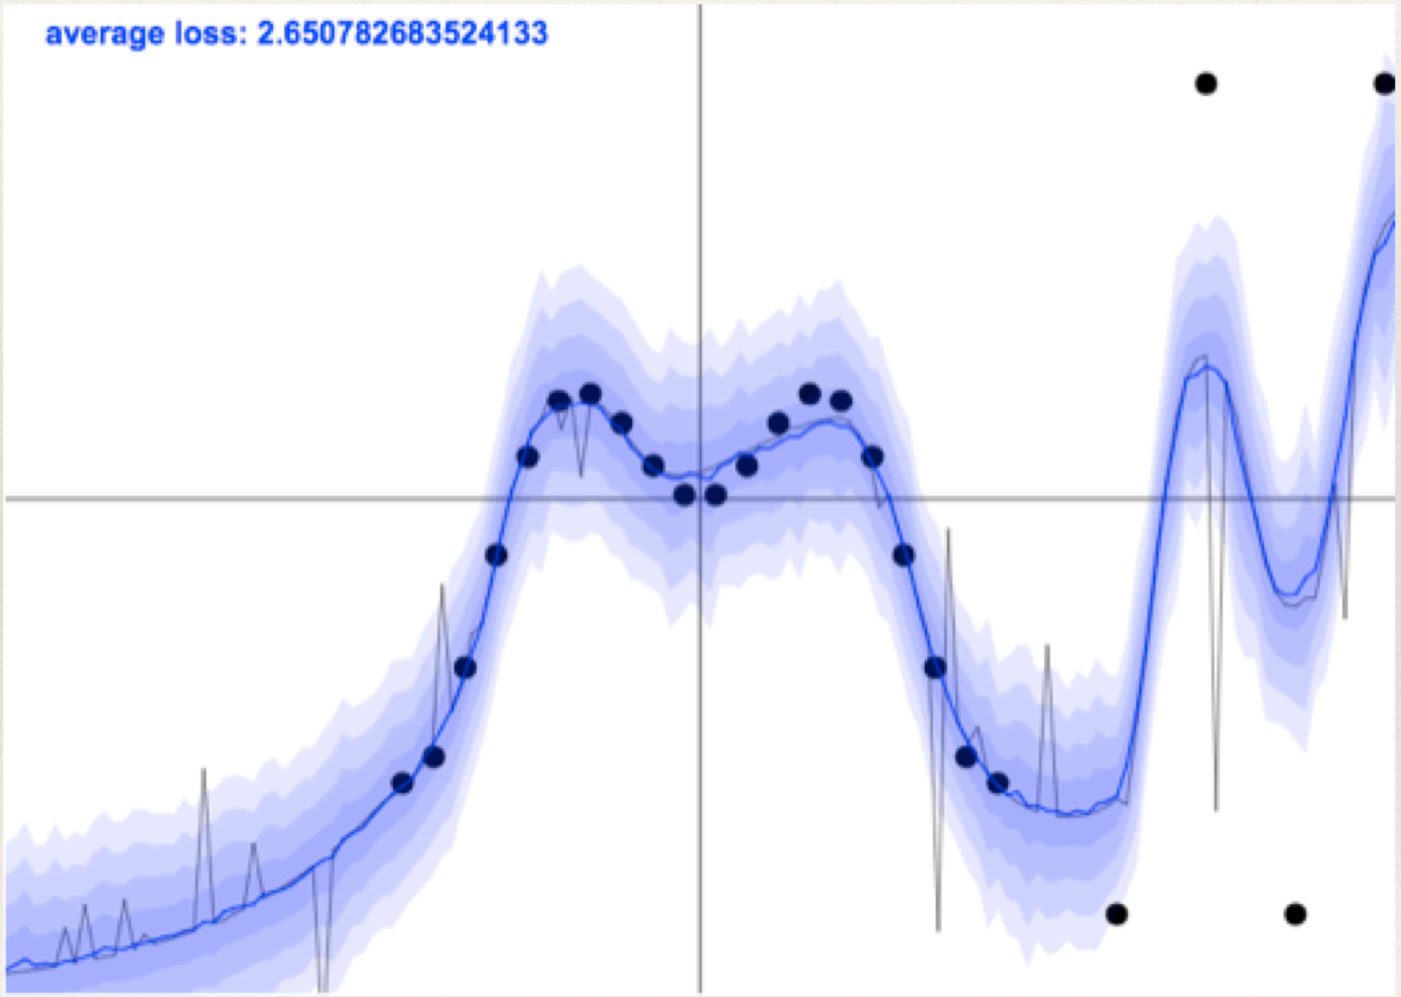

Supervised Learning

Recurrent Neural Networks: LSTMs

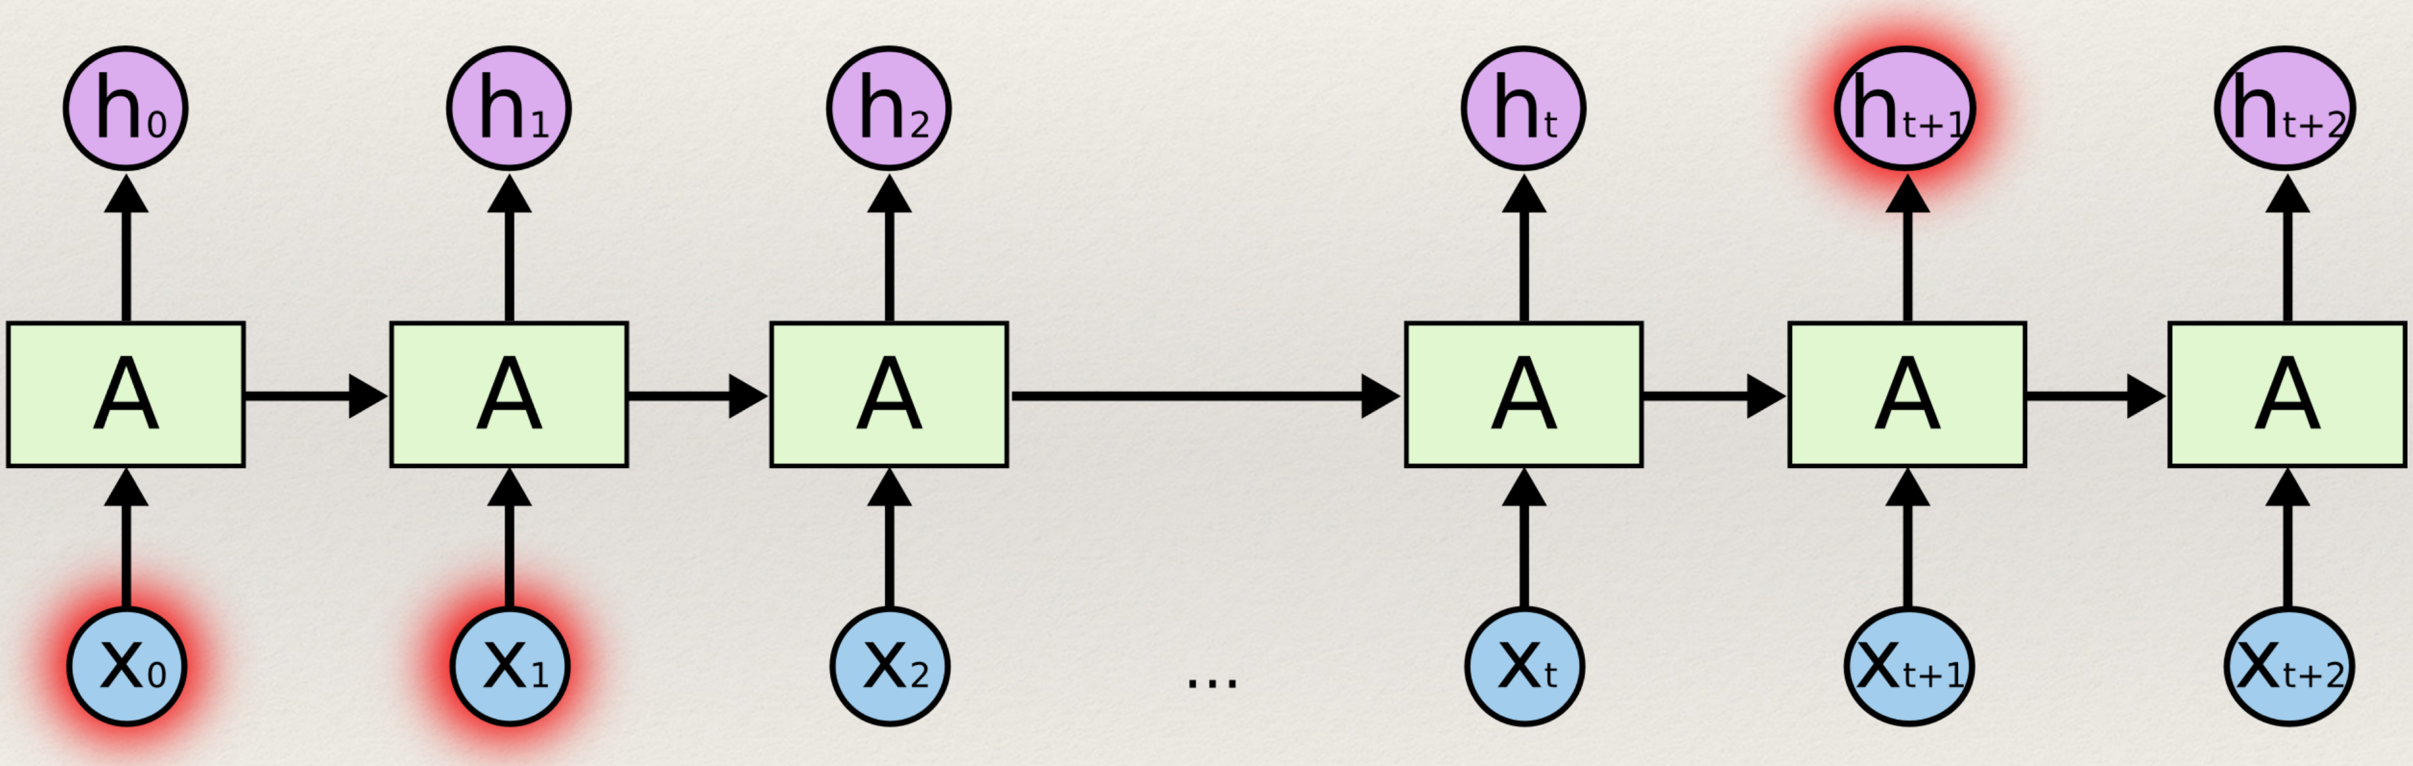

Unsupervised Learning

Generative Adversarial Networks (GANs)

CelebA dataset

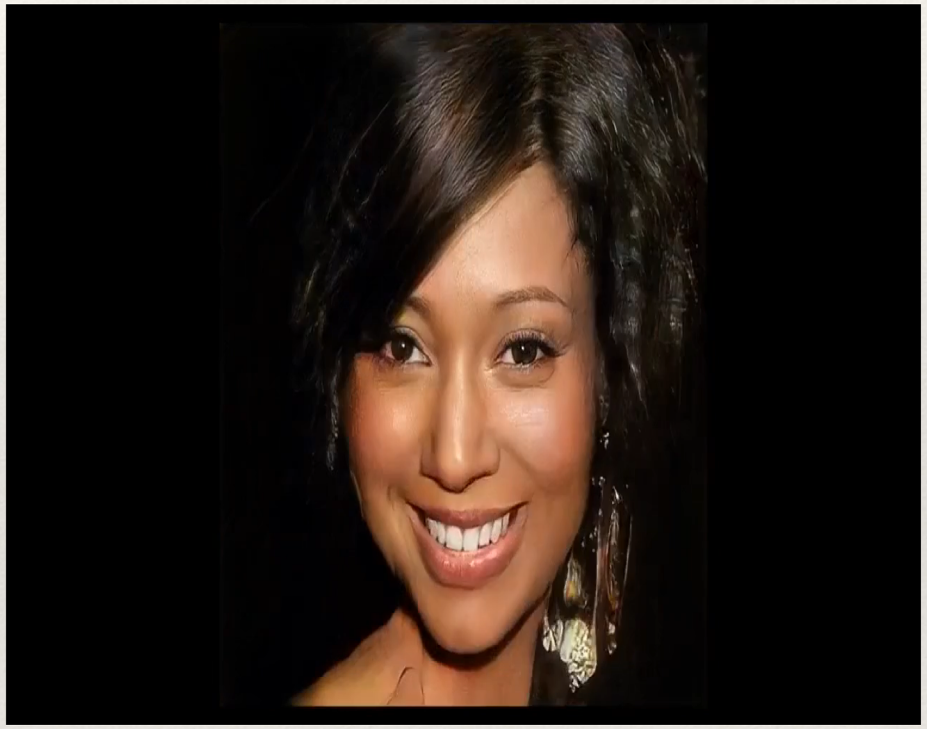

Karras et al. 2017

Formulation, Convergence, Mode-Collapse

Unsupervised Learning

Variational AutoEncoders (VAEs)

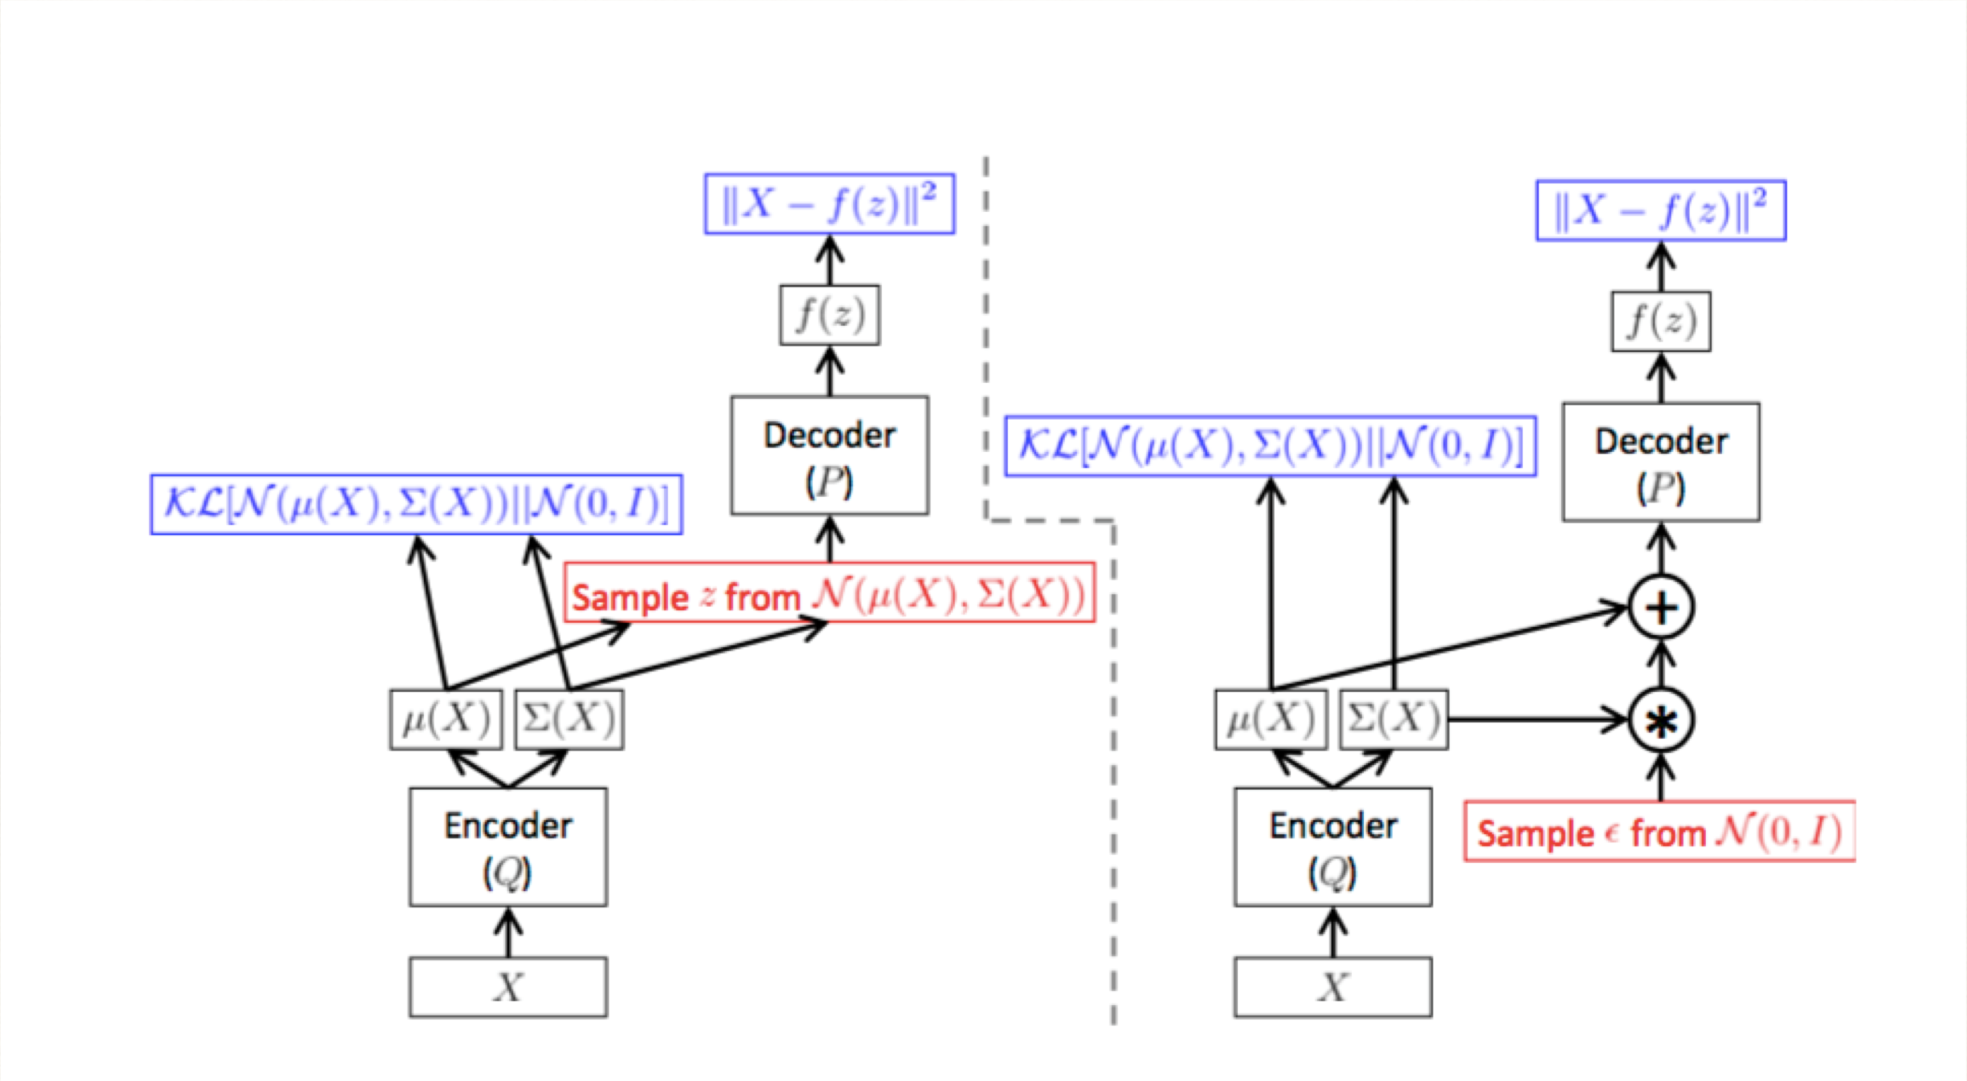
$\|X - f(z)\|^2$
$f(z)$
Decoder (P)
$\mathcal{KL}[\mathcal{N}(\mu(X), \Sigma(X)) \| \mathcal{N}(0, I)]$
Sample z from $\mathcal{N}(\mu(X), \Sigma(X))$
$\mu(X)$
$\Sigma(X)$
Encoder (Q)
X
$\|X - f(z)\|^2$
$f(z)$
$\mathcal{KL}[\mathcal{N}(\mu(X), \Sigma(X)) \| \mathcal{N}(0, I)]$
Decoder (P)
+
*
$\mu(X)$
$\Sigma(X)$
Sample ϵ from $\mathcal{N}(0, I)$
Encoder (Q)
X

Unsupervised Learning

Computing distances between distributions: optimal transport (earth-mover), divergences, etc

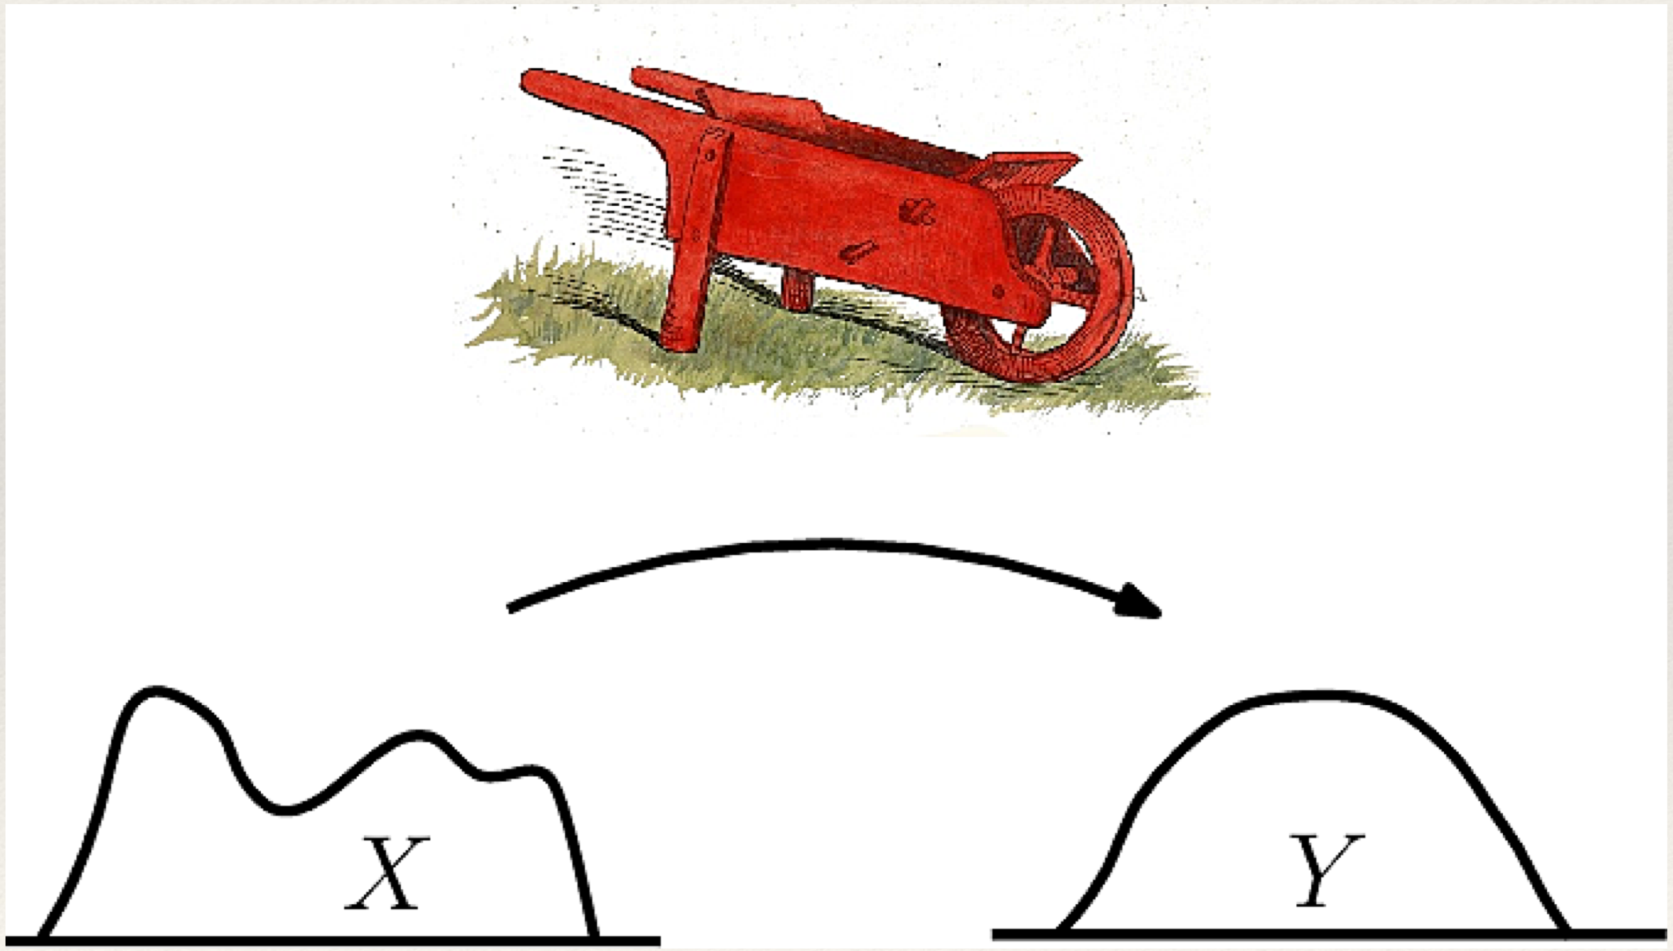

Unsupervised Learning

Nonlinear Dimensionality Reduction: TSNE

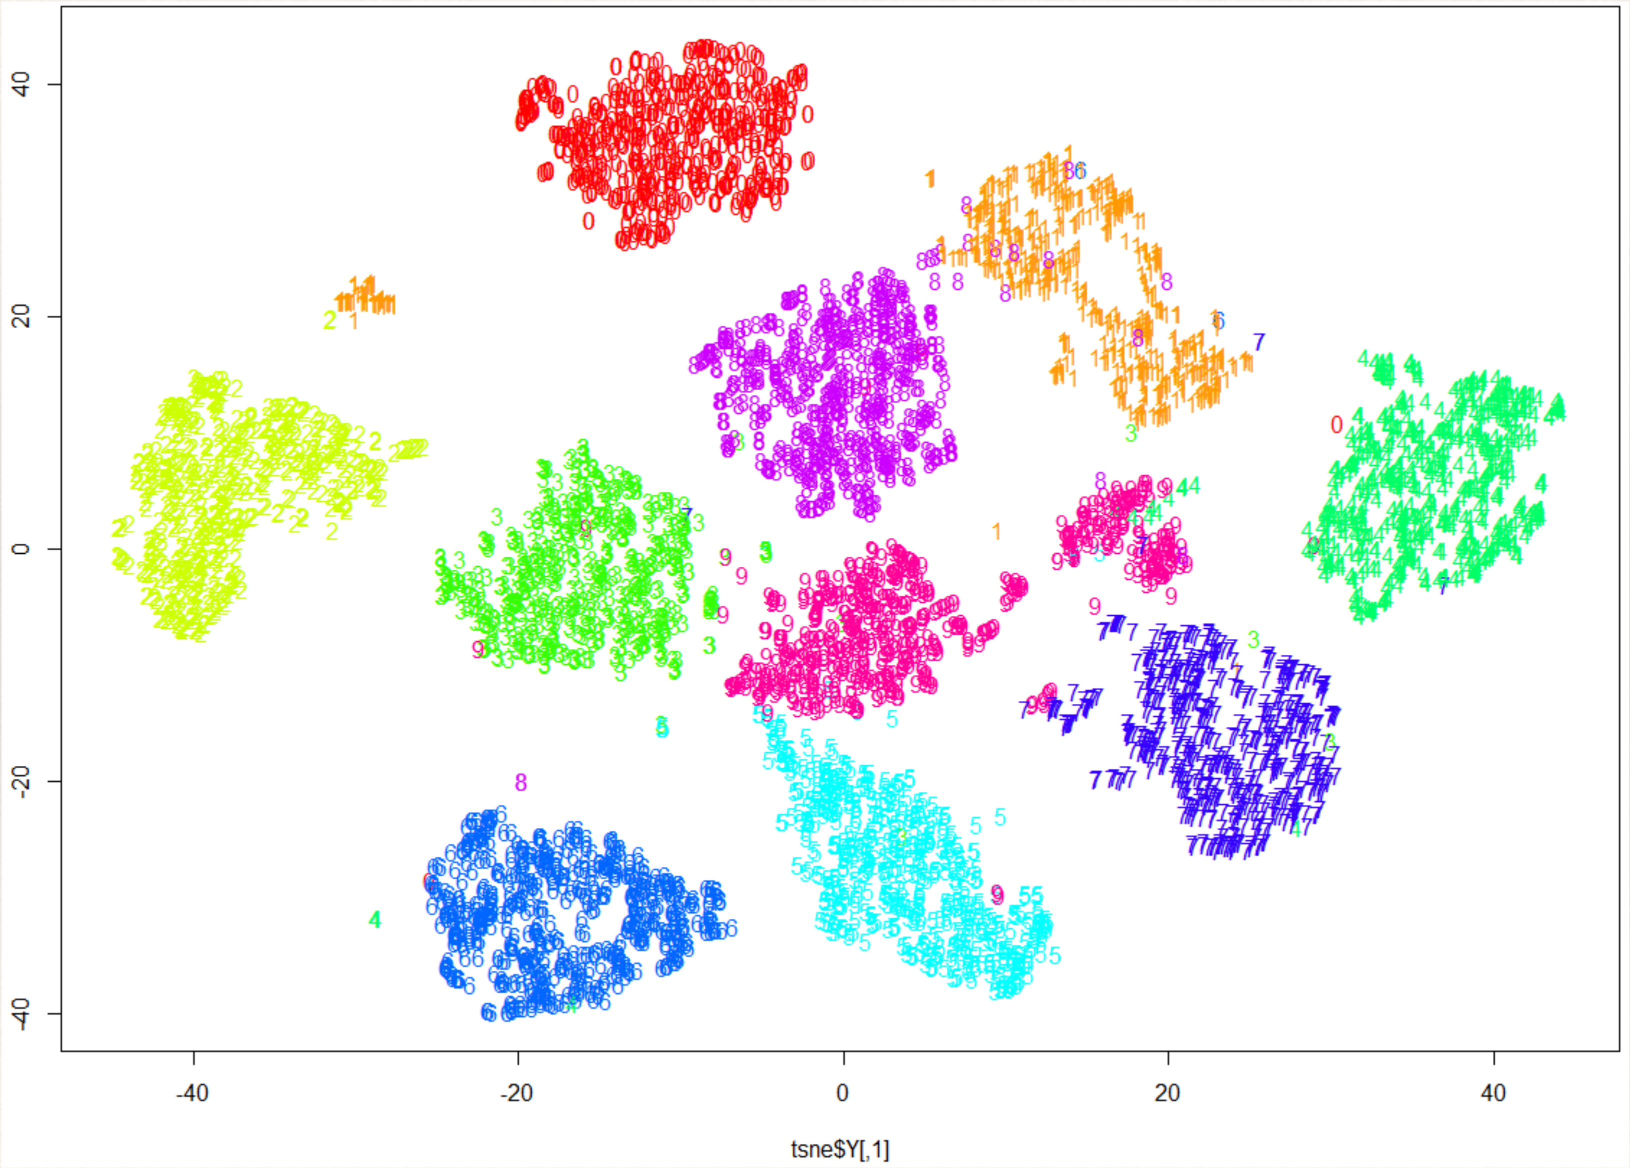

Unsupervised Learning

Nonlinear Dimensionality Reduction: Manifold Learning, Multidimensional Scaling

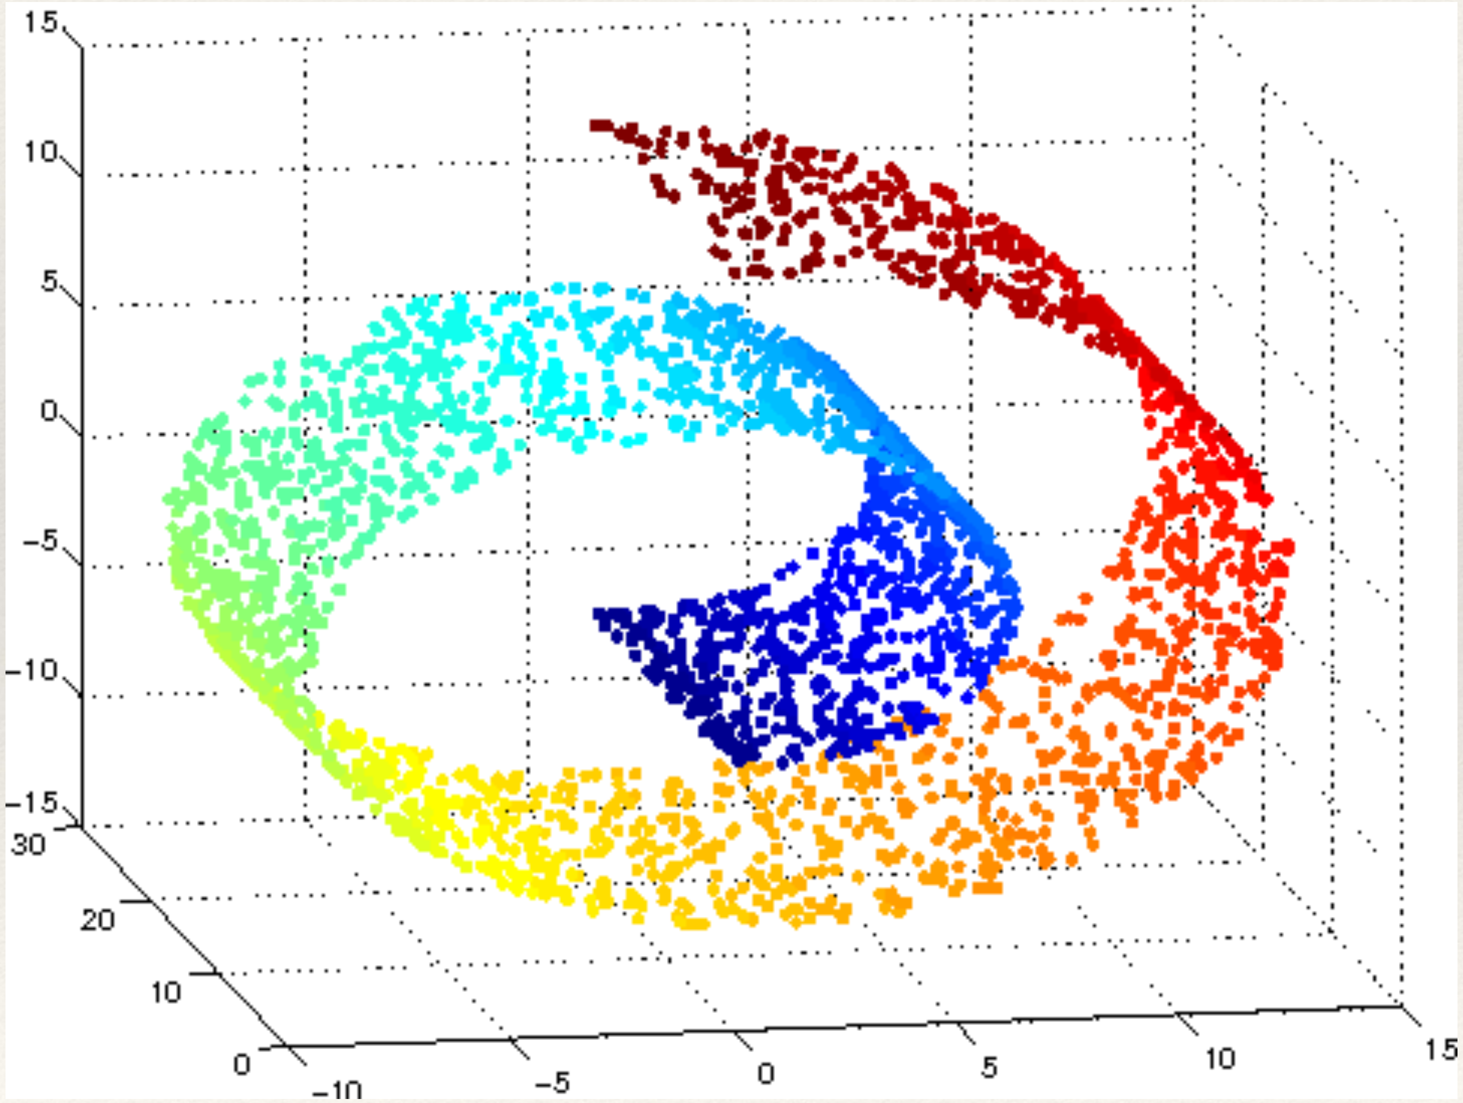

Unsupervised Learning

Embeddings: word2vec, graph2seq

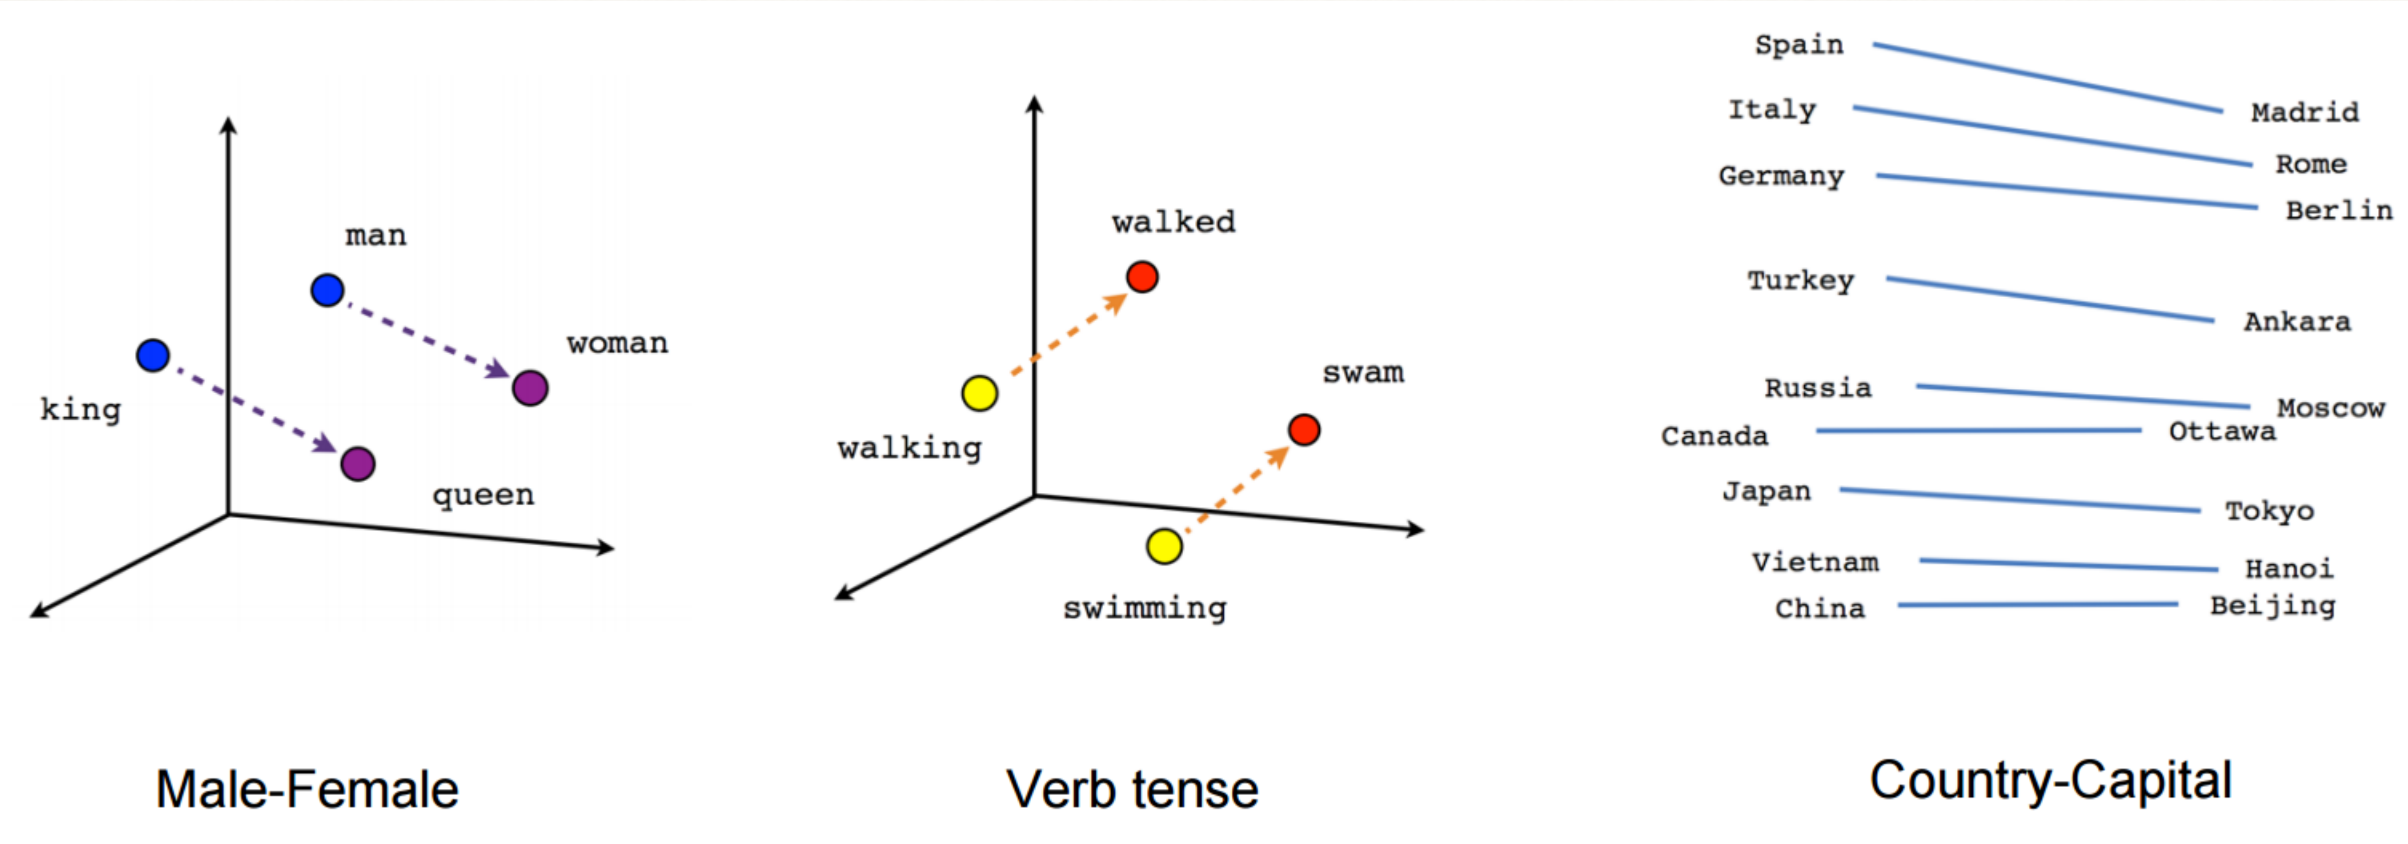

Male-Female

Verb tense

Country-Capital

Unsupervised Learning

Community Detection, Graph Clustering

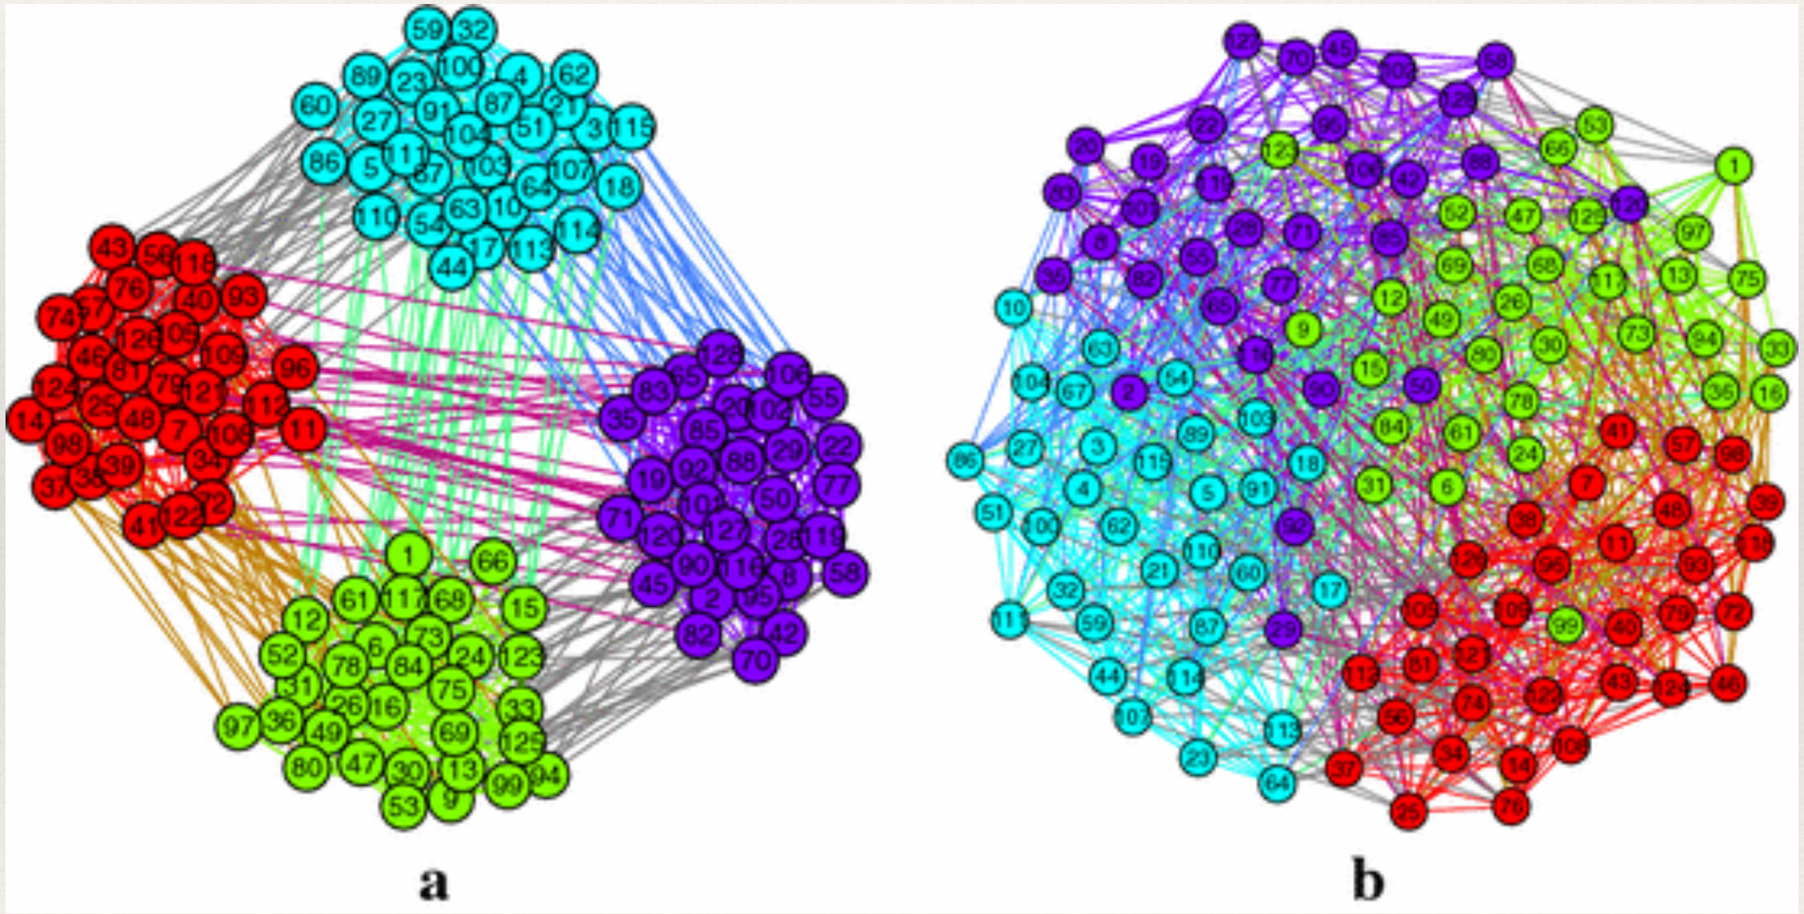

Unsupervised Learning

Domain adaptation, transfer learning

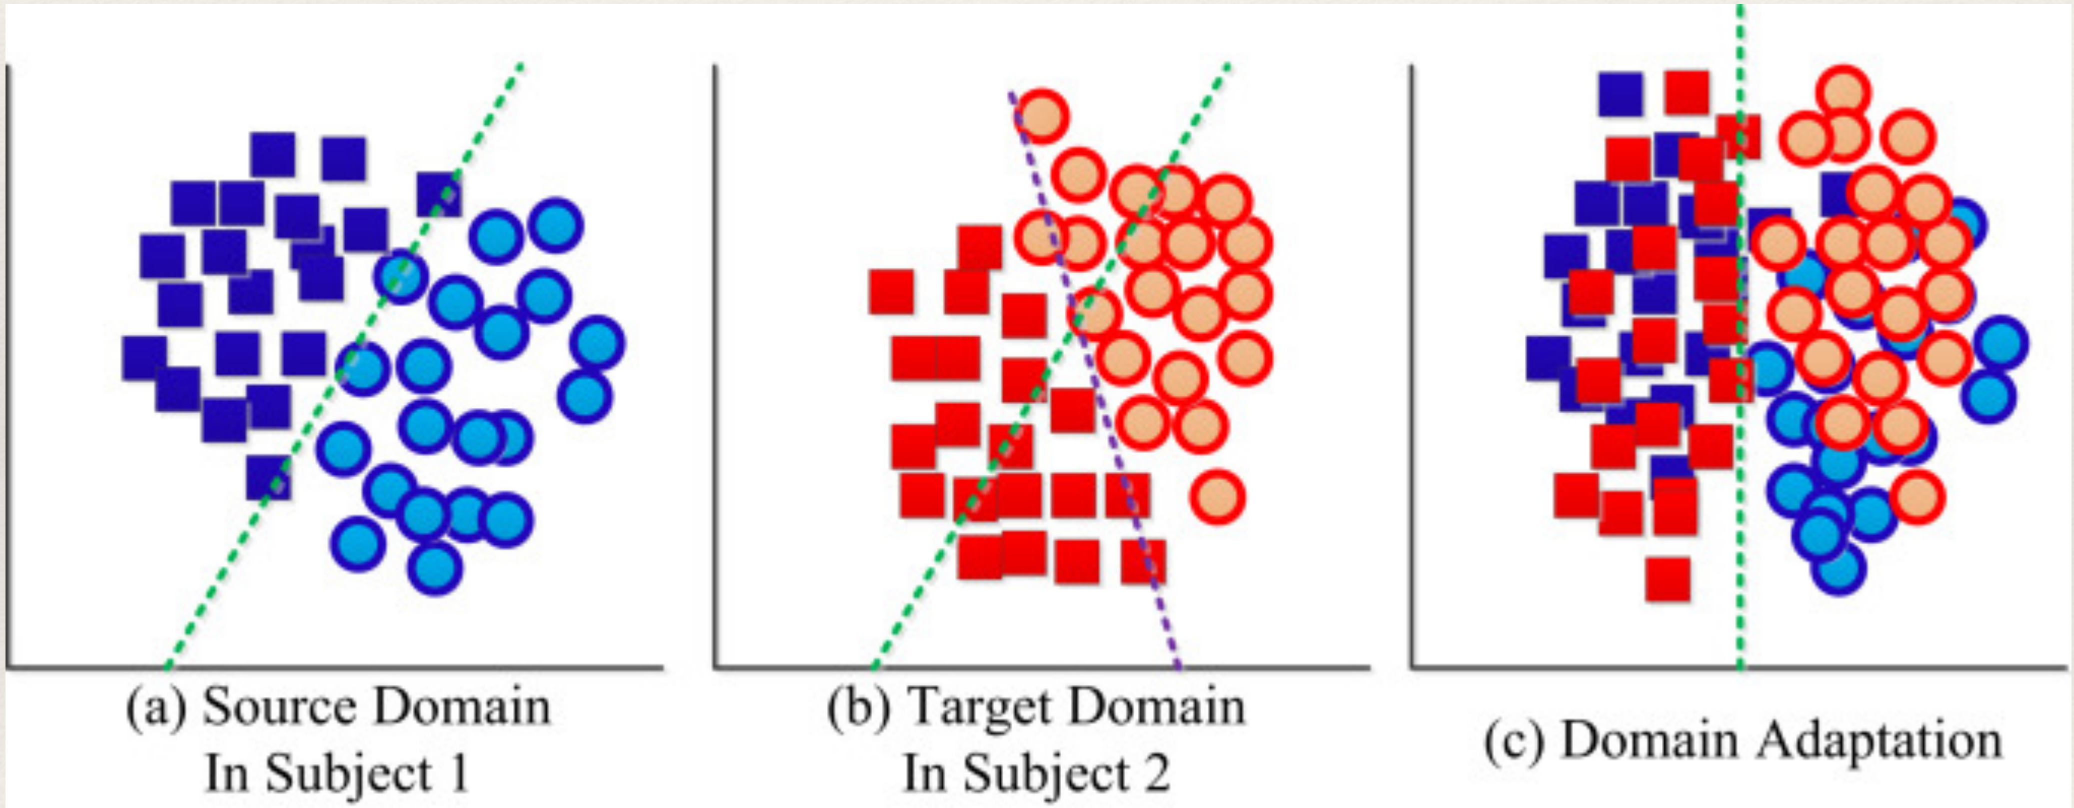

Unsupervised Learning

Topic modeling, nonnegative matrix factorization

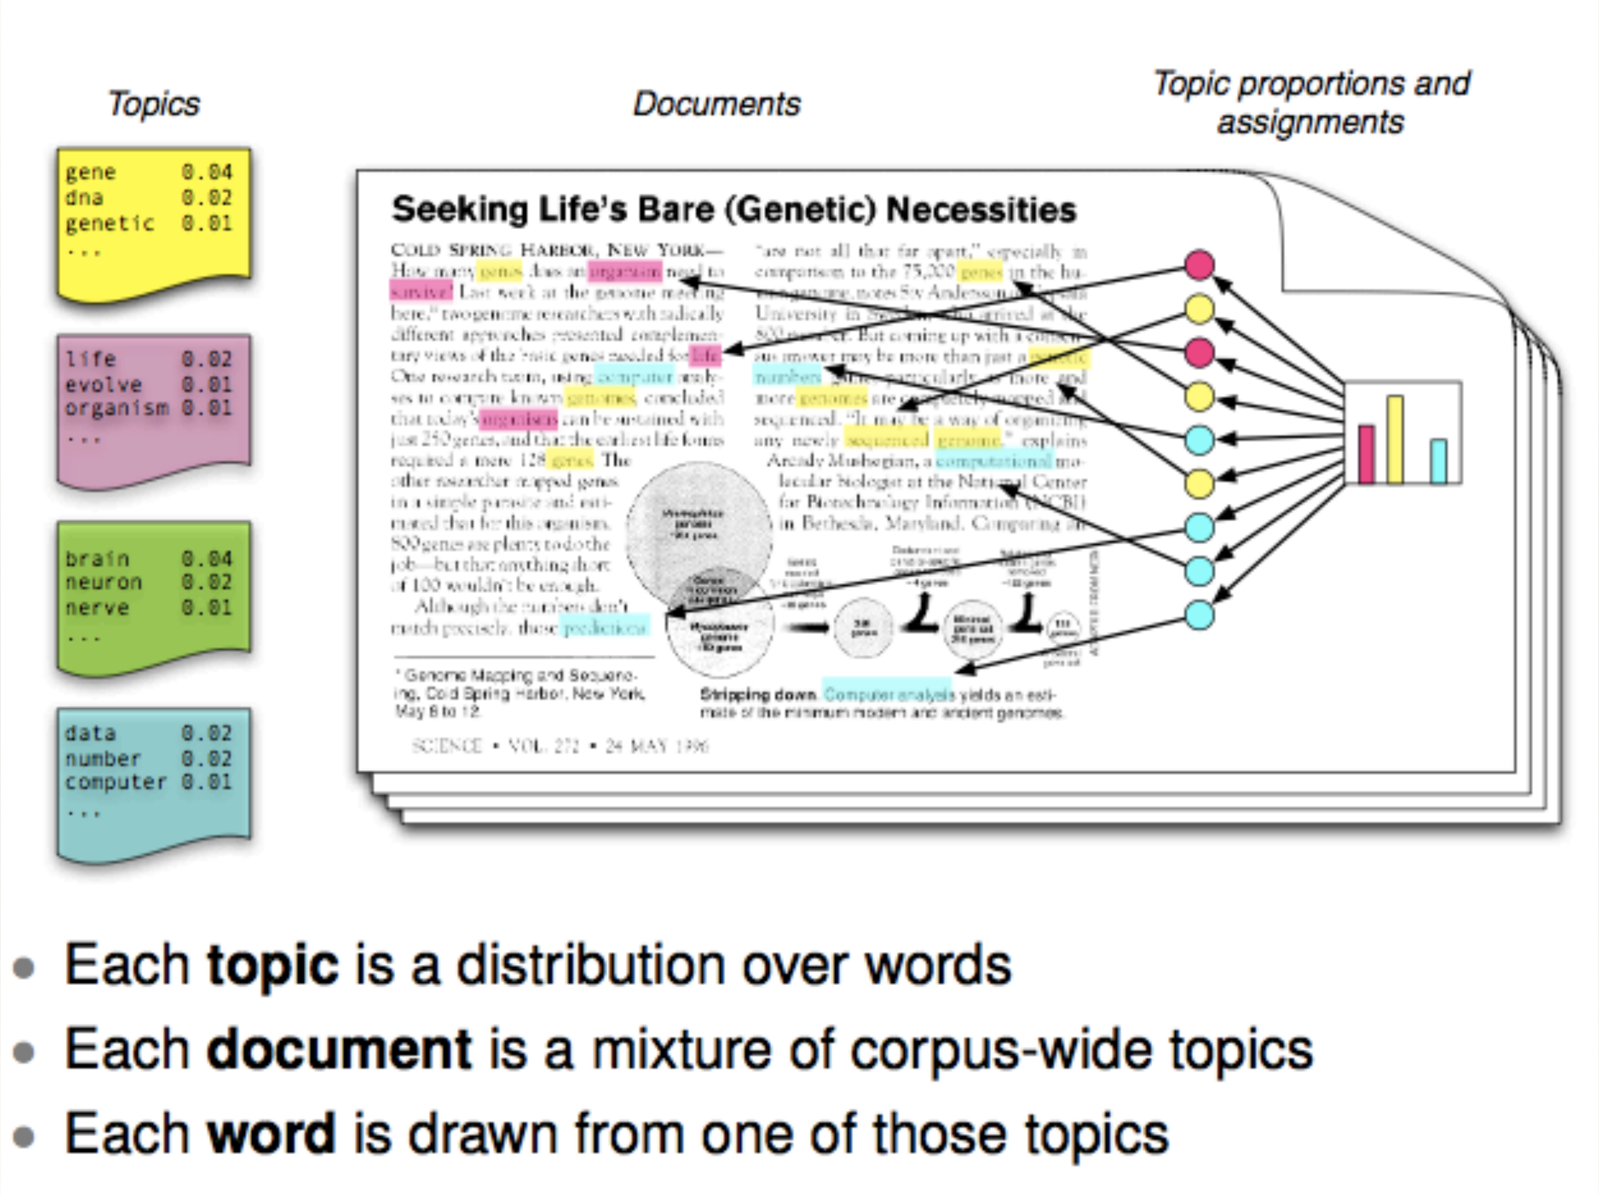

- Each **topic** is a distribution over words
- Each **document** is a mixture of corpus-wide topics
- Each **word** is drawn from one of those topics

Unsupervised Learning

Denoising: outlier detection, etc.

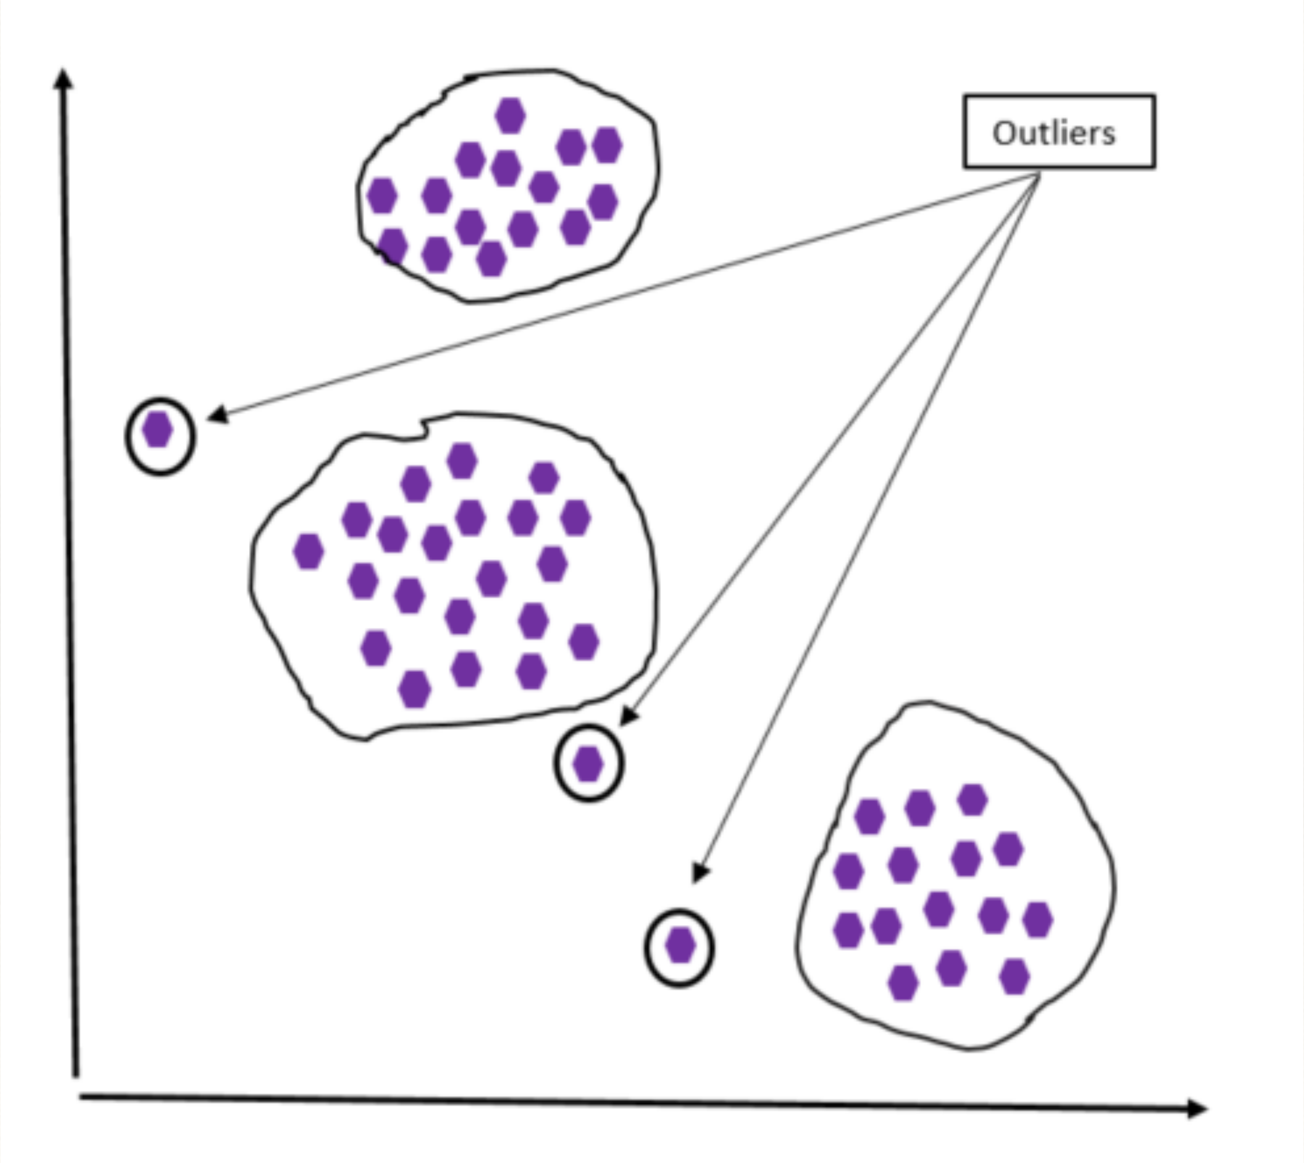

Semi-supervised Learning

Small amount of labelled data+ large amount of unlabeled data

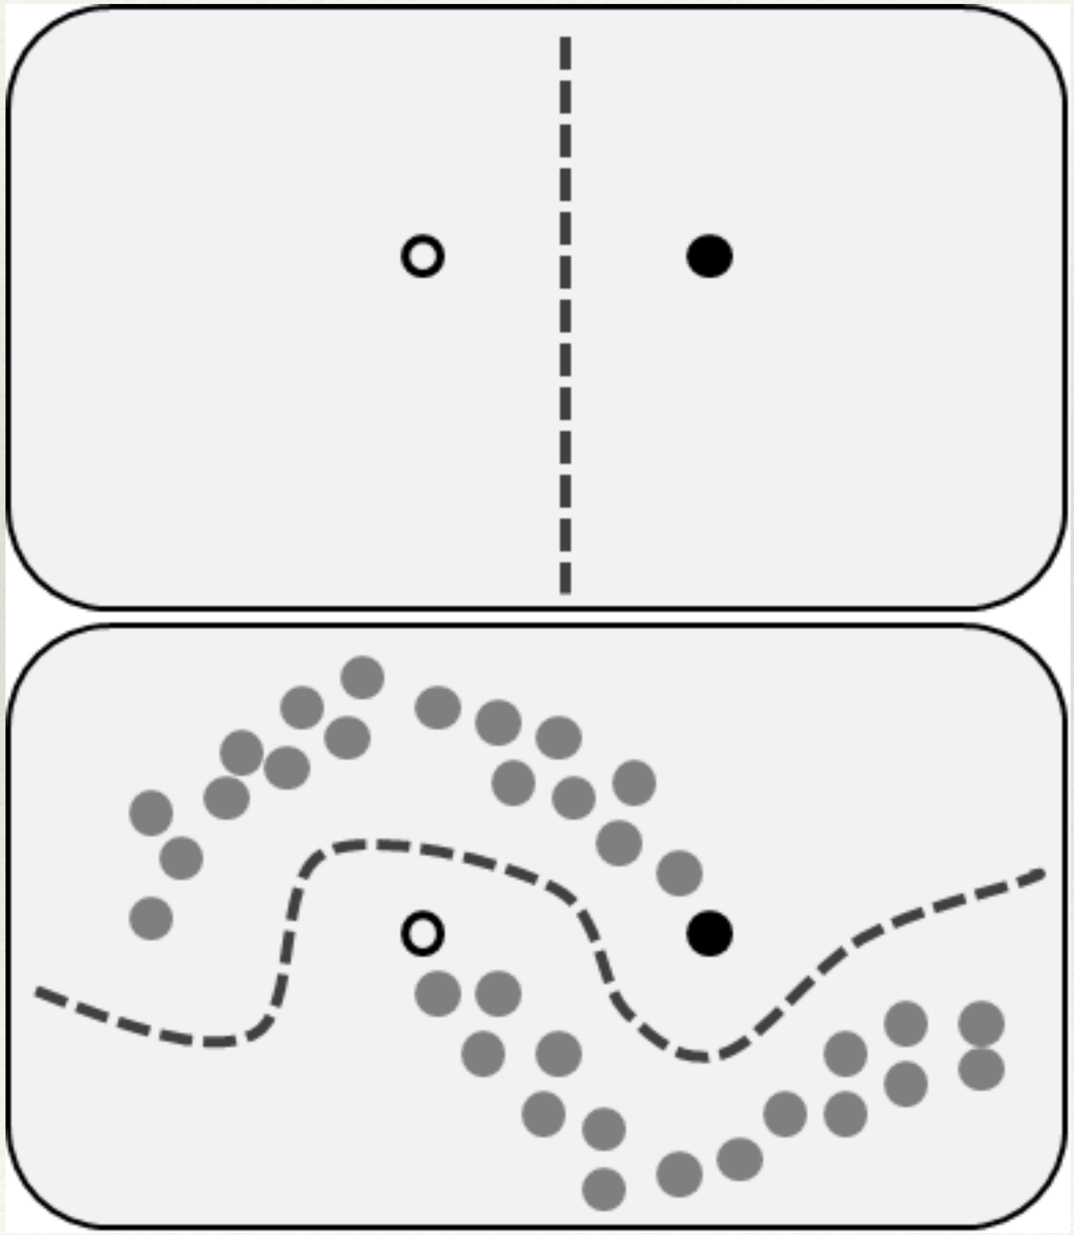

Reinforcement Learning

Approximate dynamic programming

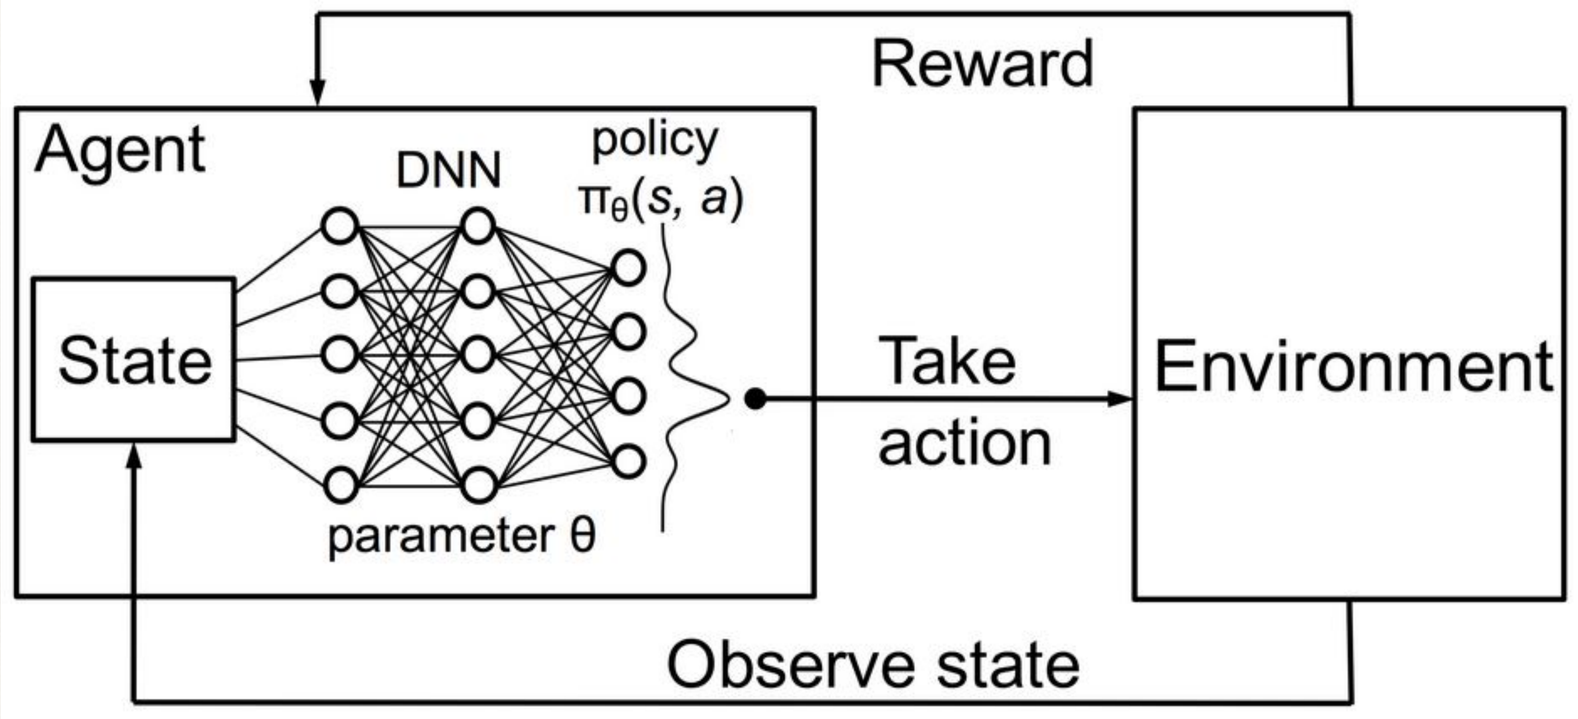

Online vs Batch Learning

Online: data becomes available in sequential order
(e.g. stock price prediction)

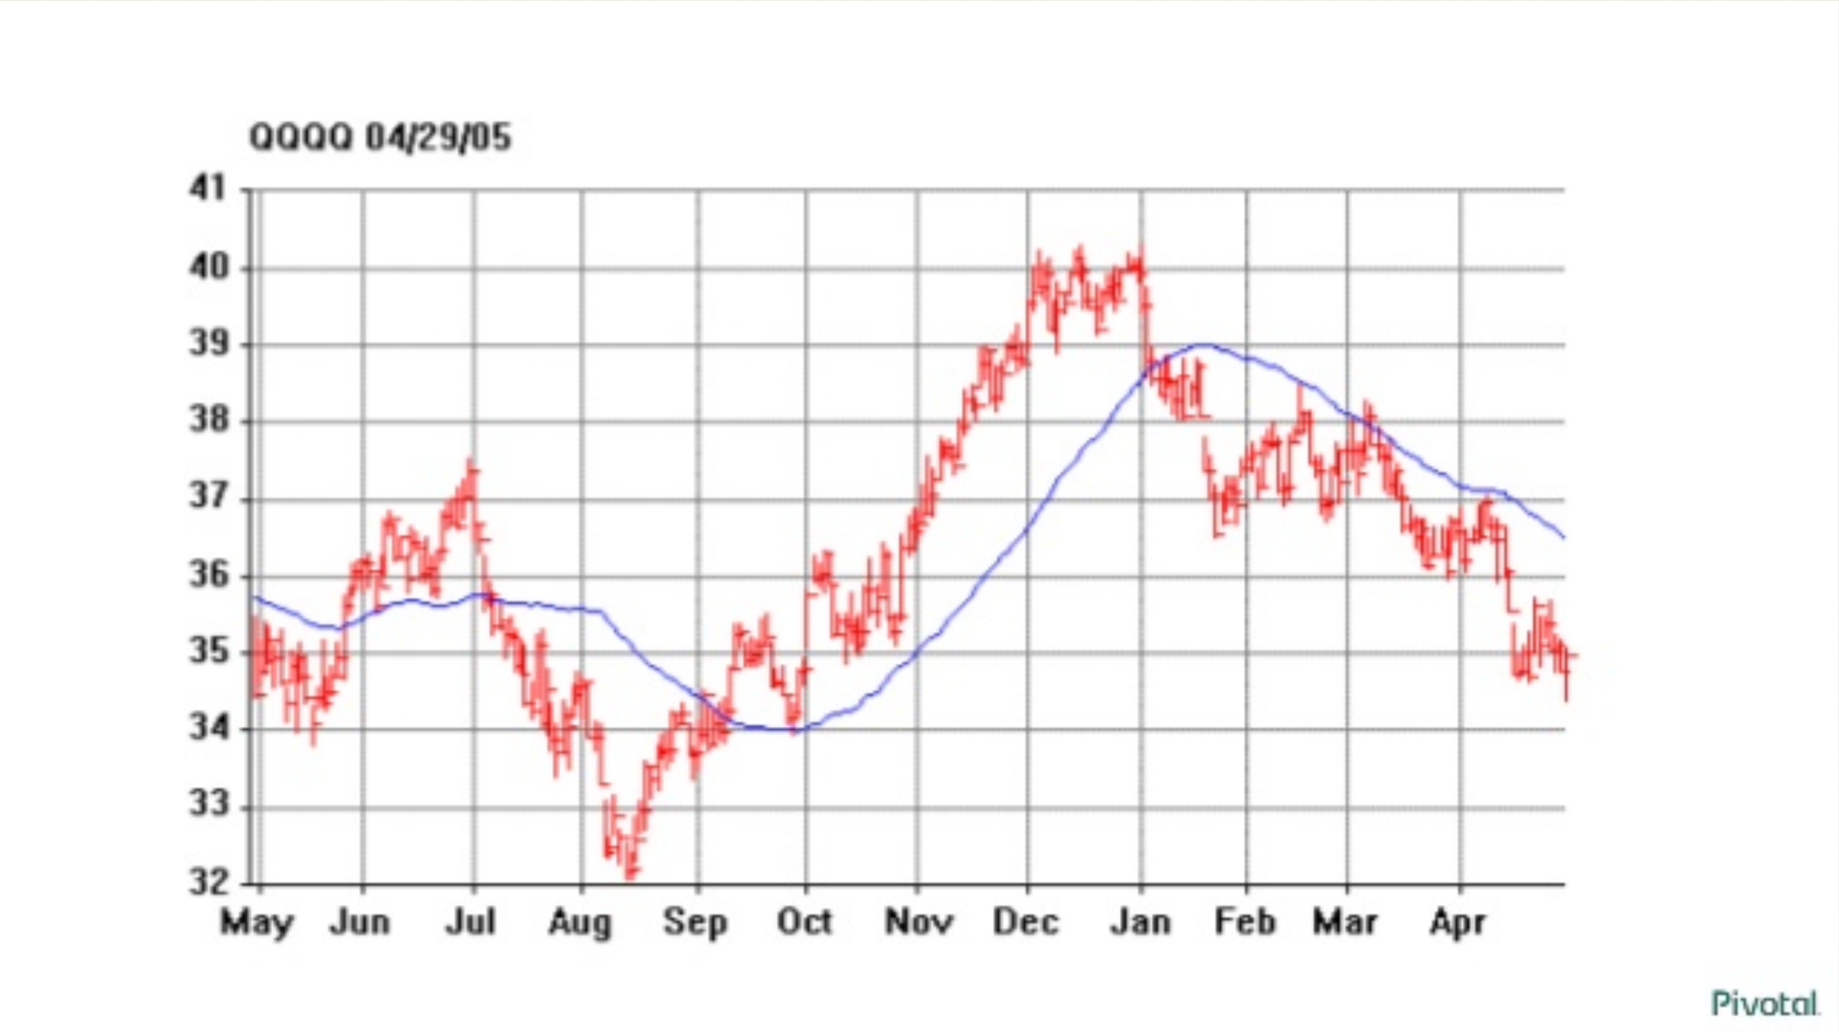

Fairness in Machine Learning

Sensitive features correlated with other features

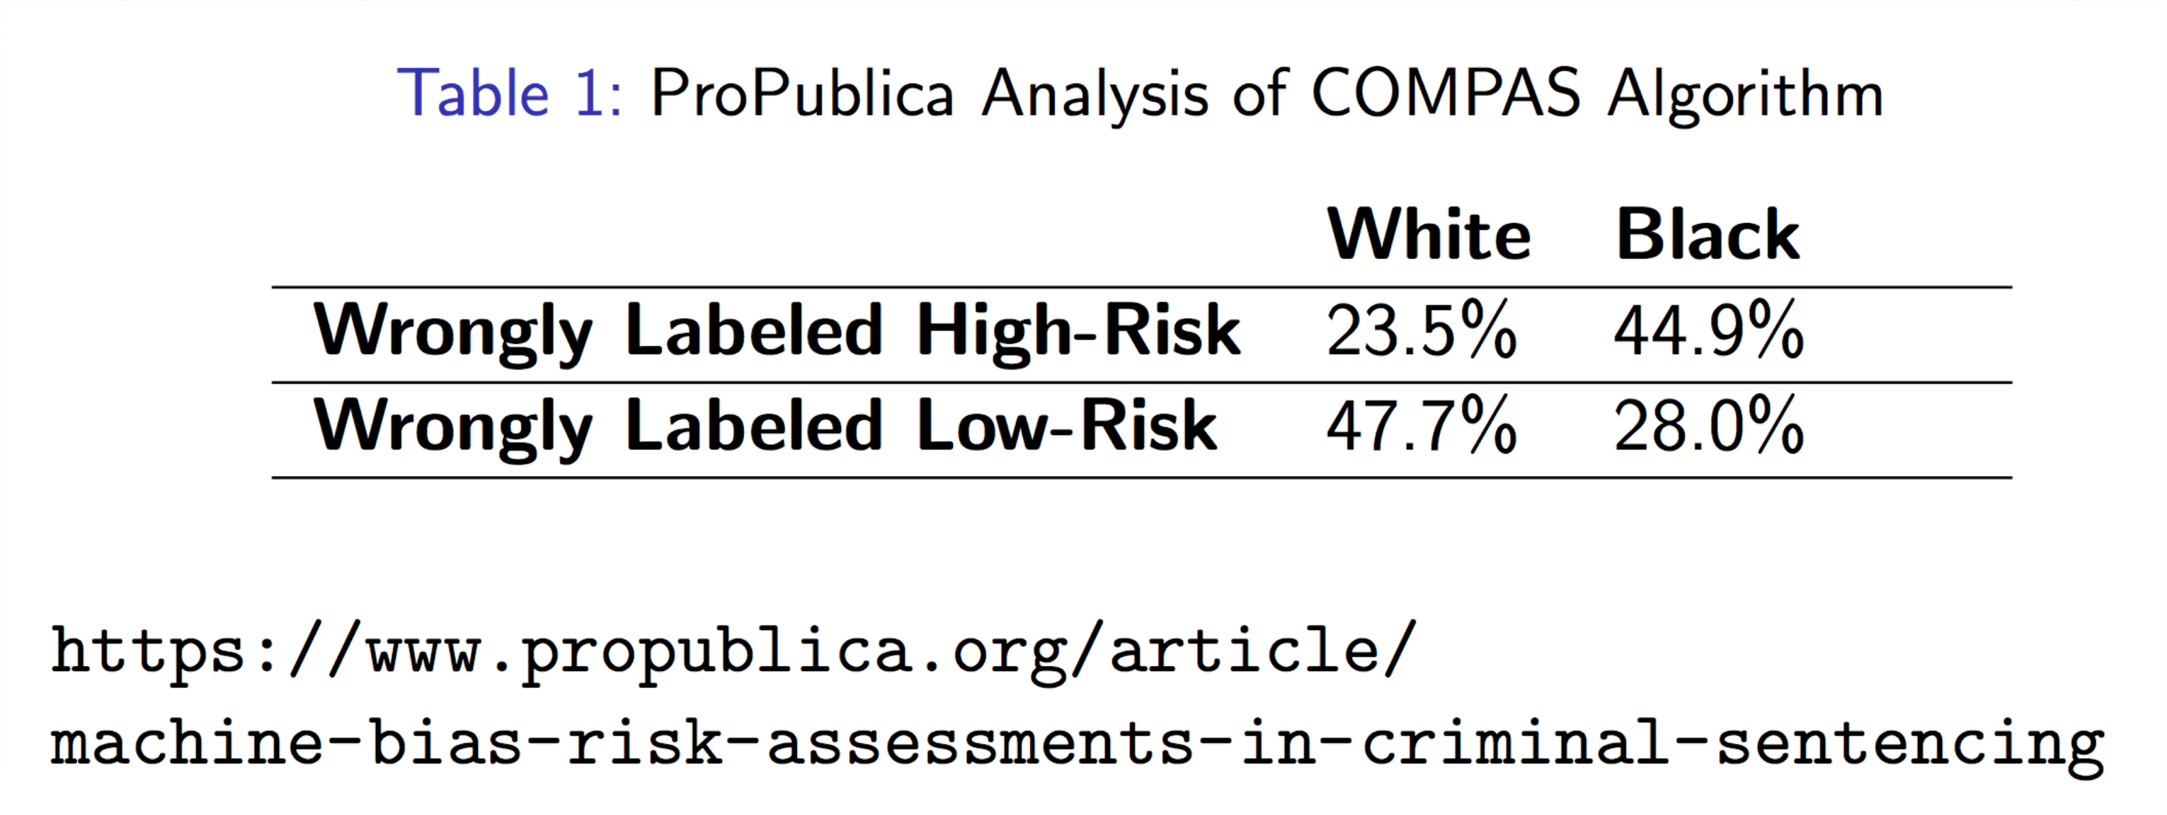

Table 1: ProPublica Analysis of COMPAS Algorithm

	White	Black
Wrongly Labeled High-Risk	23.5%	44.9%
Wrongly Labeled Low-Risk	47.7%	28.0%

https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing

Privacy in Machine Learning

Differential privacy

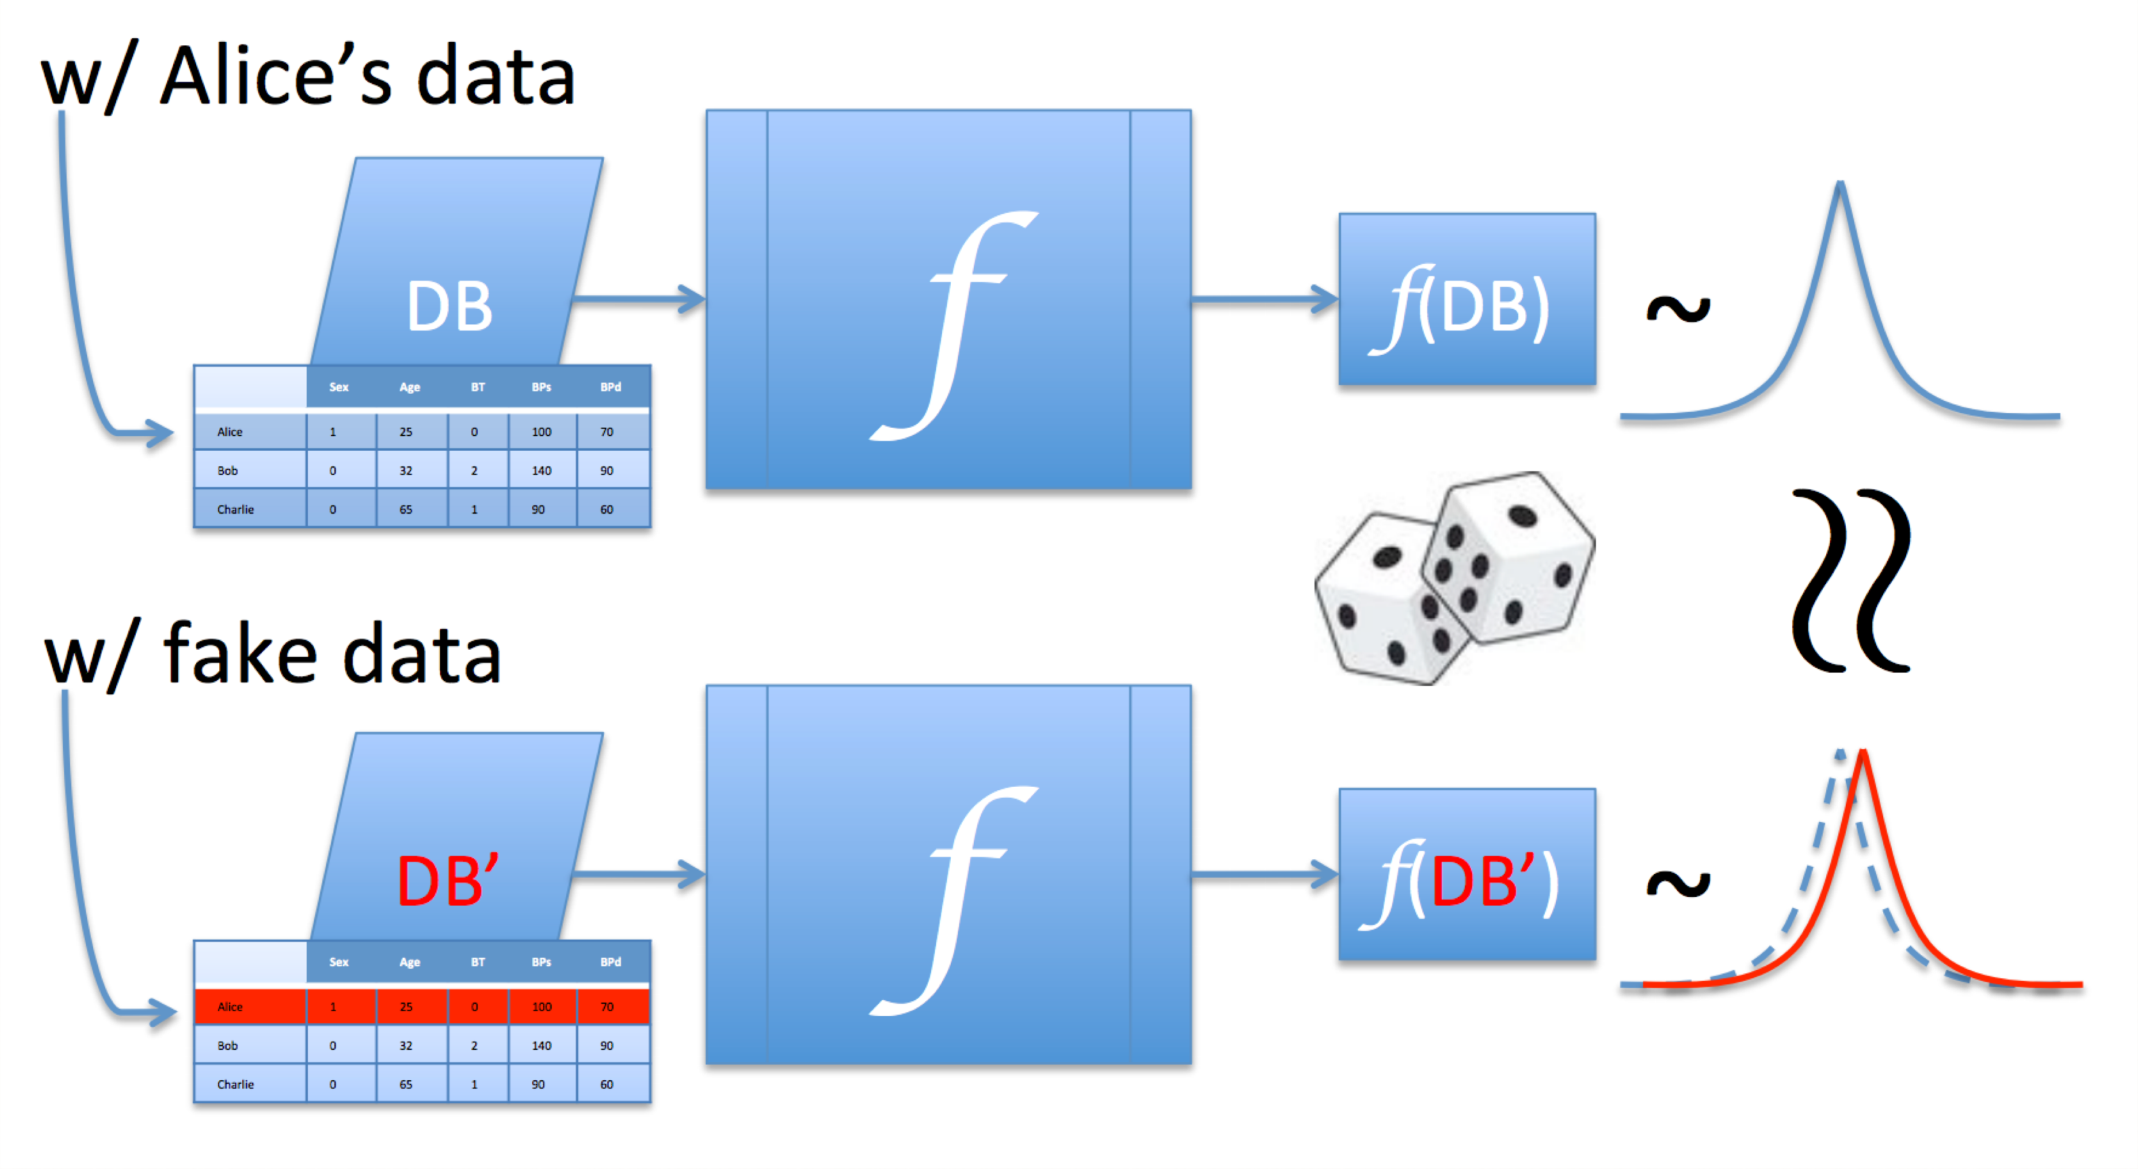

	Sex	Age	BT	BPs	BPd
Alice	1	25	0	100	70
Bob	0	32	2	140	90
Charlie	0	65	1	90	60

	Sex	Age	BT	BPs	BPd
Alice	1	25	0	100	70
Bob	0	32	2	140	90
Charlie	0	65	1	90	60