

Title: Lecture - Machine Learning, PHYS 777

Speakers: Mohamed Hibat Allah

Collection/Series: Machine Learning (Elective), PHYS 777, February 24 - March 28, 2025

Subject: Condensed Matter, Other

Date: March 05, 2025 - 9:00 AM

URL: <https://pirsa.org/25030036>

Lecture 5


Last time:

- ↳ Variants of gradient descent (SGD, Momentum)
- ↳ Backpropagation
- ↳ Intro to overfitting

Bias-variance tradeoff

E_{in} : error on the "in-samples" (training data)
 E_{out} : error on the "out-samples" (Not used for training)

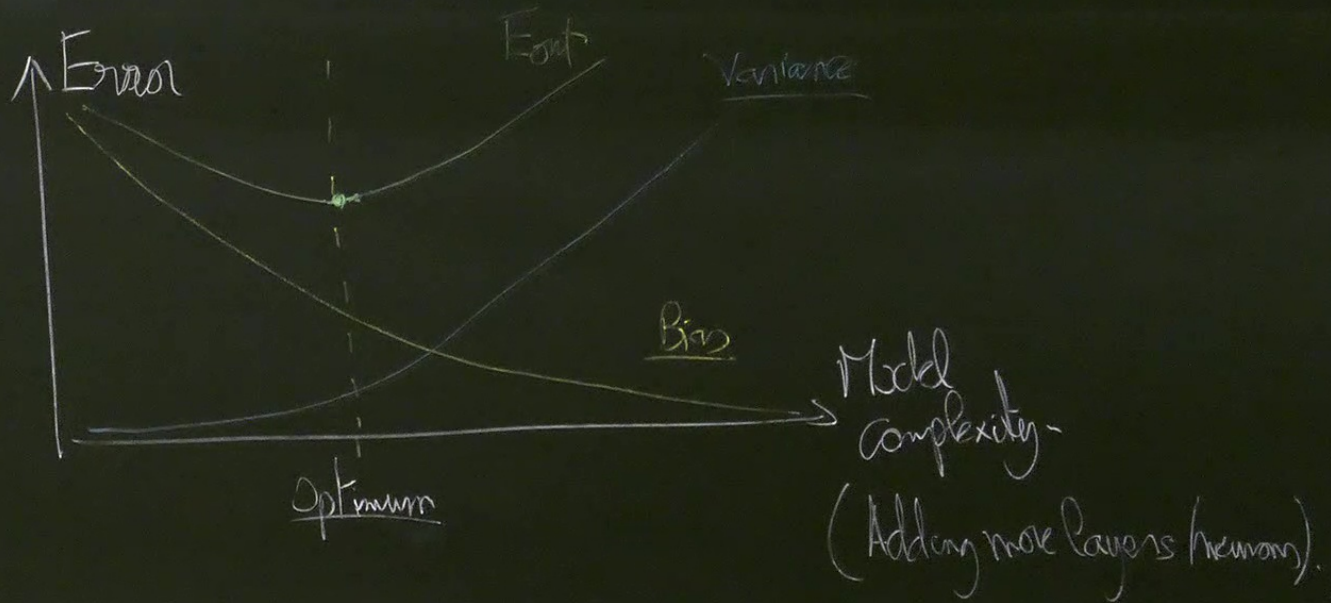
For a fixed model



$$D = \{ (\vec{x}_i, \vec{y}_i) \}$$

← "in" → "out"

↳ Intro to overfitting



→ Bias measures how well the model can do in the infinite-data limit.

↳ More complex models tend to lower the bias

→ Variance measures how dependent the fit is on the particular realization of the training data.

→ The difference between E_m and E_{out} measures the difference between "fitting" and "predicting".

↳ If $E_{out} - E_m \gg 1$, the model is overfitting.

Firefox File Edit View History Bookmarks Tools Window Help

Polynomials_fitting_illustration x A Neural Network Playground x A Neural Network Playground x A Neural Network Playground x A Neural Network Playground x Deeper Playground - Deep Learn: x +

https://colab.research.google.com/drive/1d8019_xyNKakm8axGp8_8RnDLBJ59_ux#scrollTo=4p4CU7q6m8Hf 150%

Polynomials_fitting_illustration.ipynb

File Edit View Insert Runtime Tools Help

+ Code + Text Reconnect

ground truth

- samples
- model

0s completed at 7:17 AM

The image shows a Jupyter Notebook interface in a Firefox browser. The notebook is titled 'Polynomials_fitting_illustration.ipynb'. The main content is a plot with a white background. It features several red circular data points scattered across the plot area. A smooth blue curve is drawn through the points, representing a polynomial fit. A straight orange line is also drawn across the plot, representing a linear fit. A legend in the top right corner of the plot area identifies the red dots as 'samples' and the orange line as 'model'. The notebook interface includes a menu bar with 'File', 'Edit', 'View', 'Insert', 'Runtime', 'Tools', and 'Help'. On the left side, there are icons for file operations and a sidebar with '+ Code' and '+ Text' tabs. At the bottom of the notebook, a status bar shows '0s completed at 7:17 AM'. The browser's address bar shows the URL 'https://colab.research.google.com/drive/1d8019_xyNKakm8axGp8_8RnDLBJ59_ux#scrollTo=4p4CU7q6m8Hf'. The system tray at the bottom of the screen shows various application icons, including a calendar for March 5th.

Firefox File Edit View History Bookmarks Tools Window Help

Polynomials_fitting_illustration x A Neural Network Playground x A Neural Network Playground x A Neural Network Playground x A Neural Network Playground x Deeper Playground - Deep Learn: x +

https://colab.research.google.com/drive/1d8019_xyNKakm8axGp8_8RnDLBJ59_ux#scrollTo=4p4CU7q6m8Hf 150% ☆

Polynomials_fitting_illustration.ipynb ☆ Share M

File Edit View Insert Runtime Tools Help

+ Code + Text Reconnect ^

The following notebook plots the result of using gradient descent to learn a degree 'd' Legendre polynomial on training samples. The default configuration uses training samples which are essentially a cubic with noise (feel free to change this below).

Note that the output of gradient descent starting from 0 initialization of the objective

$$\min_{\beta} \|X\beta - y\|_2^2$$

is equivalent to $\hat{\beta} = X^\dagger y$.

```
[ ] 1 import numpy as np
    2 import numpy.random as rand
    3 import numpy.linalg as la
    4 import pylab as plt
    5 import ipywidgets as widgets
    6 from ipywidgets import interact

[ ] 1 rseed = 1932
    2 rand.seed(rseed)

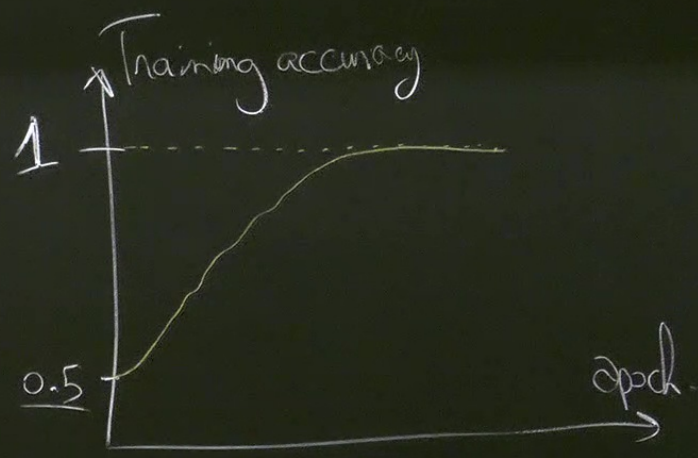

[ ] 1 n = 20
    2 sigma = 0.3 # stdev of noise
    3
```

✓ 0s completed at 7:17 AM

Underfitting

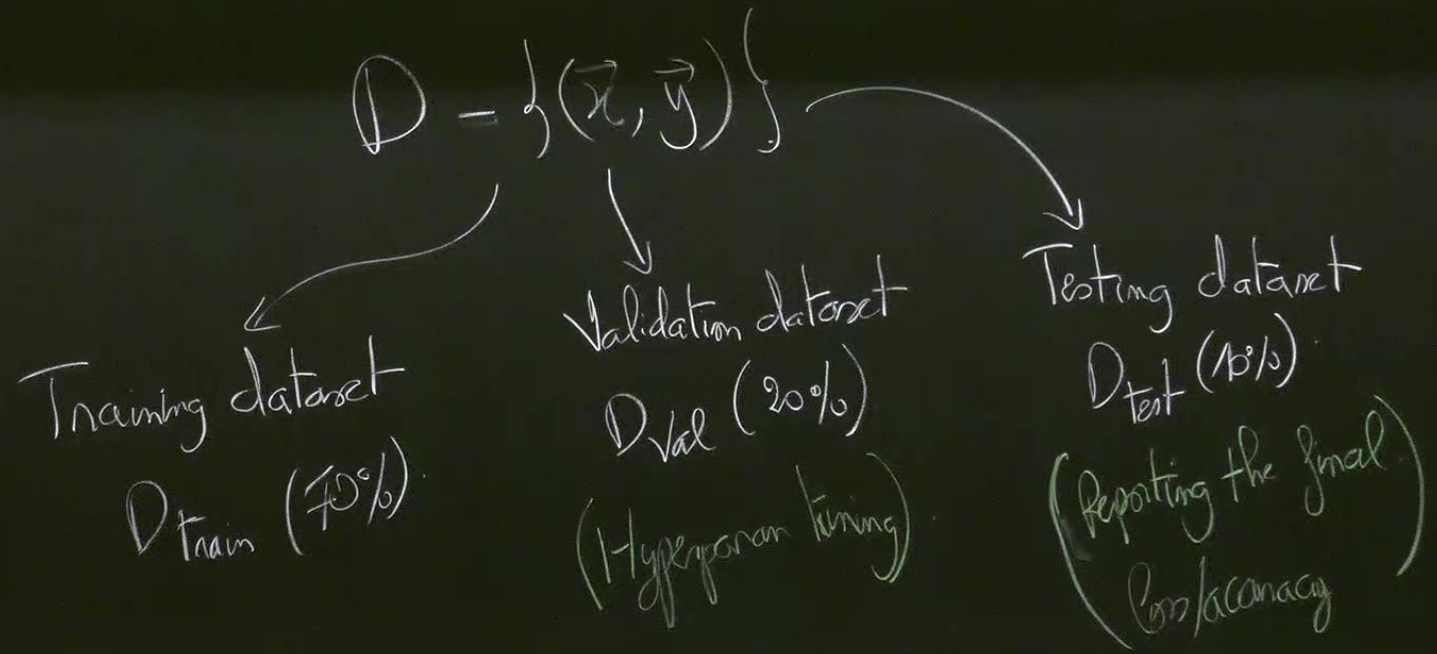
Overfitting

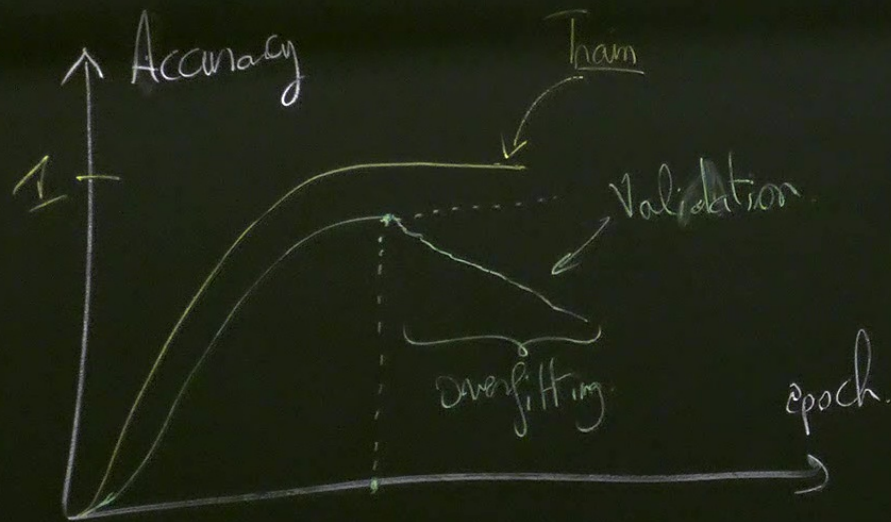
Overfitting



memory).

→ The difference between E_m and E_{out} measures the difference between
↳ If $E_{out} - E_m \gg 1$, the model is overfitting. « predicting »





Methods that prevent overfitting.

↳ Early stopping.

↳ L1 or L2 regularization.

Idea add an extra term to the cost function

$$C_{L2} = C_{\text{mit}} + \lambda \sum_{\text{weights}} |w_{ij}|^2$$


$\xrightarrow{(\lambda > 0)}$

$\underbrace{\sum_{\text{weights}} |w_{ij}|^2}_{\| \text{weights} \|_2^2}$

Hyperparameter
" L2 regularization

$$C_{L1} = C_{\text{mit}} + \lambda \sum_{\text{weights}} |w_{ij}|$$

$\rightarrow \lambda$ too small \rightarrow overfitting
 $\rightarrow \lambda$ too small \rightarrow underfitting



Don't Worry, You Can't Break It. We Promise.

Epoch
 Learning rate Activation Regularization Regularization rate Problem type

DATA

Which dataset do you want to use?



Ratio of training to test data: 50%



Noise: 50



Batch size: 10



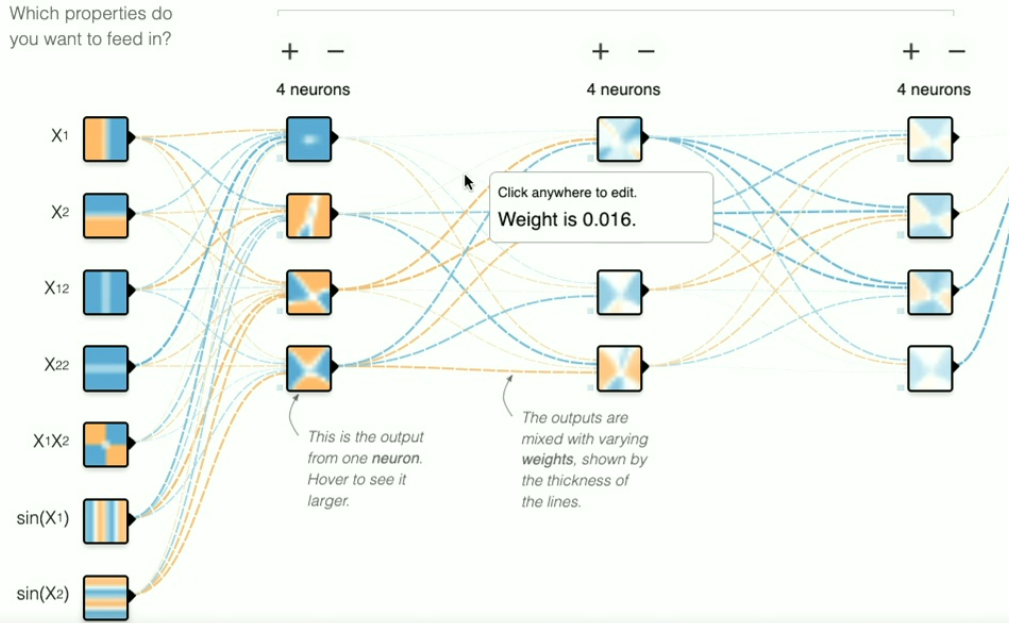
REGENERATE

FEATURES

Which properties do you want to feed in?


- X1
- X2
- X12
- X22
- X1X2
- sin(X1)
- sin(X2)

3 HIDDEN LAYERS



OUTPUT

Test loss 0.491
Training loss 0.490




Colors shows data, neuron and weight values.


IS

OUTPUT

Test loss 0.168
Training loss 0.032



+ -
4 neurons



Epoch 001,467 Learning rate 0.01 Activation Tanh Regularization L2 Regularization rate 0.1 Problem type Classification

DATA

Which dataset do you want to use?



Ratio of training to test data: 50%



Noise: 50



Batch size: 10



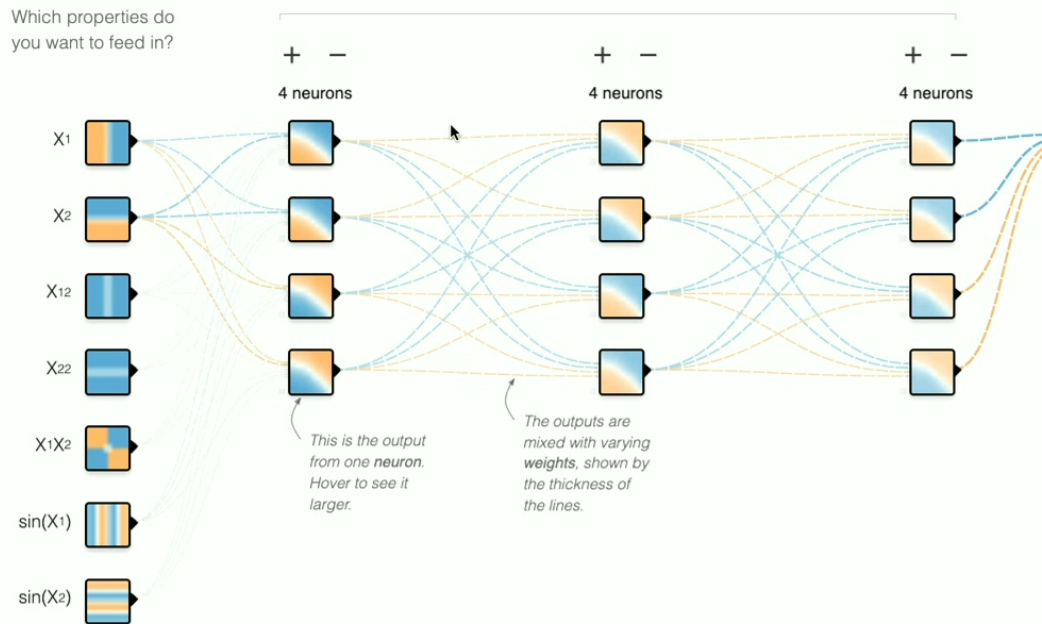
REGENERATE

FEATURES

Which properties do you want to feed in?


- X1
- X2
- X1²
- X2²
- X1X2
- sin(X1)
- sin(X2)

3 HIDDEN LAYERS



OUTPUT



Test loss 0.125
Training loss 0.140



Colors shows data, neuron and weight values. -1 0 1

Show test data Discretize output

Dropout



$P_{\text{dropout}} =$ "Prob of chopping a neuron in a layer"

Removing neurons randomly

Epoch: 000,000

Learning rate: 0.03

Activation: Tanh

Regularization: L2

Regularization rate: 0

Problem type: Classification

Animation speed: 5%

Dropout: Drop 70%

Momentum: 0


Layerwise gradient normalization: p = 0 (None)

Learning rate autotuning: None

Prevent loss increases: No

DATA

Which dataset do you want to use?



Ratio of training to test data: 50%



Noise: 50



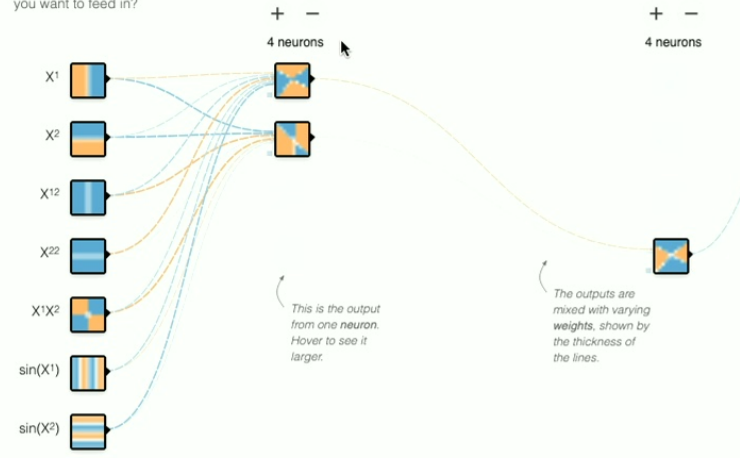
Batch size: 10



REGENERATE

FEATURES


Which properties do you want to feed in?



OUTPUT

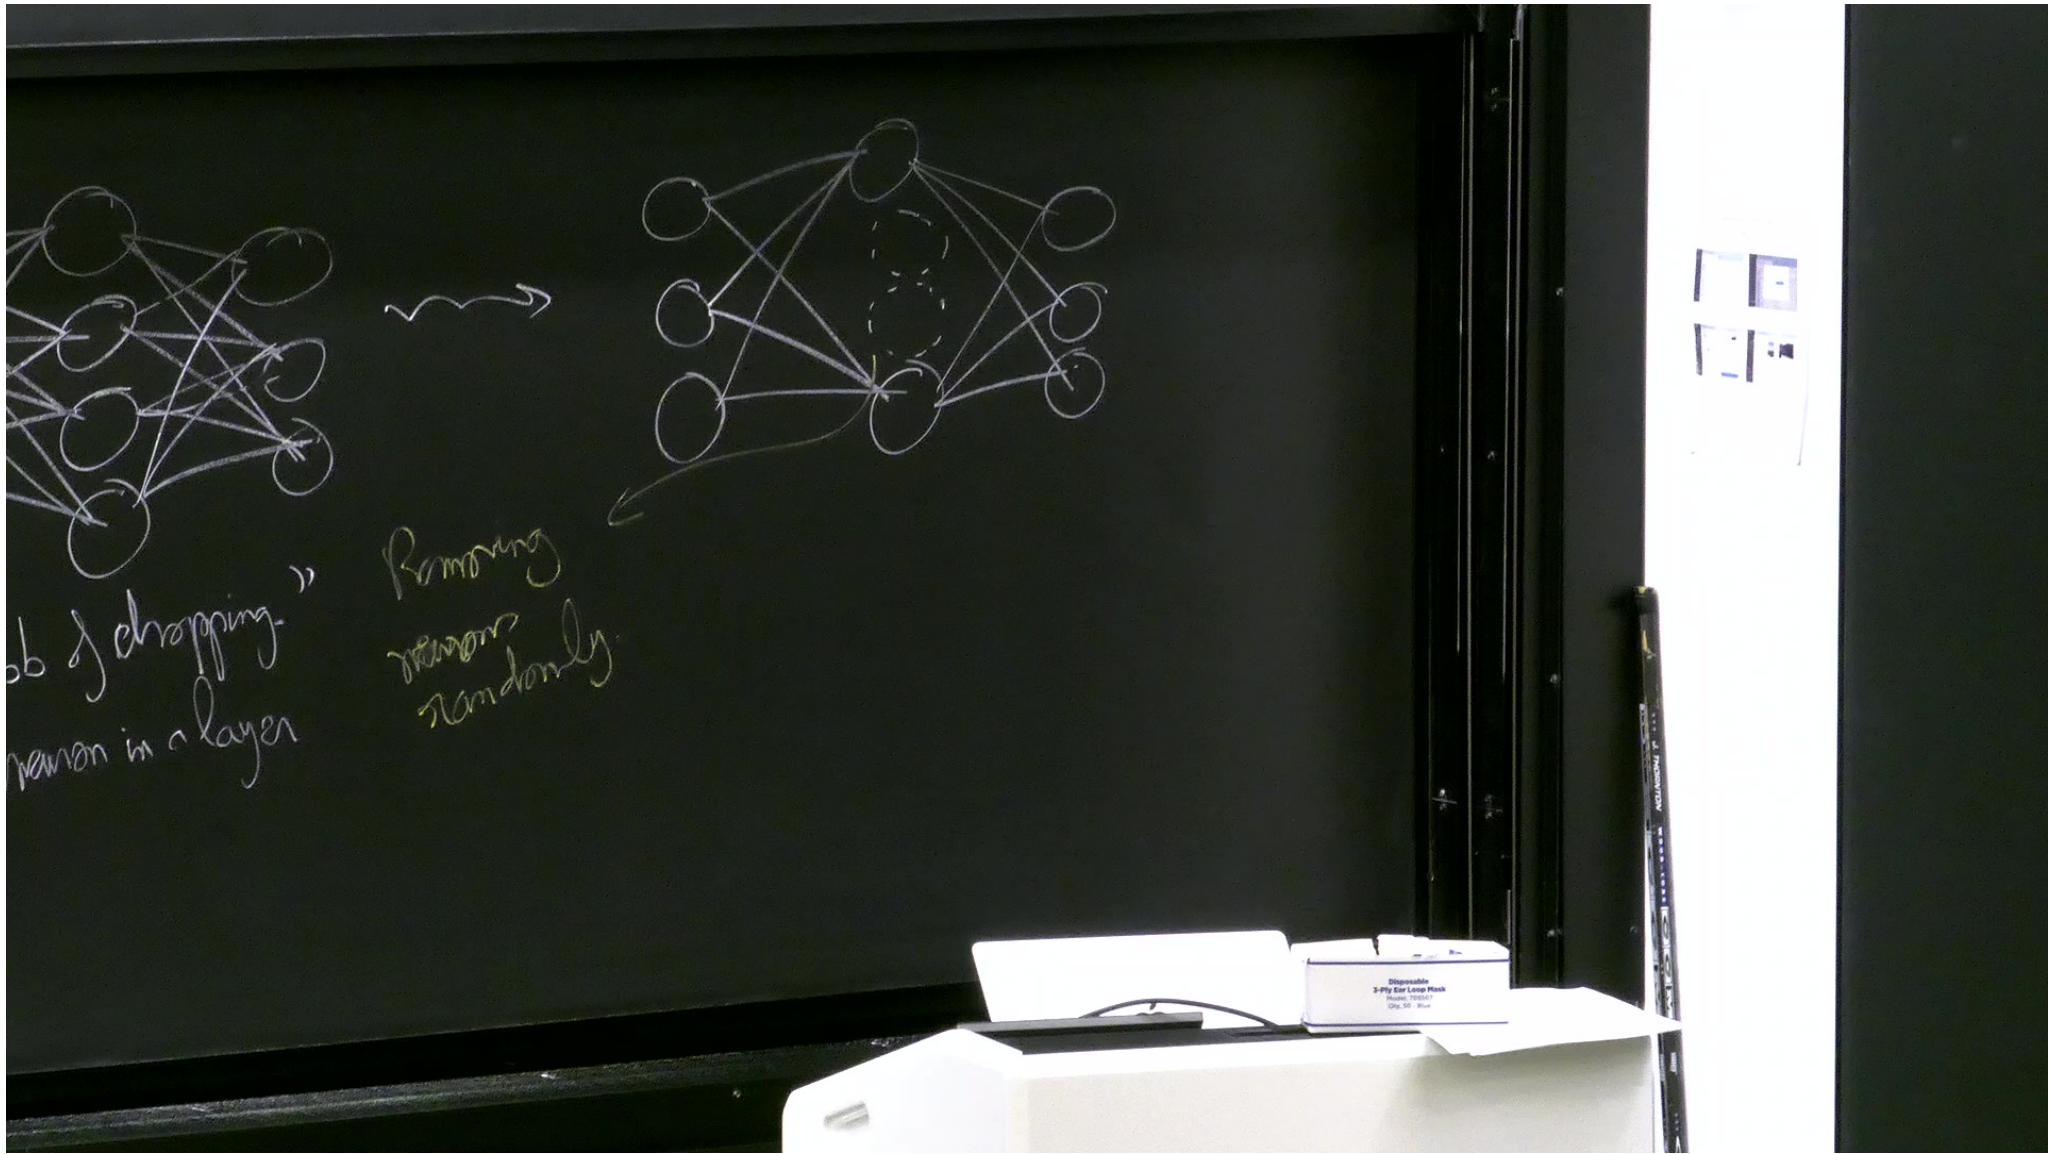
Test loss 0.708

Training loss 0.730



Colors shows data, neuron and weight values.

Show test data Discretize output





Underfitting

Overfitting

Distinguishing classical phases of matter using SL:

$$H = -J \sum_{\langle ij \rangle} S_i S_j$$

$$S_i = \pm 1$$



neurons).

→ The difference between E_m and E_{out} measures the difference between

training & predicting

↳ If $E_{out} - E_m \gg 1$, the model is overfitting.

$$\vec{x} = \begin{pmatrix} +1 & -1 & \dots & +1 \\ s_1 & s_2 & \dots & s_N \end{pmatrix}$$

Labels $y = \begin{matrix} 0 & 1 \\ \downarrow & \downarrow \\ FM & PM \end{matrix}$