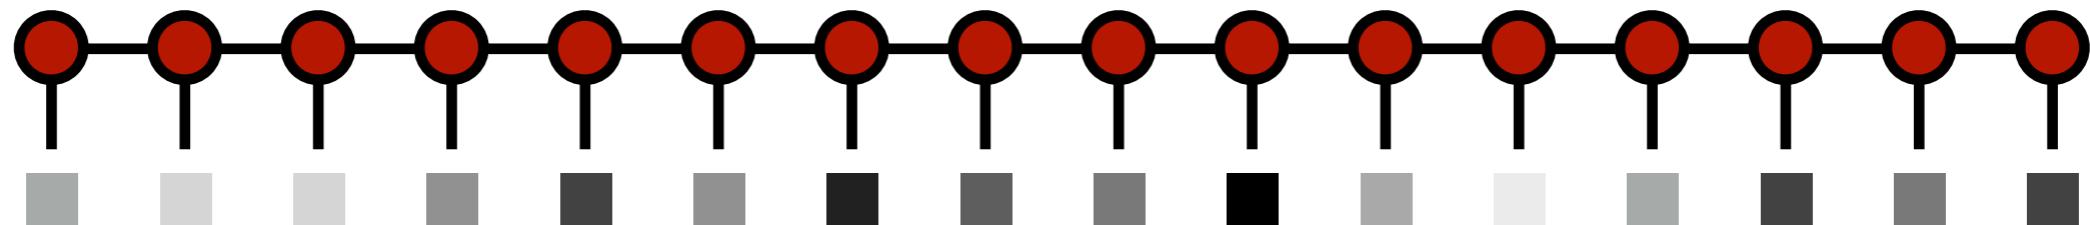


# Introduction to Tensor Networks for Machine Learning



# Plan of Lectures

Tensor networks are a natural way to parameterize interesting and powerful machine learning models

Today:

- Intro to Machine Learning
- Intro to Tensor Networks
- Tensor Network Machine Learning

Thursday:

- Architectures beyond MPS
- Potential of new algorithms

# Machine learning galvanizing industry & science



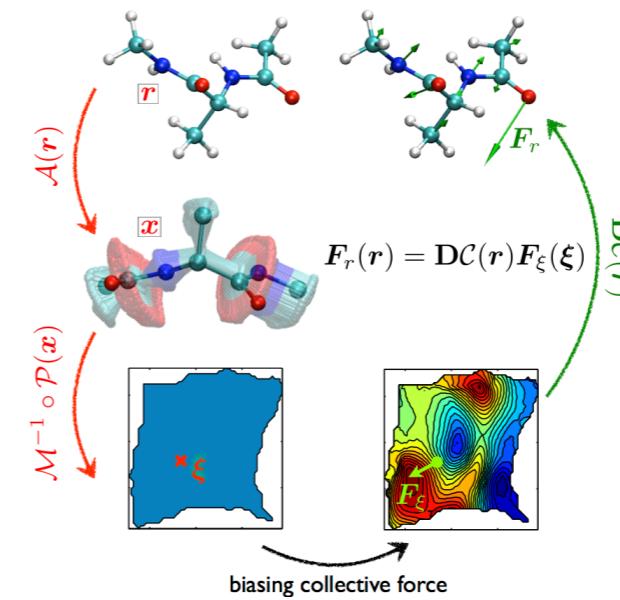
Language Processing



Self-driving cars



Medicine



Materials Science / Chemistry

# Google rebranded a "machine learning first company"



Neural nets replace linguistic approach to Google Translate

## HOW GOOGLE IS REMAKING ITSELF AS A "MACHINE LEARNING FIRST" COMPANY

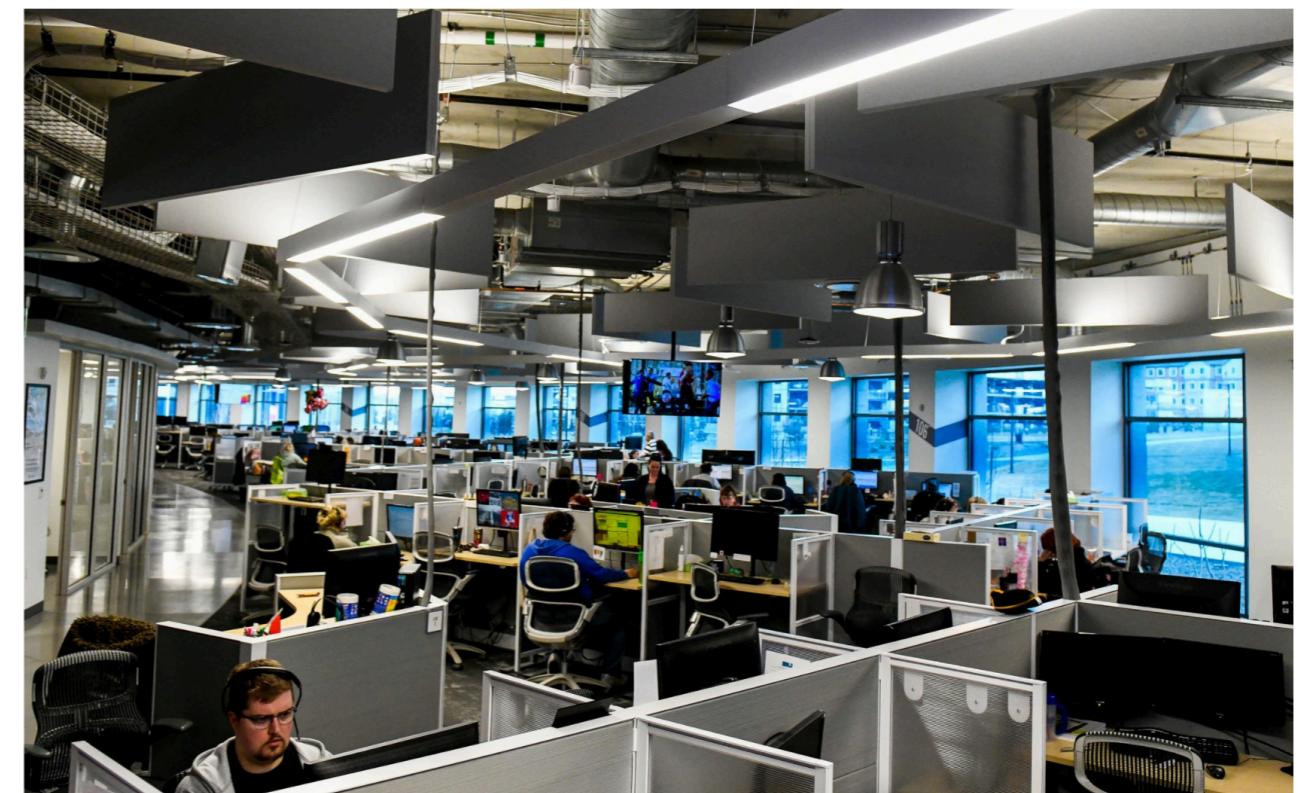
BY STEVEN LEVY DEC. 14, 2016

### How Google is Remaking Itself as a "Machine Learning First" Company

TECHNOLOGY

#### *Why A.I. Researchers at Google Got Desks Next to the Boss*

By CADE METZ FEB. 19, 2018



arXiv.org > quant-ph > arXiv:1802.06002

Quantum Physics

#### Classification with Quantum Neural Networks on Near Term Processors

Edward Farhi, Hartmut Neven

(Submitted on 16 Feb 2018)

Quantum machine learning

# **Examples of Machine Learning**

# Image recognition

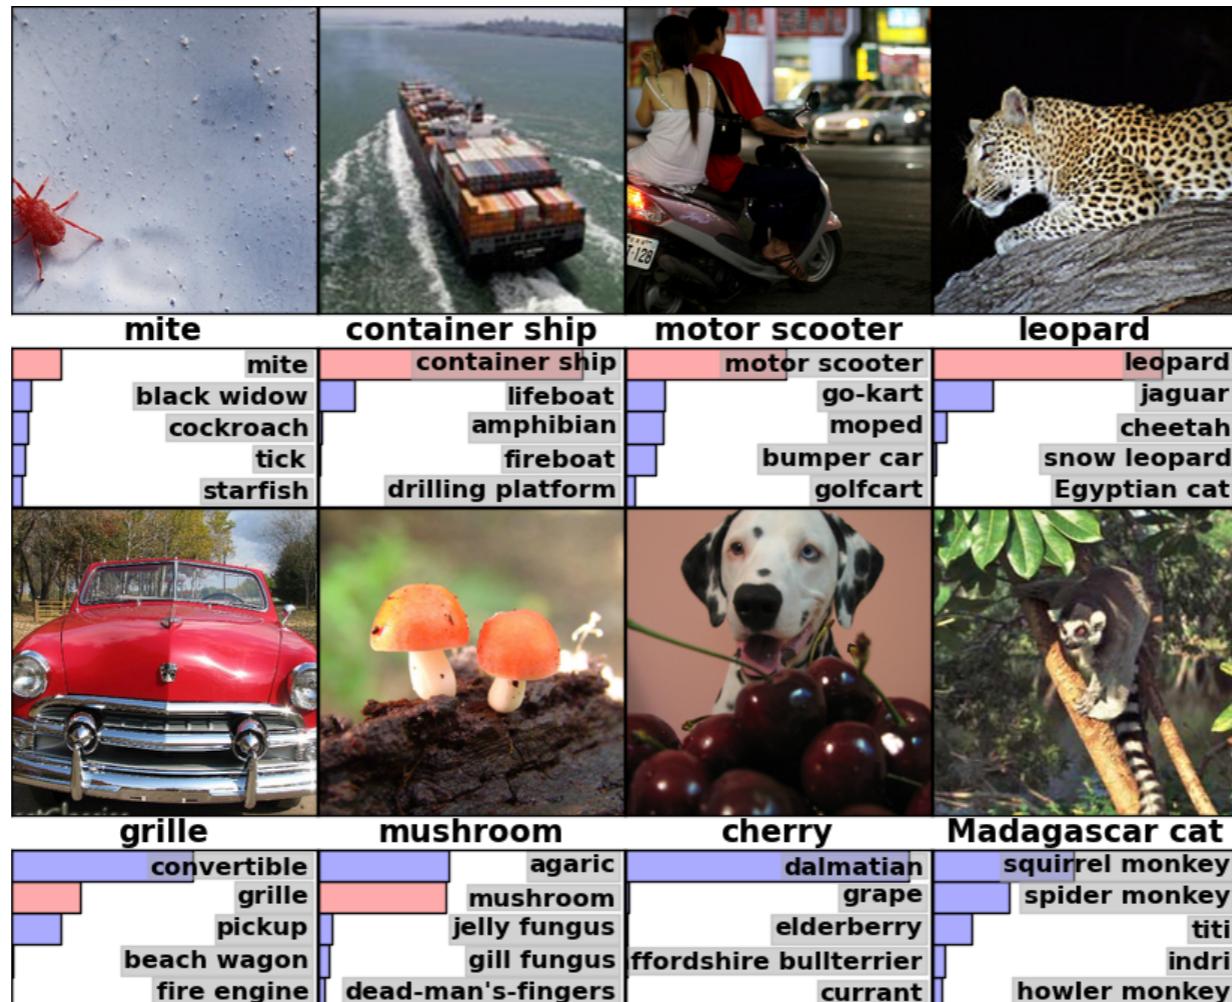
## ImageNet Classification with Deep Convolutional Neural Networks

Alex Krizhevsky  
University of Toronto  
kriz@cs.utoronto.ca

Ilya Sutskever  
University of Toronto  
ilya@cs.utoronto.ca

Geoffrey E. Hinton  
University of Toronto  
hinton@cs.utoronto.ca

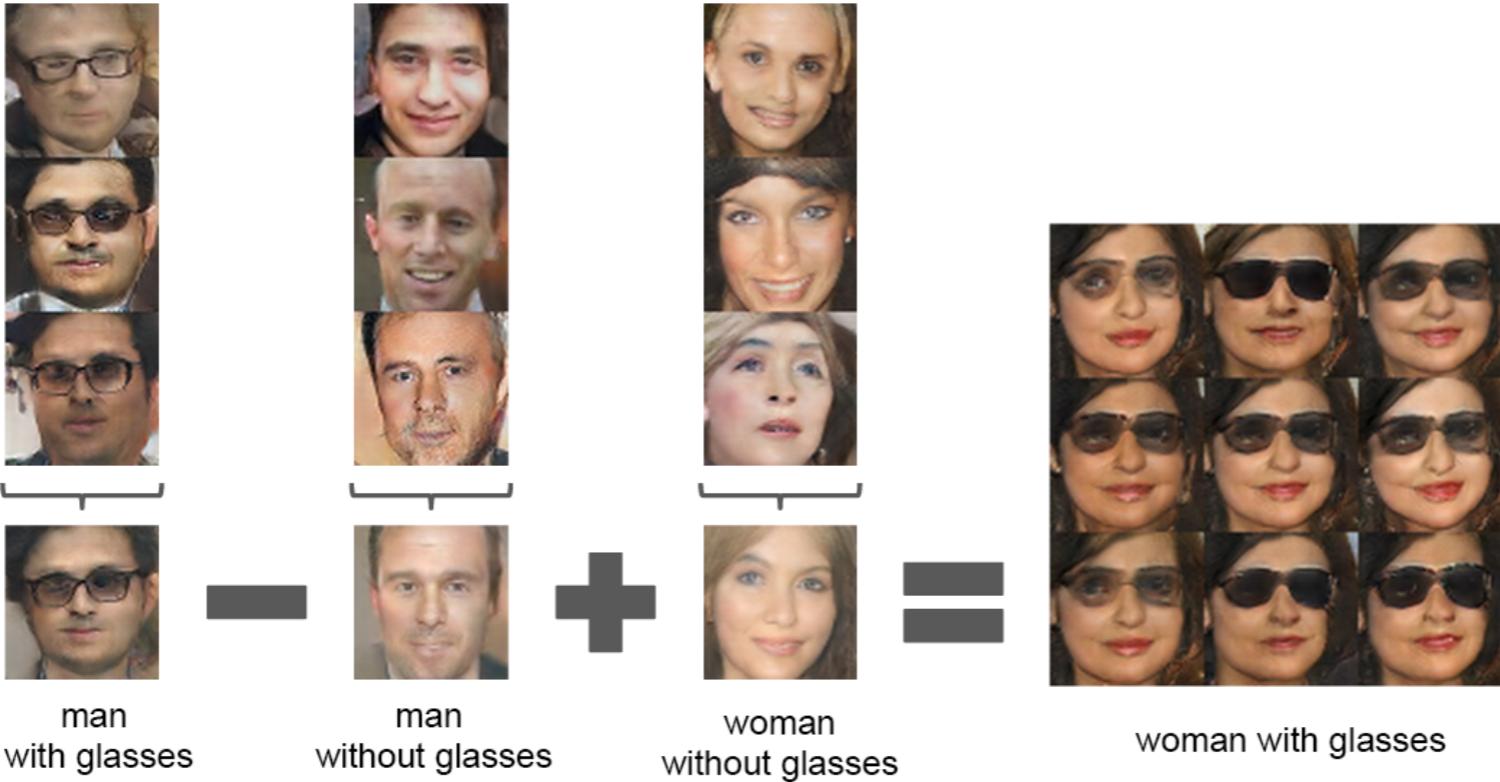
2012 paper that launched recent deep learning era (20k citations)



ImageNet:

- 1.2 million training images (150k test)
- 1000 categories
- 15% neural net error
- 26% next best error

# Image Generation

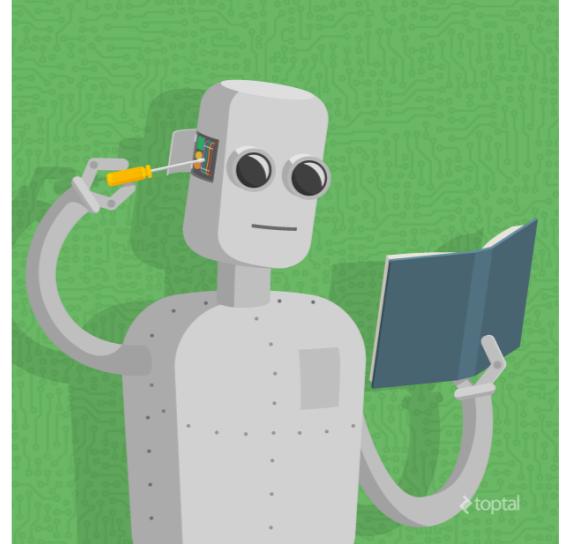


## UNSUPERVISED REPRESENTATION LEARNING WITH DEEP CONVOLUTIONAL GENERATIVE ADVERSARIAL NETWORKS

**Alec Radford & Luke Metz**  
indico Research  
Boston, MA  
[{alec,luke}@indico.io](mailto:{alec,luke}@indico.io)

**Soumith Chintala**  
Facebook AI Research  
New York, NY  
[soumith@fb.com](mailto:soumith@fb.com)

# What is machine learning?



*Data driven* problem solving

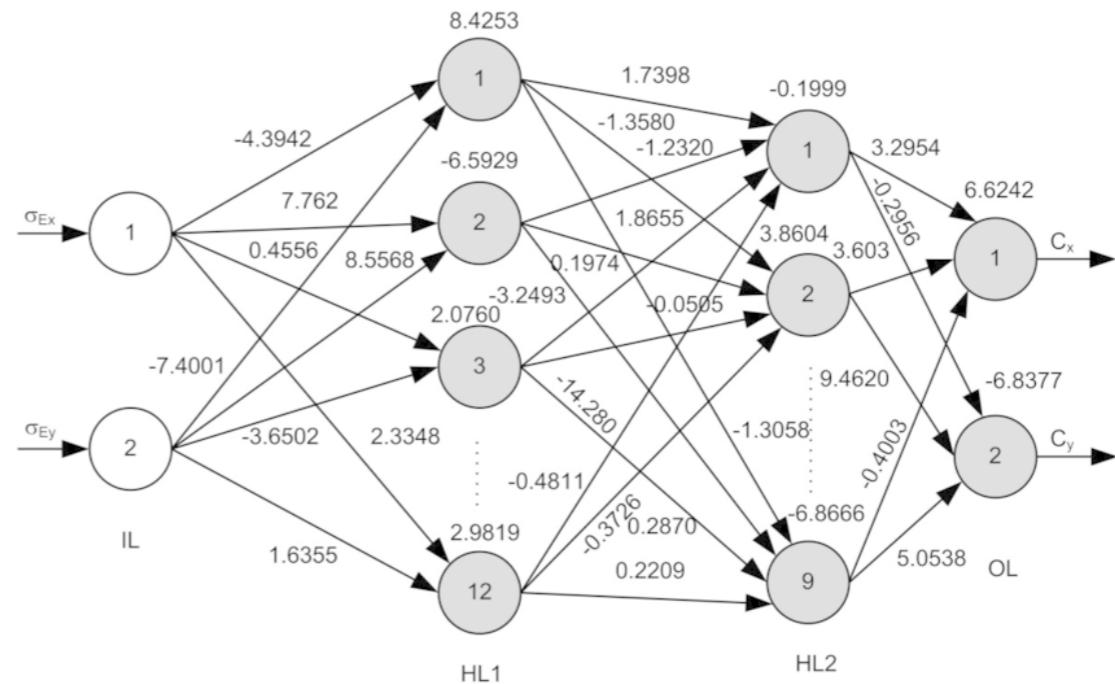
Any system that, given more data, performs  
increasingly better at some task

Framework / philosophy, not single method

# Software 1.0



# Software 2.0



Andrej Karpathy [Follow](#)

Director of AI at Tesla. Previously Research Scientist at OpenAI and PhD student at Stanford. I like to train deep neural nets on large datasets.

Nov 11, 2017 · 7 min read

<https://medium.com/@karpathy/software-2-0-a64152b37c35>

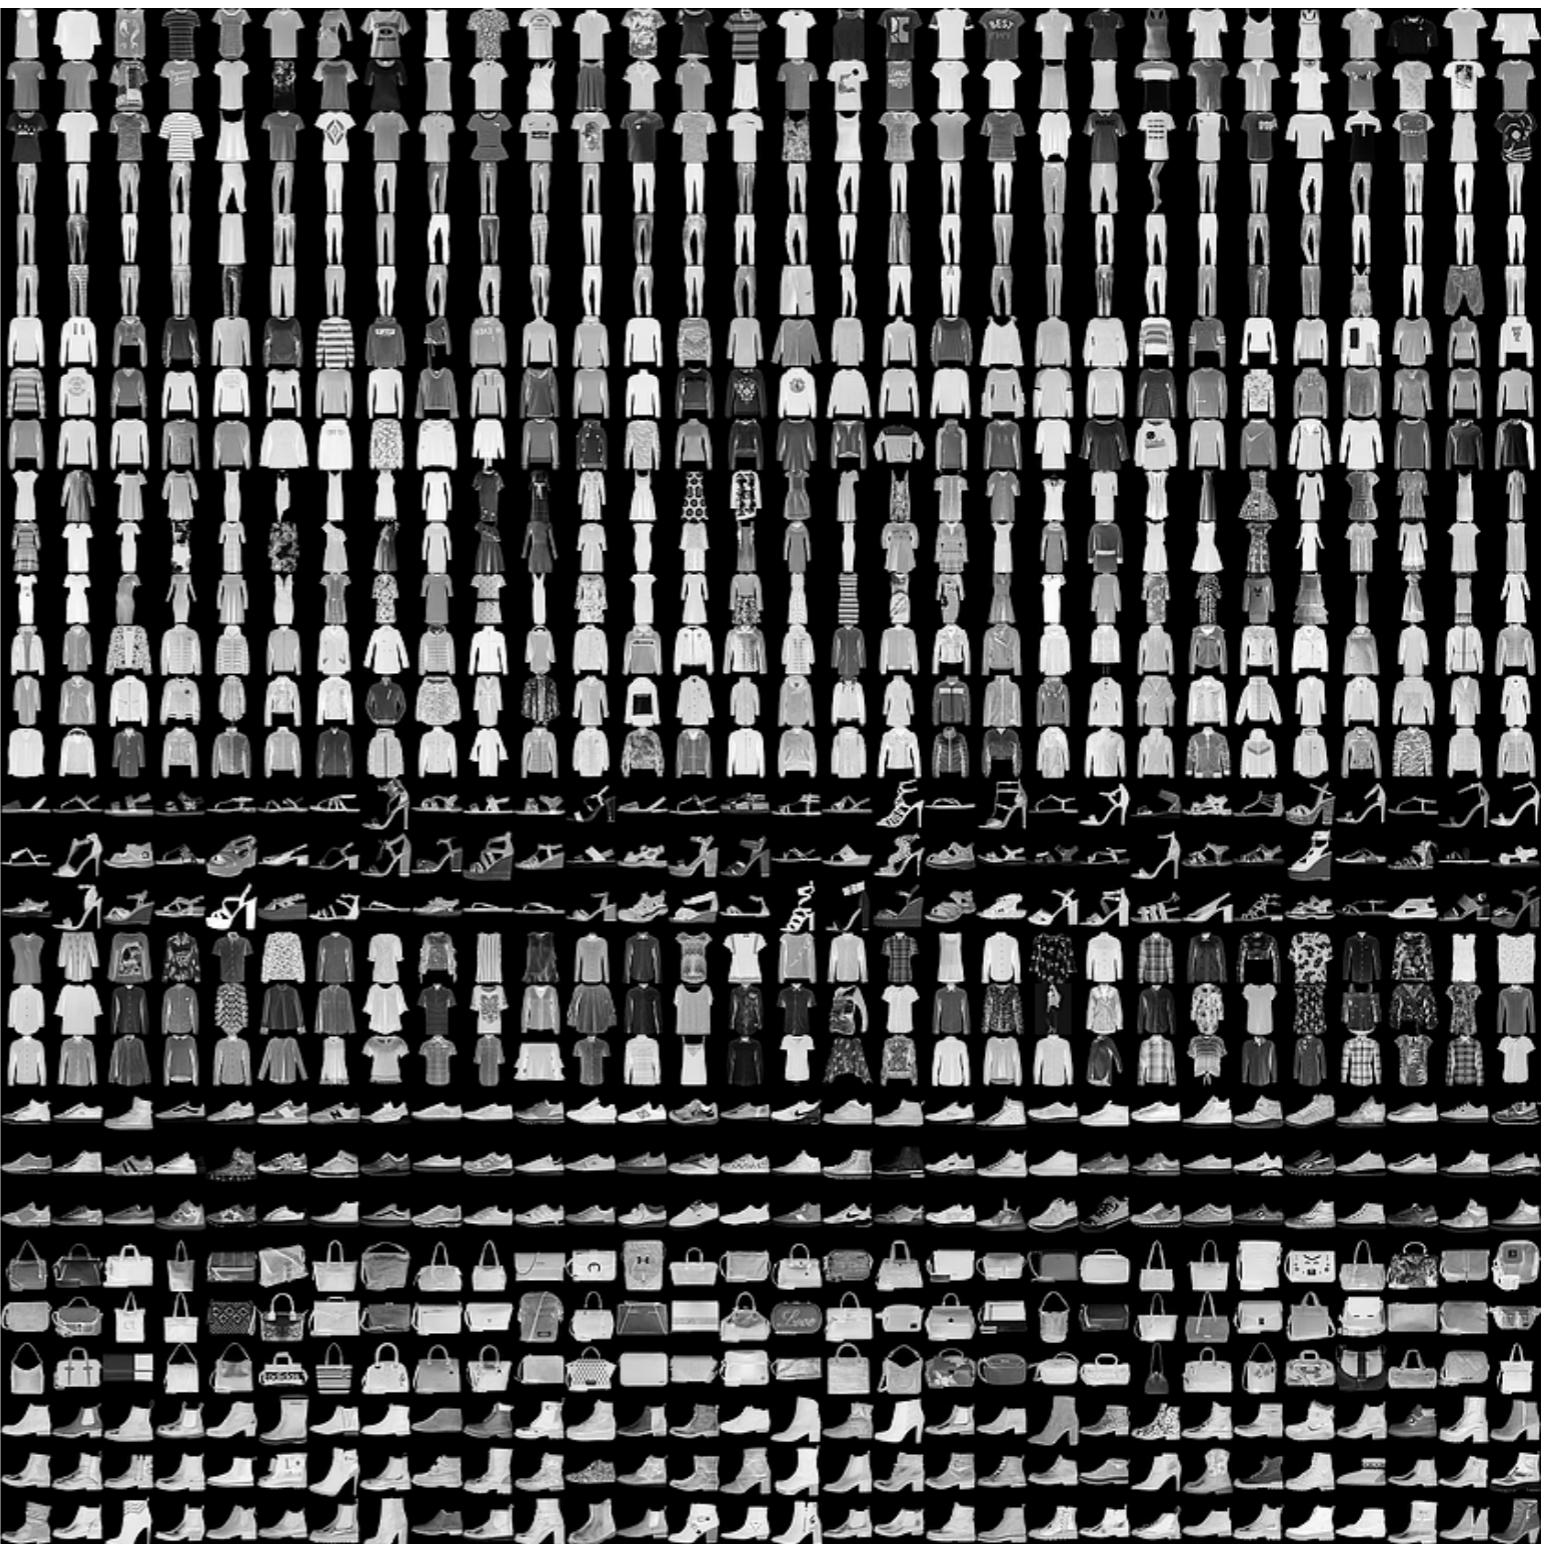
# **Basics of Machine Learning**

# Example of a Dataset – Fashion MNIST

10 categories (labels)

28x28 grayscale

70000 labeled images

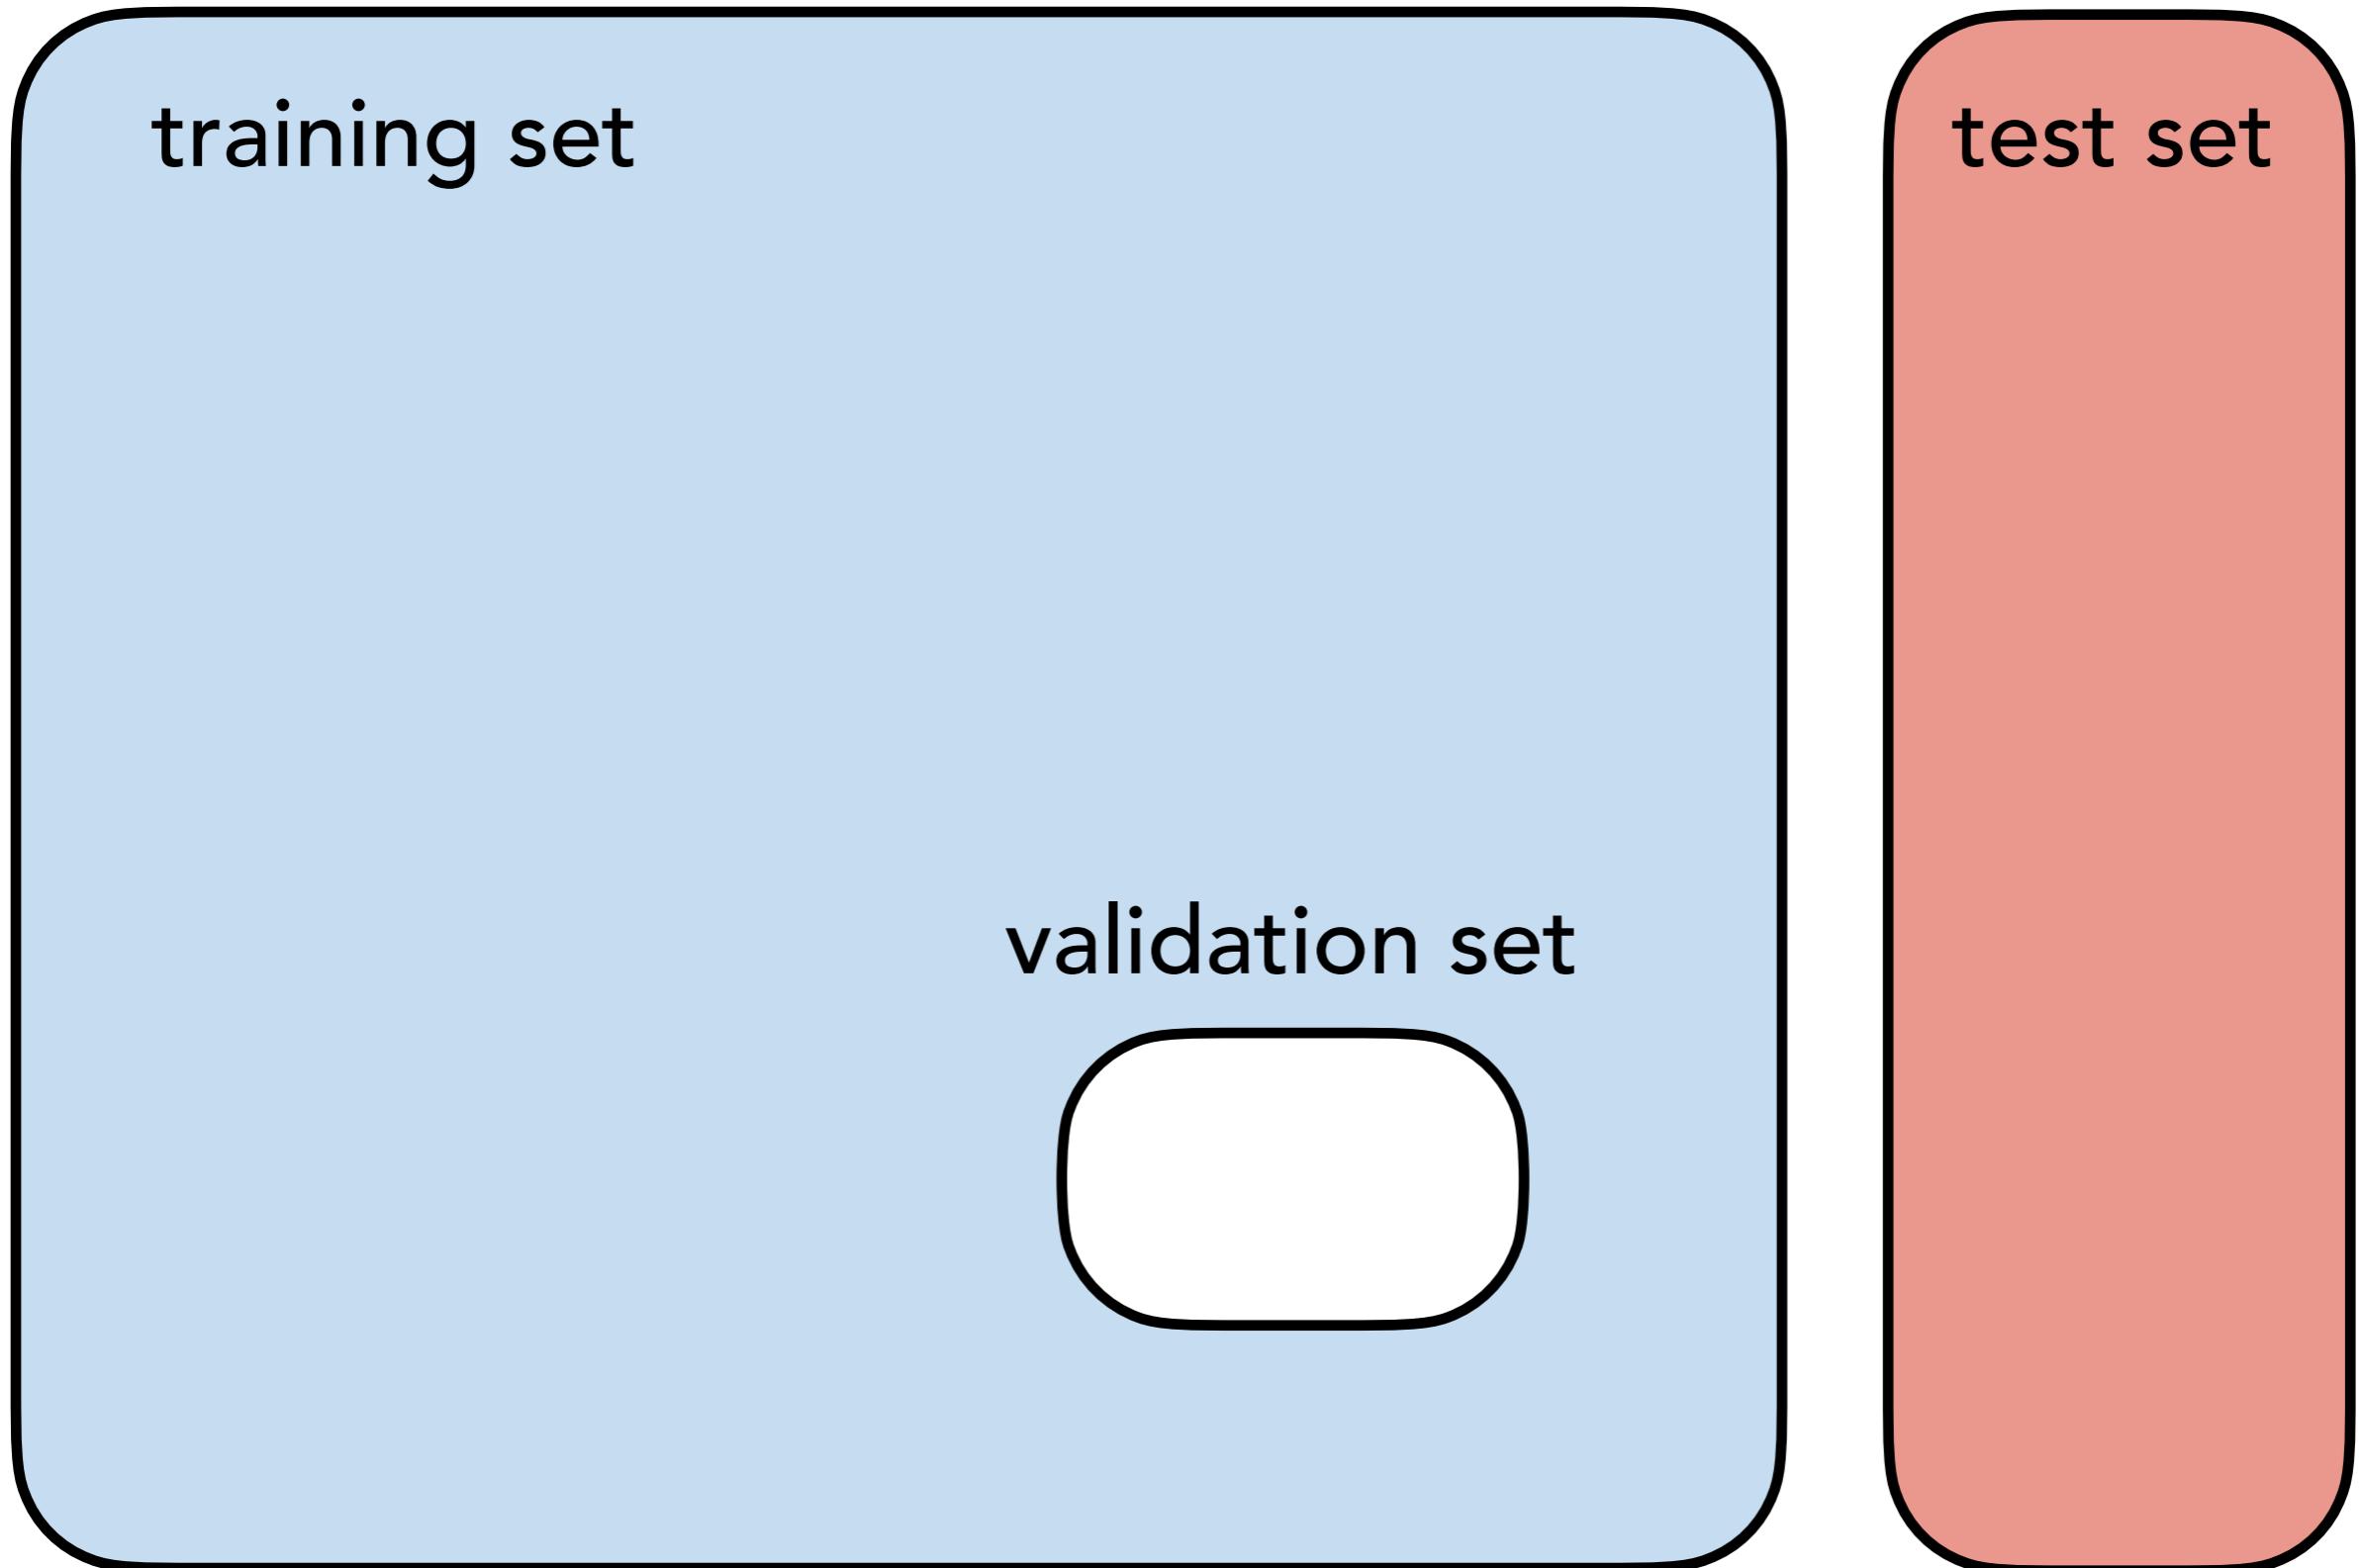


# Anatomy of a Dataset

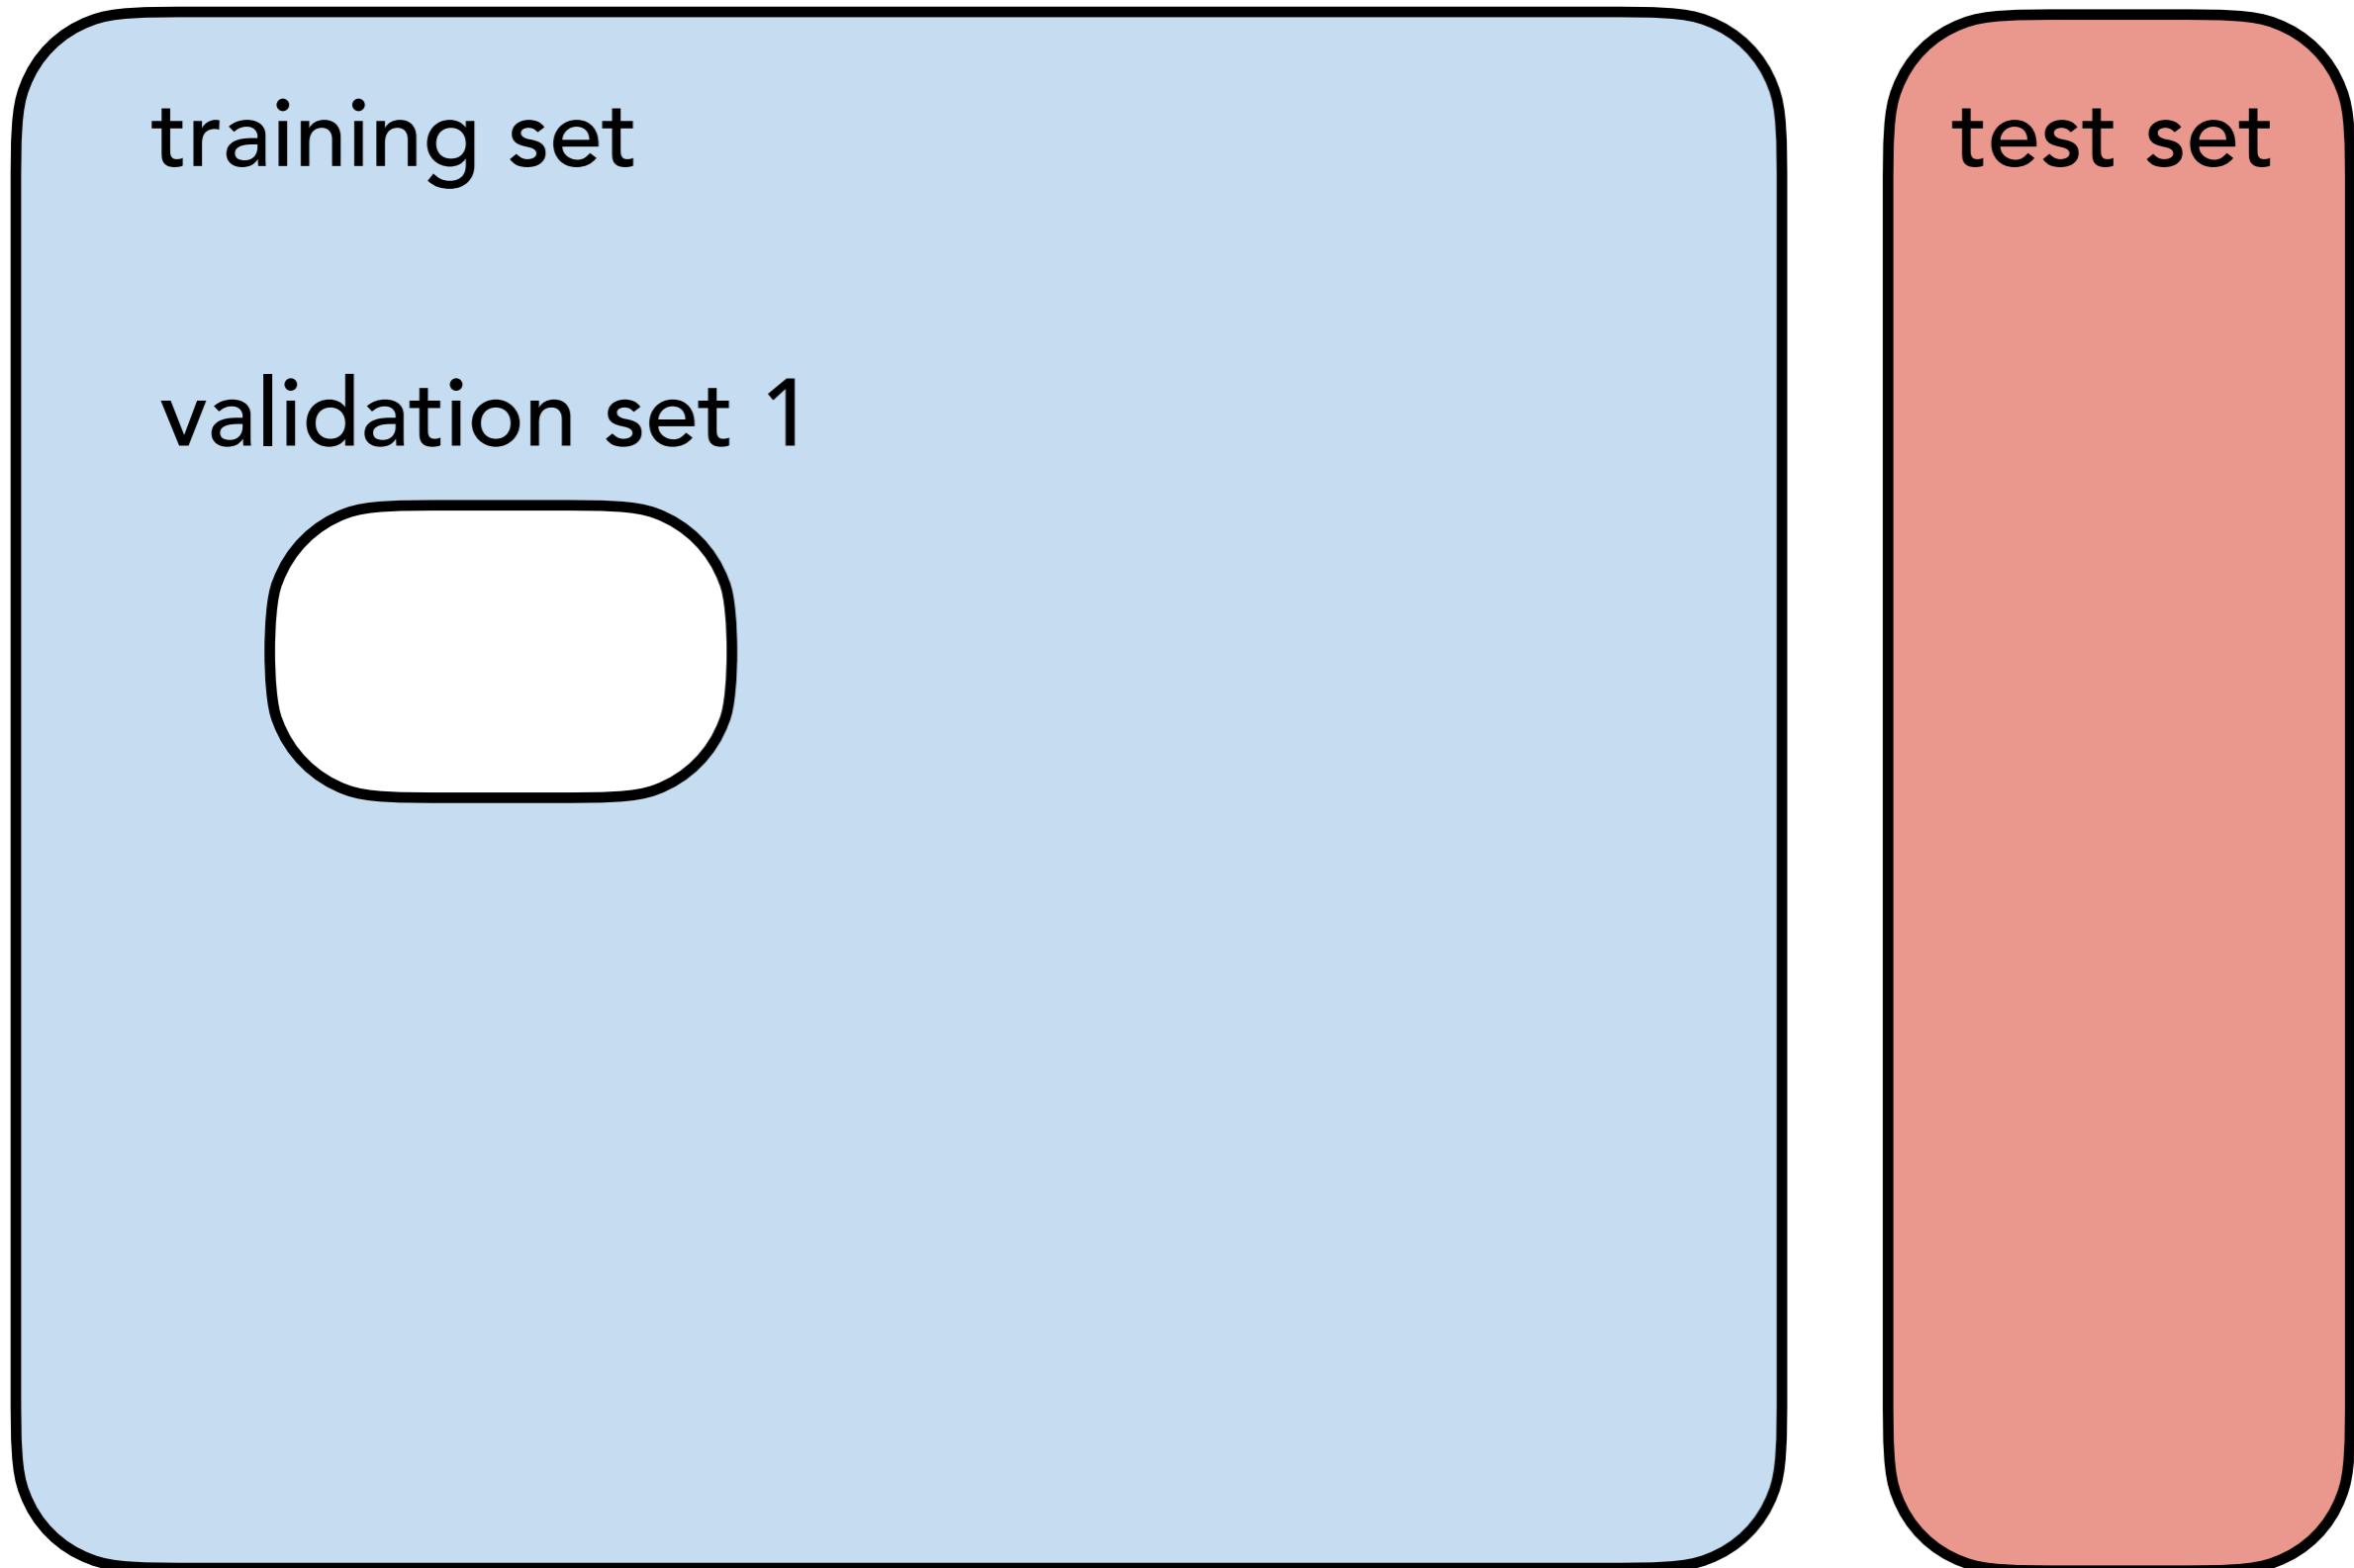
training set

test set

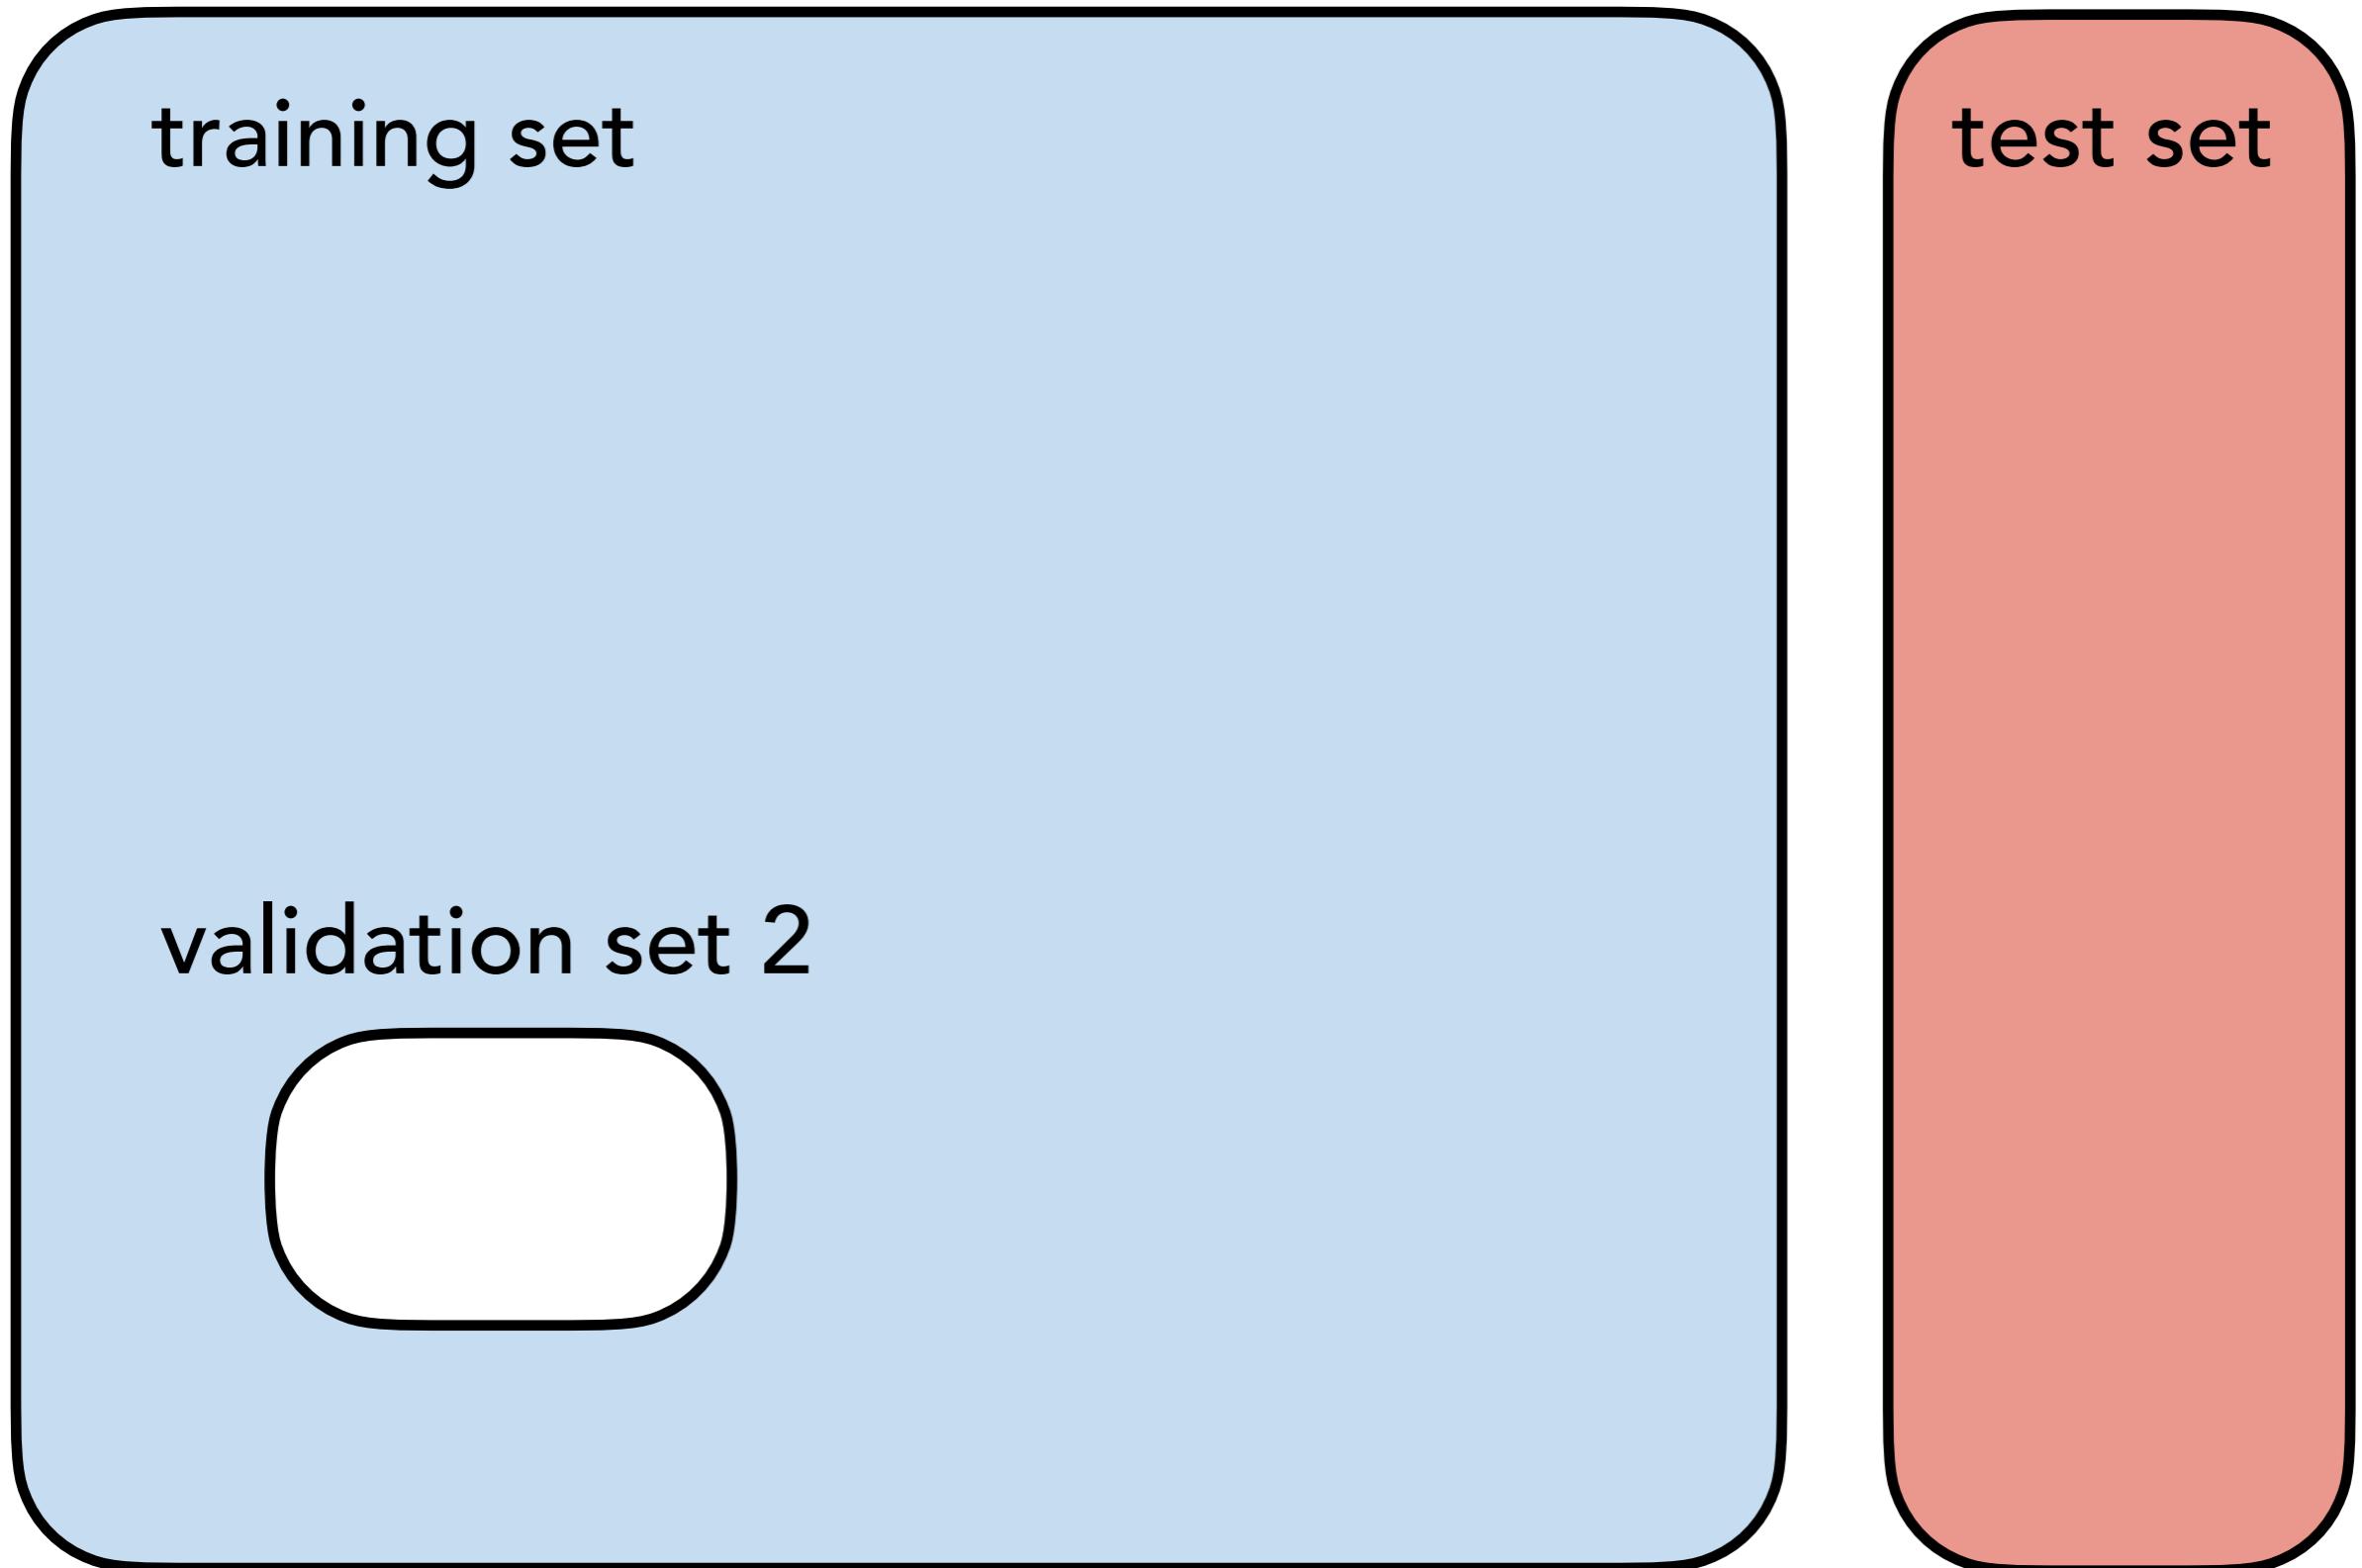
# Anatomy of a Dataset



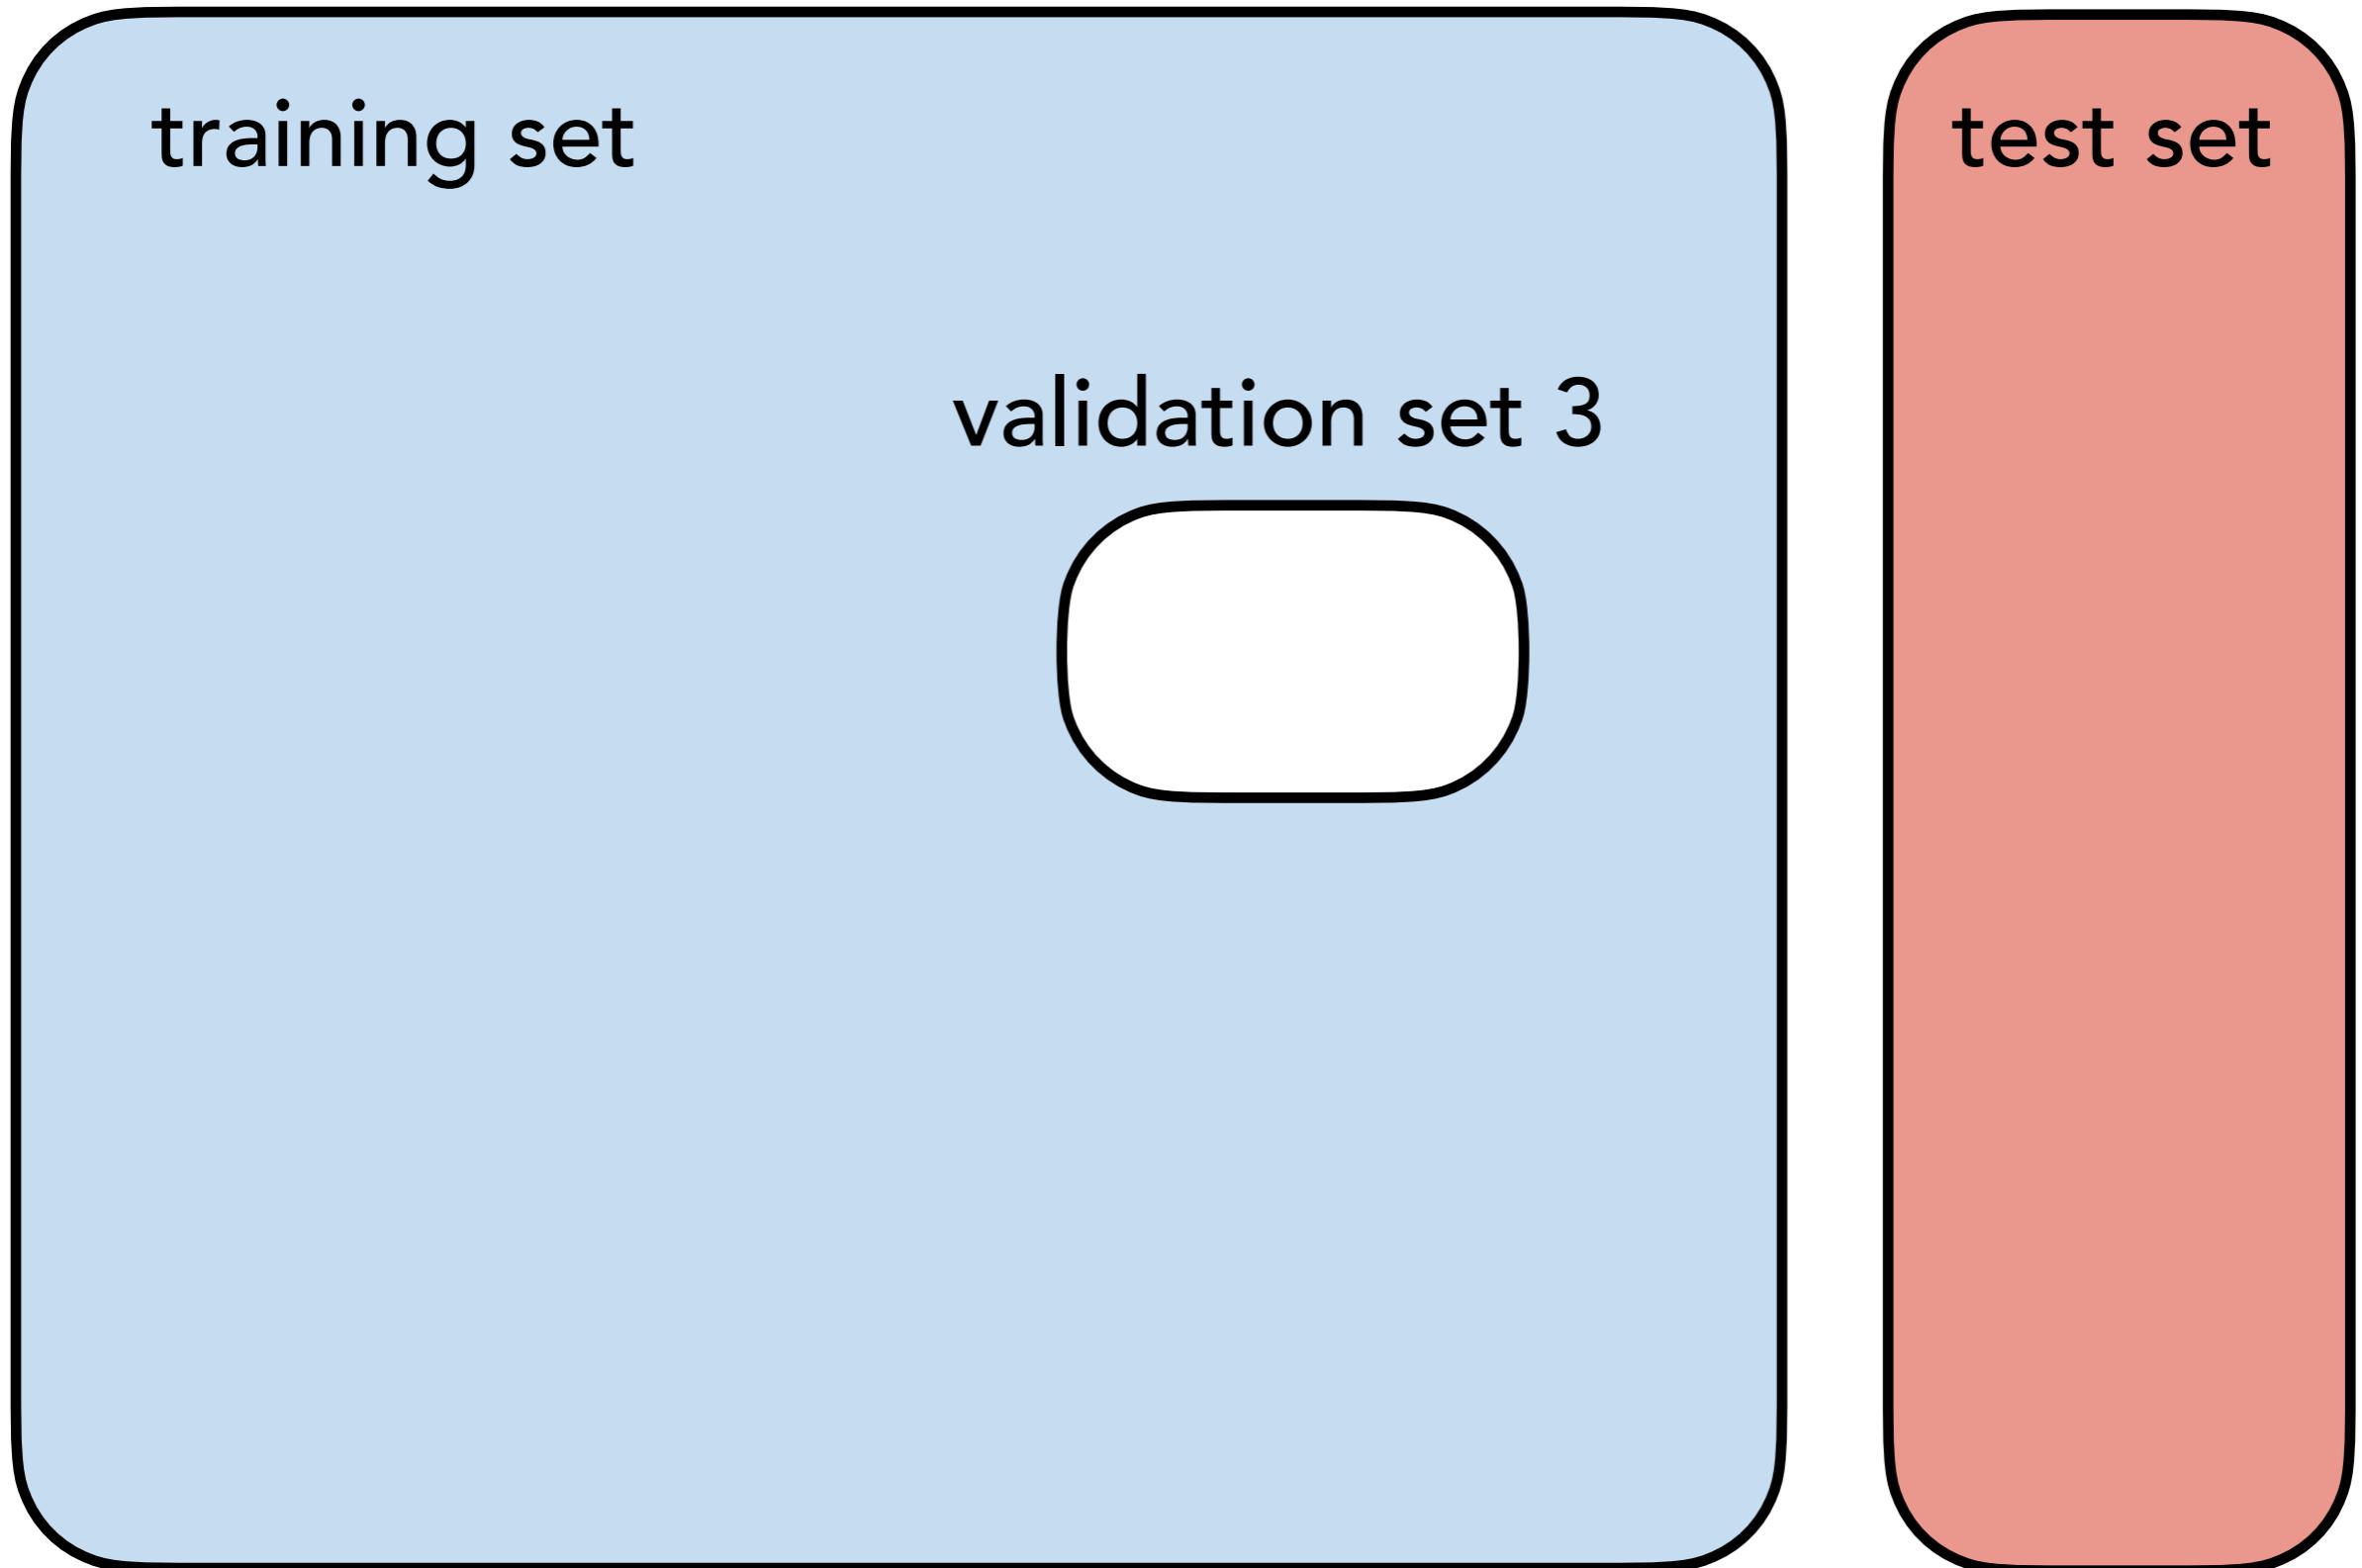
# Anatomy of a Dataset



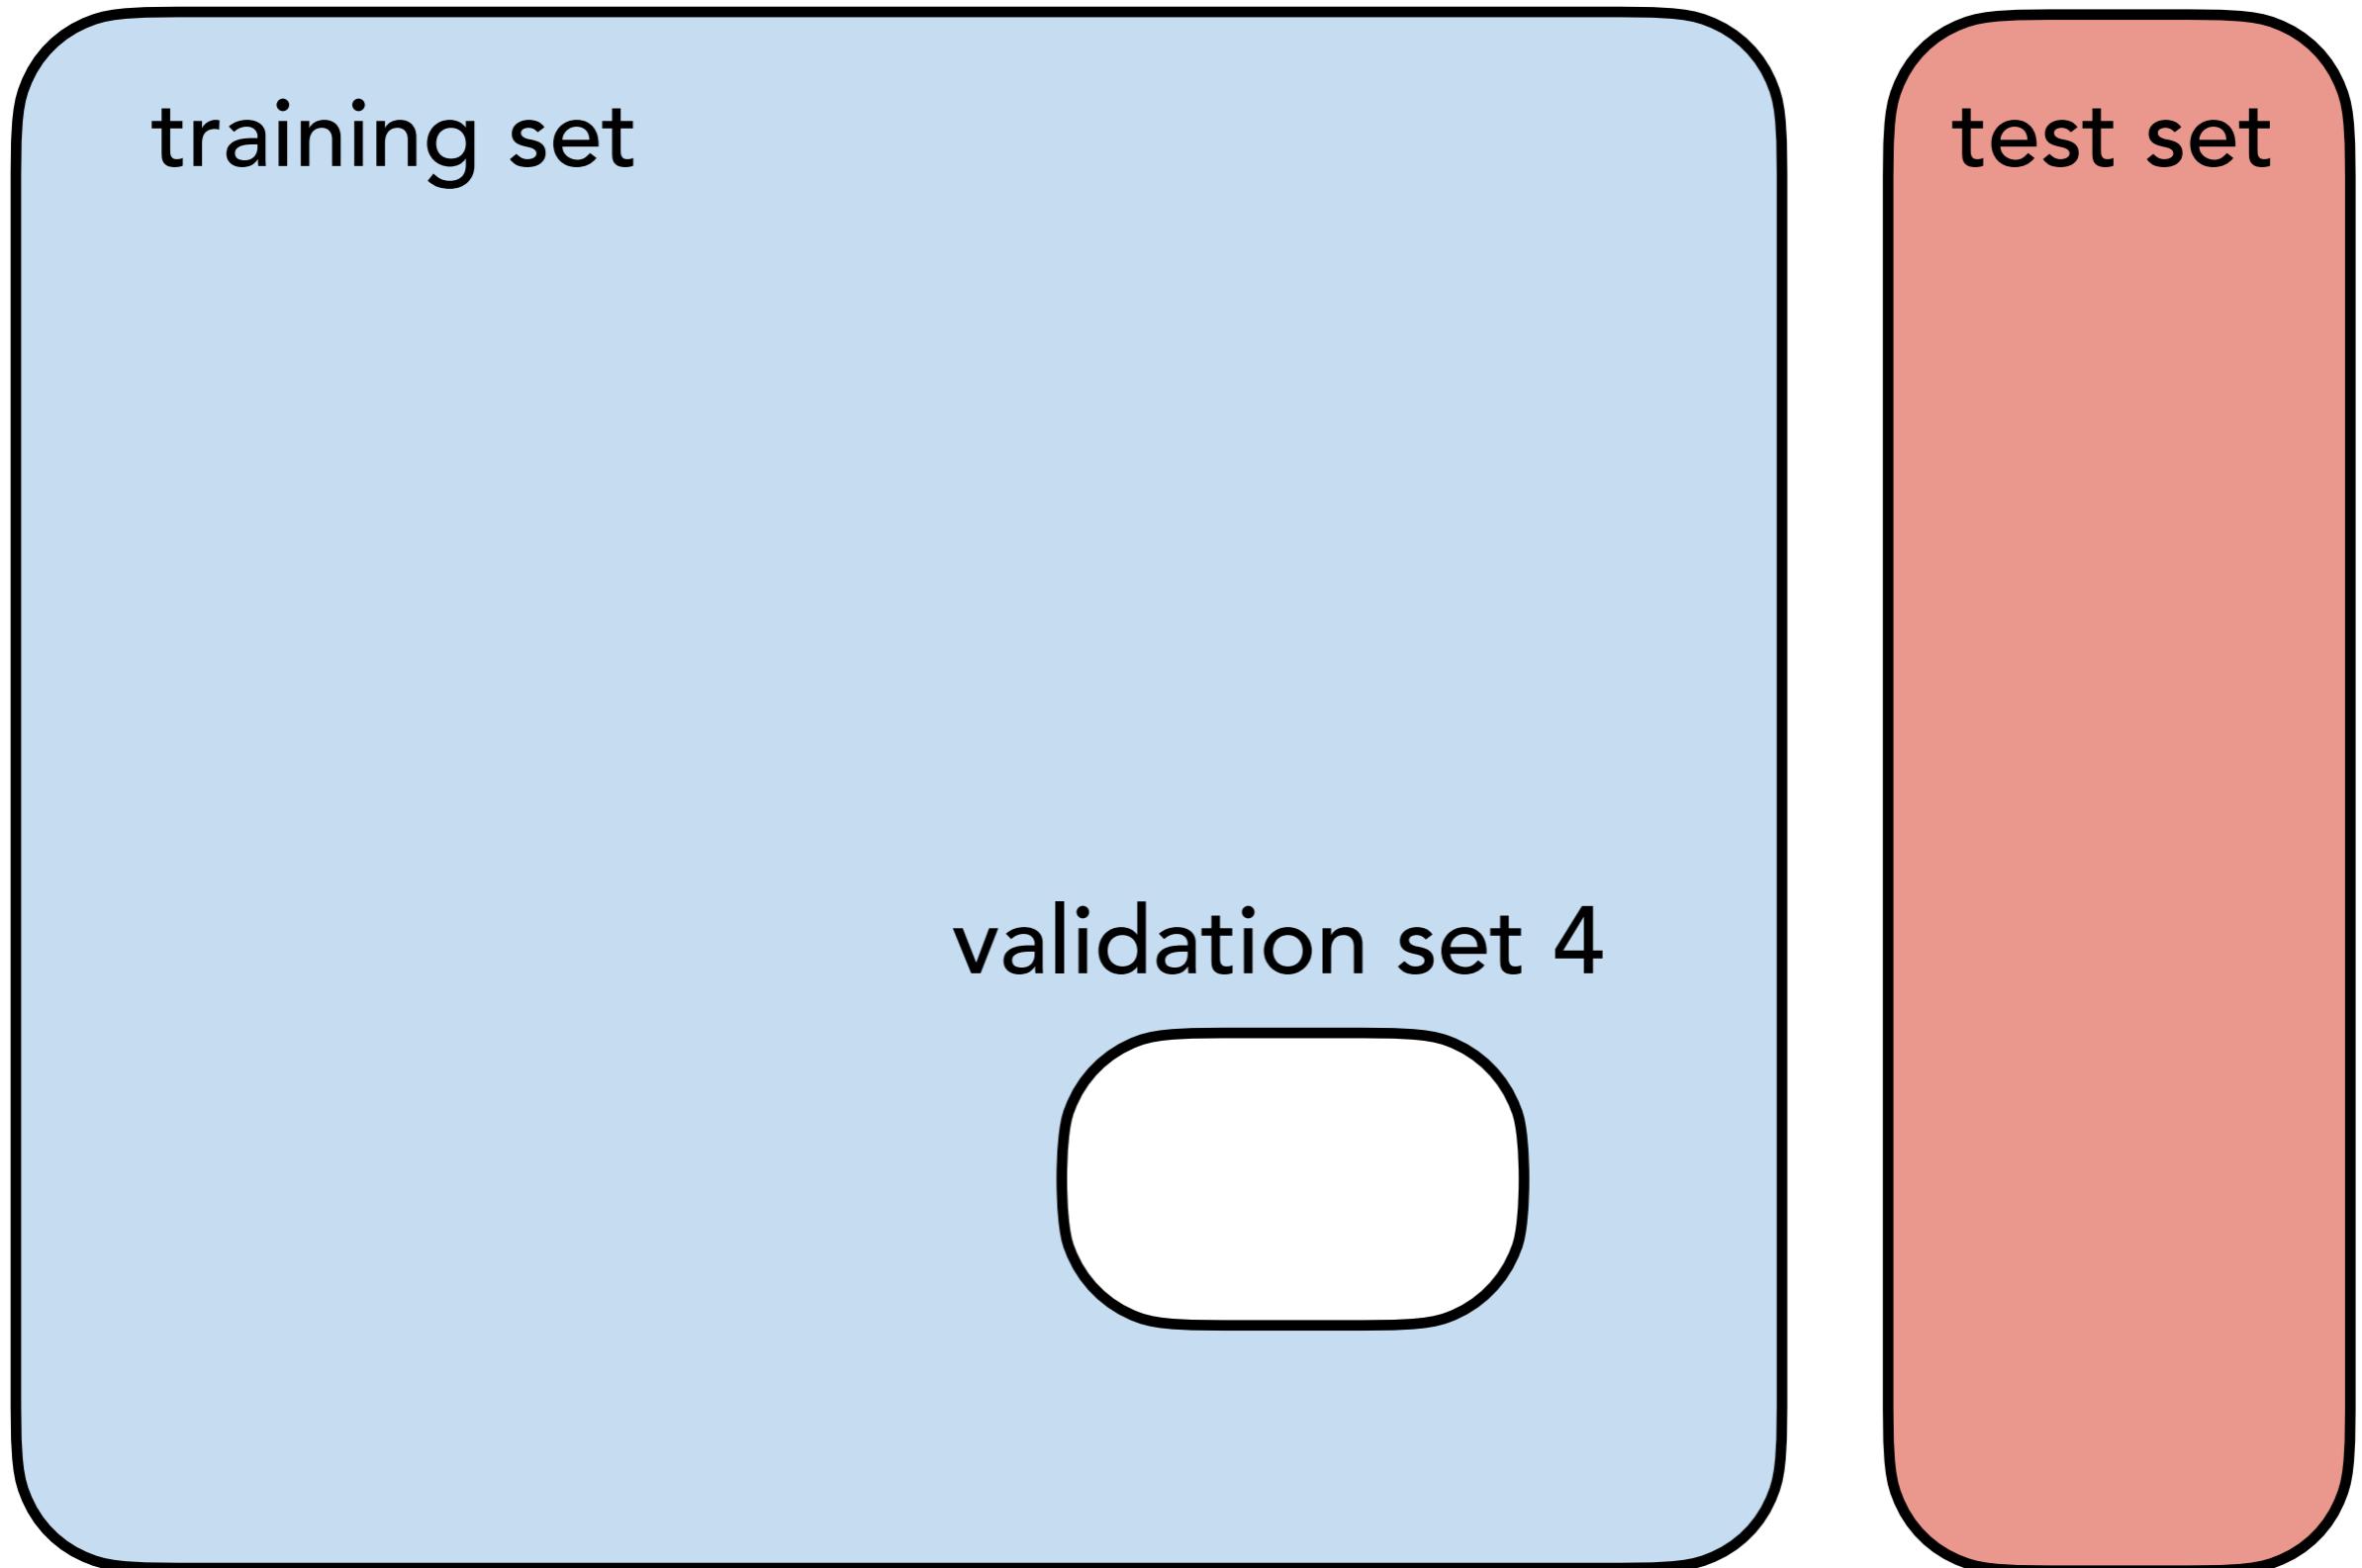
# Anatomy of a Dataset



# Anatomy of a Dataset



# Anatomy of a Dataset



# Types of learning tasks:

- Supervised learning (labeled data) *a priori* knowledge  
  
high
- Unsupervised learning (unlabeled data)
- Reinforcement learning (must collect data) low

# Supervised Learning

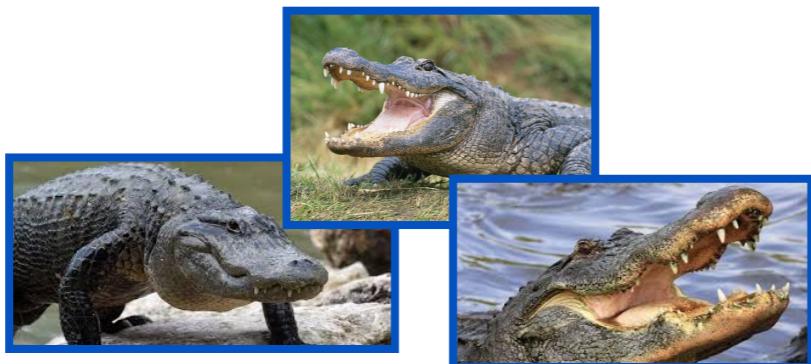
Given labeled training data (labels  $A$  and  $B$ )

Find *decision function*  $f(\mathbf{x})$

$$f(\mathbf{x}) > 0 \quad \mathbf{x} \in A$$

$$f(\mathbf{x}) < 0 \quad \mathbf{x} \in B$$

Example: identify photos of **alligators** and **bears**



# Supervised Learning

Typical strategy:

given training set  $\{\mathbf{x}_j, y_j\}$ , minimize cost function

$$C = \frac{1}{N_T} \sum_j (f(\mathbf{x}_j) - y_j)^2$$

$$y_j = \begin{cases} +1 & \mathbf{x}_j \in A \\ -1 & \mathbf{x}_j \in B \end{cases}$$

by varying adjustable params of  $f$

Cost function measures distance of trial function  $f(\mathbf{x}_j)$  from idealized "indicator" function  $y_j$

# Unsupervised Learning

Given unlabeled training data  $\{\mathbf{x}_j\}$

- Find function  $f(\mathbf{x})$  such that  $f(\mathbf{x}_j) \simeq p(\mathbf{x}_j)$
- Find function  $f(\mathbf{x})$  such that  $|f(\mathbf{x}_j)|^2 \simeq p(\mathbf{x}_j)$
- Find data clusters and which data belongs to each cluster
- Discover reduced representations of data for other learning tasks (e.g. supervised)

# General Philosophy of Machine Learning

- Solution to problem just some function  $y(\mathbf{x})$
- Parameterize very flexible functions  $f(\mathbf{x})$   
(prefer convenient over "correct")
- Of all  $f$  that come closest to  $y$  for training data,  
prefer the simplest  $f$



# Tensors in Machine Learning

*Where can tensors appear in  
machine learning applications?*

# Multi-Dimensional Data

## *Image Data*



leopard



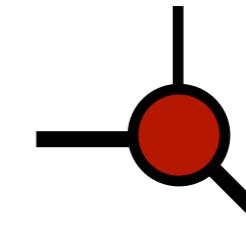
mushroom



grille

r,g,b value

x coord



y coord

## *Medical Data*

age

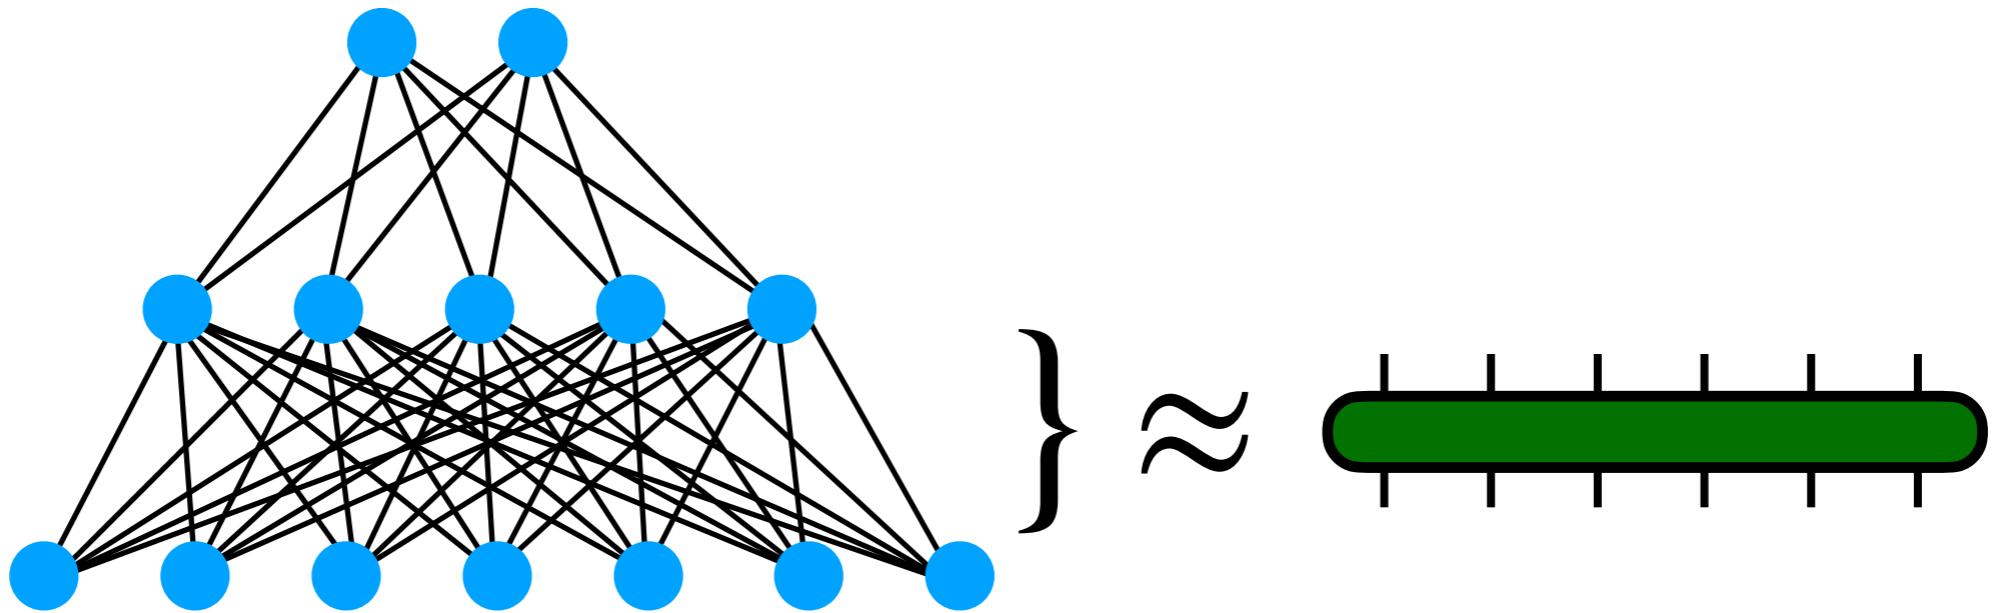
height

weight

symptom



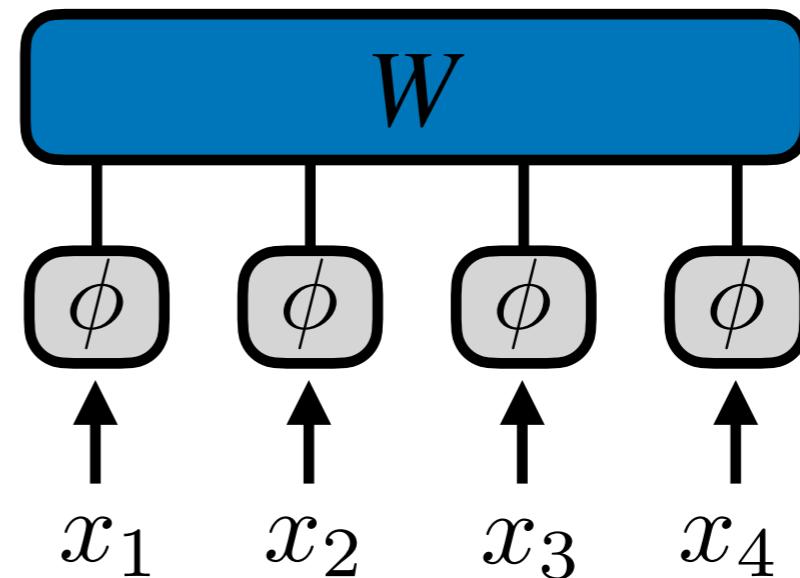
# Neural Network Weight Layers



Possible to interpret as a very high-order tensor  
(not just a matrix)

# Weights of Discriminative, Generative Models

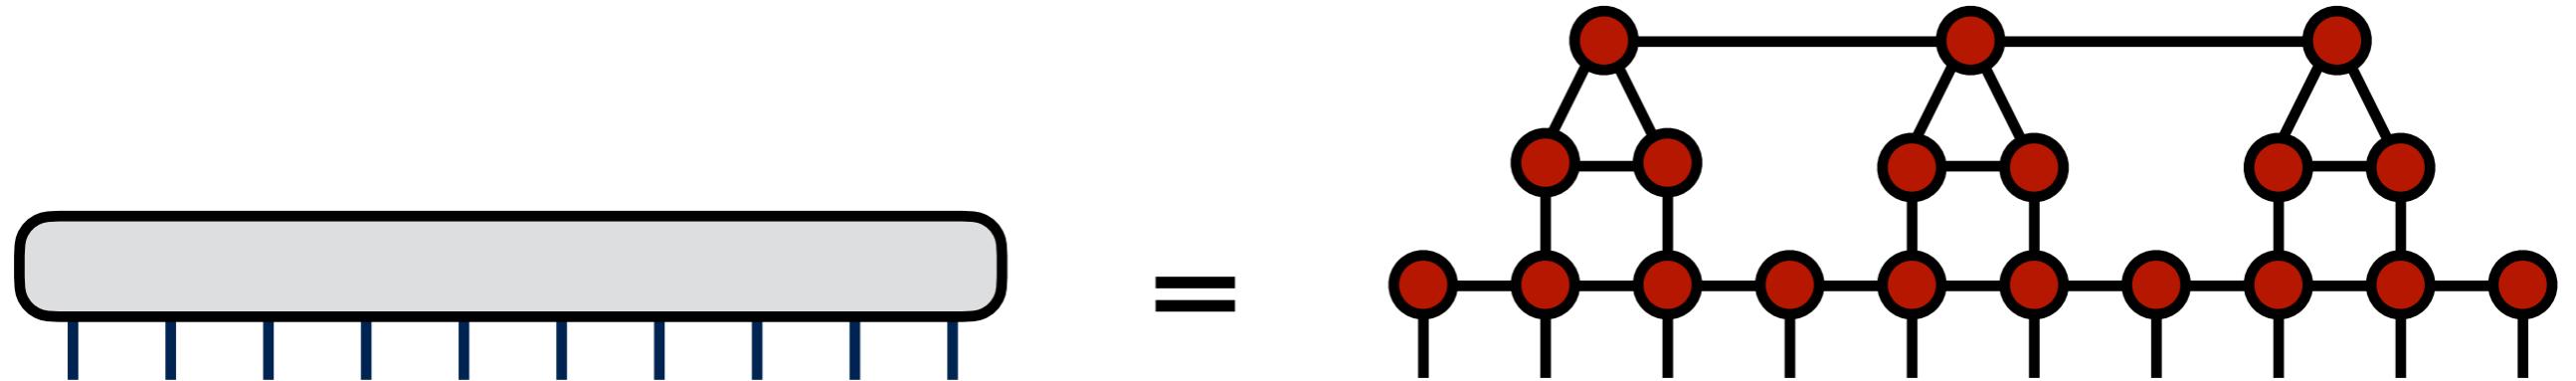
$$f(\mathbf{x}) = W \cdot \Phi(\mathbf{x})$$



For certain cases of kernel learning and Gaussian processes, weights are naturally a *high-order tensor*

# **Why Tensor Networks?**

Tensor network = factorization of huge tensor into contracted product of smaller tensors



Benefits:

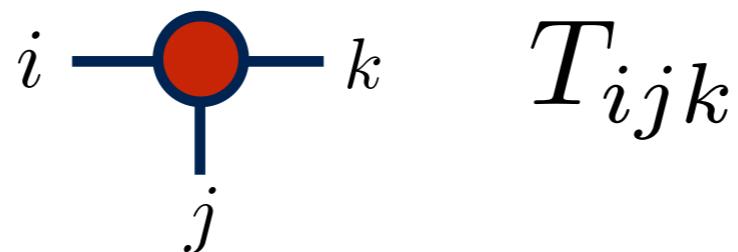
- exponential reduction in **memory** needed
- exponential speedup of **computations** (addition, product)
- theoretical **insight** and **interpretation**
- **estimation** of missing or corrupted entries
- many optimization **algorithms** & strategies

# Notation – Tensor Diagrams

N-index tensor represented as shape with N lines

$$T^{s_1 s_2 s_3 \cdots s_N} = \begin{array}{ccccccccc} s_1 & s_2 & s_3 & s_4 & \cdot & \cdot & \cdot & \cdot & s_N \\ | & | & | & | & | & | & | & | & | \\ \text{---} & \text{---} \end{array}$$


# Diagrams for low-order tensors



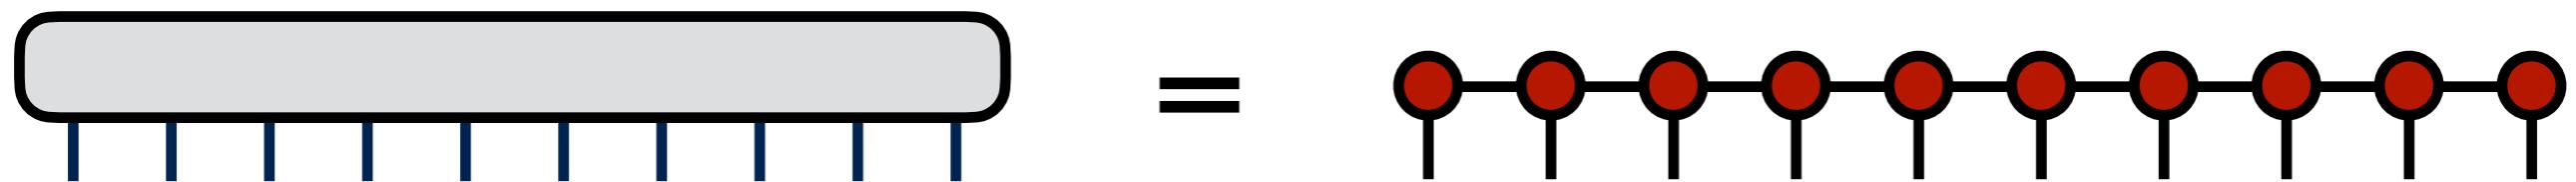
Joining lines implies contraction, can omit names

$$\begin{array}{c} \text{---} \\ | \quad | \\ i \quad j \end{array} \longleftrightarrow \sum_j M_{ij} v_j$$

$$\begin{array}{c} \text{---} \\ | \quad | \\ \text{green} \quad \text{orange} \end{array} \longleftrightarrow A_{ij} B_{jk} = AB$$

$$\begin{array}{c} \text{---} \\ | \quad | \\ \text{red} \quad \text{blue} \end{array} \longleftrightarrow A_{ij} B_{ji} = \text{Tr}[AB]$$

Best understood tensor network in physics is the  
***matrix product state (MPS)***<sup>1,2</sup>



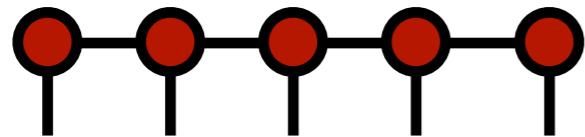
Known as ***tensor train*** in applied math literature<sup>3</sup>

[1] Östlund, Rommer, PRL 75, 3537 (1995)

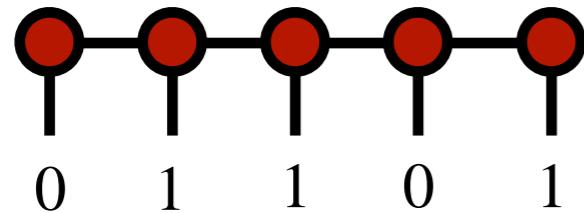
[2] Vidal, PRL 91, 147902 (2003)

[3] Oseledets, SIAM J. Sci. Comp. 33, 2295 (2011)

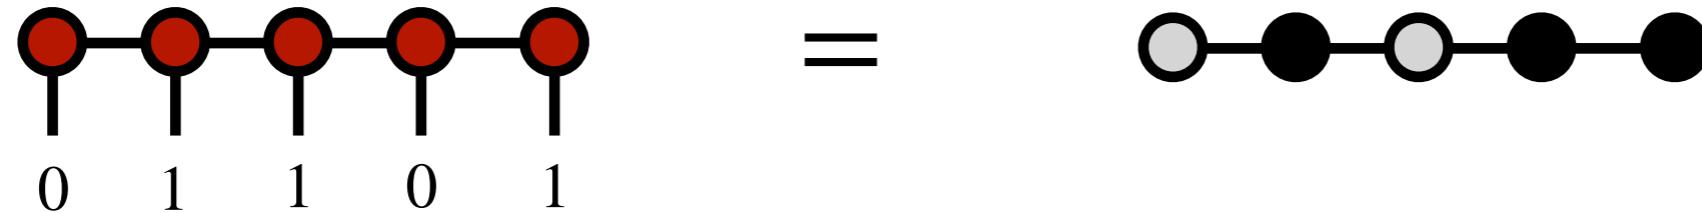
Name matrix product state refers to retrieving elements:



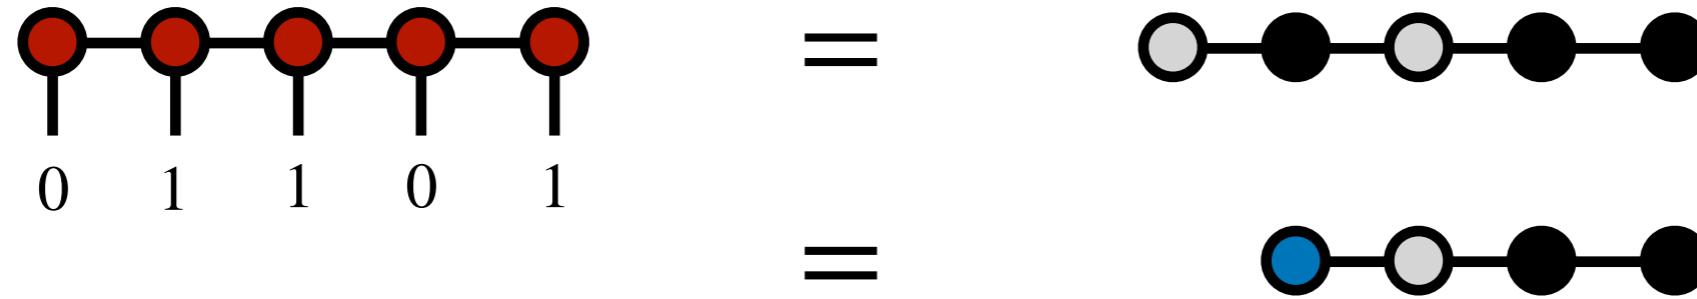
Name matrix product state refers to retrieving elements:



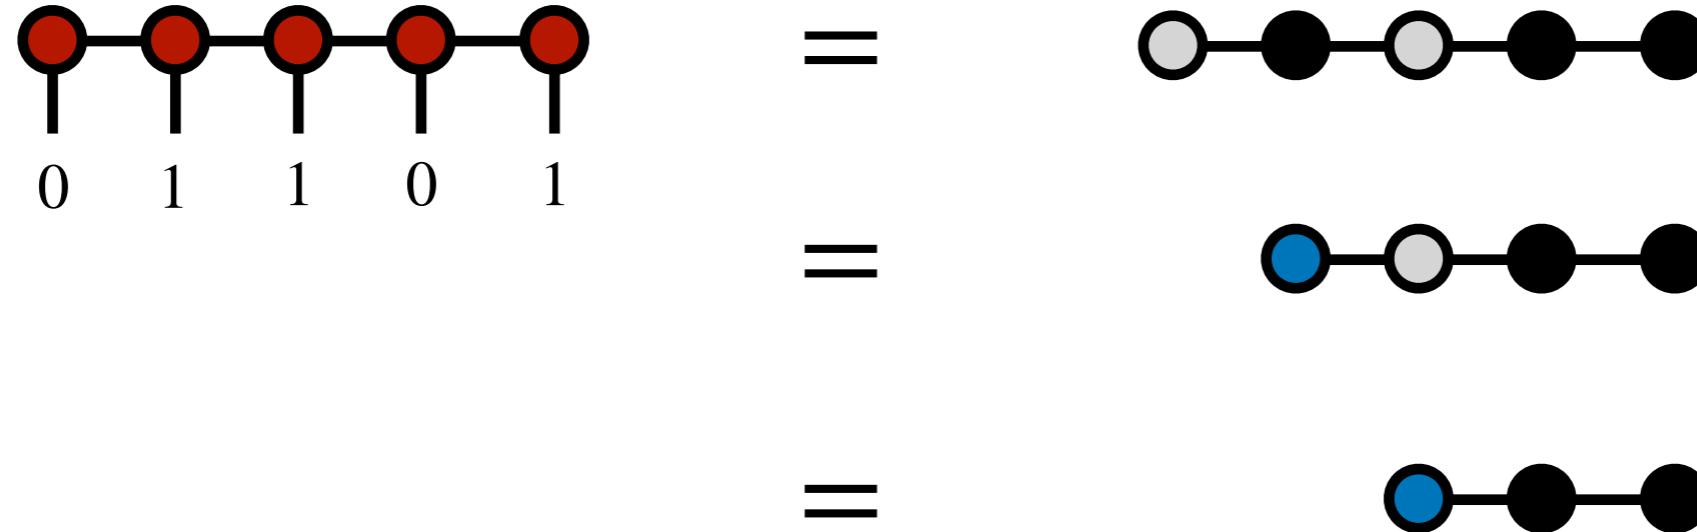
Name matrix product state refers to retrieving elements:



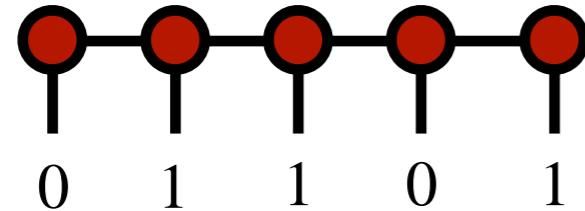
Name matrix product state refers to retrieving elements:



Name matrix product state refers to retrieving elements:



Name matrix product state refers to retrieving elements:



=



=



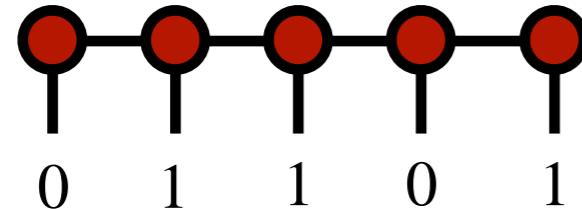
=



=



Name matrix product state refers to retrieving elements:



=



=



=



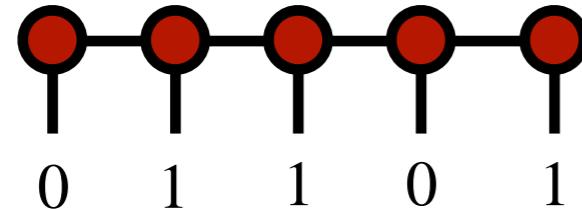
=



=



Name matrix product state refers to retrieving elements:



=



=



=



=

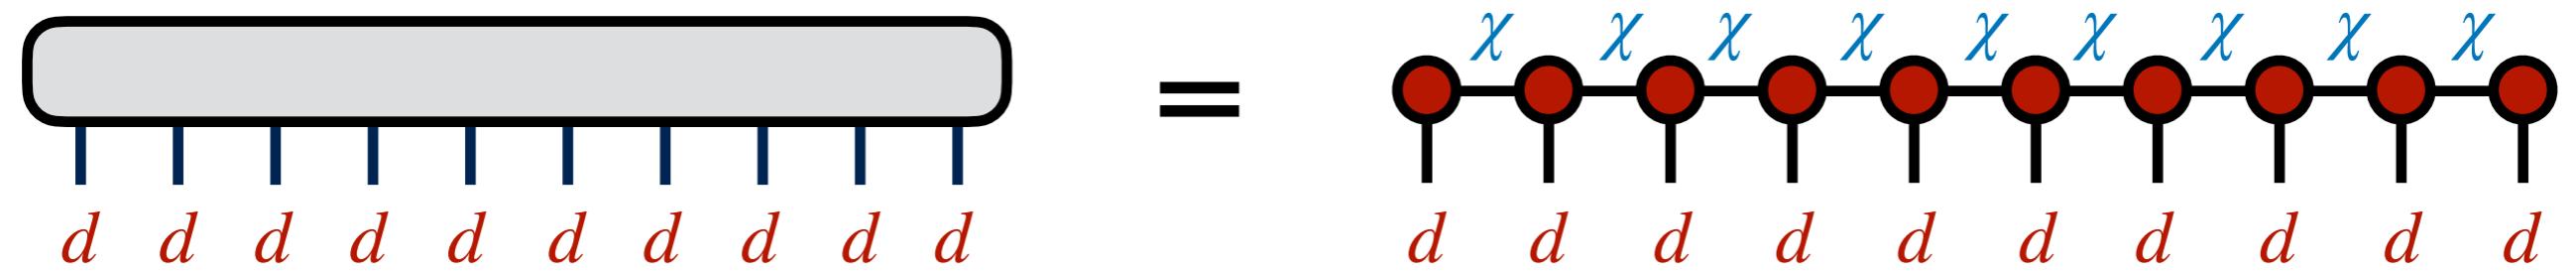


=



$\approx T^{01101}$

Adjustable parameter of matrix product state (MPS) is bond dimension  $\chi$



If modest  $\chi$  yields good approximation,  
obtain massive compression:

$$d^N \longrightarrow N d \chi^2$$

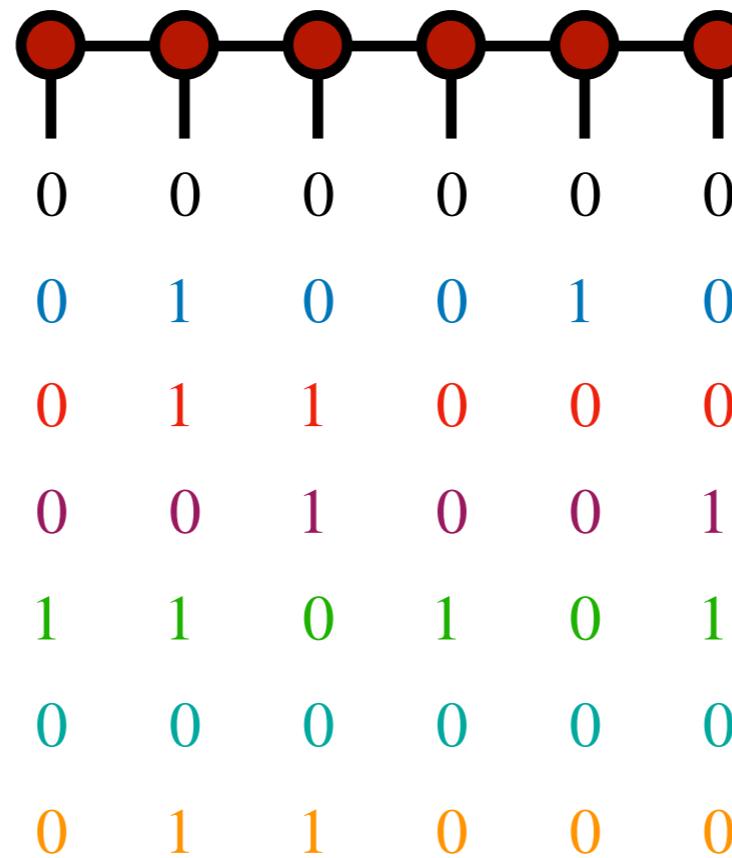
Can efficiently sum MPS compressed form:

$$\begin{array}{c} \text{Red circles} \\ \text{Blue circles} \end{array} + = \begin{array}{c} \text{Purple circles} \end{array}$$

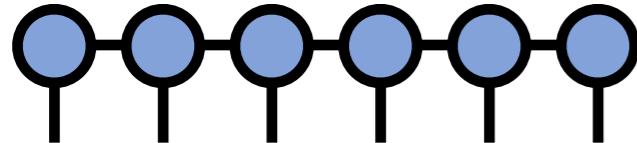
Or multiply by other networks:

$$\begin{array}{c} \text{Red circles} \\ \text{Grey circles} \end{array} = \begin{array}{c} \text{Teal circles} \end{array}$$

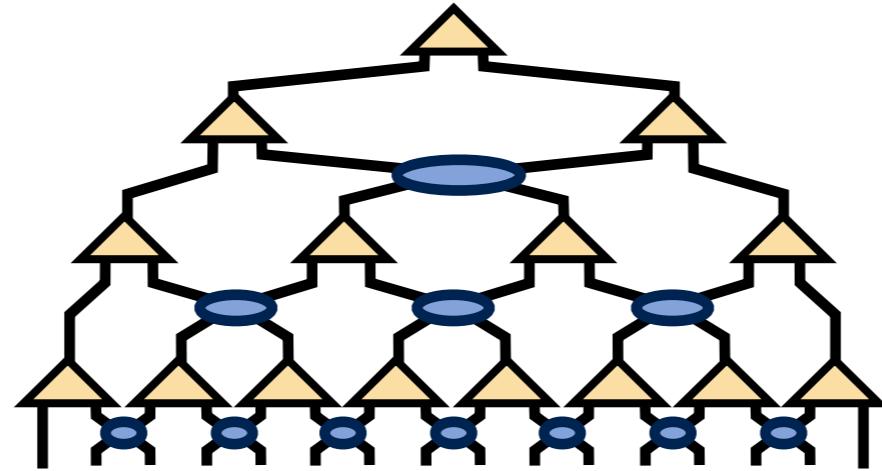
MPS can be perfectly sampled (no autocorrelation):



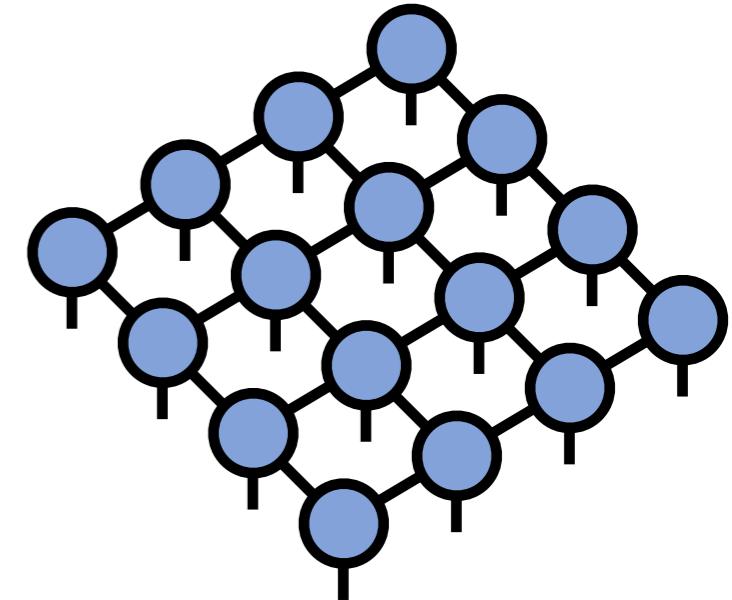
In quantum physics context, have rich theory of which tensor networks are suited for particular "data"



1D system



1D *critical* system

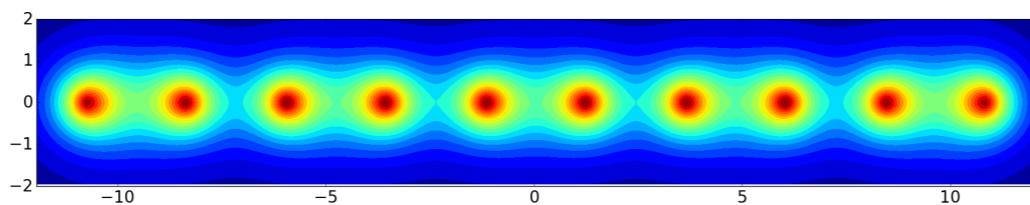


2D system

(Here "data" = samples/measurements of a quantum wavefunction)

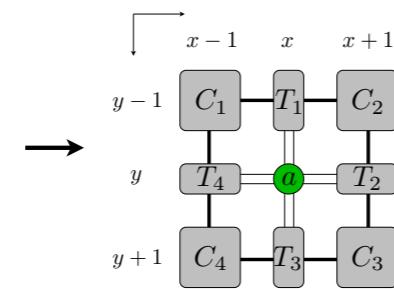
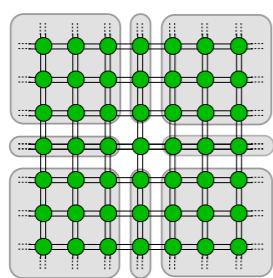
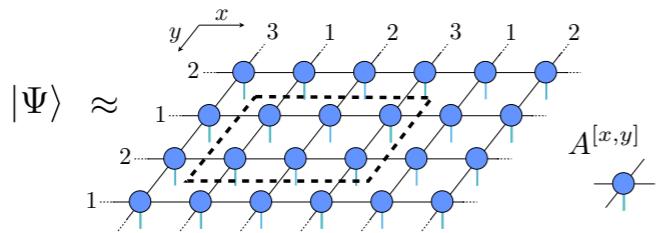
Last but not least, can pull off impressive calculations

*capture wavefunctions of 1000's of atoms*

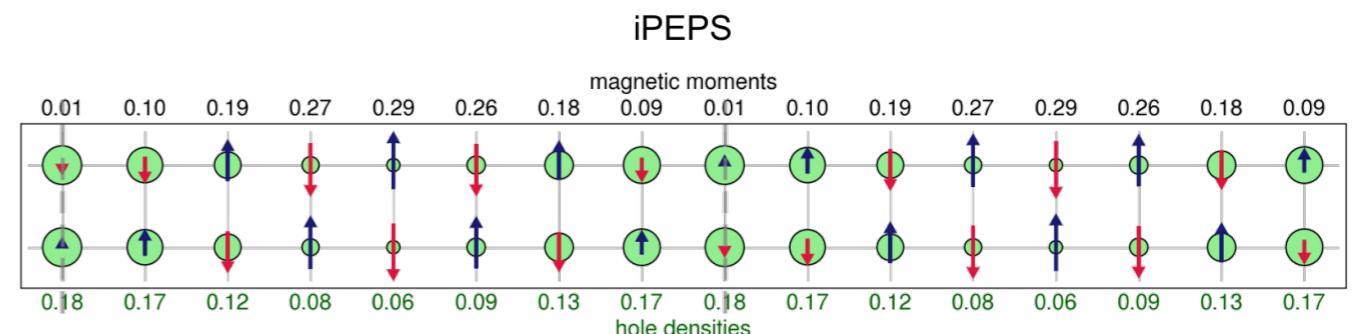
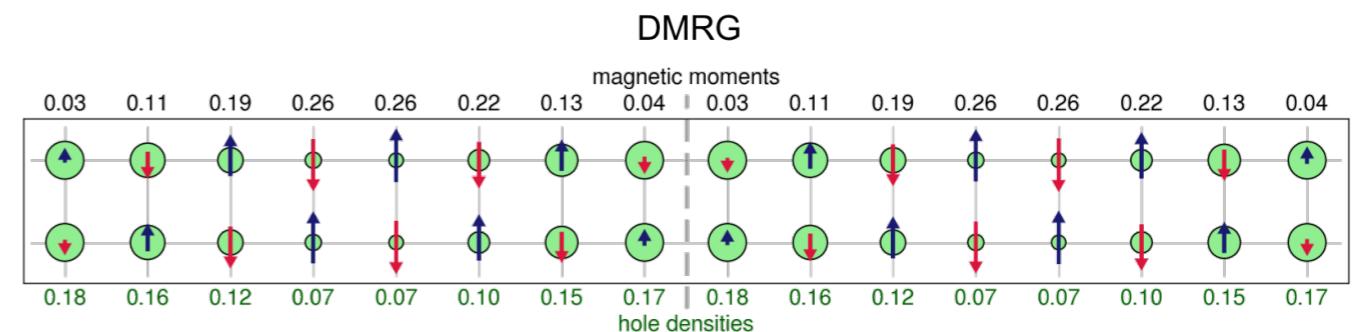


Stoudenmire, White, PRL 119, 046401 (2017)

*study challenging, correlated electron models in 2D*



Corboz, PRB 94, 035133 (2016)



Zheng et al., Science 358, 1155 (2017)

# Are tensor networks only useful for wavefunctions? Why not other tensors too?

Actually, since late 2000's, applied math community exploring tensor networks (Hackbusch, Oseledets,...)

Comparative study of Discrete Wavelet Transforms and Wavelet Tensor Train decomposition to feature extraction of FTIR data of medicinal plants

Pavel Kharyuk<sup>a,b</sup>, Dmitry Nazarenko<sup>c,\*</sup>, Ivan Oseledets<sup>a,d</sup>

Principal Component Analysis with Tensor Train Subspace

Wenqi Wang, Vaneet Aggarwal, and Shuchin Aeron

A Tensor-Train Deep Computation Model for Industry Informatics Big Data Feature Learning

4 Author(s)

Qingchen Zhang ; Laurence T. Yang  ; Zhikui Chen  ; Peng Li [View All Authors](#)

# Tensor networks a general tool for linear algebra in exponentially high-dimensional spaces



Notions like entanglement entropy correspond to multilinear tensor rank

# MPS and other tensor networks already being explored for machine learning:

## *supervised learning*

0 0 0 0 0 0 0 0 0  
1 1 1 1 1 1 1 1 1  
2 2 2 2 2 2 2 2 2  
3 3 3 3 3 3 3 3 3  
4 4 4 4 4 4 4 4 4  
5 5 5 5 5 5 5 5 5  
6 6 6 6 6 6 6 6 6  
7 7 7 7 7 7 7 7 7  
8 8 8 8 8 8 8 8 8  
9 9 9 9 9 9 9 9 9



Novikov, Trofimov, Oseledets, arxiv:1605.03795

Stoudenmire, Schwab, *Advances in N.I.P.S.*, **29**, 4799

I. Glasser, N. Pancotti, J.I. Cirac, arxiv:1806.05964

Efthymiou, et al., arxiv:1906.06329

Selvan, et al., arxiv:2004.10076

## *unsupervised learning*

$$p(\mathbf{x}) = \left| \begin{array}{cccccc} x_1 & x_2 & x_3 & x_4 & x_5 & x_6 \end{array} \right|^2$$

Bengua, Phien, Tuan, 10.1109/BigDataCongress.2015.105 (2015)

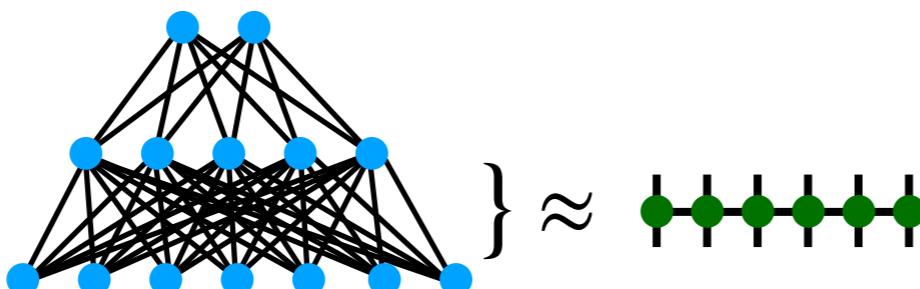
Cichocki et al, [dx.doi.org/10.1561/2200000067](https://dx.doi.org/10.1561/2200000067) (2017)

Han, et al, *Phys. Rev. X* 8, 031012 (2018)

Stokes, Terilla, arxiv:1902.06888

Miller, et al., arxiv:2003.01039

## *compressing neural nets*



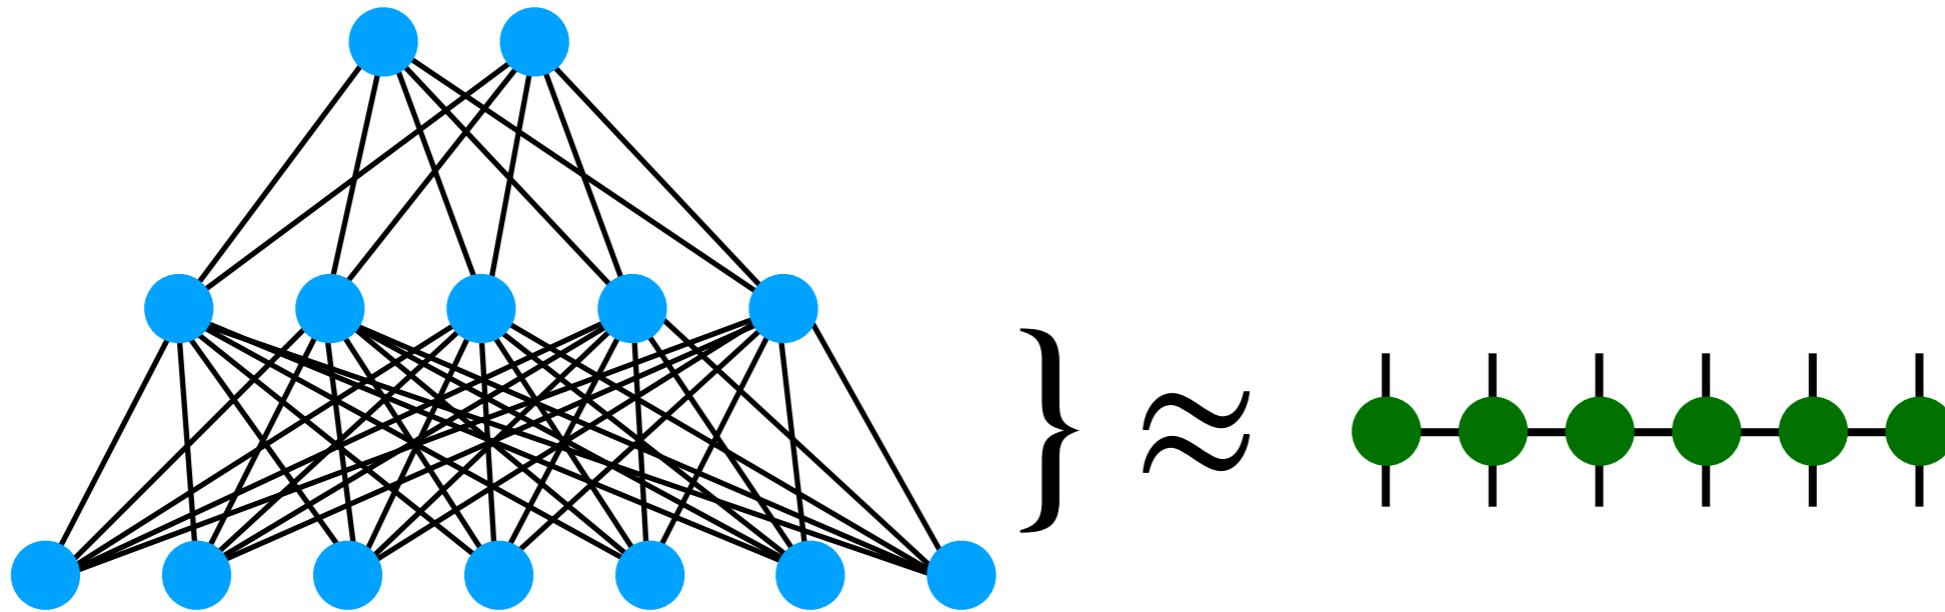
Novikov et al., *Advances in NIPS* (2015), arxiv:1509.06569

Garipov, Podoprikhin, Novikov, arxiv:1611.03214

Yu, Zheng, Anandkumar, , arxiv:1711.00073

*...and many others*

# Compressing Neural Network Weight Layers



Novikov et al., *Advances in Neural Information Processing* (2015) ([arxiv:1509.06569](https://arxiv.org/abs/1509.06569))

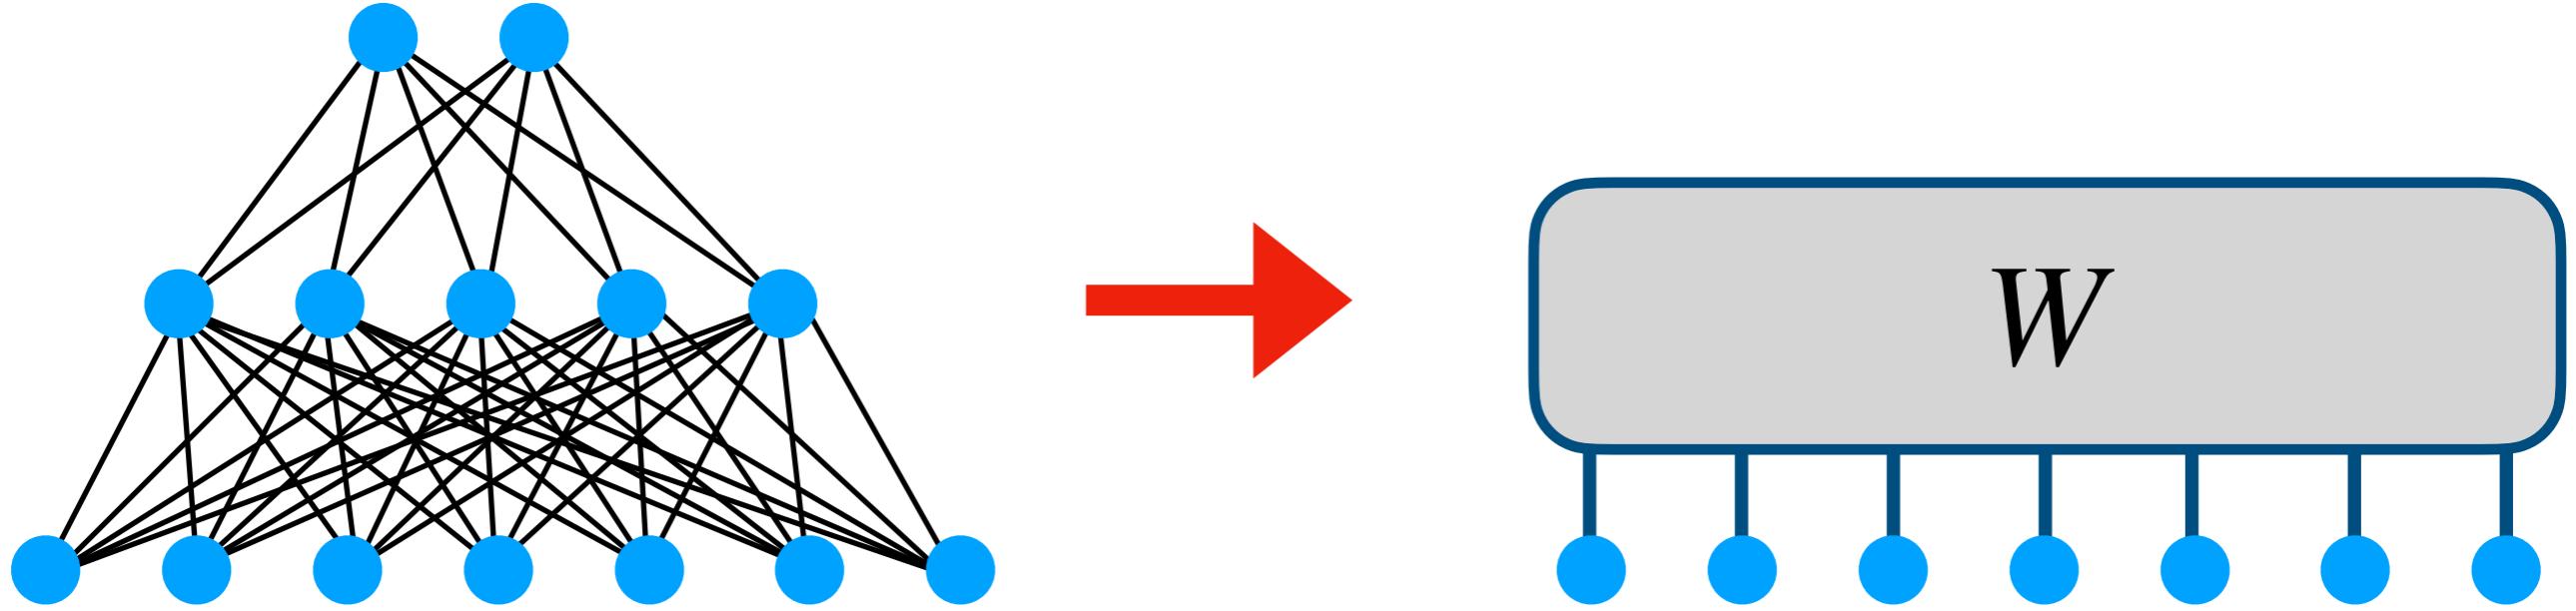
Garipov, Podoprikin, Novikov, [arxiv:1611.03214](https://arxiv.org/abs/1611.03214)

Hallam, Andrew, et al. "Compact neural networks based on the multiscale entanglement renormalization ansatz." [arXiv:1711.03357](https://arxiv.org/abs/1711.03357)

Rose Yu, Stephan Zheng, Anima Anandkumar, Yisong Yue, "Long-term Forecasting using Tensor-Train RNNs", [arXiv:1711.00073](https://arxiv.org/abs/1711.00073)

- Train very "wide" model: 262,144 hidden units
- Achieve 80x compression, only 1% accuracy loss

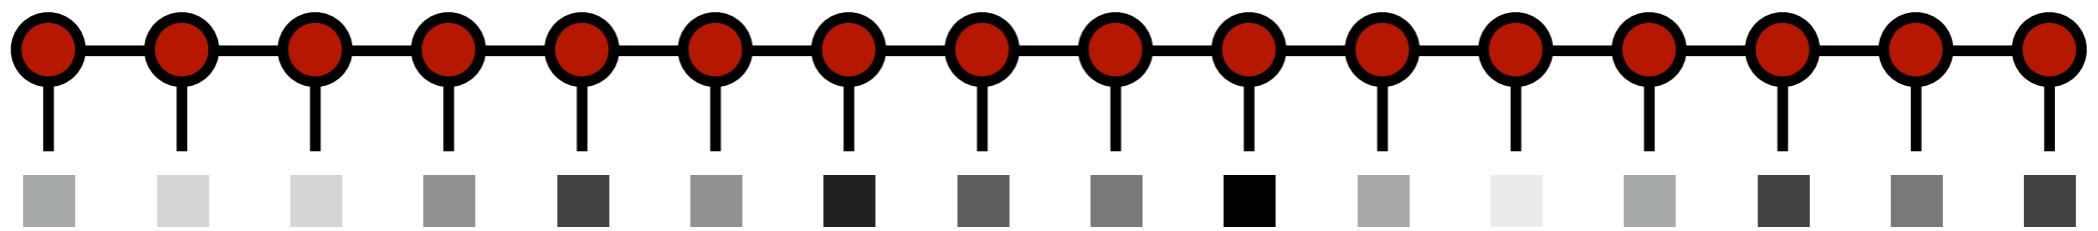
Framework where tensor network plays central role?



Motivation:

- Can natural images be more complex than wavefunctions?
- Import many ideas, algorithms from physics
- Improve tensor network methods

# MPS and Tensor Networks for General Data



# Ingredients for machine learning:

Data ( + labels):

$$\{\mathbf{x}_j, y_j\}$$

Objective / cost function:

$$\min_f \frac{1}{|\mathcal{T}|} \sum_{j=1}^{|\mathcal{T}|} (f(\mathbf{x}_j) - y_j)^2$$

Class of model functions:  $f(\mathbf{x})$

How to get a class of **functions** where a huge (order-N) **tensor** appears?

Consider a polynomial over N variables:

$$f(x_1, x_2, \dots, x_N) = \sum_{\mathbf{n}} W_{n_1 n_2 \dots n_N} x_1^{n_1} x_2^{n_2} \cdots x_N^{n_N}$$

How to get a class of **functions** where a huge (order-N) **tensor** appears?

Consider a polynomial over N variables:

$$f(x_1, x_2, \dots, x_N) = \sum_{\mathbf{n}} W_{n_1 n_2 \dots n_N} x_1^{n_1} x_2^{n_2} \cdots x_N^{n_N}$$

$$= W_{00000} + W_{10000} x_1 + W_{01000} x_2 + \dots$$

$$+ W_{11000} x_1 x_2 + W_{10100} x_1 x_3 + \dots$$

$$+ W_{11111} x_1 x_2 x_3 x_4 x_5$$

More generally, use any basis of products of functions

$$f(x_1, x_2, \dots, x_N) = \sum_{\mathbf{n}} W_{n_1 n_2 \dots n_N} \phi^{n_1}(x_1) \phi^{n_2}(x_2) \cdots \phi^{n_N}(x_N)$$

More generally, use any basis of products of functions

$$f(x_1, x_2, \dots, x_N) = \sum_{\mathbf{n}} W_{n_1 n_2 \dots n_N} \phi^{n_1}(x_1) \phi^{n_2}(x_2) \cdots \phi^{n_N}(x_N)$$

$$\phi(x) = \begin{bmatrix} 1 \\ x \end{bmatrix}$$

gives polynomials

More generally, use any basis of products of functions

$$f(x_1, x_2, \dots, x_N) = \sum_{\mathbf{n}} W_{n_1 n_2 \dots n_N} \phi^{n_1}(x_1) \phi^{n_2}(x_2) \cdots \phi^{n_N}(x_N)$$

$$\phi(x) = \begin{bmatrix} 1 \\ x \end{bmatrix} \quad \text{gives polynomials}$$

$$\phi(x) = \begin{bmatrix} \cos(\pi x/2) \\ \sin(\pi x/2) \end{bmatrix} \quad \text{gives symmetry} \quad x \leftrightarrow (1-x)$$

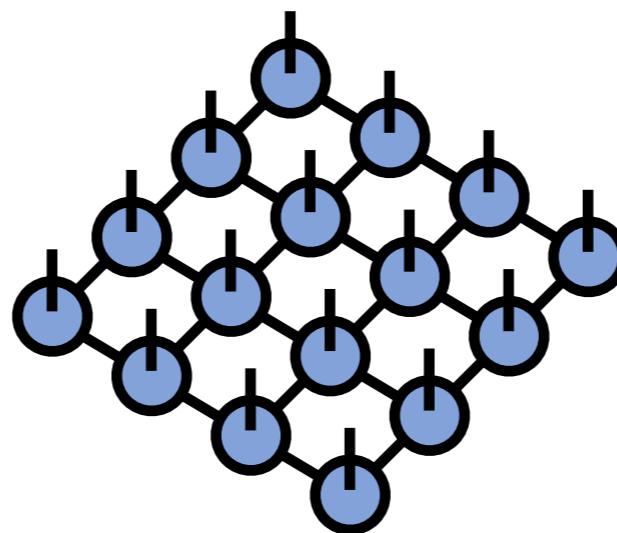
# Main idea: factorize weight tensor

$$W_{n_1 n_2 n_3 n_4 n_5 n_6} = \text{Diagram}$$

$$= \text{---} \circ \text{---} \circ \text{---} \circ \text{---} \circ \text{---} \circ \text{---}$$

or

—



or other  
tensor  
networks

Can use as a model class for supervised or unsupervised learning

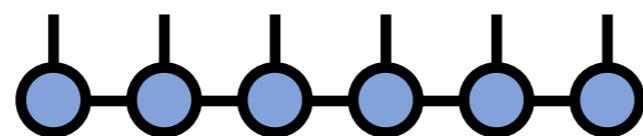
$$f(x_1, x_2, \dots, x_6) = \sum_{\mathbf{n}} \text{Diagram} \quad x_1^{n_1} x_2^{n_2} \cdots x_6^{n_6}$$

The diagram shows a sequence of six blue circles connected by horizontal lines. Above each circle is a label  $n_1, n_2, n_3, n_4, n_5, n_6$  respectively. Below the sequence is the label  $\mathbf{n}$ .

$$\min_{\bullet} \frac{1}{|\mathcal{T}|} \sum_{j=1}^{|\mathcal{T}|} \left( f(\mathbf{x}_j) - y_j \right)^2$$

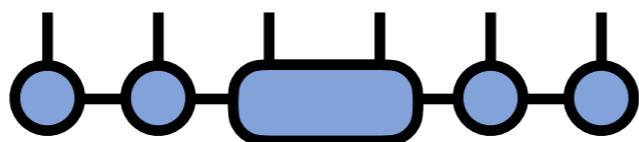
# Immediate payoff: adaptivity of weights

1) merge (contract) a pair of tensors



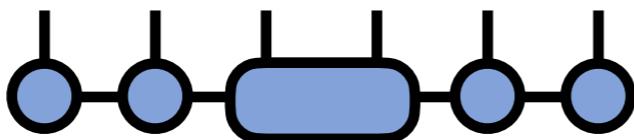
# Immediate payoff: adaptivity of weights

1) merge (contract) a pair of tensors



# Immediate payoff: adaptivity of weights

1) merge (contract) a pair of tensors

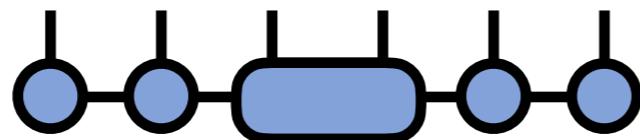


2) optimize parameters (e.g. gradient steps)

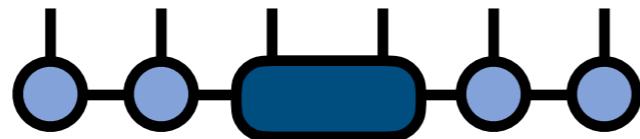


# Immediate payoff: adaptivity of weights

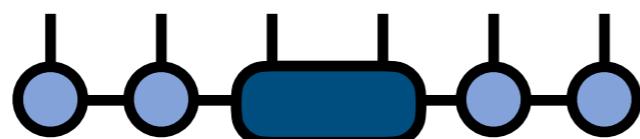
1) merge (contract) a pair of tensors



2) optimize parameters (e.g. gradient steps)

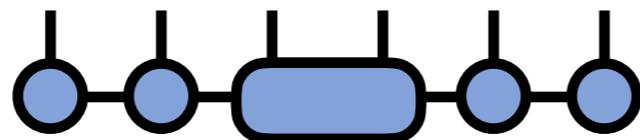


3) SVD factorization to adapt size of bond index



# Immediate payoff: adaptivity of weights

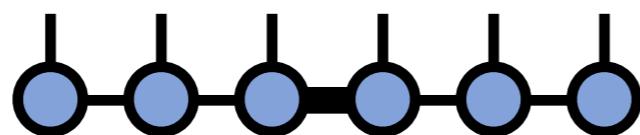
1) merge (contract) a pair of tensors



2) optimize parameters (e.g. gradient steps)



3) SVD factorization to adapt size of bond index



How well do tensor networks perform for supervised  
& unsupervised learning?

What applications can be done?

# Supervised Learning

MPS have been used as proof-of-principle for tensor networks in supervised learning

## 1. MNIST data set (99% test accuracy)

*Stoudenmire, Schwab, NIPS 29, 4799 (2016)*

*Efthymiou, Hidary, Leichenauer, arxiv:1906.06329*

0 6 0 8 1 1  
9 3 2 2 5 9  
7 6 5 5 3 1  
3 6 0 3 8 9  
8 7 0 6 0 0  
7 4 1 7 4 9

## 2. Fashion MNIST data set (89% test accuracy)

*Stoudenmire, Quant. Sci. Tech. 3, 034003 (2018)*

*Glasser, Pancotti, Cirac, arxiv:1806.05964 (2018)*



# Generative Modeling

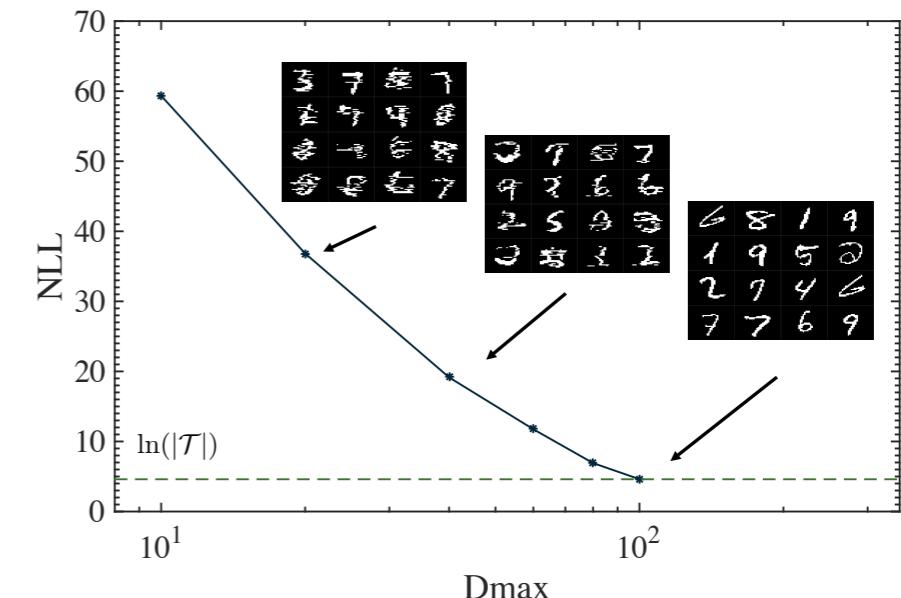
MPS have been used to parameterize generative models

## 1. Trained by gradient optimization of log-likelihood

*MNIST data set*

Han, Wang, Fan, Wang, Zhang, PRX 8, 031012 (2018)

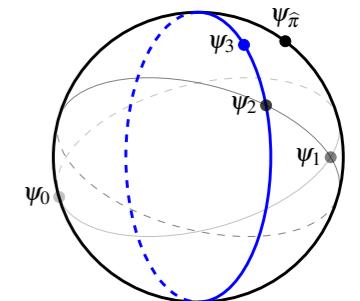
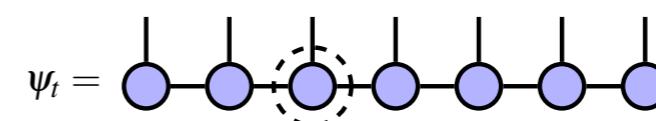
Cheng, Wang, Xiang, Zhang, PRB 99, 155131 (2019)



## 2. Trained by local-exact-solution algorithm

*parity ( $N=20$ ) data set*

Stokes, Terilla, arxiv:1902.06888



Concatenating multiple MPS (string-bond) and pre-processing with a CNN layer gives state-of-the-art on Fashion MNIST:

Method	Test Accuracy
AlexNet	88.90%
Tree tensor net. <sup>1</sup>	88.97%
String bond state <sup>2</sup>	89.2%
CNN-String bond state <sup>2</sup>	92.3%
GoogLeNet	93.7%



# Extension #1: hierarchical optimization of tensor network unsupervised / supervised hybrid

## Data feature map:

$$x \longrightarrow \Phi(x) = \text{[six blue circles with black outlines]}.$$

## Deterministic tensor learning:

$$(a) \quad \rho_{12} = \sum_j s'_1 s'_2 = \begin{array}{c} \text{Diagram showing two coupled spins } s_1 \text{ and } s_2 \text{ with coupling strength } j. \\ \text{The spins are coupled through a central interaction region.} \end{array}$$

$$(b) \quad \rho_{12} = \begin{array}{c} s'_1 \quad s'_2 \\ \text{---} \text{---} \\ \text{blue oval} \\ \text{---} \text{---} \\ s_1 \quad s_2 \end{array} = \begin{array}{c} s'_1 \quad s'_2 \\ \text{---} \text{---} \\ \text{yellow triangle} \\ \text{---} \text{---} \\ s_1 \quad s_2 \end{array} U_{12} \quad \begin{array}{c} \text{---} \text{---} \\ \text{blue circle} \\ \text{---} \text{---} \\ \text{yellow triangle} \\ \text{---} \text{---} \\ s_1 \quad s_2 \end{array} U_{12}^\dagger P_{12}$$

Resulting model – train top tensors for supervised objective:

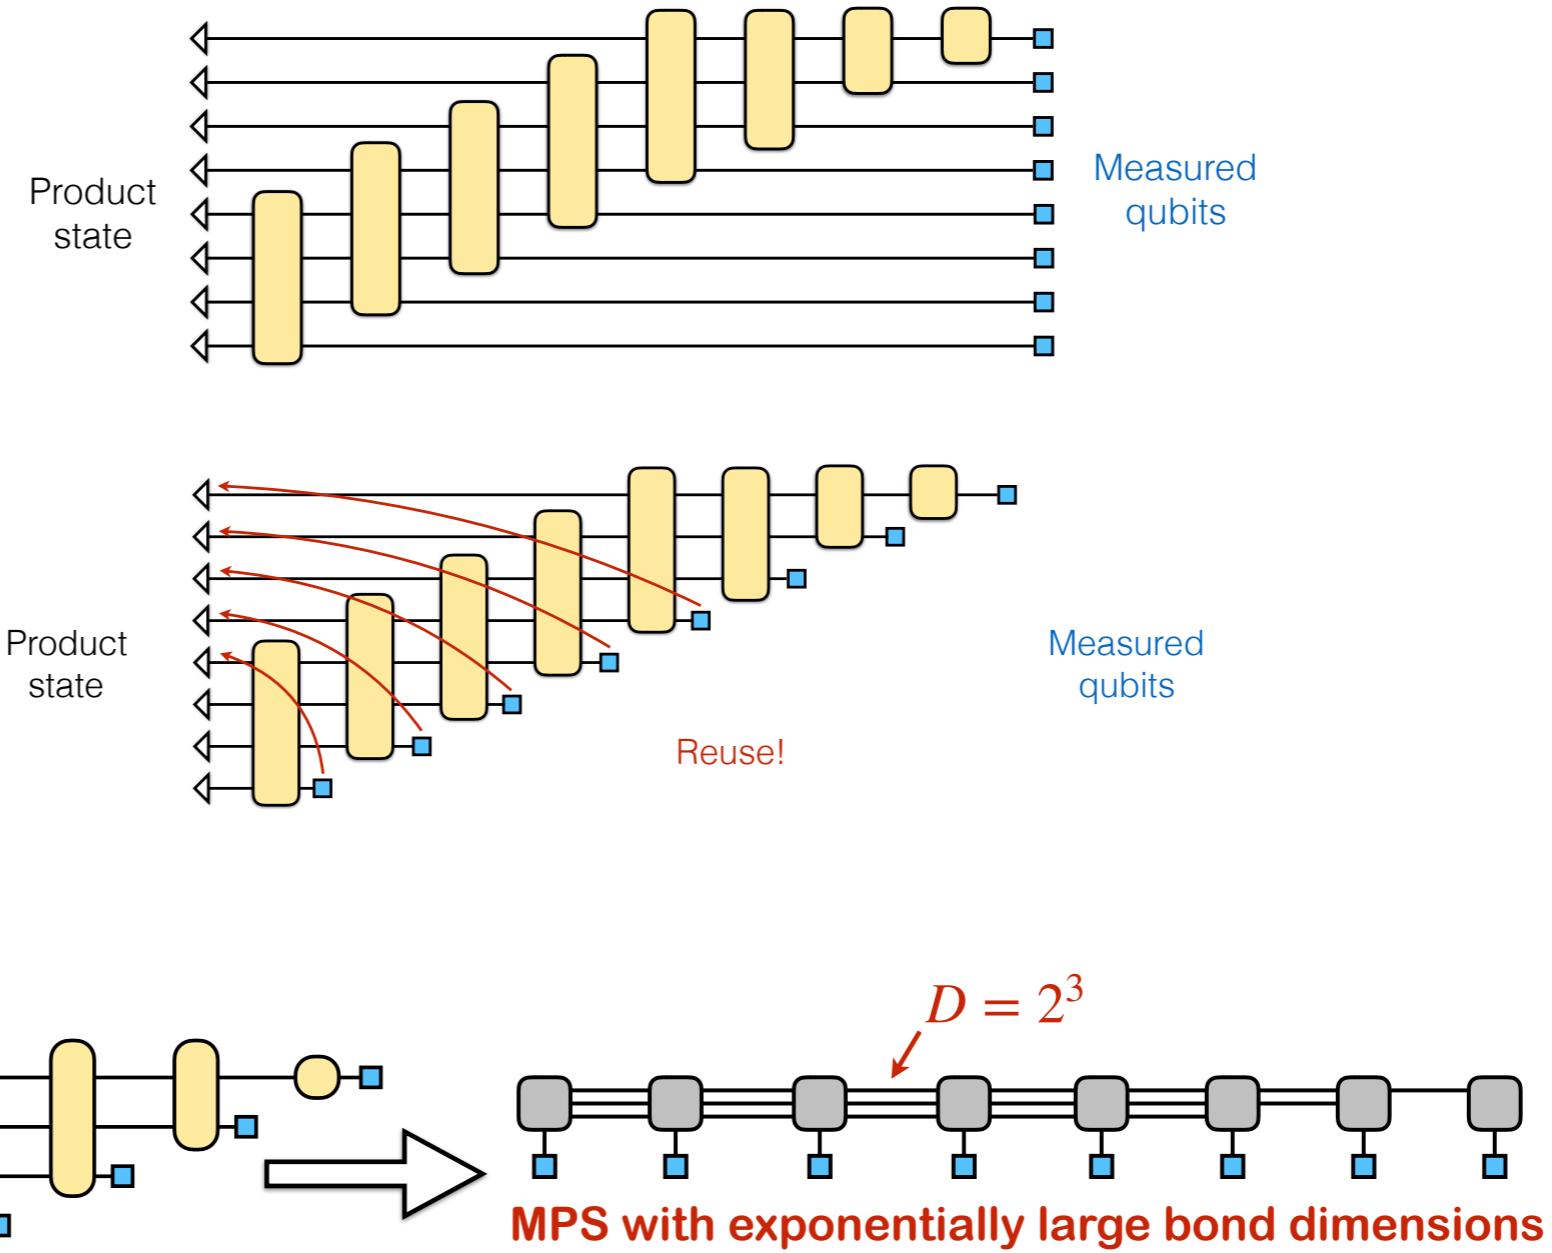
The diagram illustrates a neural network layer  $f^\ell(\mathbf{x}) = \Phi(\mathbf{x})$ . It features a sequence of four hidden units, each represented by a grey circle at the top. Below each unit is a tree structure representing a function  $u$ , with orange triangles as internal nodes and blue circles as leaf nodes. The entire set of trees is grouped by a brace on the right labeled  $\{\Phi(\mathbf{x})\}$ . Above the first three hidden units is a label  $\ell$ , indicating the layer index.



*89% accuracy on  
Fashion MNIST data set*

# Extension #2: algorithm for quantum machine learning identical to optimization of a tensor network!

qubit efficient  
scheme:



Related idea: infinite DMRG on quantum computers  
with finite number of physical qubits

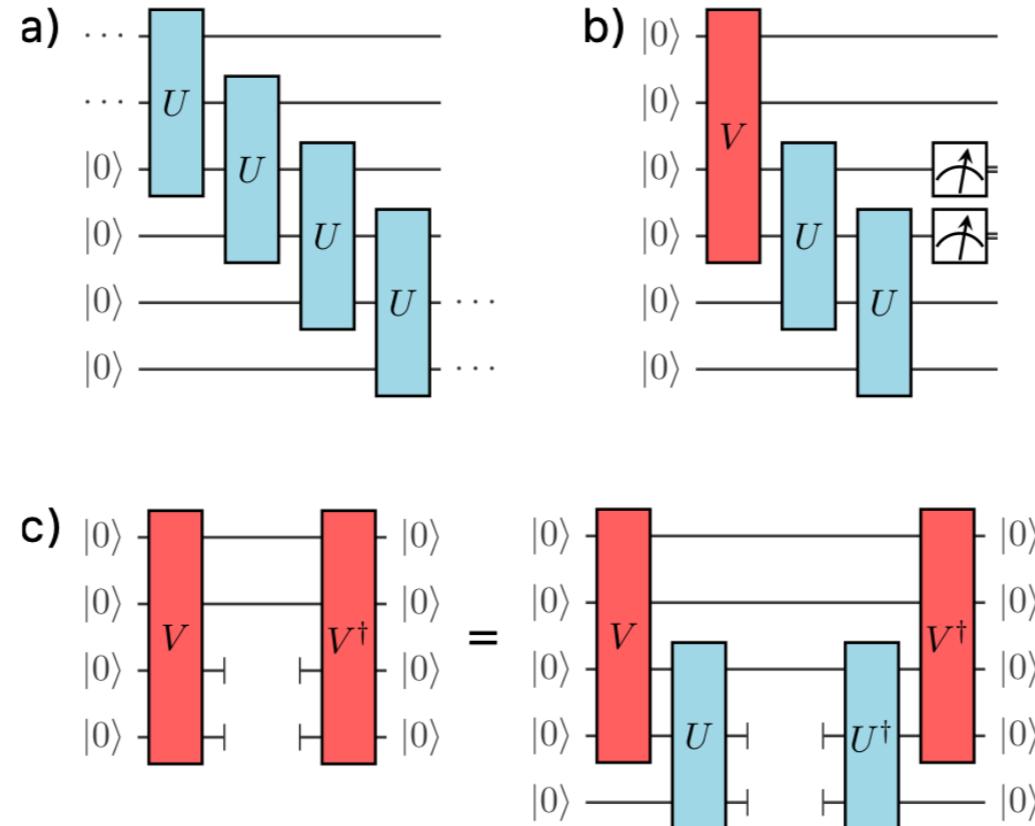


FIG. 2. Quantum circuits for translationally invariant states and their local measurement. a) An infinite

Some of these techniques could be used for  
machine learning too

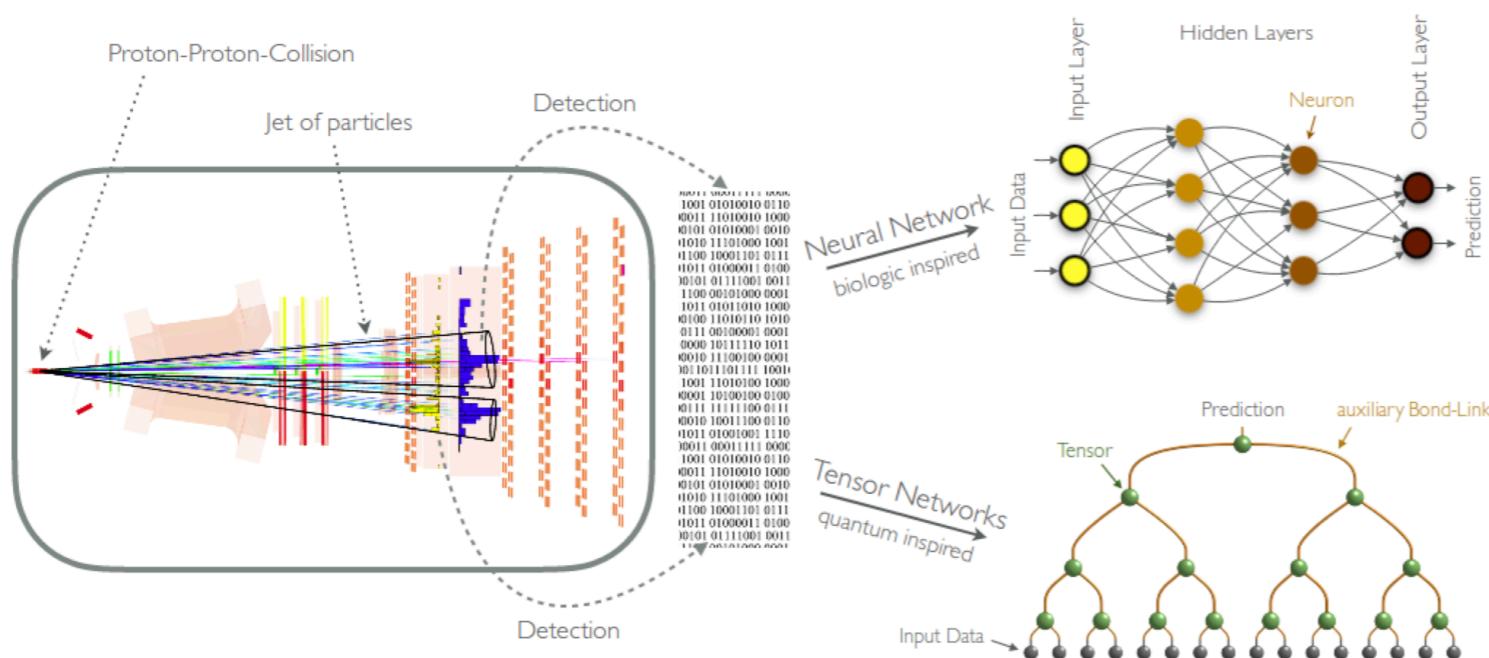
# Review: Notable Recent Works Using *Tensor Network Machine Learning*

# Quantum-Inspired Machine Learning on High-Energy Physics Data



*Marco Trenti, Lorenzo Sestini, Alessio Gianelle, Davide Zuliani, Timo Felser, Donatella Lucchesi, Simone Montangero, arxiv:2004.13747*

- Used tree tensor network
- Determine charge of quark from events
- Compared to neural networks and gave excellent results



**Schematic Learning Process**

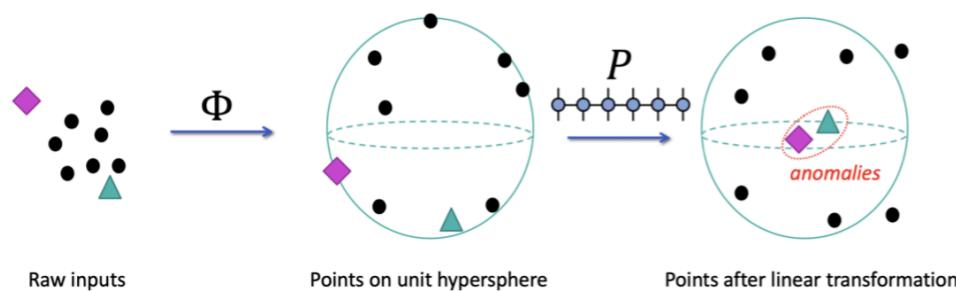
# Anomaly Detection with Tensor Networks



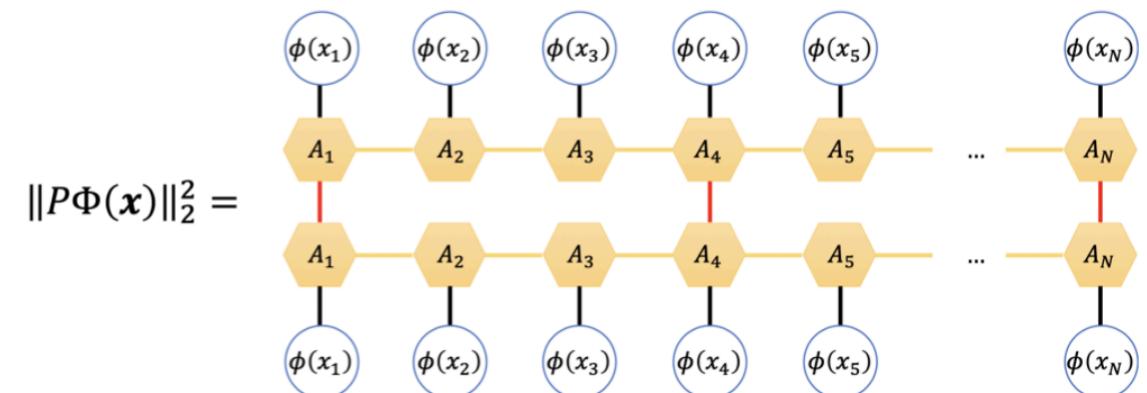
Jinhui Wang, Chase Roberts, Guifre Vidal, Stefan Leichenauer, arxiv:2006.02516



- Used squared tensor network (locally purified state)
- Perform anomaly detection task
- Strongly competes with, and for tabular data **exceeds performance** of state-of-art neural net approaches



**Geometric Anomaly  
Detection**



**Model Architecture**

Table 3: Mean AUROC scores (in %) and standard errors on ODDS datasets.

Dataset	OC-SVM	IF	GOAD	DAGMM	TNAD
<i>Wine</i>	<b>60.0</b>	$46.0 \pm 8.4$	$48.2 \pm 24.7$	$51.7 \pm 19.3$	<b><math>97.3 \pm 4.5</math></b>
<i>Glass</i>	<b>62.0</b>	$57.2 \pm 1.6$	$53.5 \pm 13.6$	$52.5 \pm 12.9$	<b><math>81.8 \pm 7.3</math></b>
<i>Thyroid</i>	98.8	<b><math>99.0 \pm 0.1</math></b>	$95.8 \pm 1.3$	$88.8 \pm 6.8$	<b><math>99.0 \pm 0.1</math></b>
<i>Satellite</i>	<b>79.9</b>	$77.2 \pm 0.9$	$60.6 \pm 5.3$	$72.1 \pm 4.7$	<b><math>81.3 \pm 0.5</math></b>
<i>Forest</i>	<b>97.7</b>	$71.7 \pm 2.6$	$64.6 \pm 4.7$	$60.9 \pm 8.9$	<b><math>98.8 \pm 0.6</math></b>

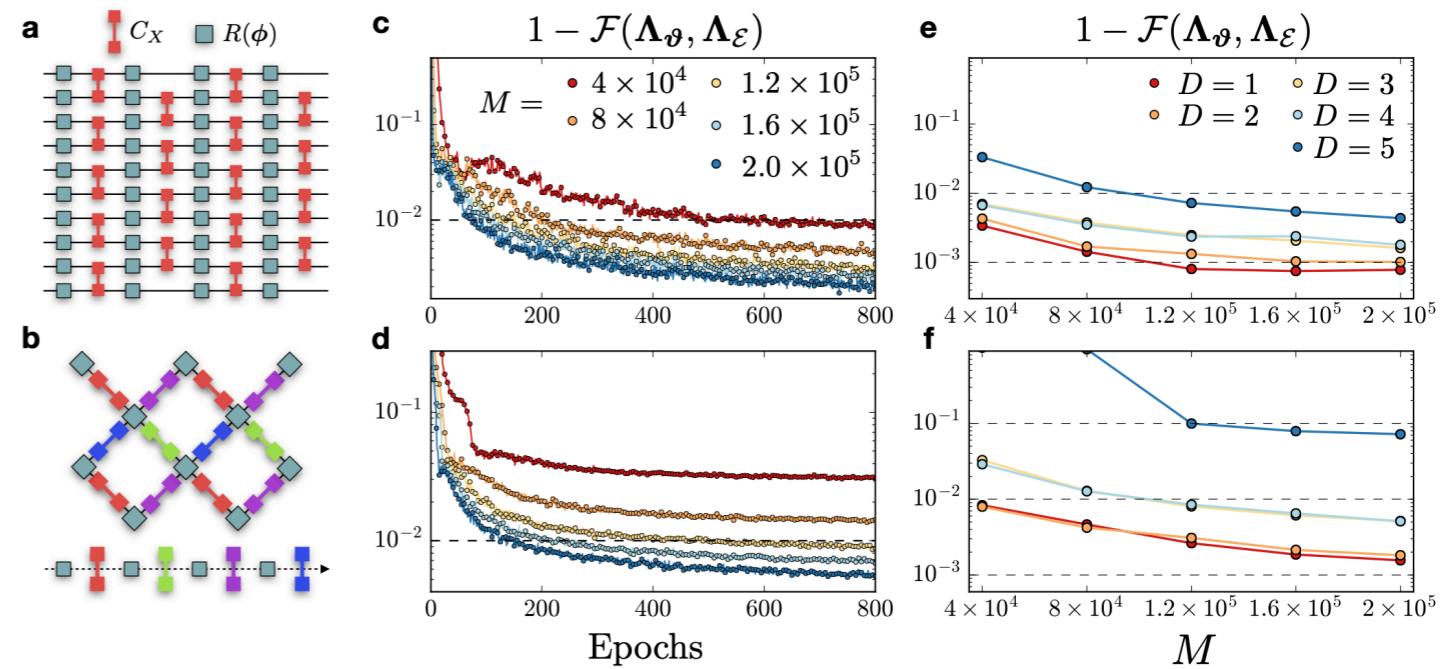
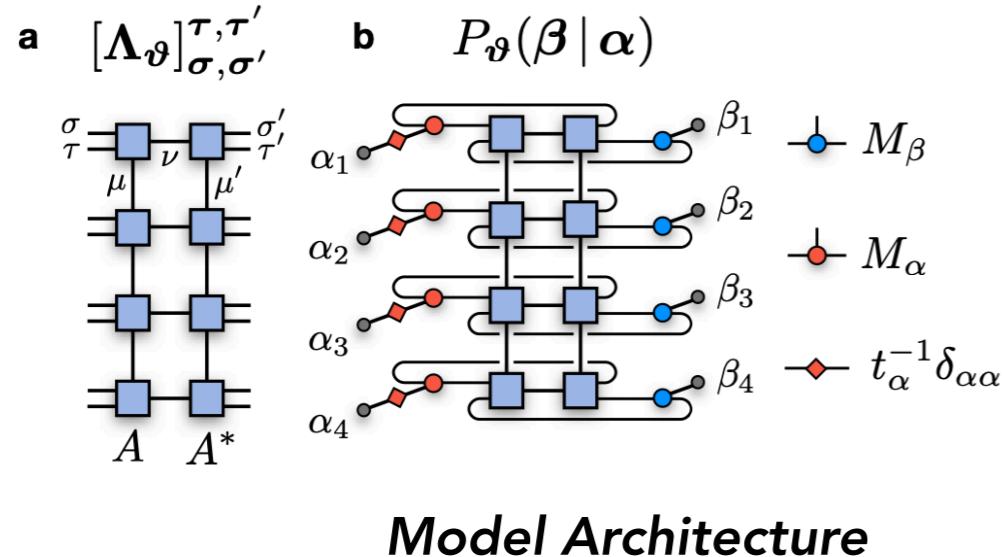
**Results on Tabular Data**

# Quantum process tomography with unsupervised learning and tensor networks



Giacomo Torlai, Christopher J. Wood, Atithi Acharya, Giuseppe Carleo, Juan Carrasquilla, Leandro Aolita, arxiv: 2006.02424

- Also used squared tensor network (locally purified state)
- Deduce noisy circuit (actually CPTP map) that explains observed data
- First scalable approach to quantum process tomography, as far as I know



**Learning Random Quantum Circuits**

# Summary & Future Directions

Tensor networks are a natural way to parameterize interesting and powerful machine learning models

Benefits:

- Linear scaling
- Adaptive weights
- Learning data "features"
- Interpretability & theory
- Better algorithms
- Quantum computing

Much work to be done:

- best approach for various data set types
- theory of learning of tensor networks (it is possible!)
- developing new learning algorithms