

- Which function best fits this data? Why?
- What are some of the choices that we can make when designing a neural network?
- What might we infer about the differences in the neural networks used for each function?
- What can we conclude about the neural network weights in the third figure from looking at the highlighted region?

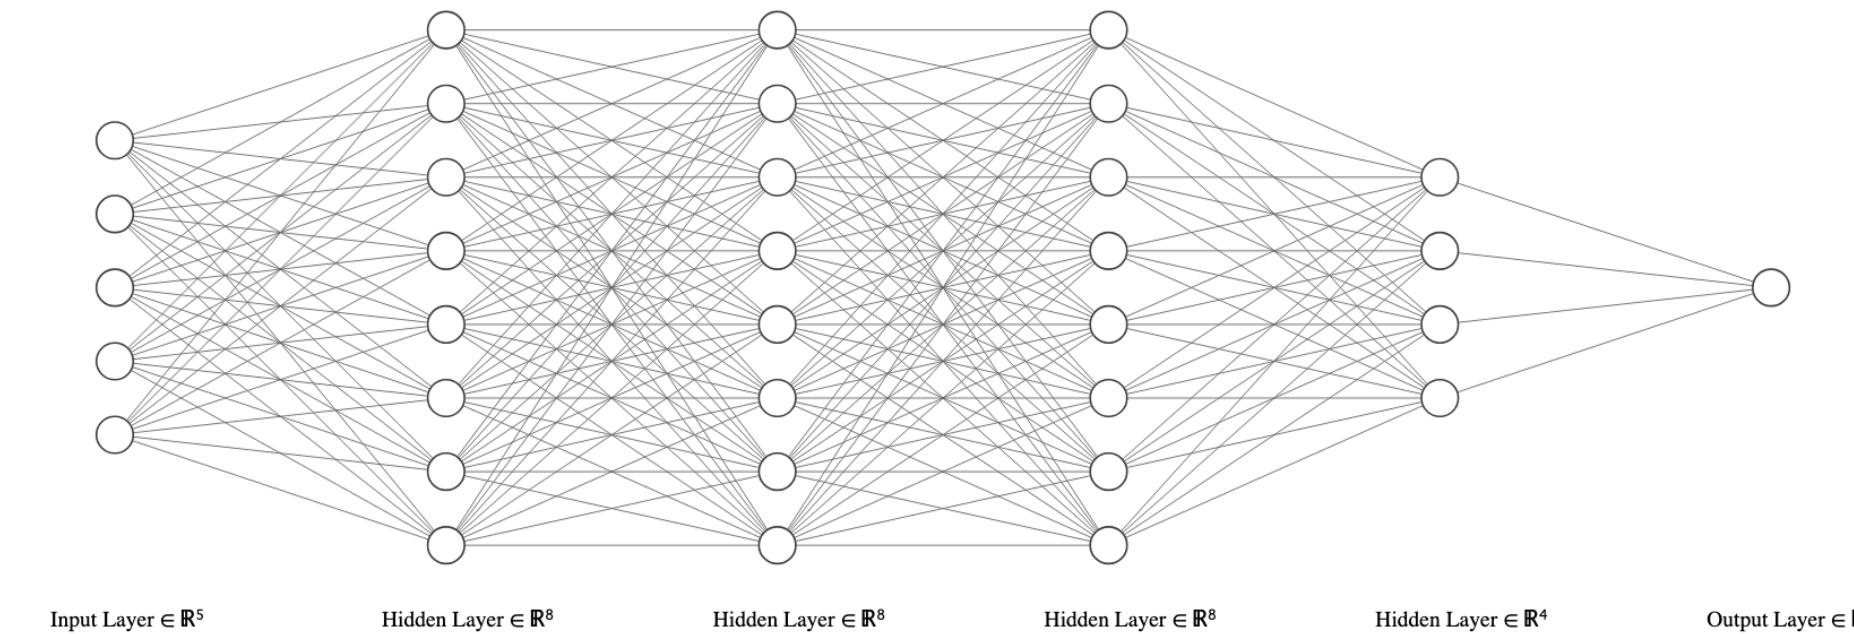
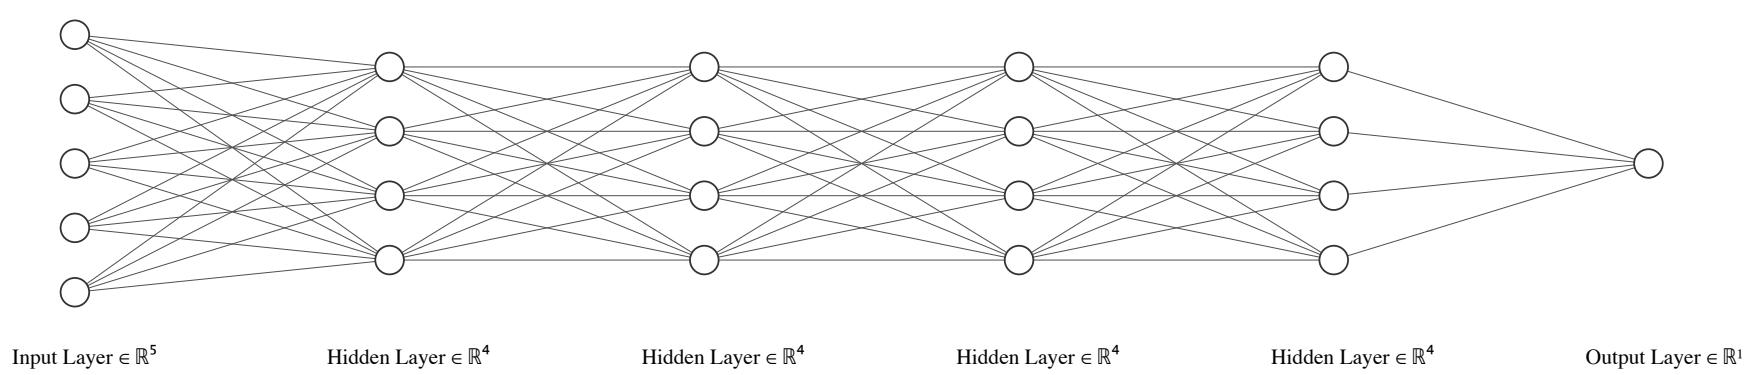
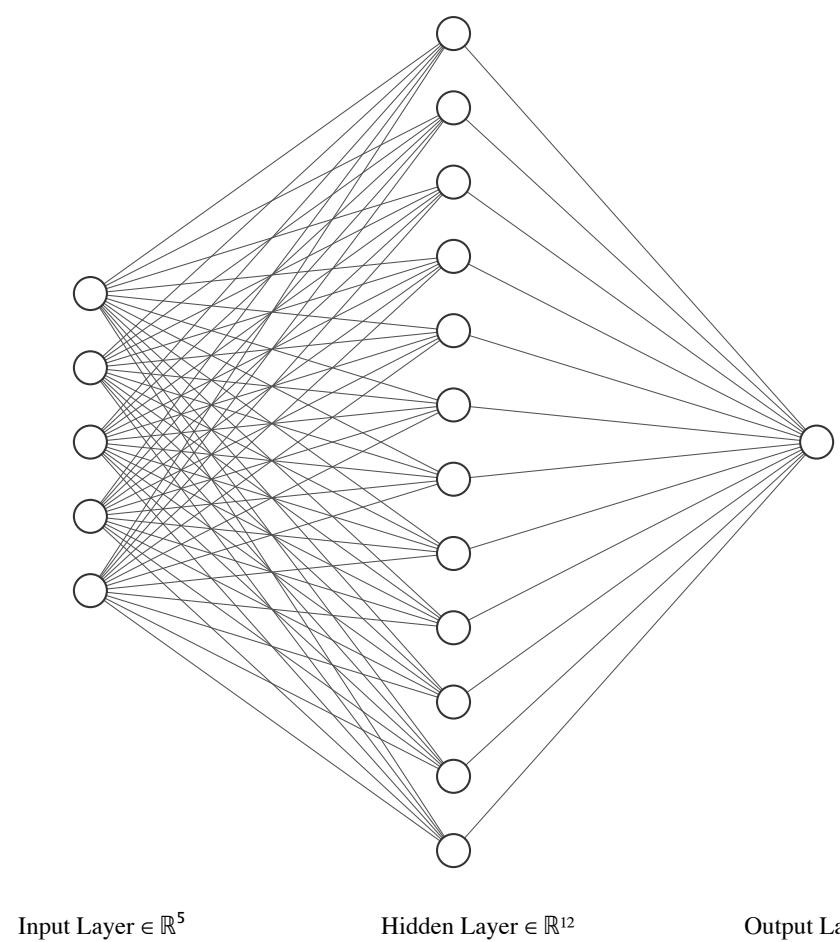
Evaluation

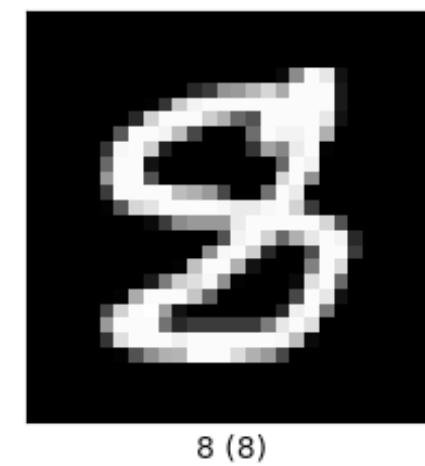


8 (8)

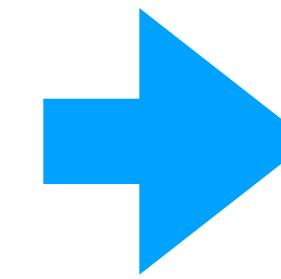
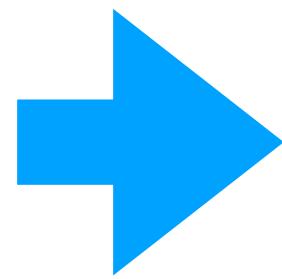


8

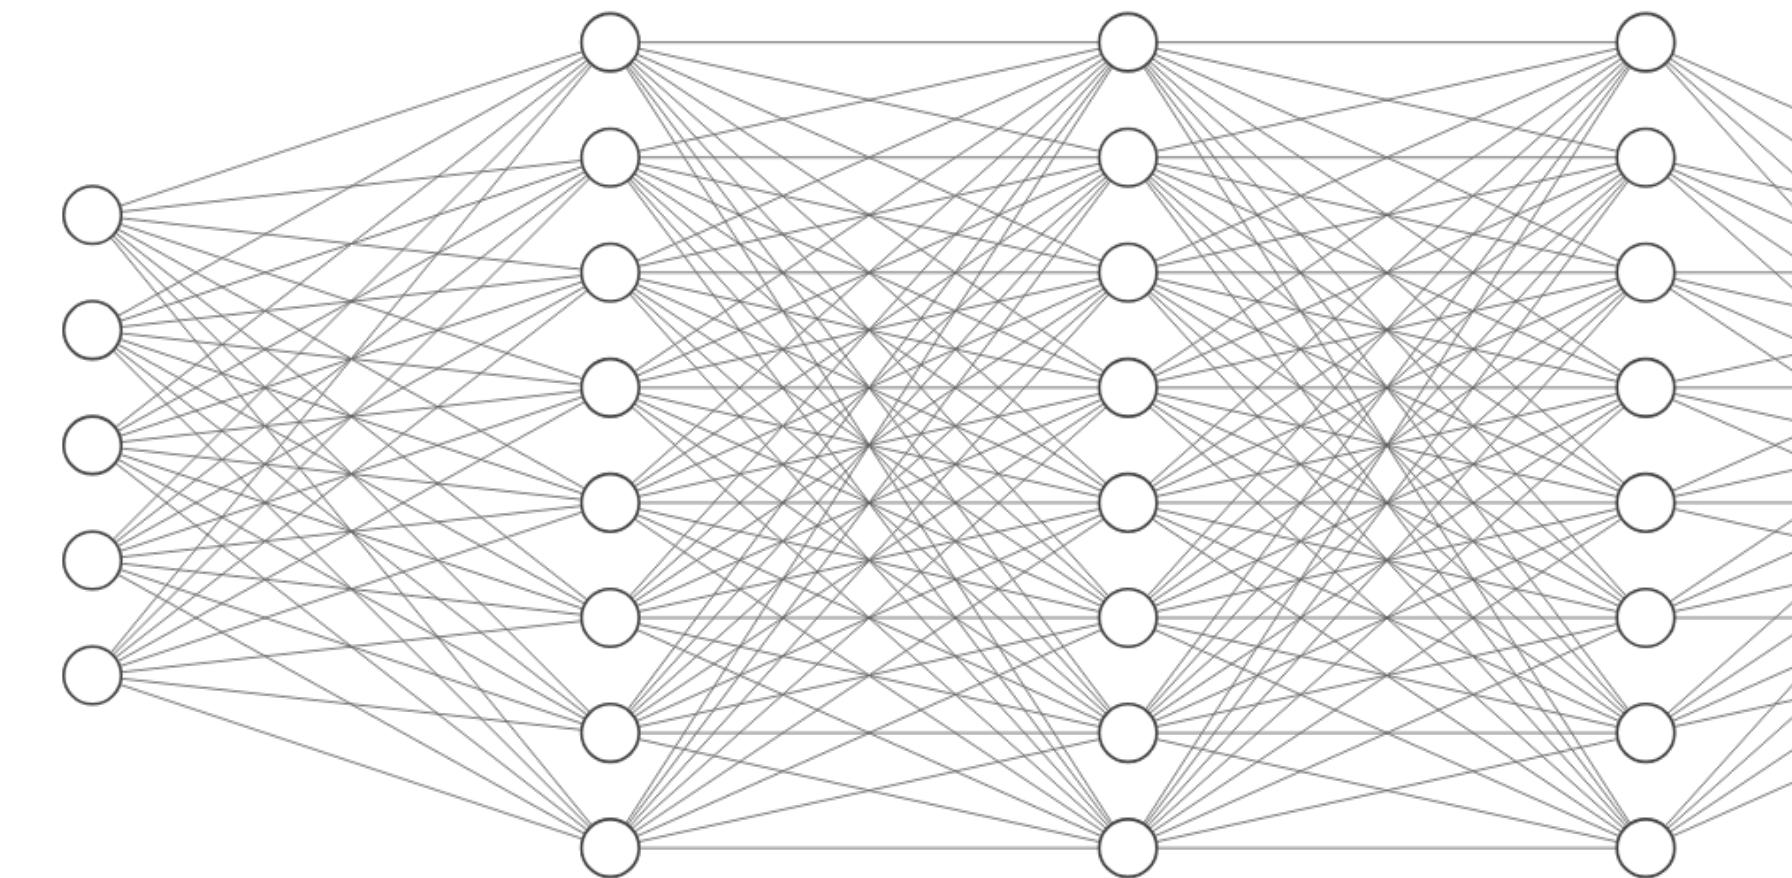




8 (8)



Input Layer $\in \mathbb{R}^5$



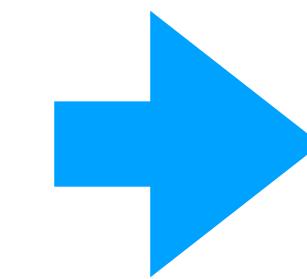
Hidden Layer $\in \mathbb{R}^8$

Hidden Layer $\in \mathbb{R}^8$

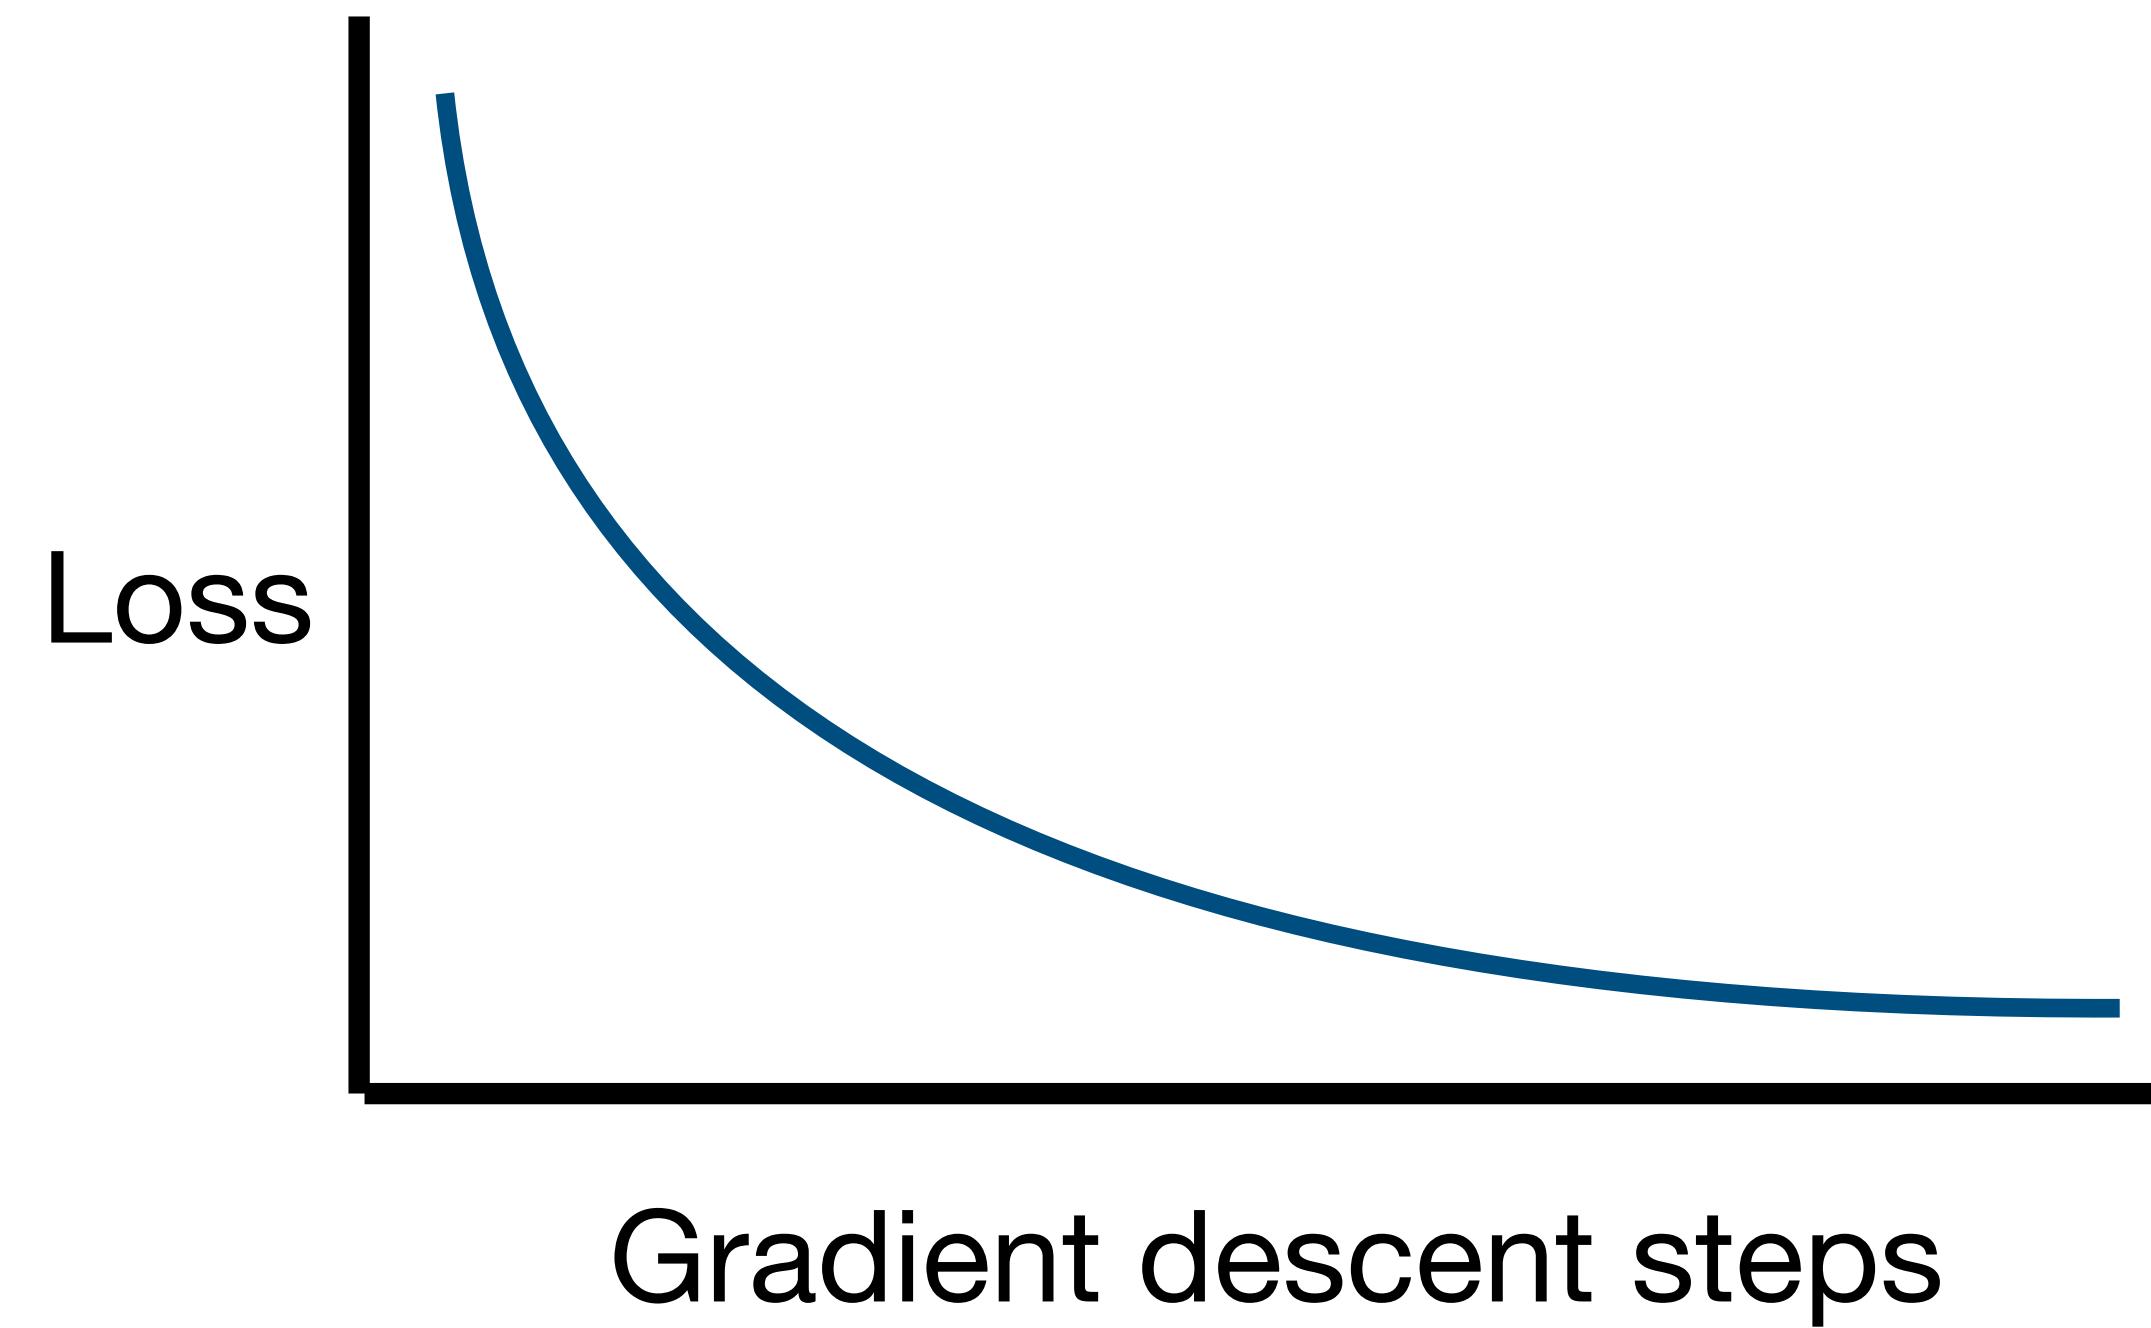
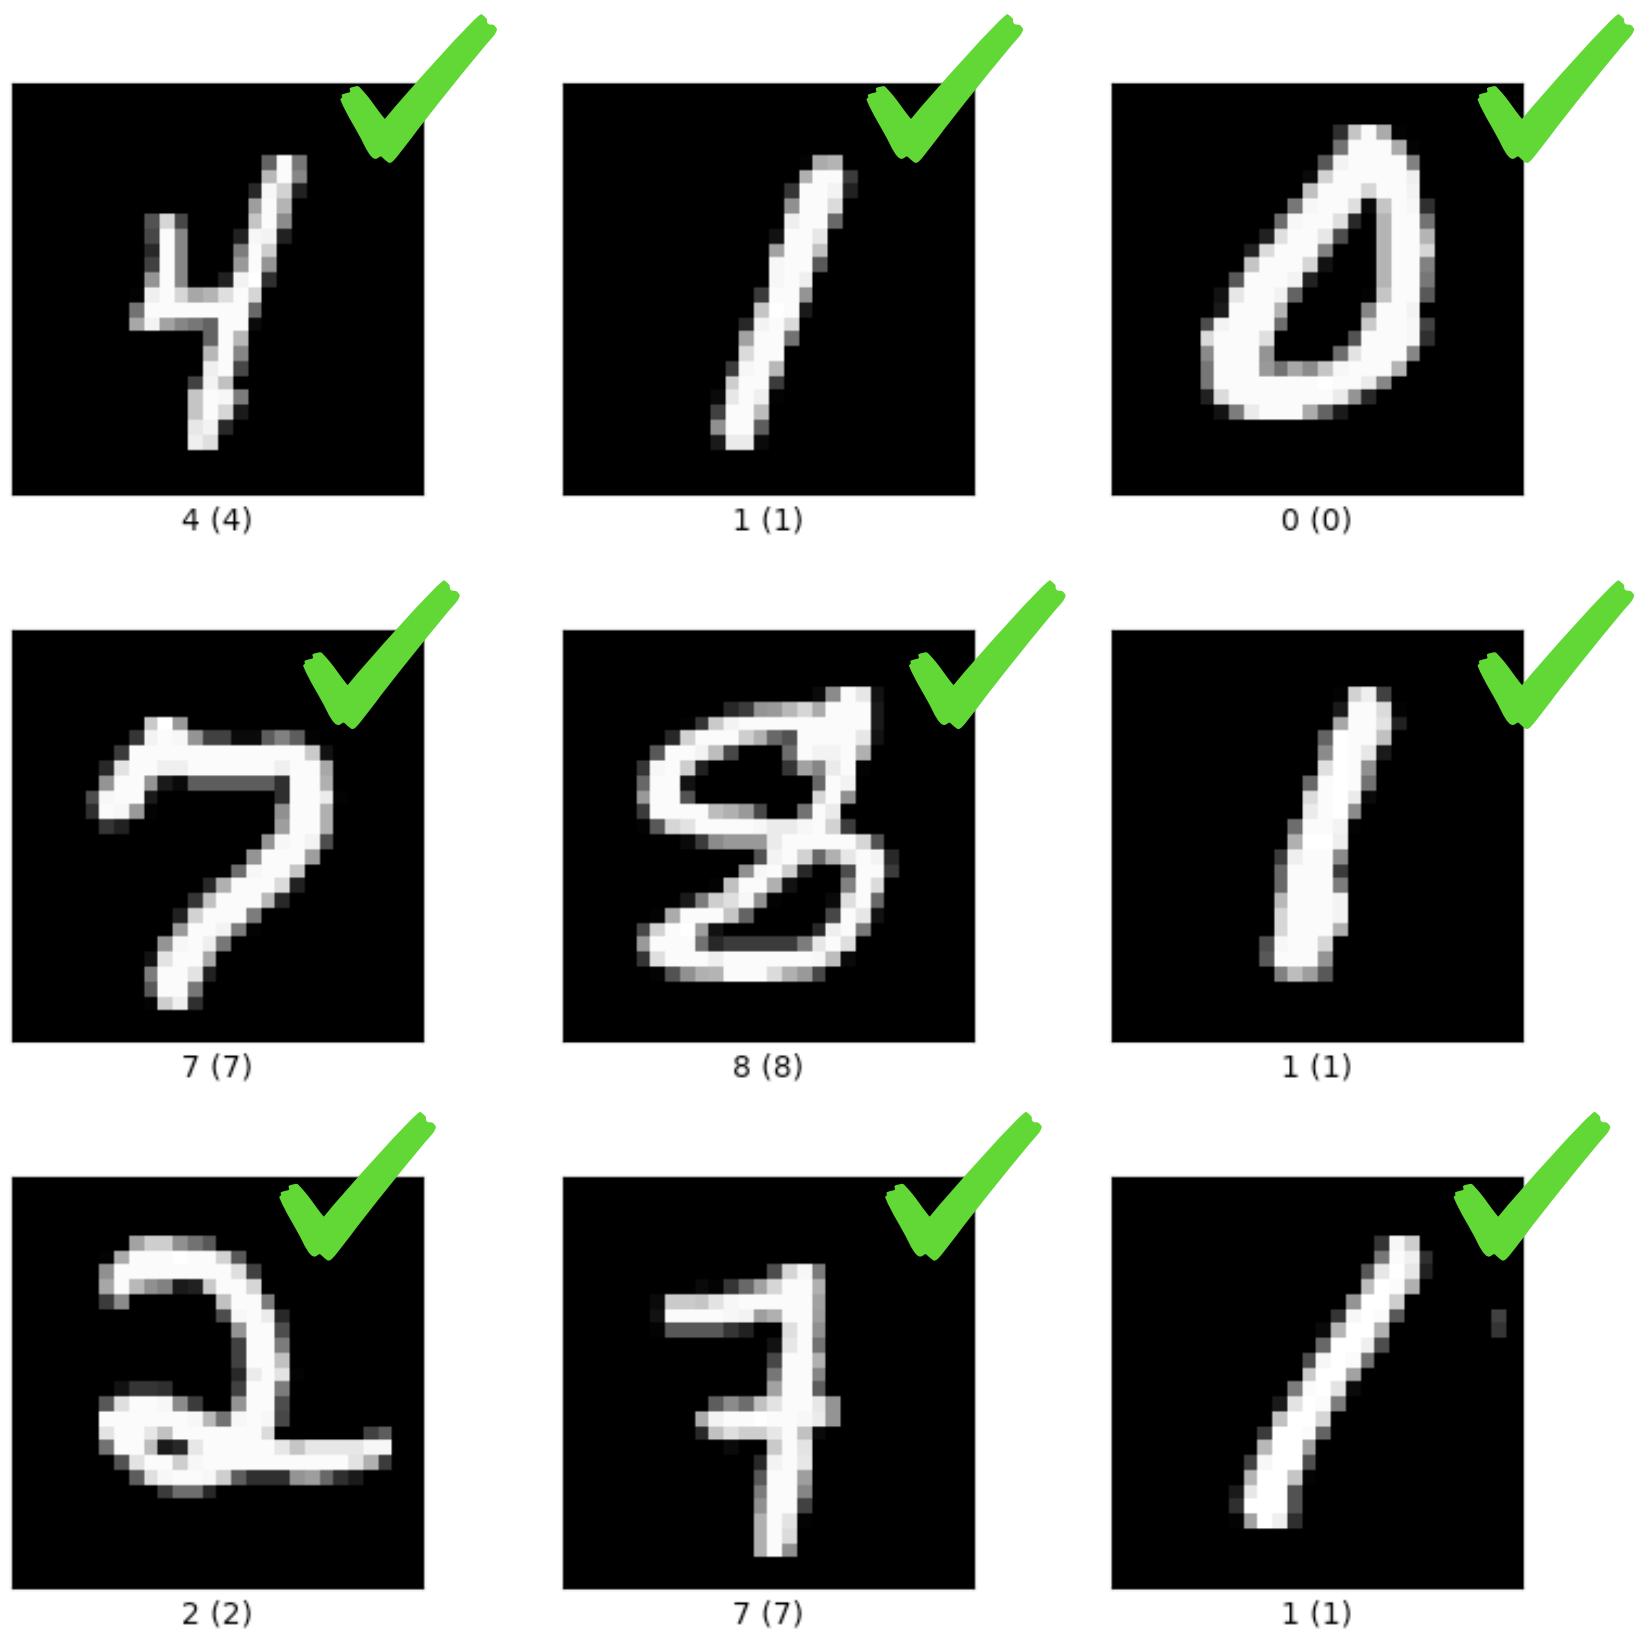
Hidden Layer $\in \mathbb{R}^8$

Hidden Layer $\in \mathbb{R}^4$

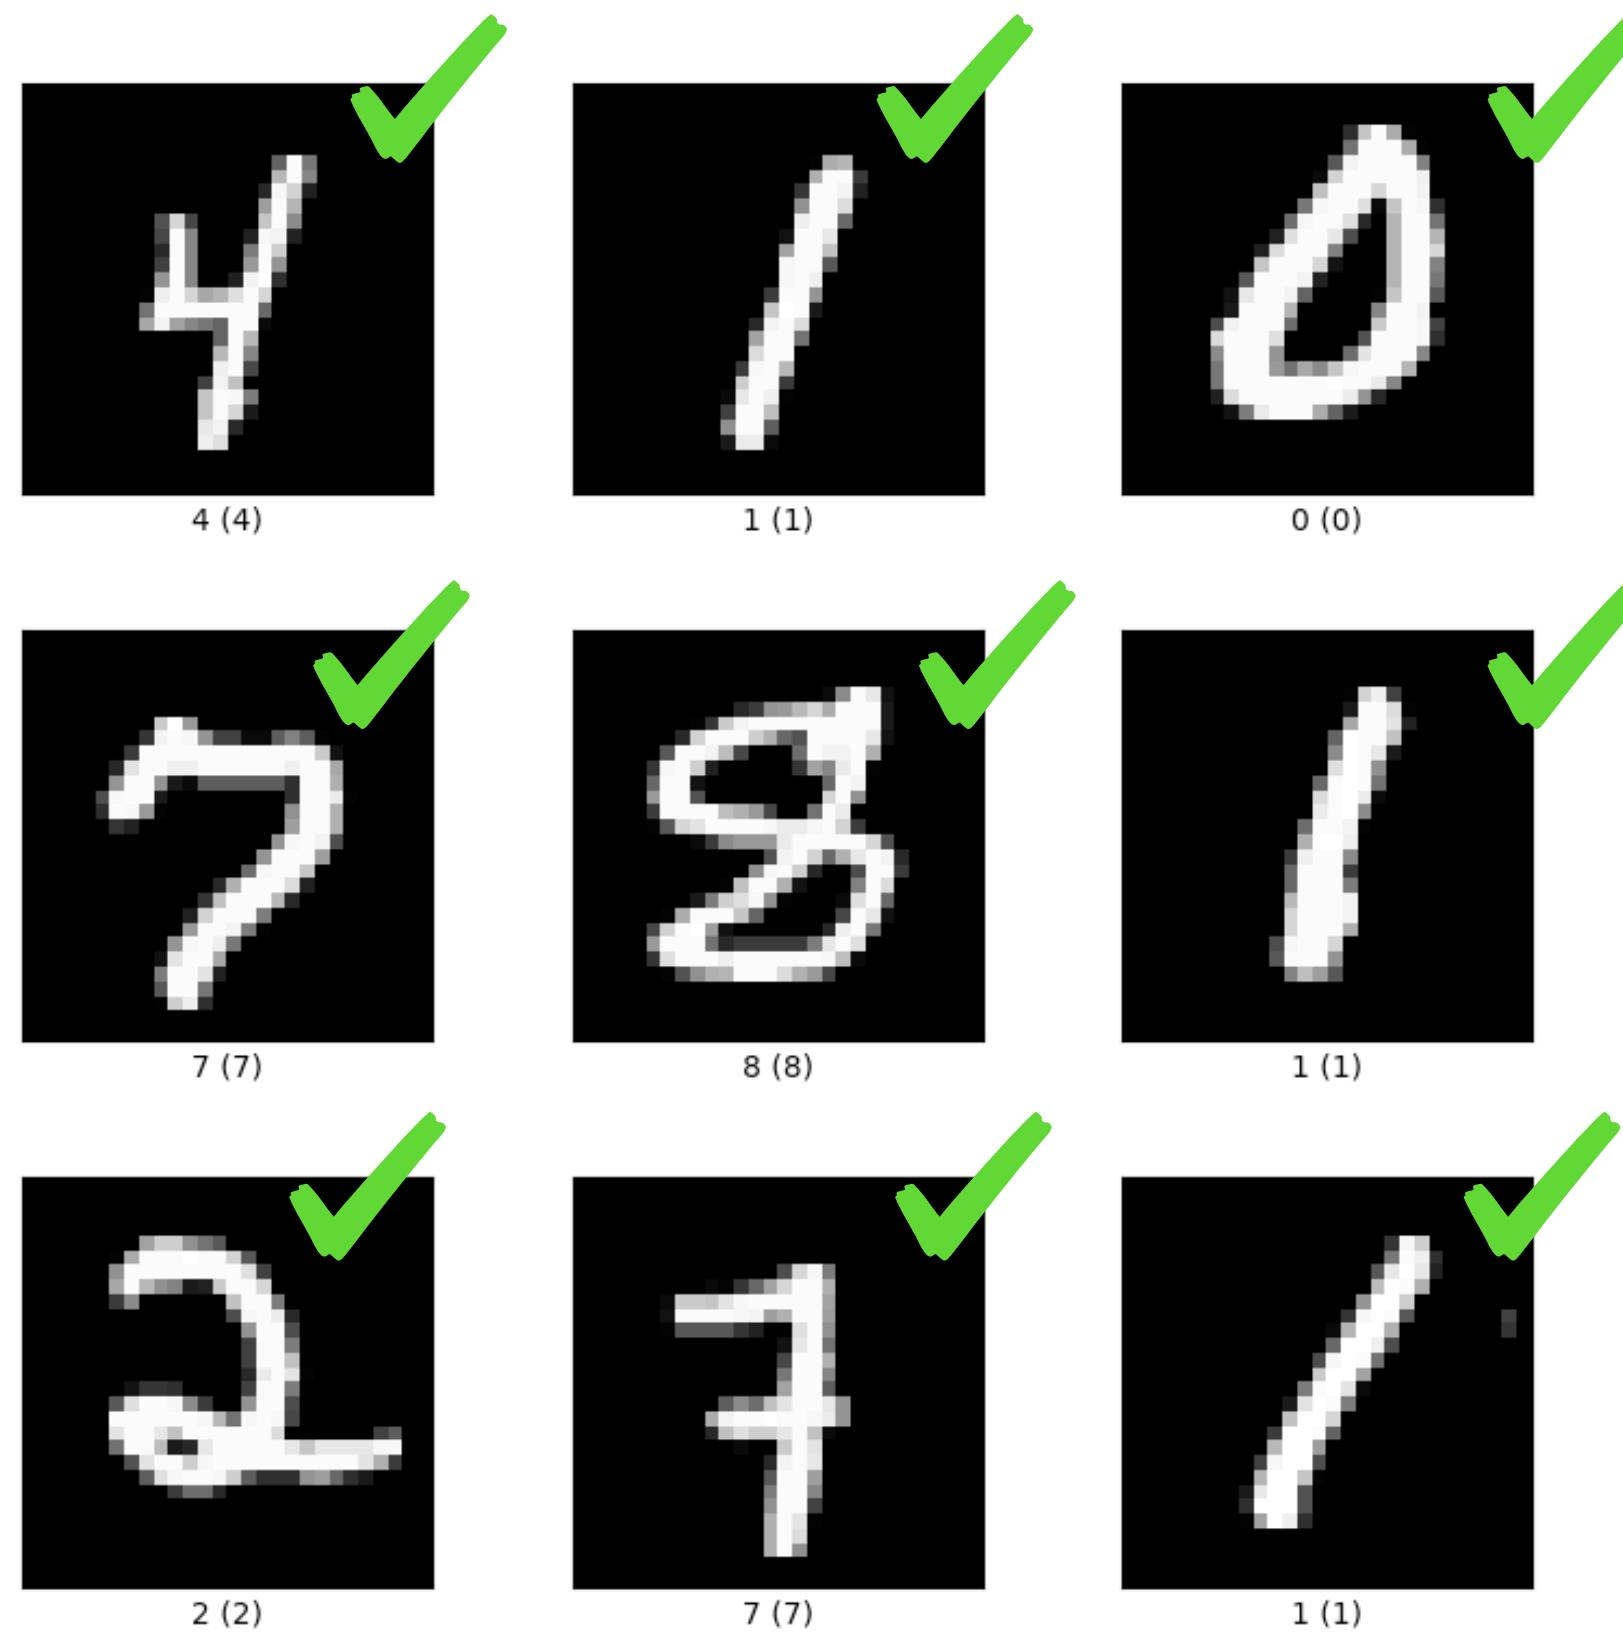
Output Layer $\in \mathbb{R}^1$



8

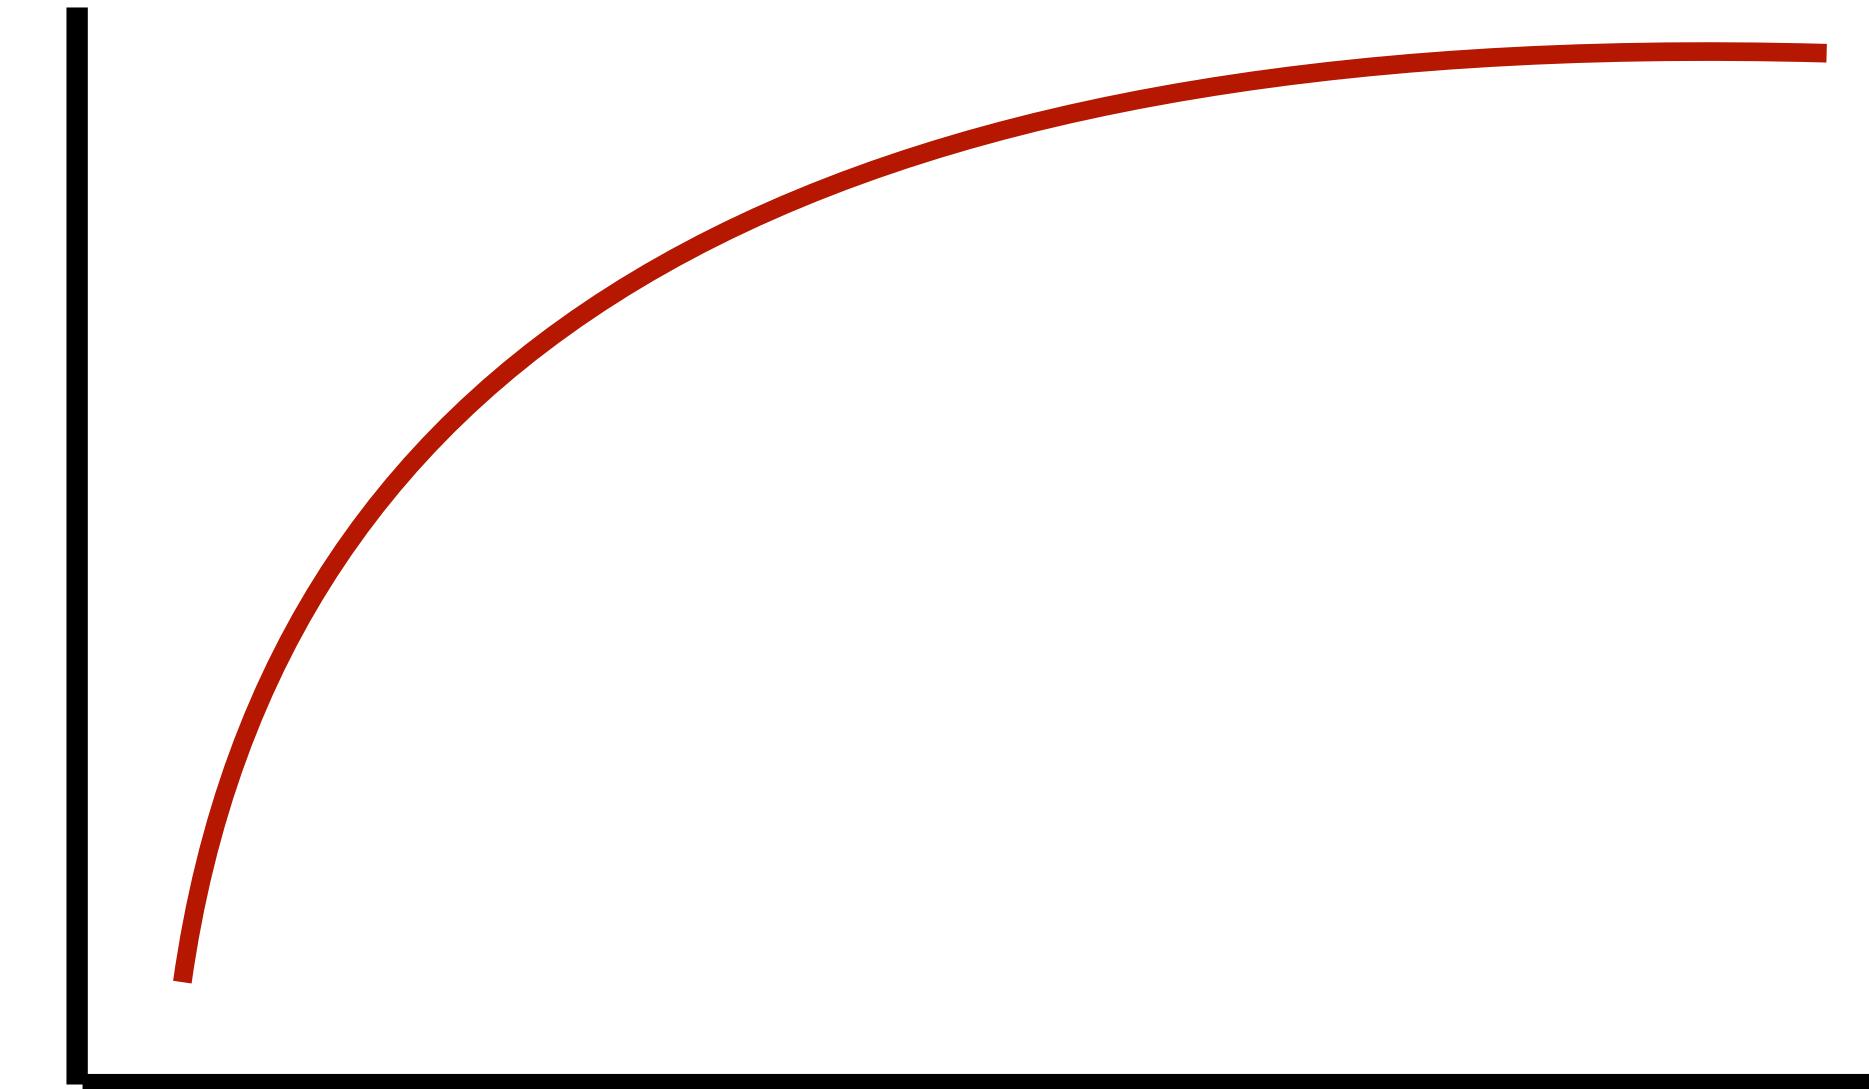


$$\text{Loss}(\mathbf{w}) = \text{NLL}(\mathbf{w}, \mathbf{X}, \mathbf{y}) = - \sum_{i=1}^N \log p(y_i \mid \mathbf{x}_i, \mathbf{w})$$



Accuracy

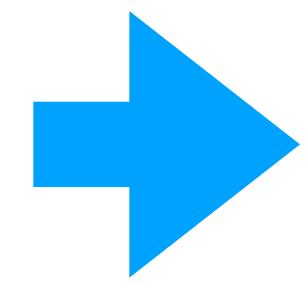
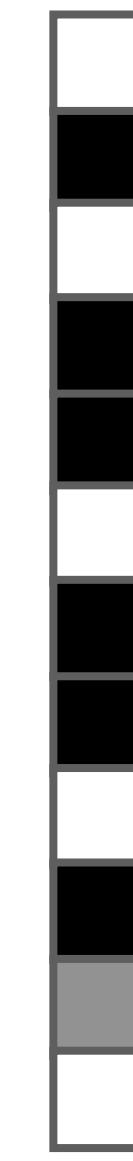
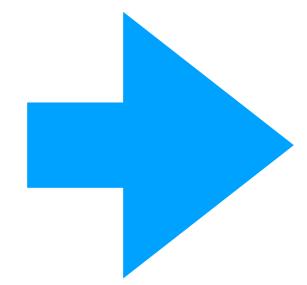
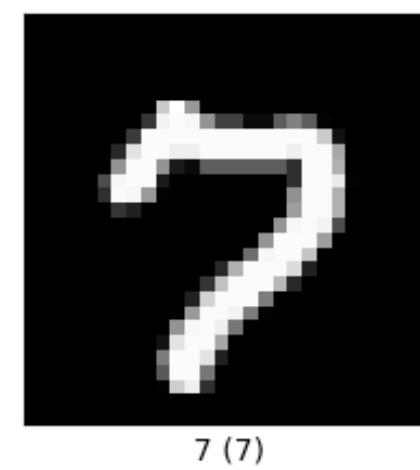
Accuracy:



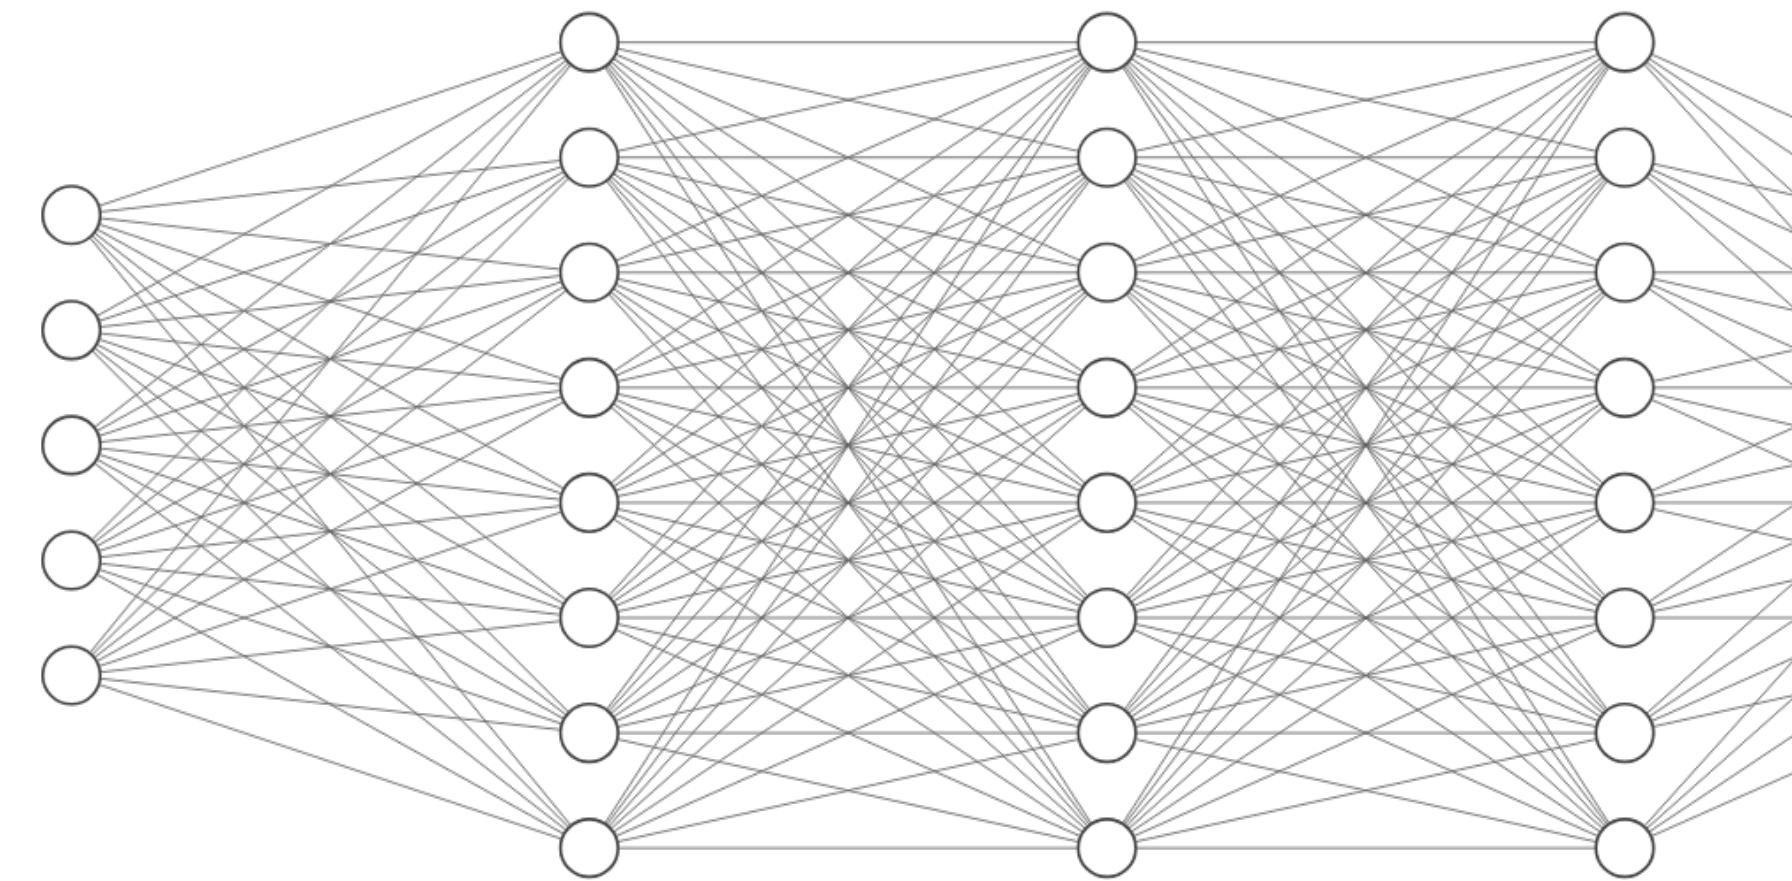
Gradient descent steps

$$\frac{\text{\# of correct predictions}}{\text{Total predictions}} = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(f(\mathbf{x}_i) = y_i)$$

Are we done?



Input Layer $\in \mathbb{R}^5$



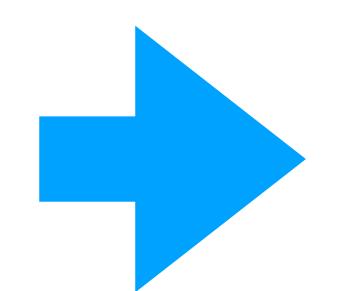
Hidden Layer $\in \mathbb{R}^8$

Hidden Layer $\in \mathbb{R}^8$

Hidden Layer $\in \mathbb{R}^8$

Hidden Layer $\in \mathbb{R}^4$

Output Layer $\in \mathbb{R}^1$

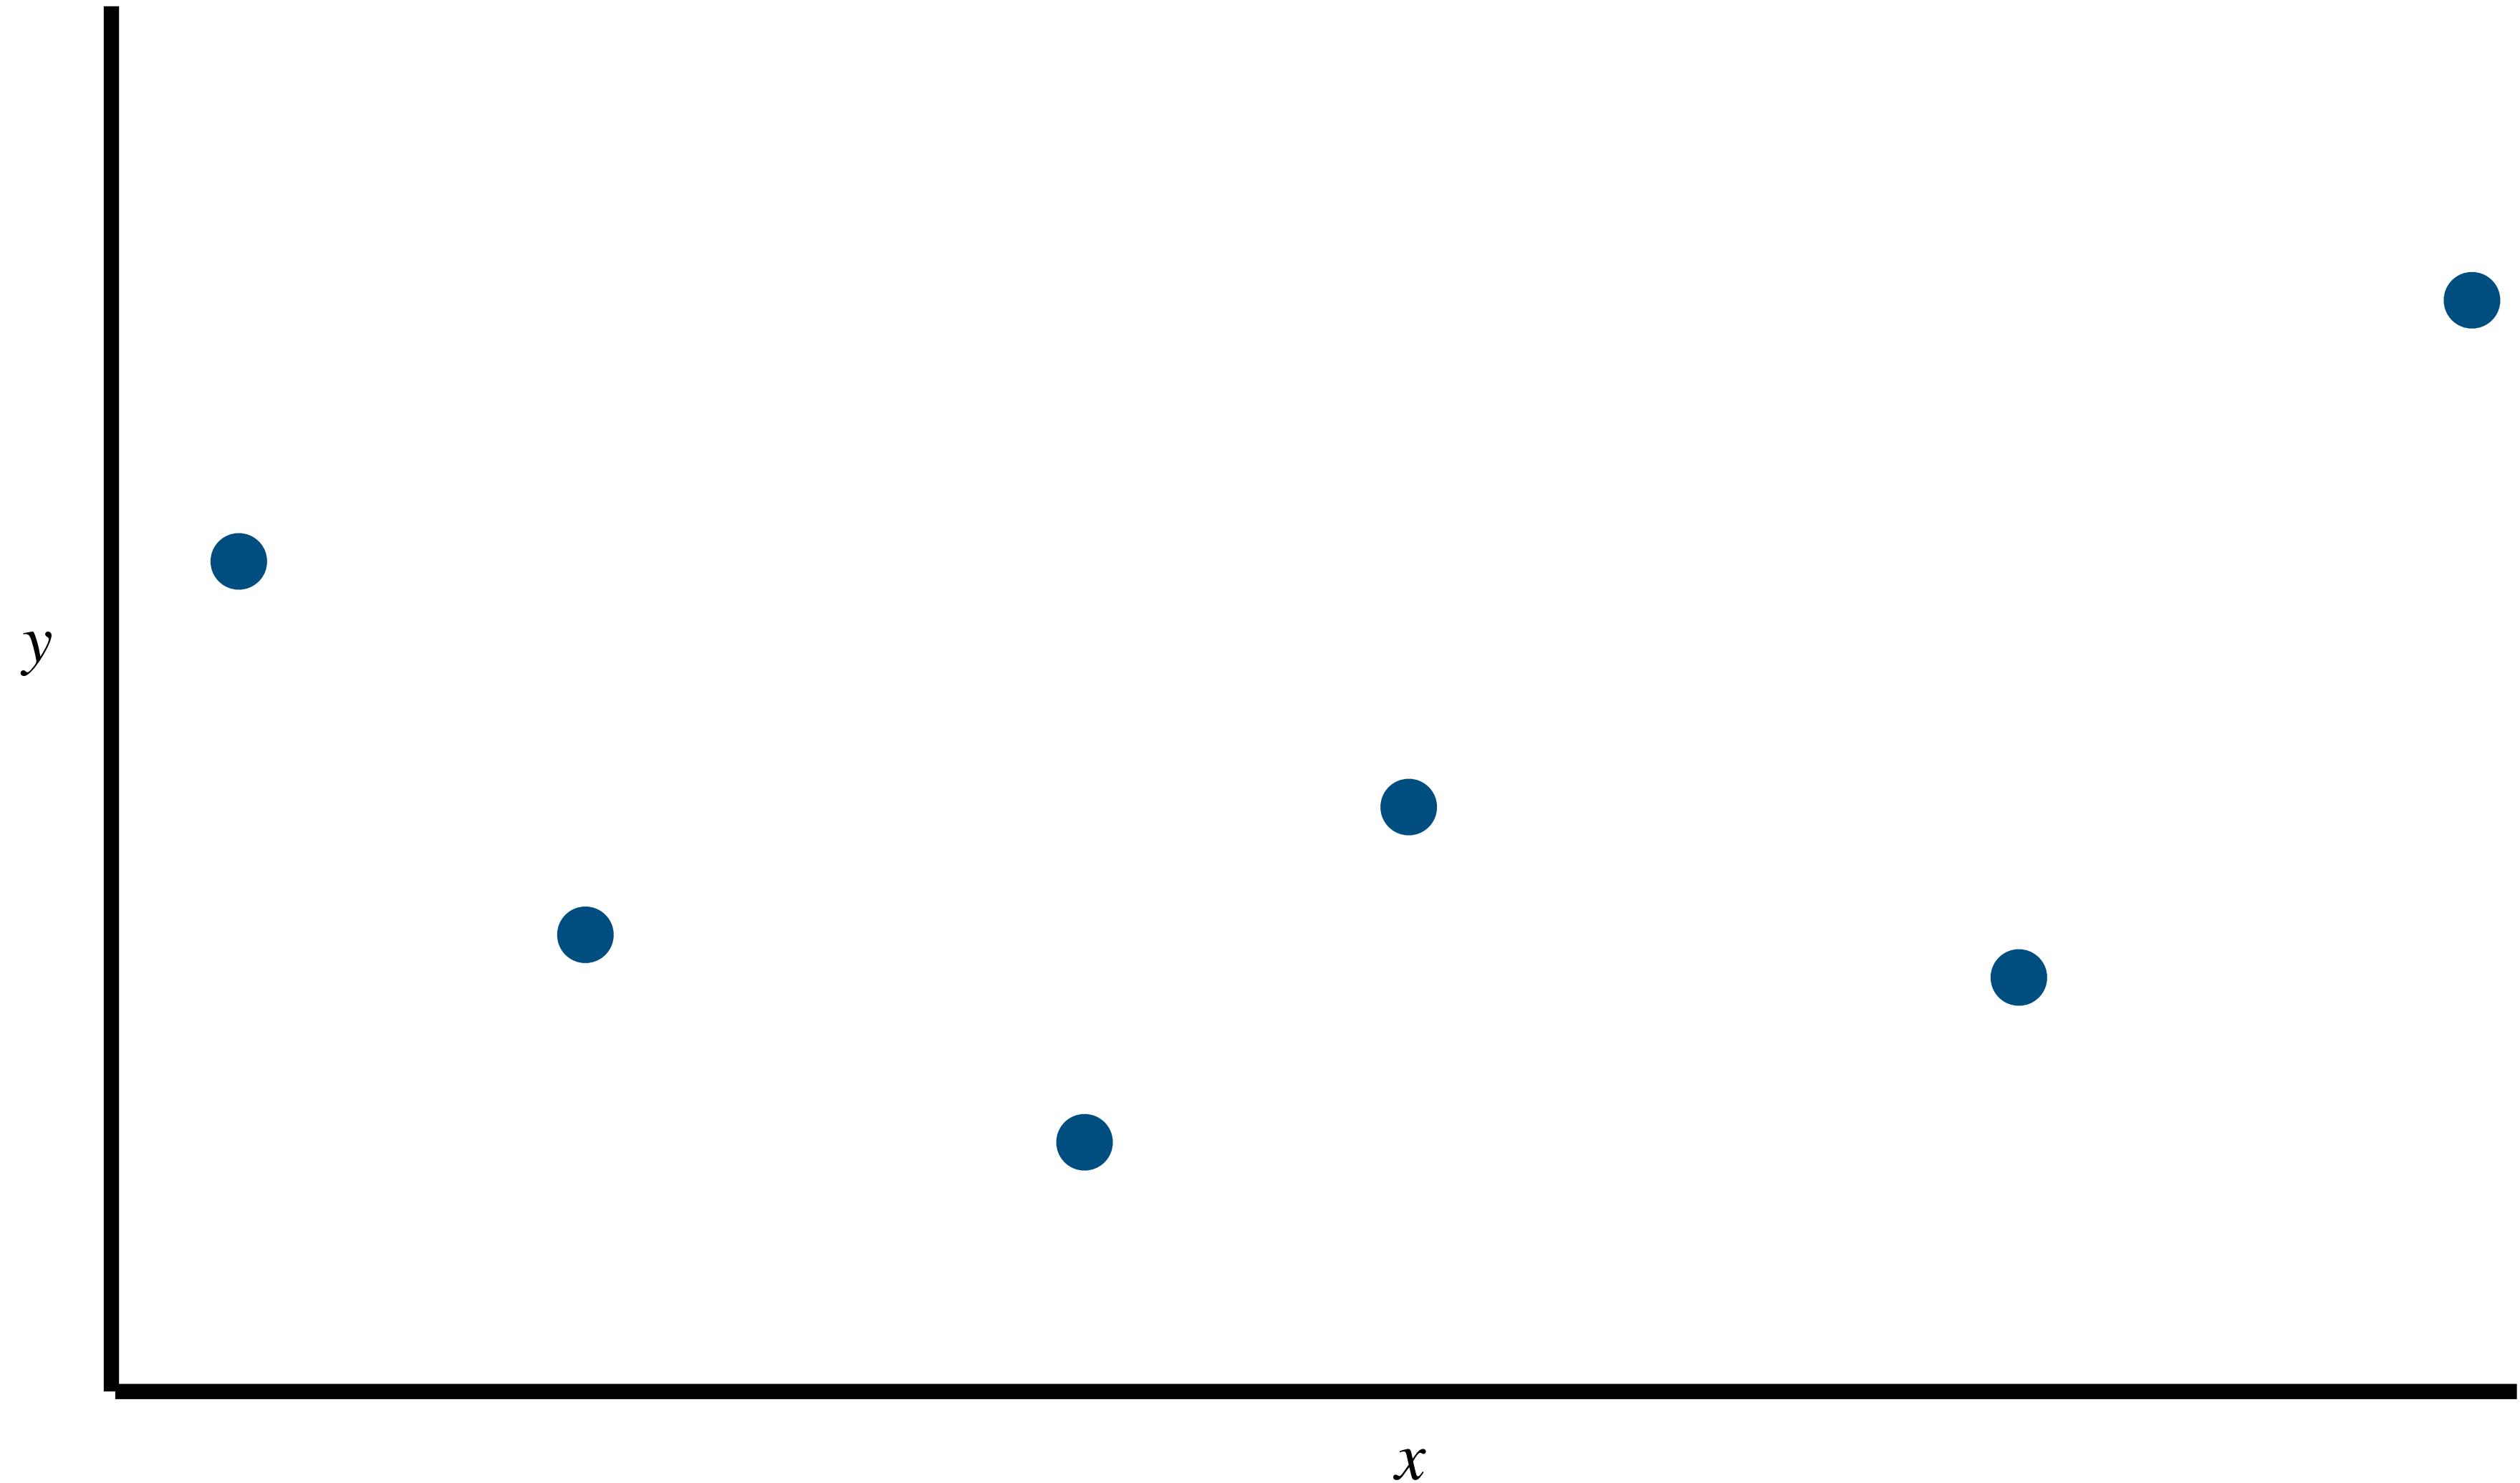


3

?

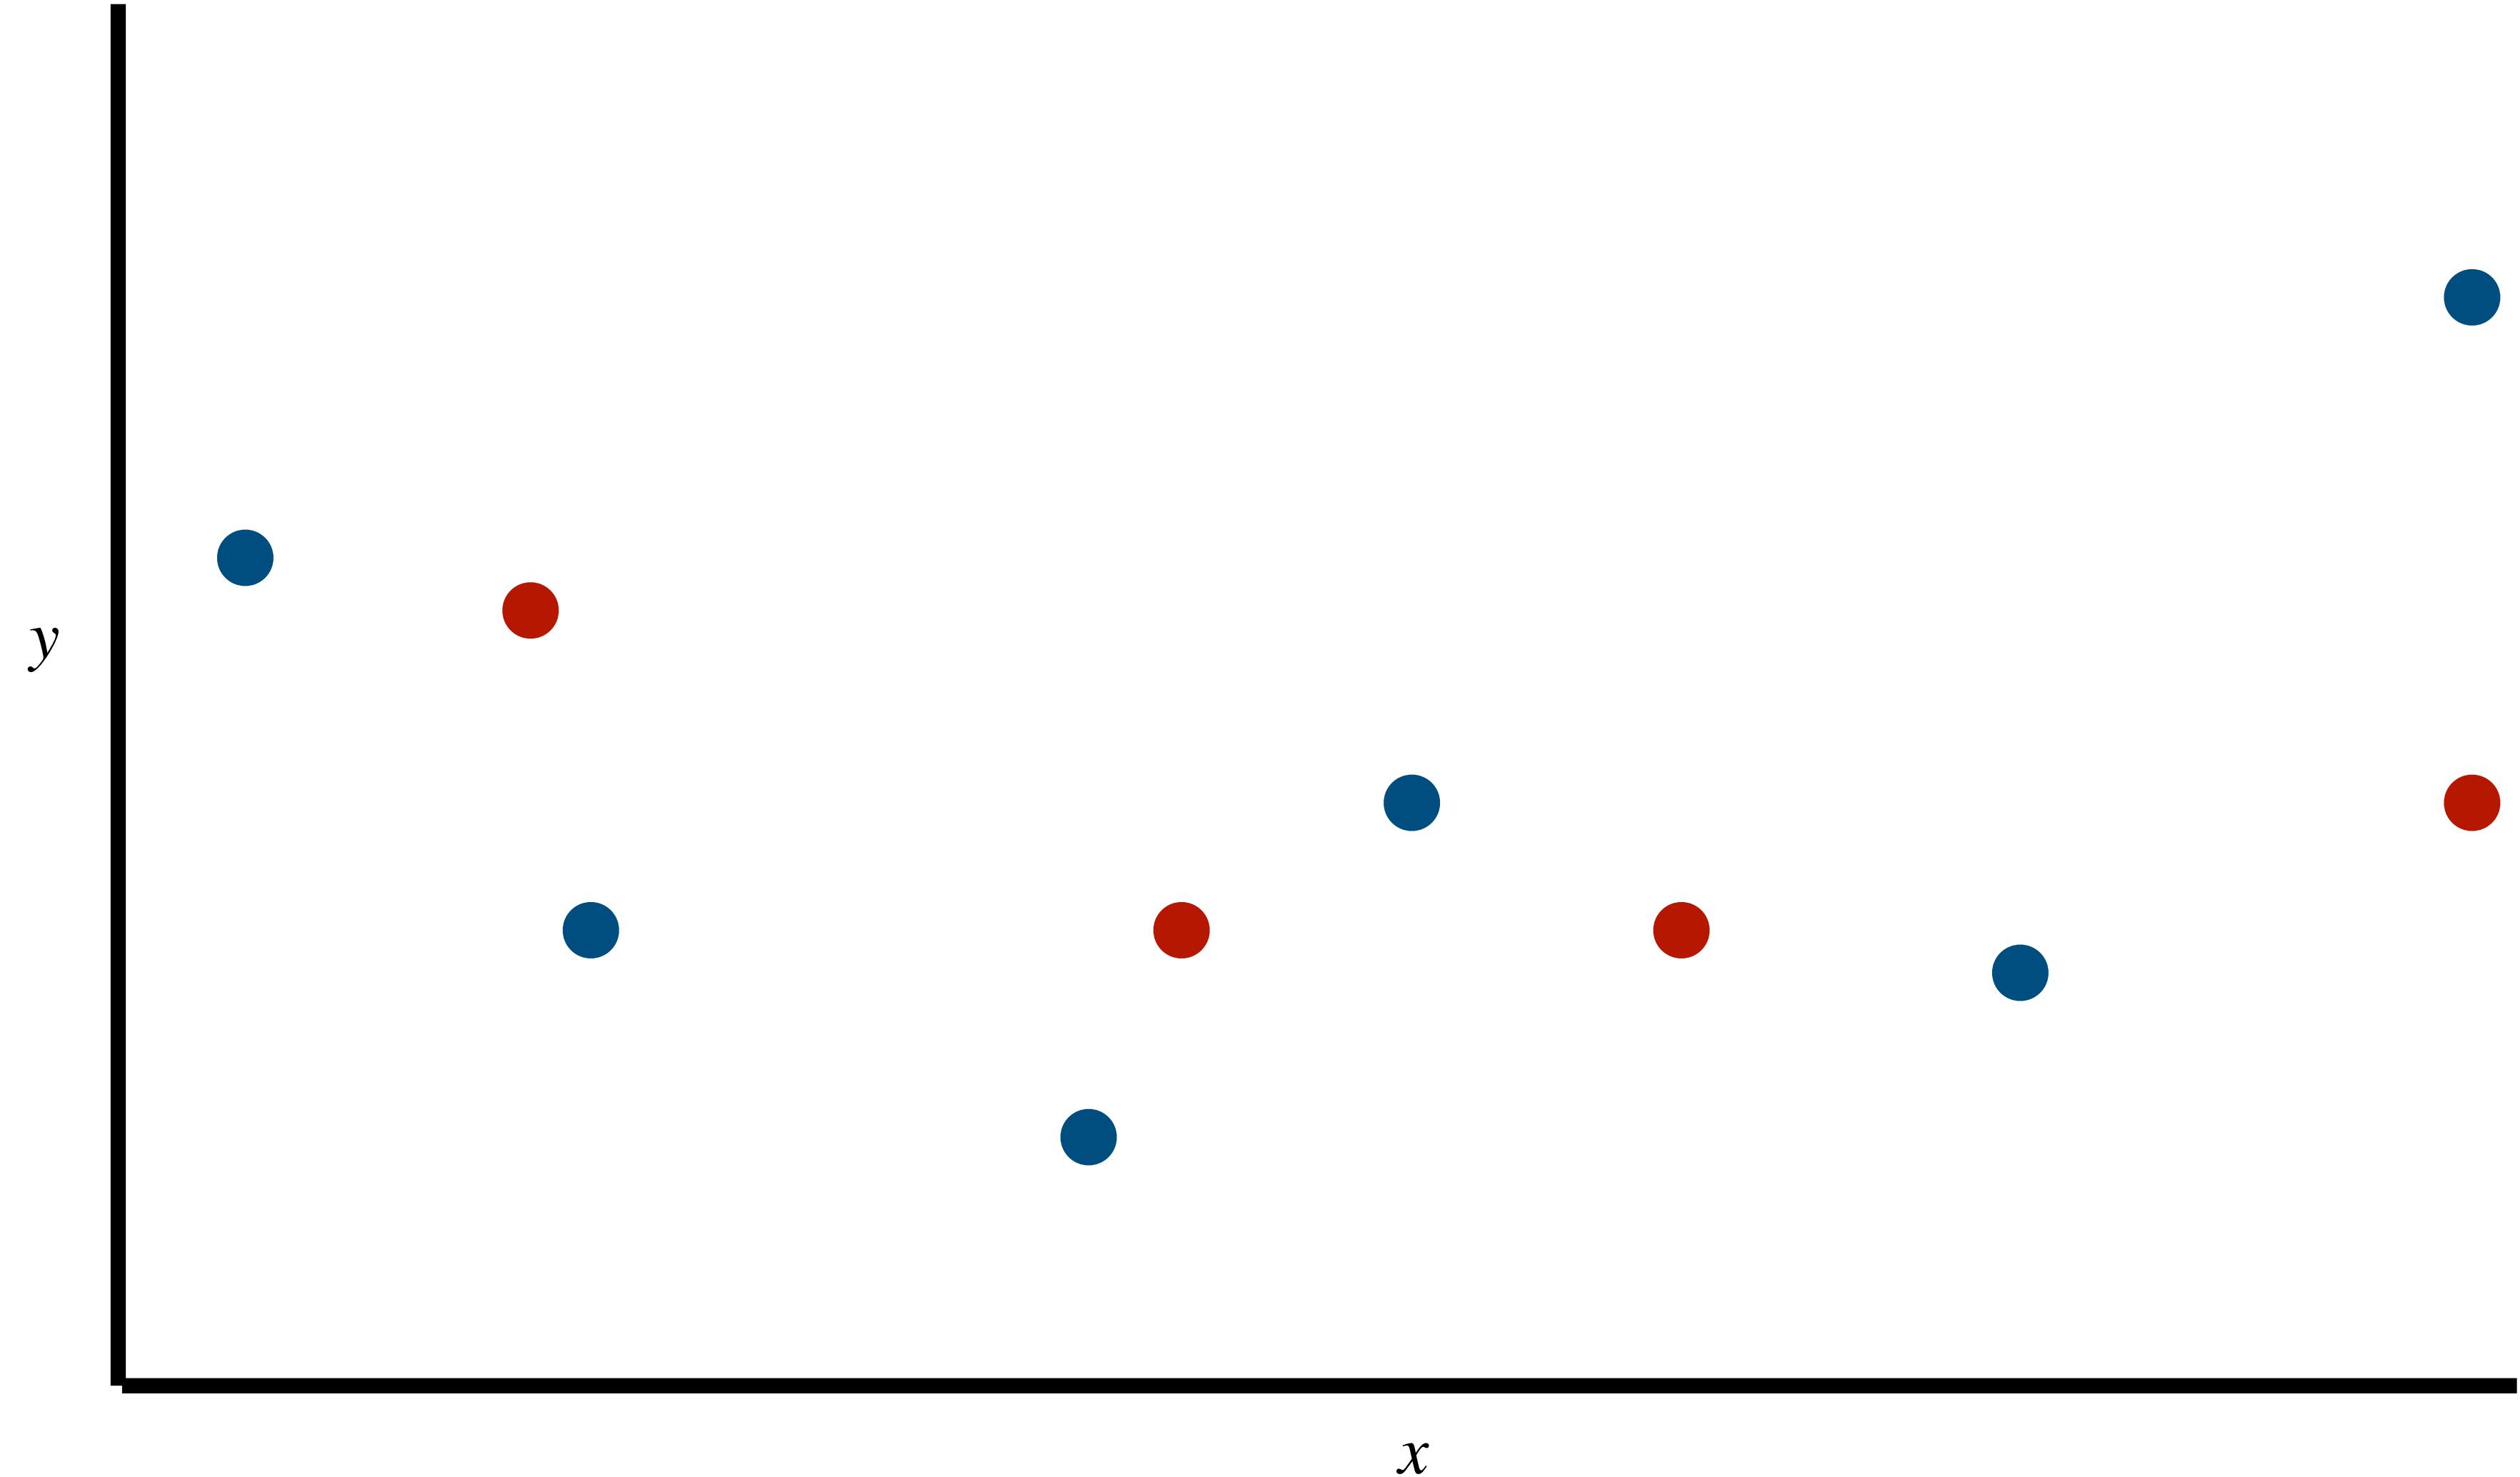
?

How many functions have 0 loss?



$$\text{MSE Loss} = \frac{1}{N} \sum_{i=1}^N (f(x_i) - y_i)^2$$

What if we then see new data?



$$\text{MSE Loss} = \frac{1}{N} \sum_{i=1}^N (f(x_i) - y_i)^2$$

Data = \mathbf{X}, \mathbf{y}



Training data = $\mathbf{X}_{train}, \mathbf{y}_{train}$



 Fit model

Split

Test data = $\mathbf{X}_{test}, \mathbf{y}_{test}$

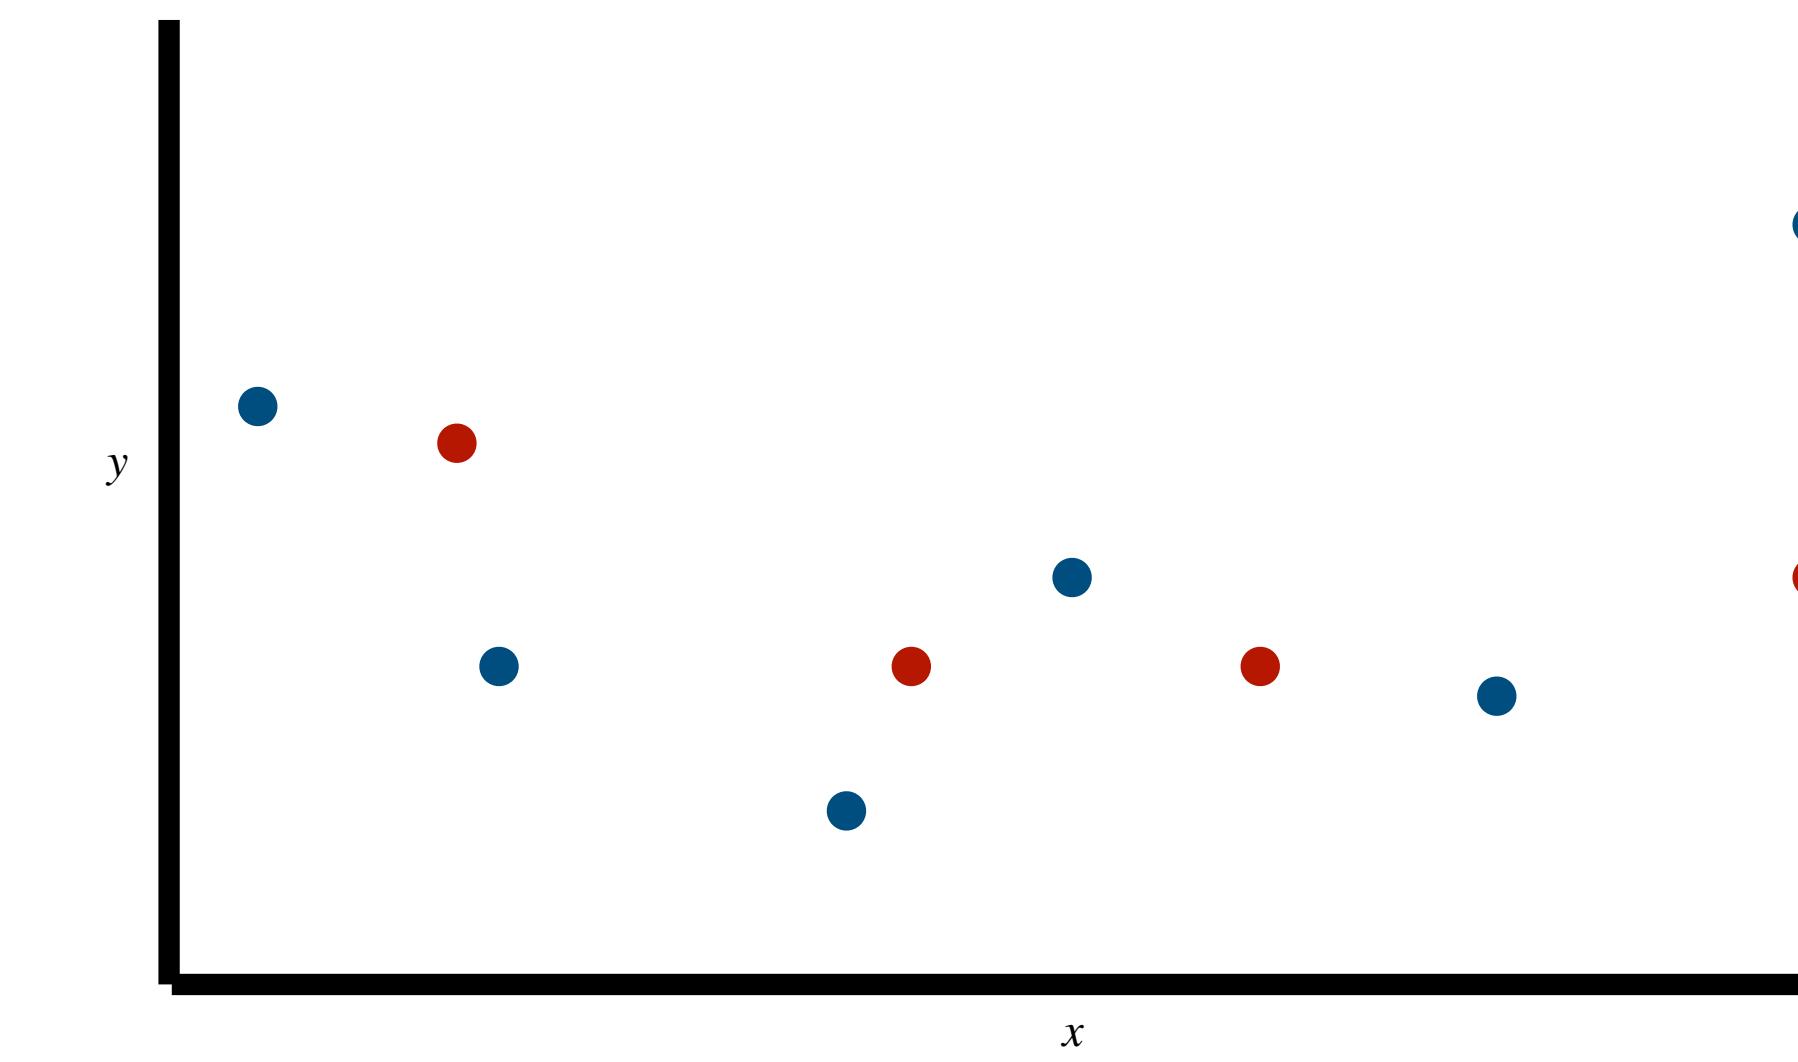
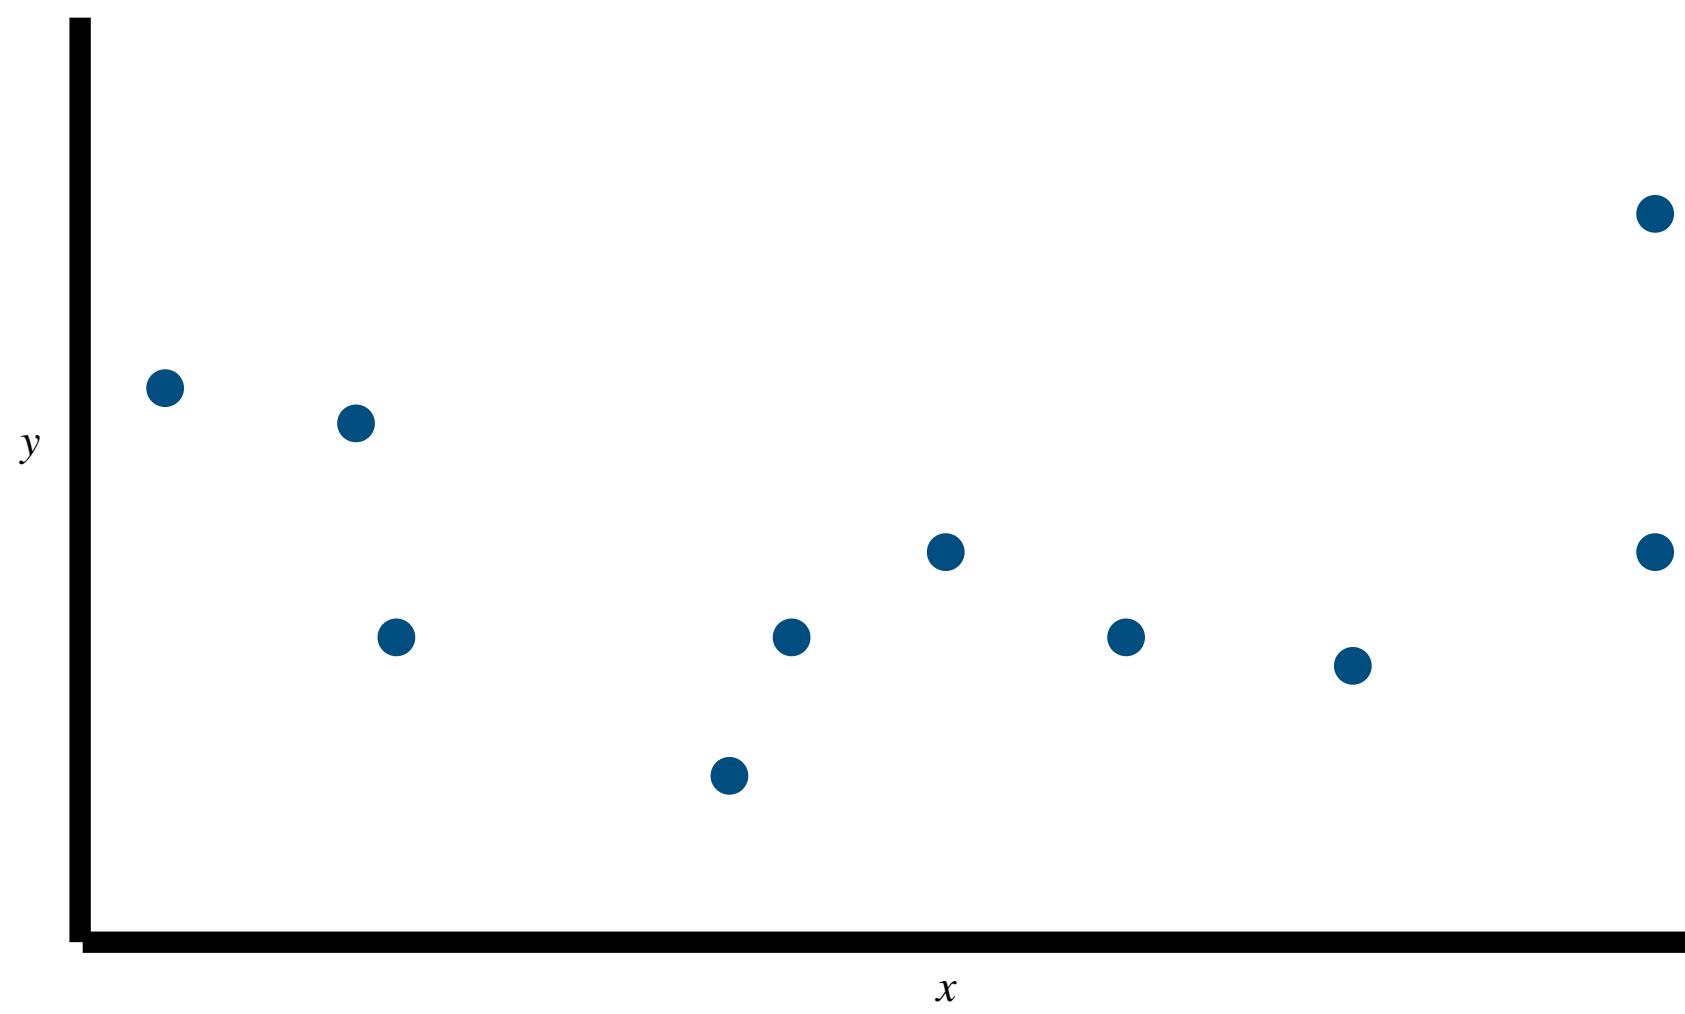


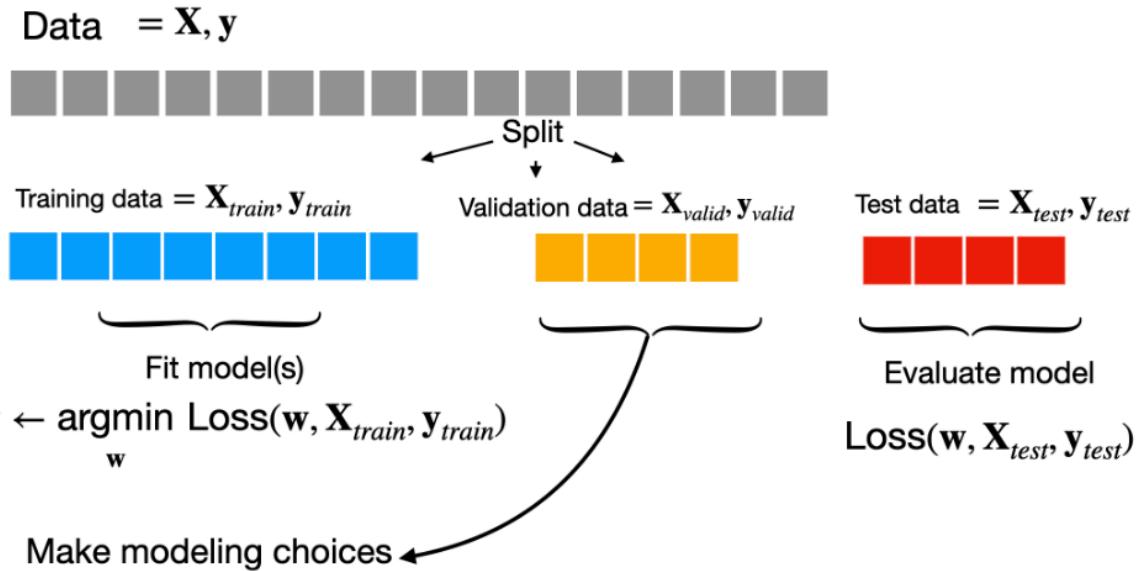
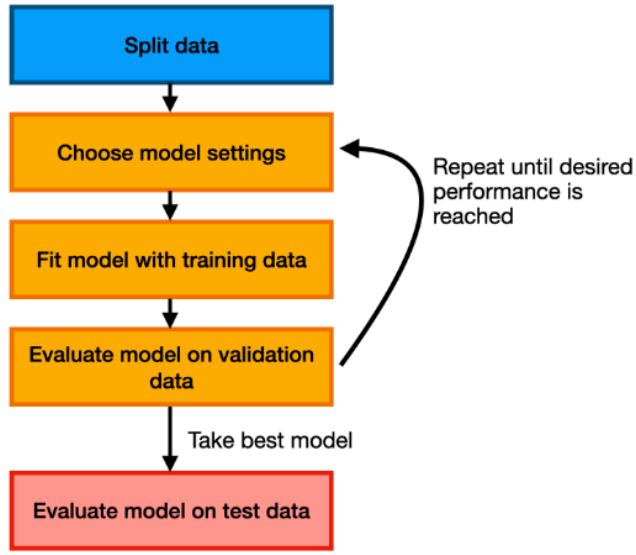
 Evaluate model

$\mathbf{w} \leftarrow \underset{\mathbf{w}}{\operatorname{argmin}} \text{Loss}(\mathbf{w}, \mathbf{X}_{train}, \mathbf{y}_{train})$

Loss($\mathbf{w}, \mathbf{X}_{test}, \mathbf{y}_{test}$)

Train-test split

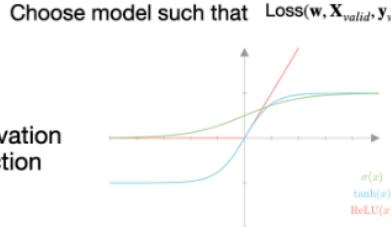




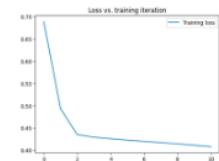
Number of layers and neurons

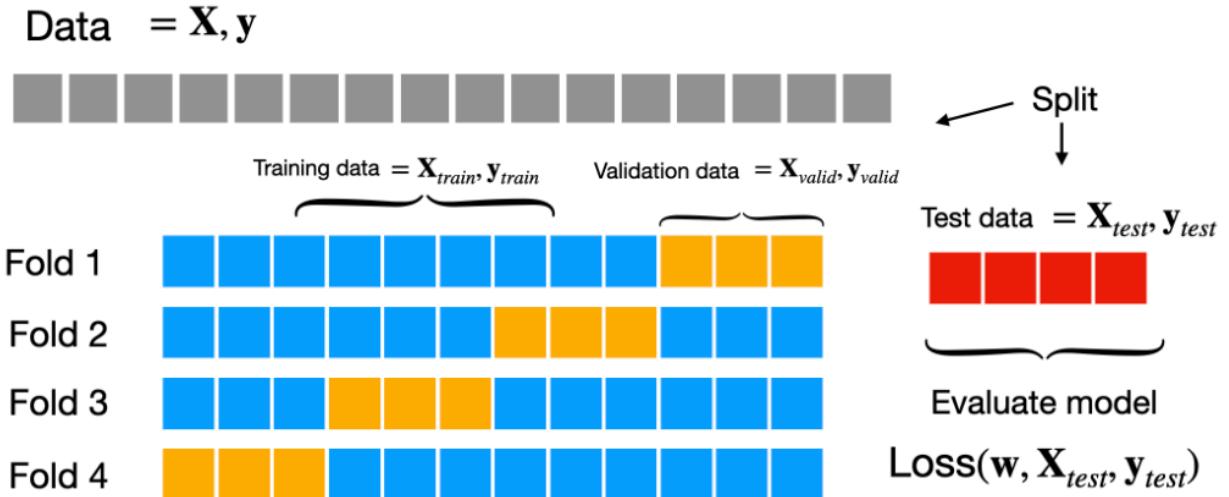
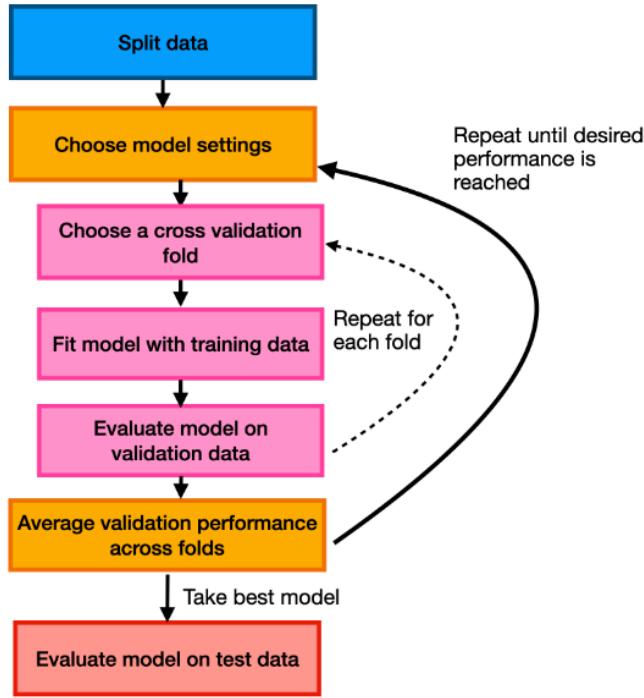


Activation function



Learning rate and gradient descent steps





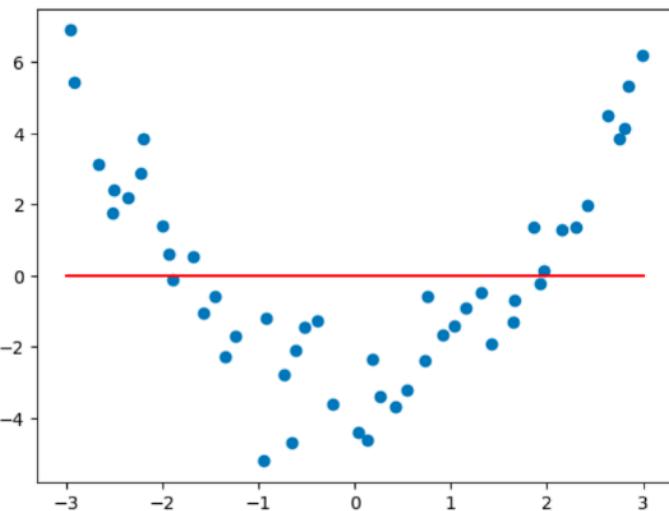
For each fold:

Fit model(s)

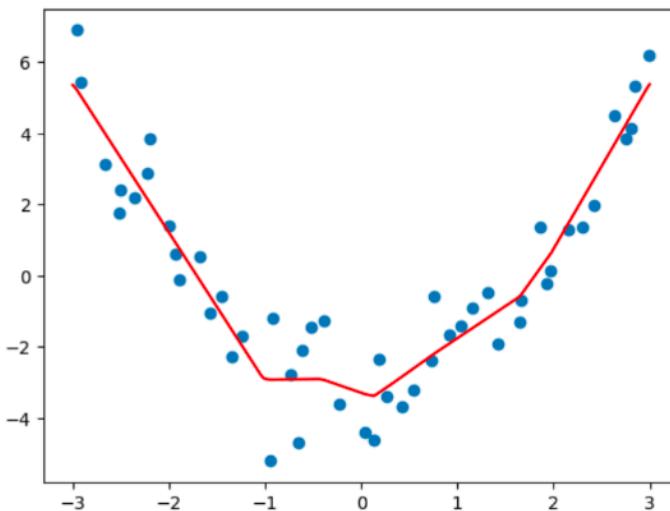
$$\mathbf{w} \leftarrow \underset{\mathbf{w}}{\operatorname{argmin}} \text{Loss}(\mathbf{w}, \mathbf{X}_{train}, \mathbf{y}_{train})$$

Evaluate model

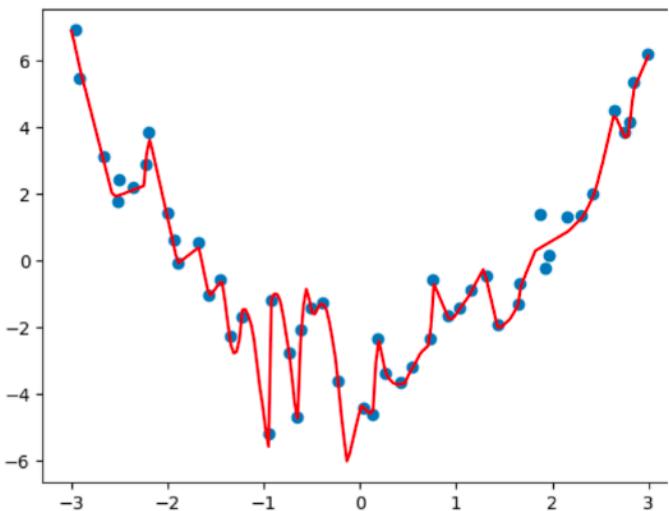
$$\text{Loss}(\mathbf{w}, \mathbf{X}_{\text{valid}}, \mathbf{y}_{\text{valid}})$$



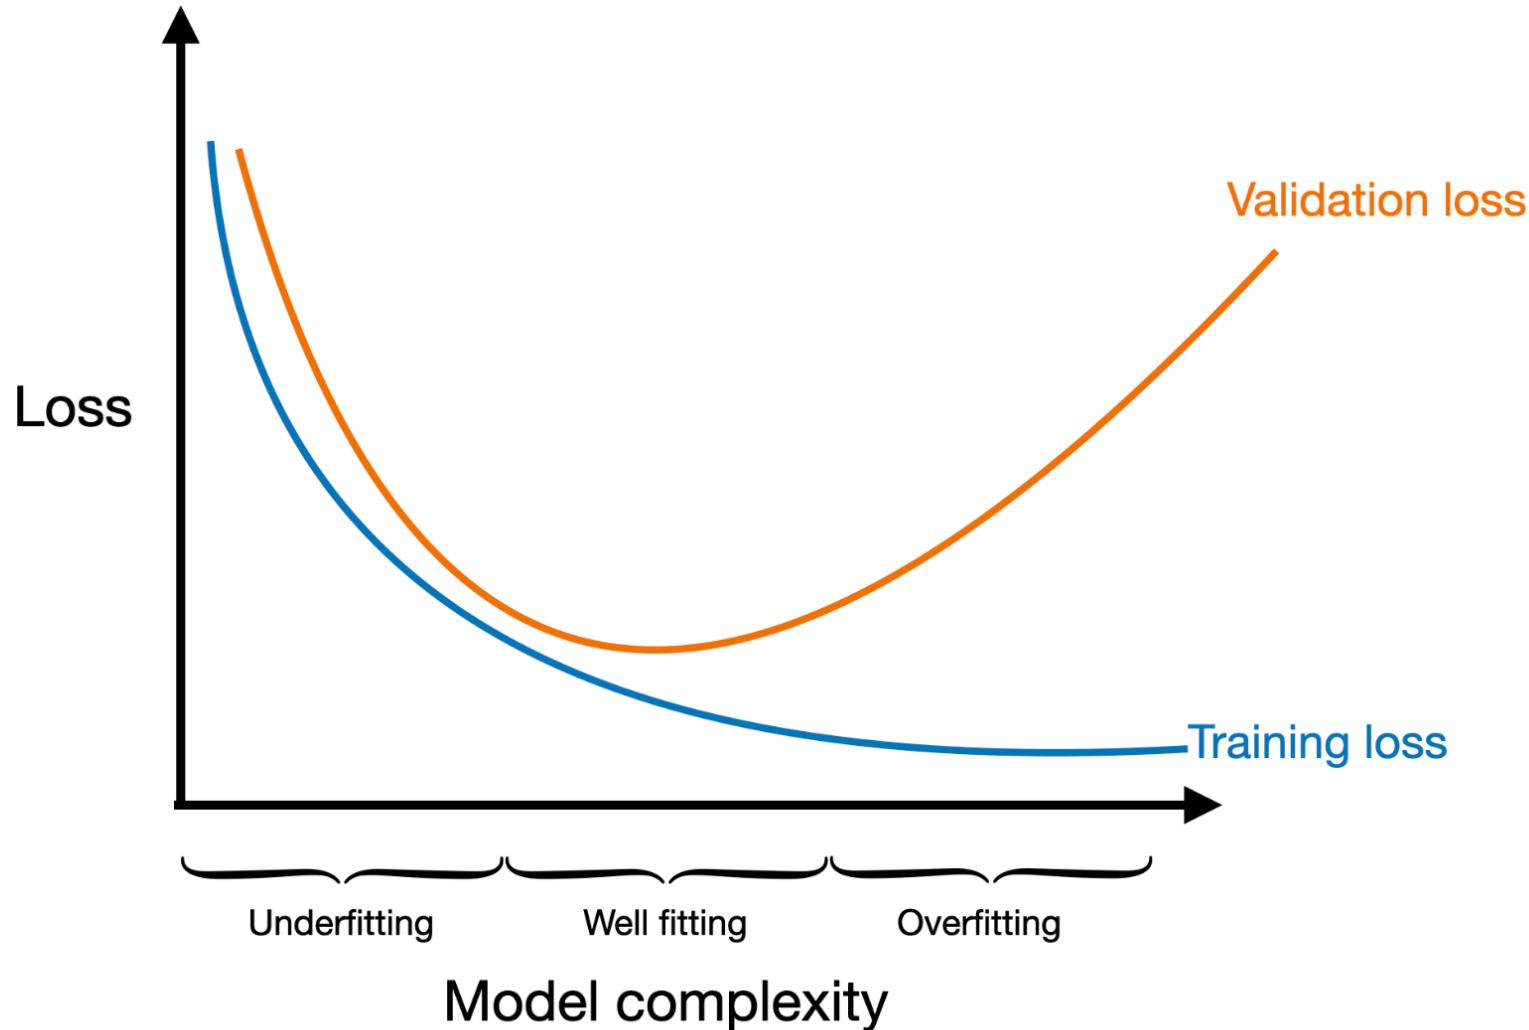
Underfitting

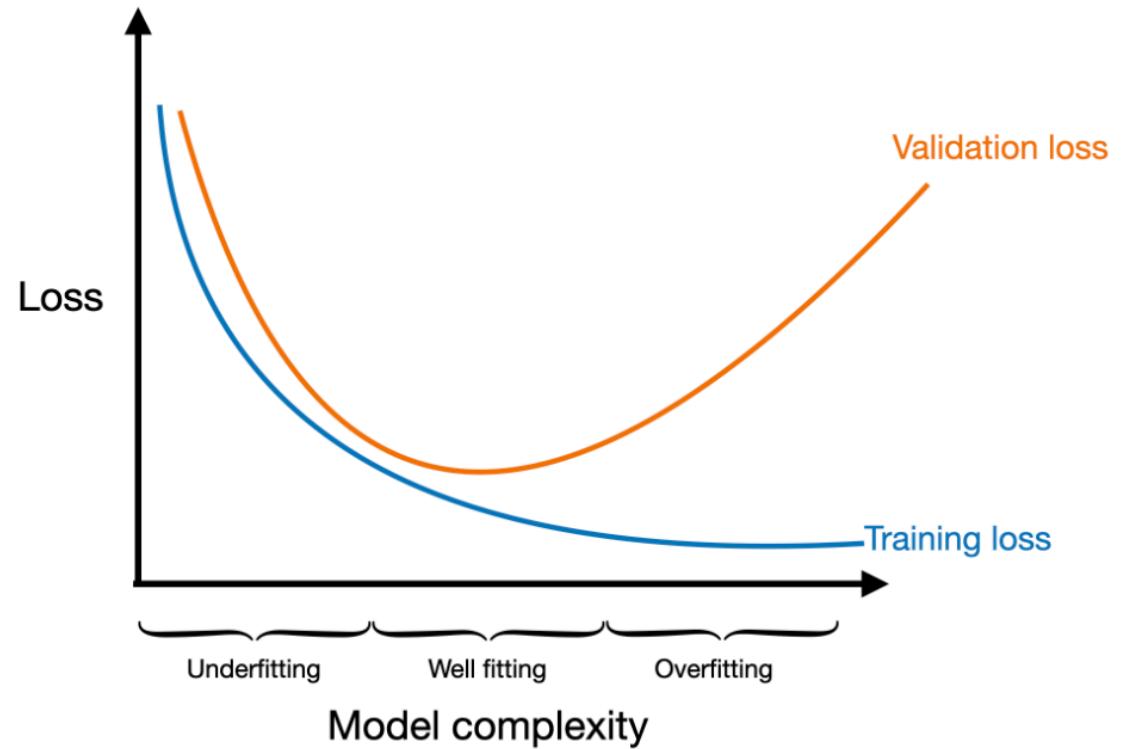
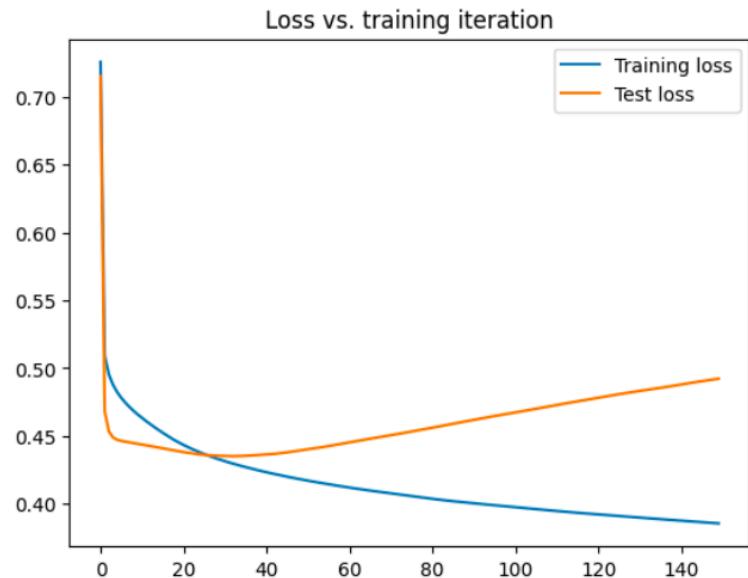


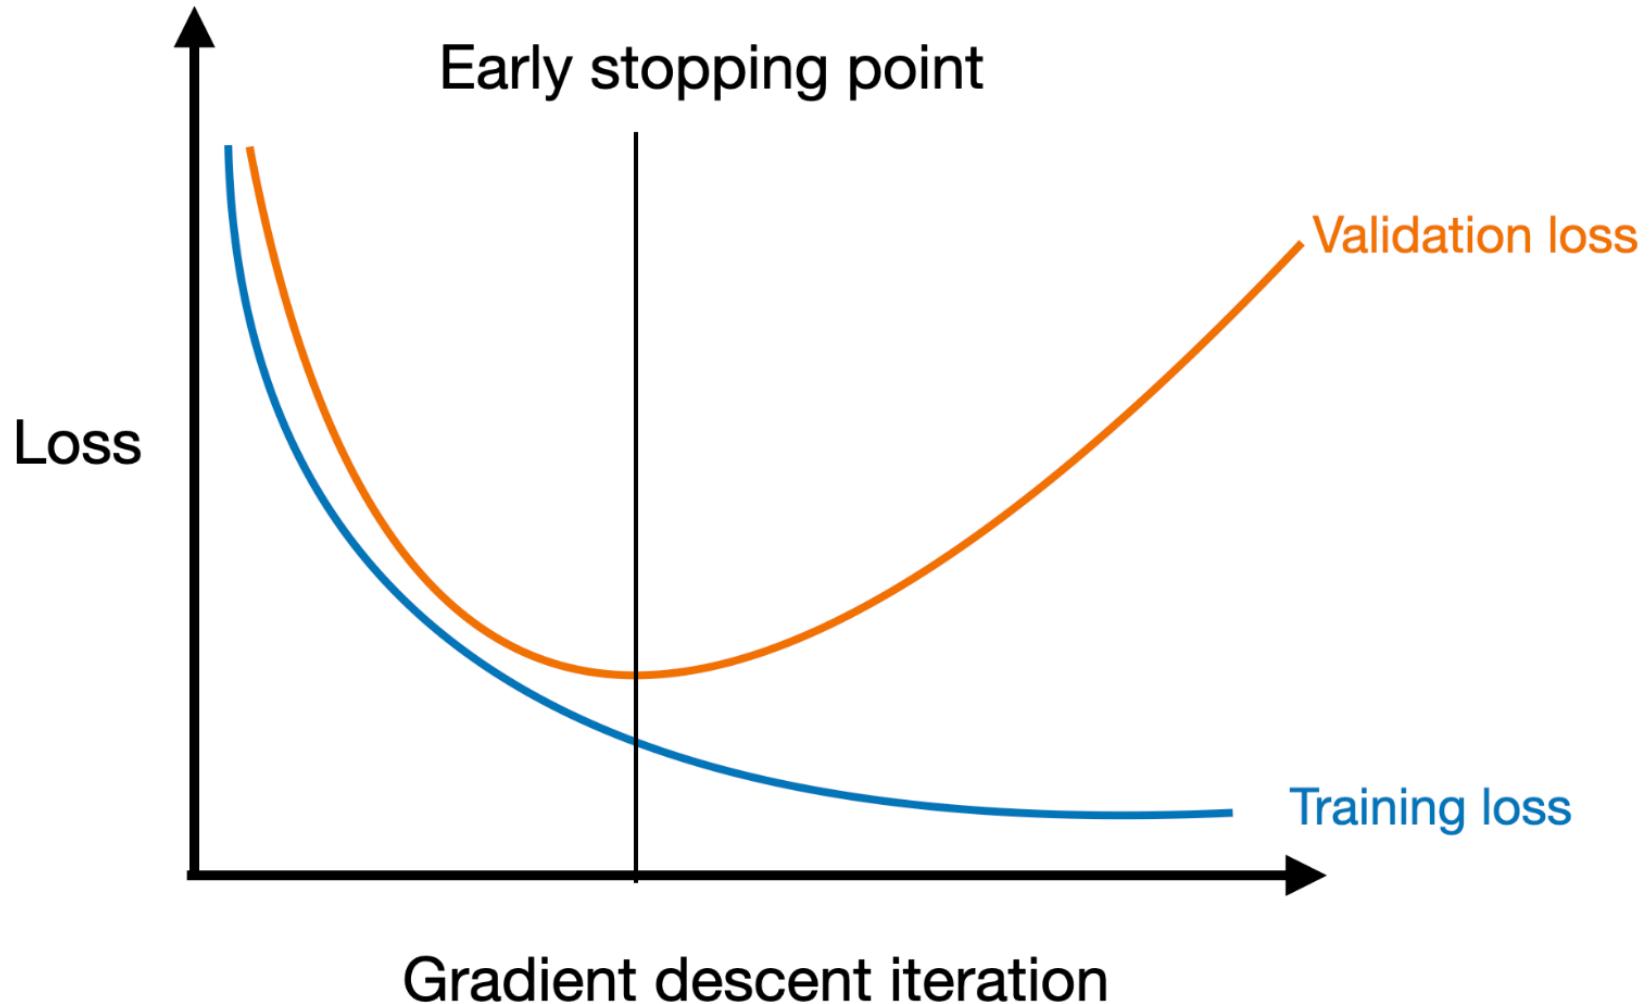
Well-fit

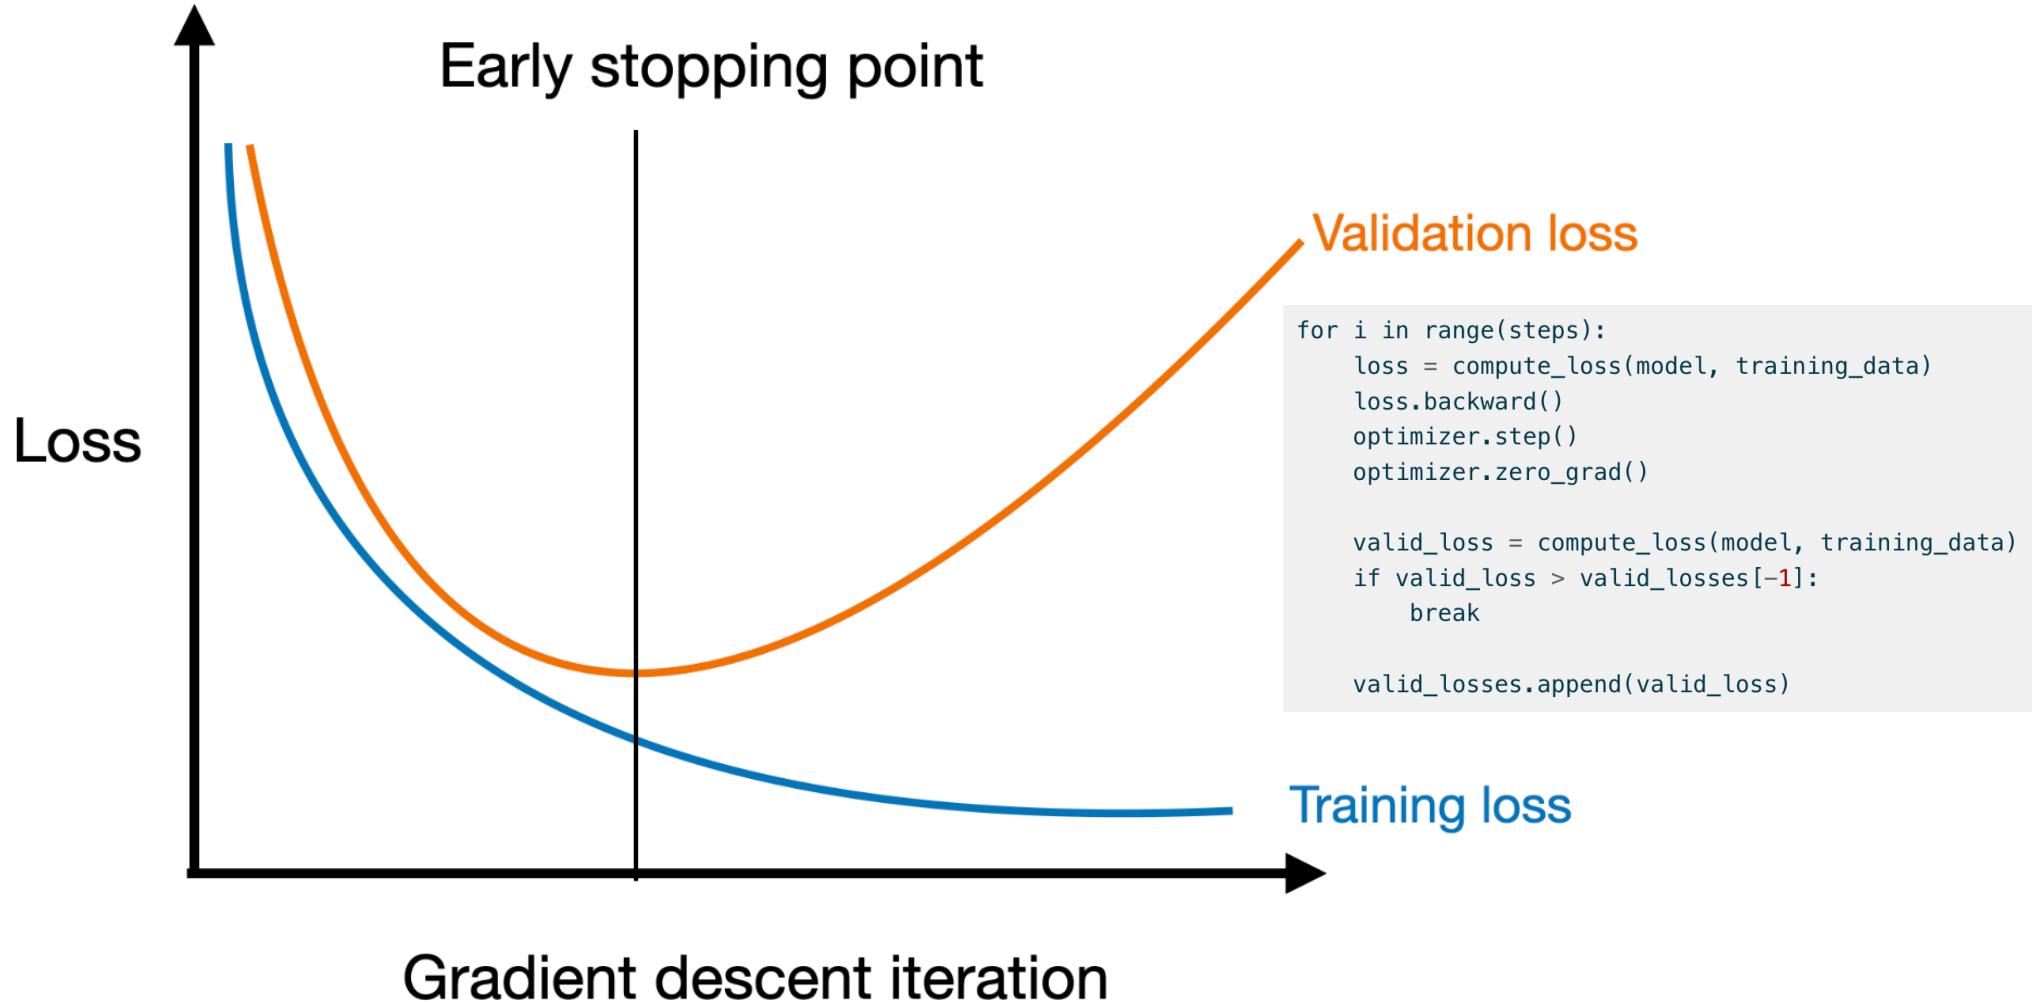


Overfit

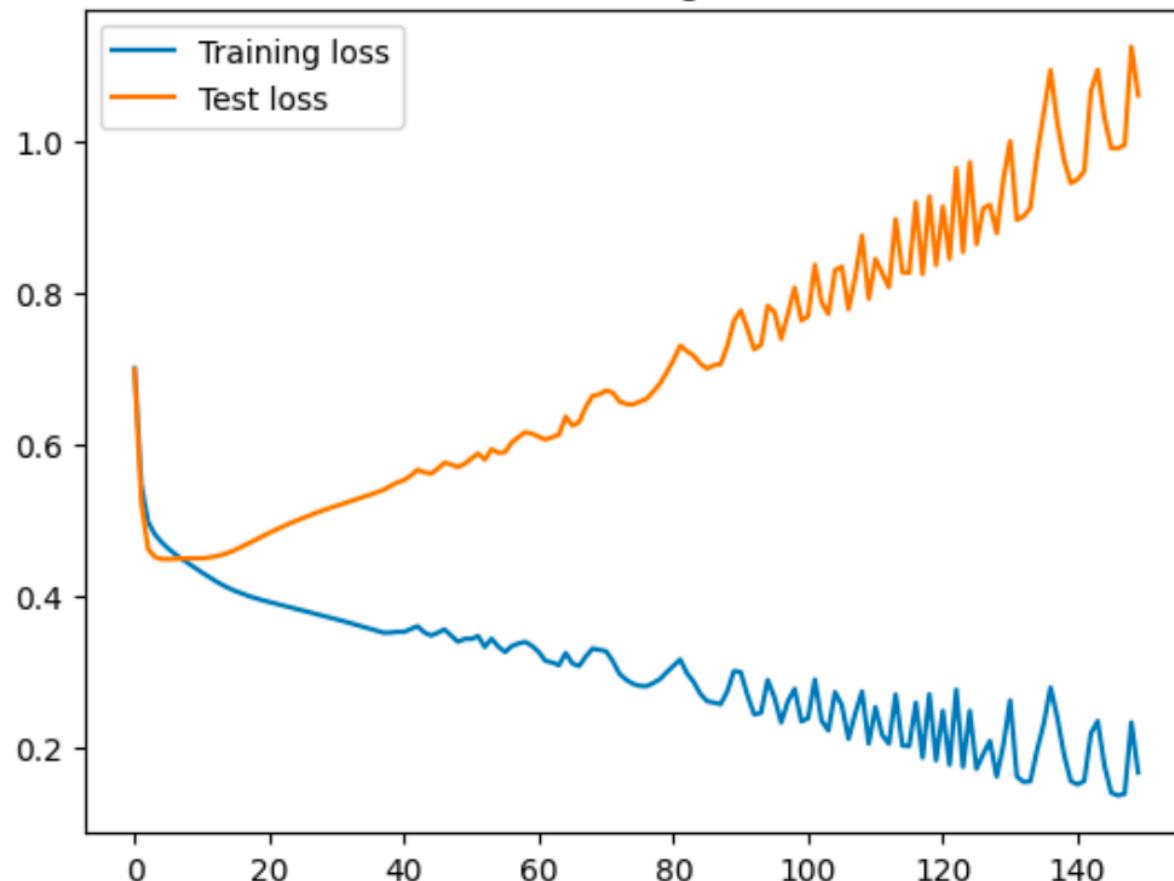




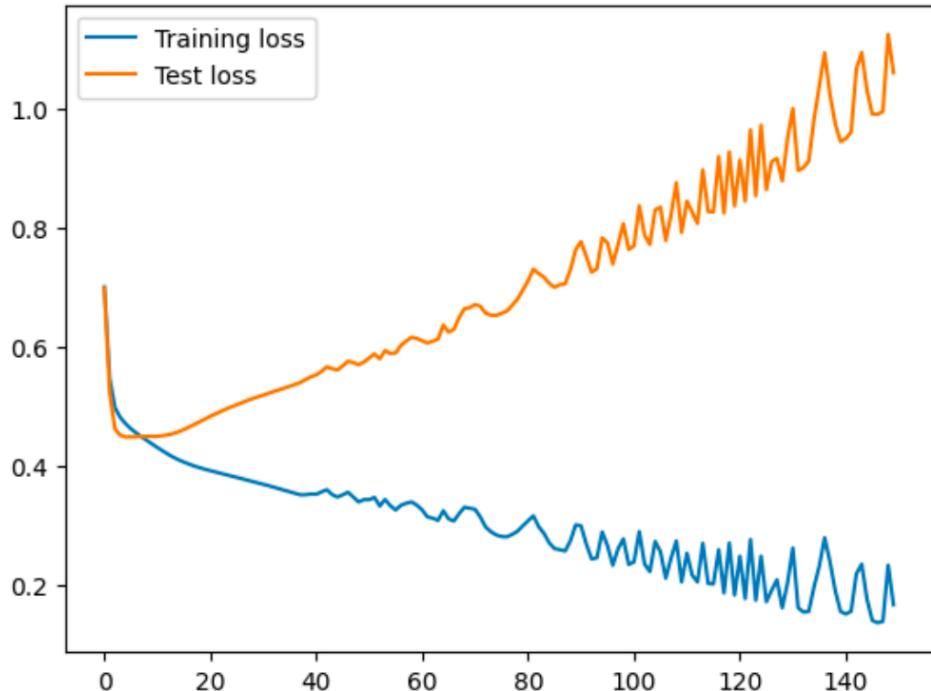




Loss vs. training iteration



Loss vs. training iteration



```
patience = 5          # Number of steps to wait before stopping
steps_since_improvement = 0 # Steps since validation loss improved
min_loss = 1e8          # Minimum loss seen so far (start large)

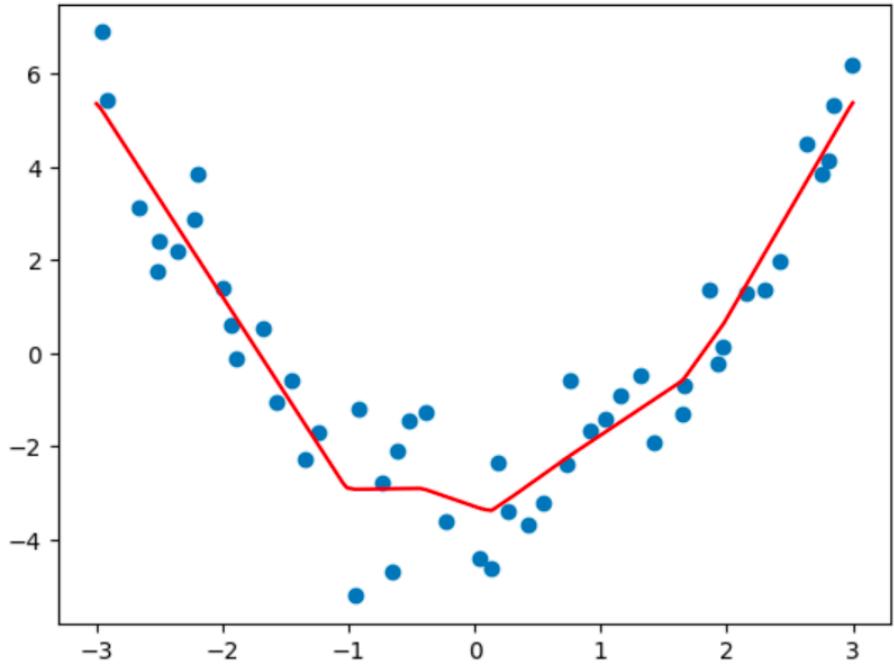
for i in range(steps):
    ...

    valid_loss = compute_loss(model, training_data)

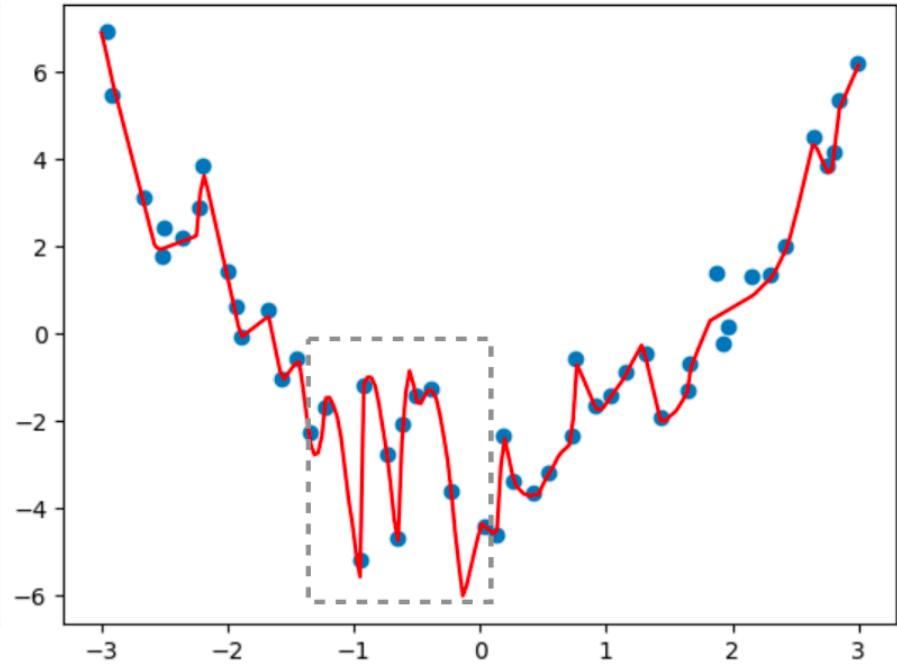
    # If the validation loss improves reset the counter
    if valid_loss < min_loss:
        steps_since_improvement = 0
        min_loss = valid_loss

    # Otherwise increment the counter
    else:
        steps_since_improvement += 1

    # If its been patience steps since the last improvement, stop
    if steps_since_improvement == patience:
        break
```

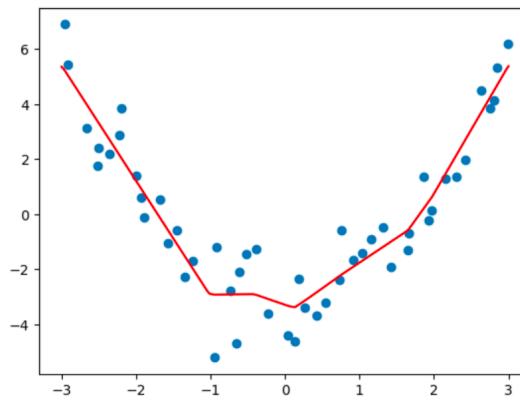


Well-fit

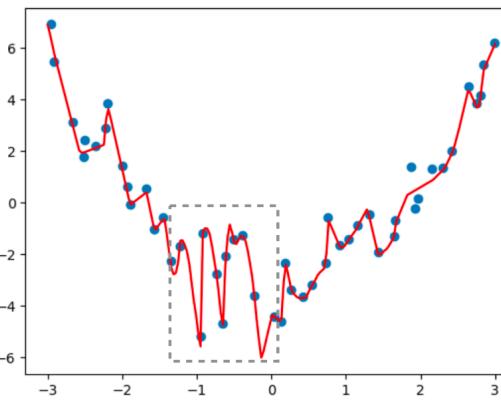


Overfit

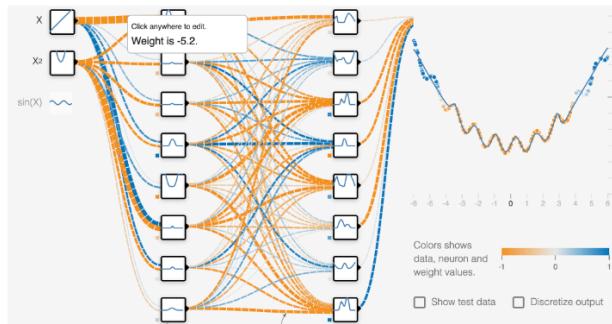
$$\text{MSE}(\mathbf{X}, \mathbf{y}, \mathbf{w}) = \frac{1}{N} \sum_{i=1}^N ((f(\mathbf{x}_i, \mathbf{w}) - y_i)^2)$$



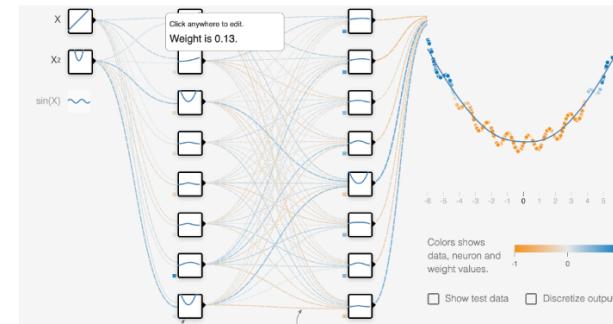
Well-fit



Overfit



An overfit network will have large weights to encode large slopes.



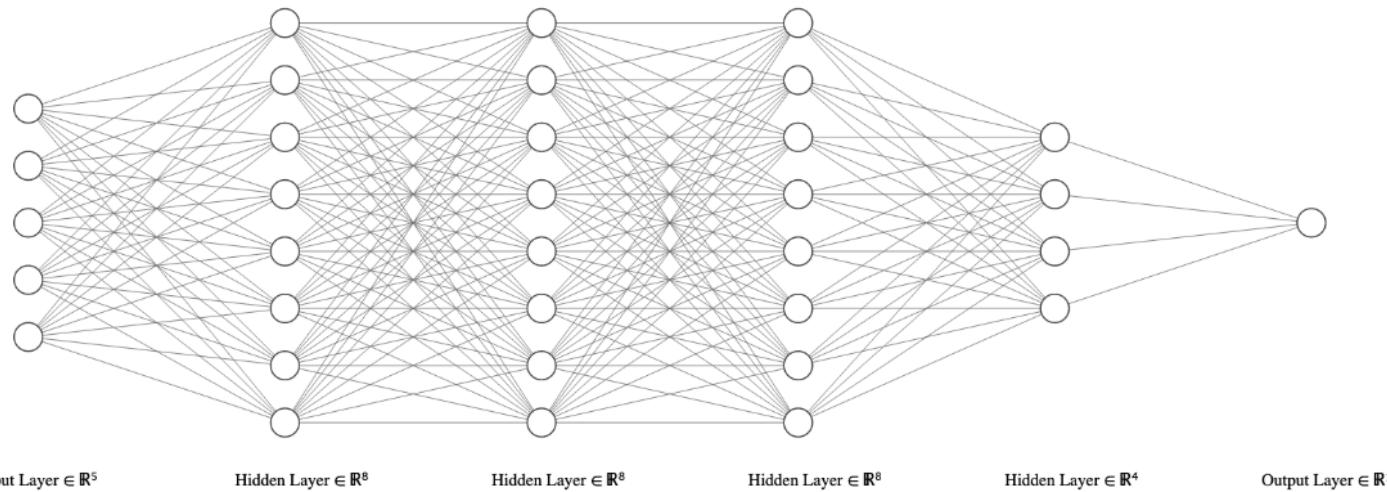
A regularized network will have smaller weights encoding a smooth function.

$$\mathbf{L}_2(\mathbf{w}) = \|\mathbf{w}\|_2^2 = \sum_{i=1}^d w_i^2$$

$$\mathbf{L}_2(\mathbf{W}) = \|\mathbf{W}\|_2^2 = \sum_{i=1}^d \sum_{j=1}^e w_{ij}^2$$

$$\mathbf{Loss}(\mathbf{X}, \mathbf{y}, \mathbf{w}) = \mathbf{MSE}(\mathbf{X}, \mathbf{y}, \mathbf{w}) + \lambda \mathbf{L}_2(\mathbf{w})$$

$$f(\mathbf{x}, \mathbf{w}_0, \dots) = \sigma(\sigma(\sigma(\sigma(\mathbf{x}^T \mathbf{W}_4)^T \mathbf{W}_3)^T \mathbf{W}_2)^T \mathbf{W}_1)^T \mathbf{w}_0$$



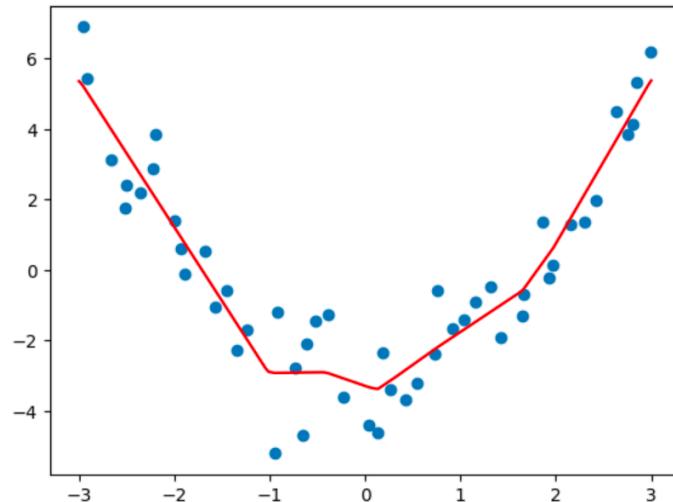
$$\mathbf{L}_2(\mathbf{w}_0, \mathbf{W}_1, \dots, \mathbf{W}_L) = \sum_{l=0}^L \|\mathbf{W}_l\|_2^2$$

In practice most networks also incorporate bias terms, so each linear function in our network can be written as:

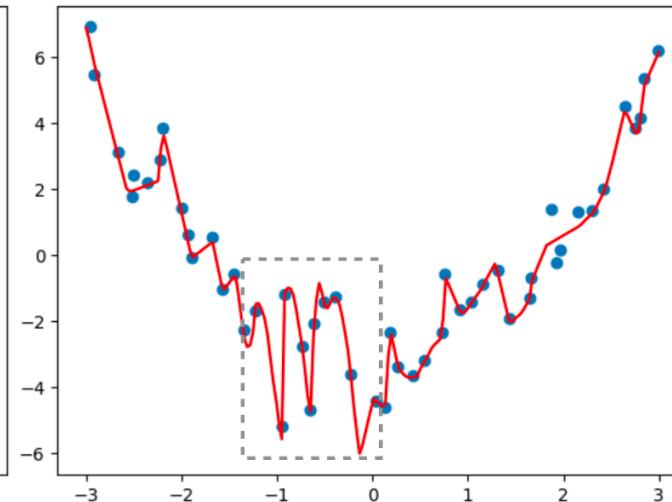
$$\mathbf{x}^T \mathbf{W} + \mathbf{b}$$

And the full prediction function for a sigmoid-activation network might be:

$$f(\mathbf{x}, \mathbf{w}_0, \dots) = \sigma(\sigma(\sigma(\sigma(\mathbf{x}^T \mathbf{W}_4 + \mathbf{b}_4)^T \mathbf{W}_3 + \mathbf{b}_3)^T \mathbf{W}_2 + \mathbf{b}_2)^T \mathbf{W}_1 + \mathbf{b}_1)^T \mathbf{w}_0 + \mathbf{b}_0$$



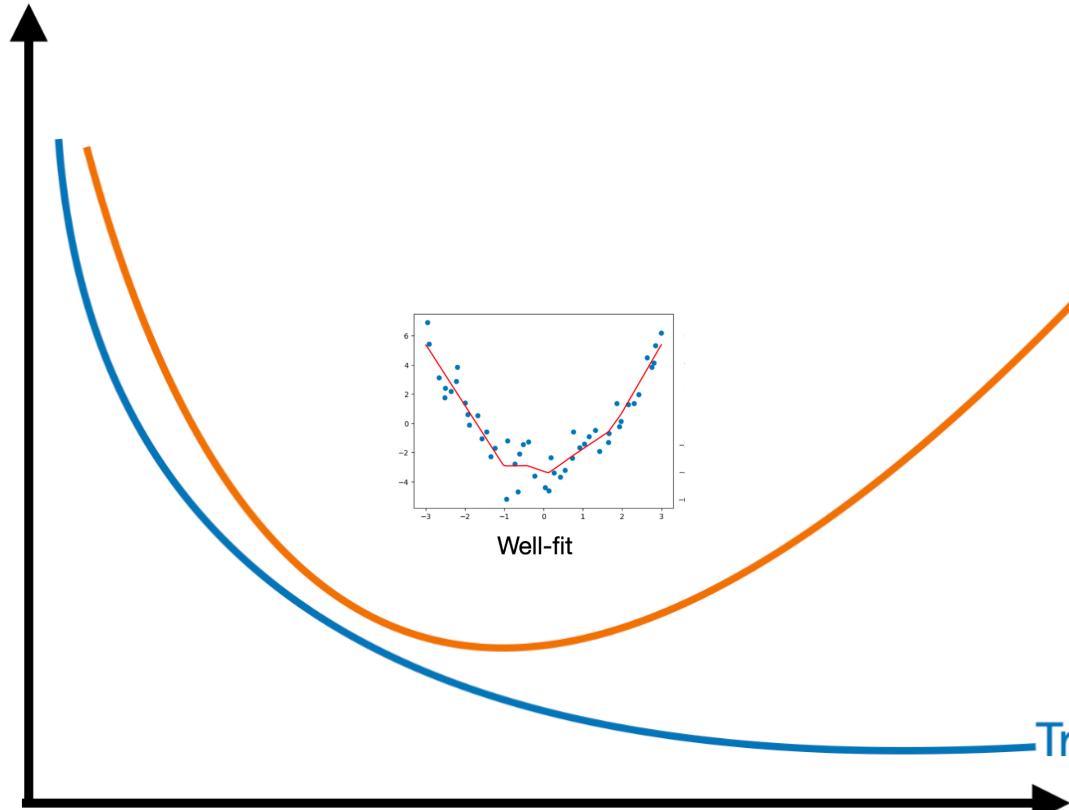
Well-fit



Overfit

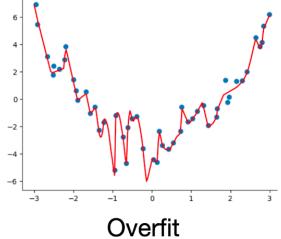
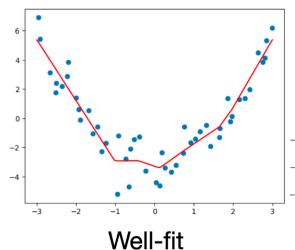
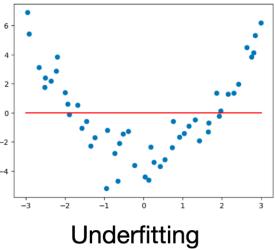
Does the bias contribute to overfitting?

Loss

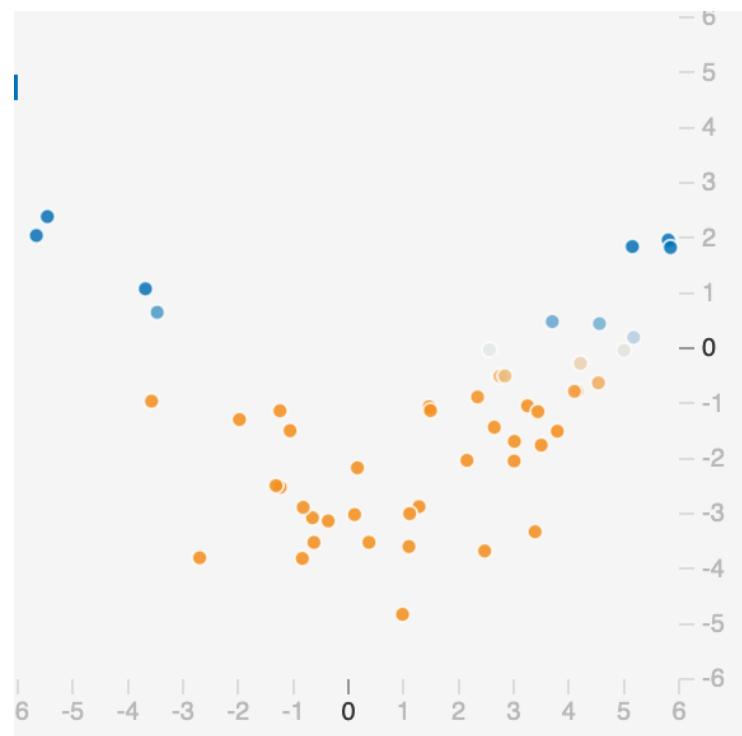
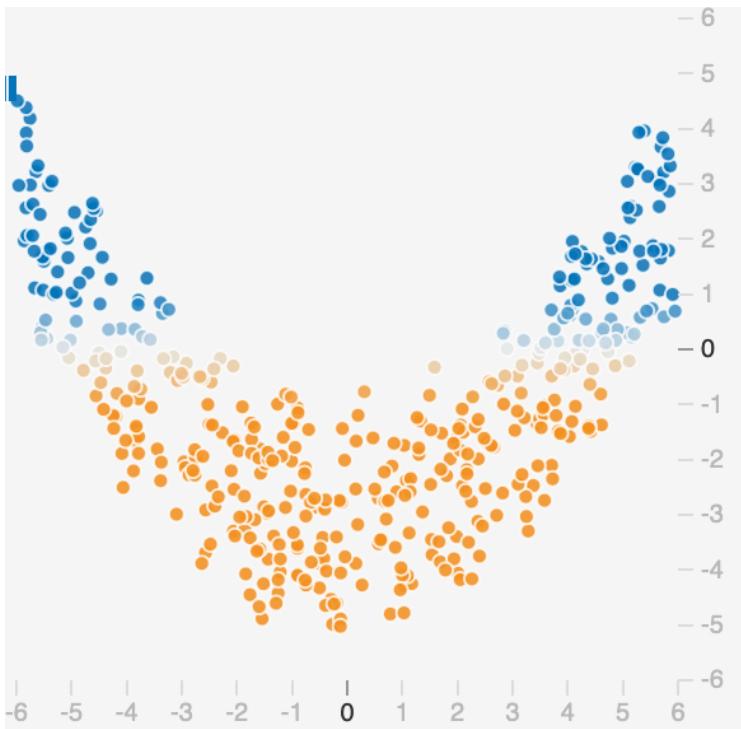


Validation loss

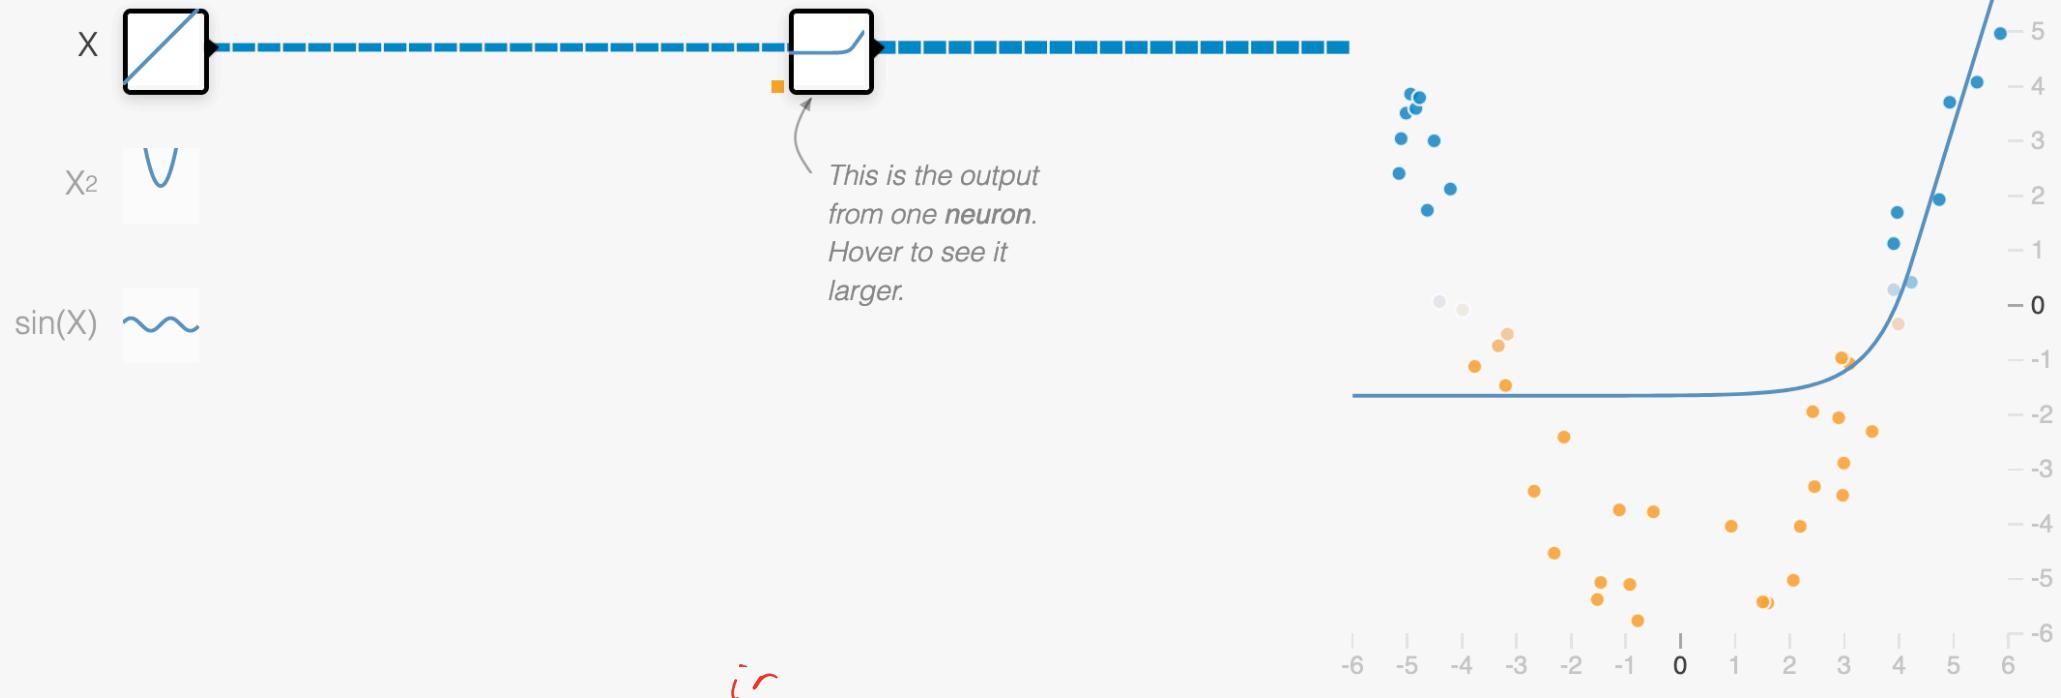
Overfit



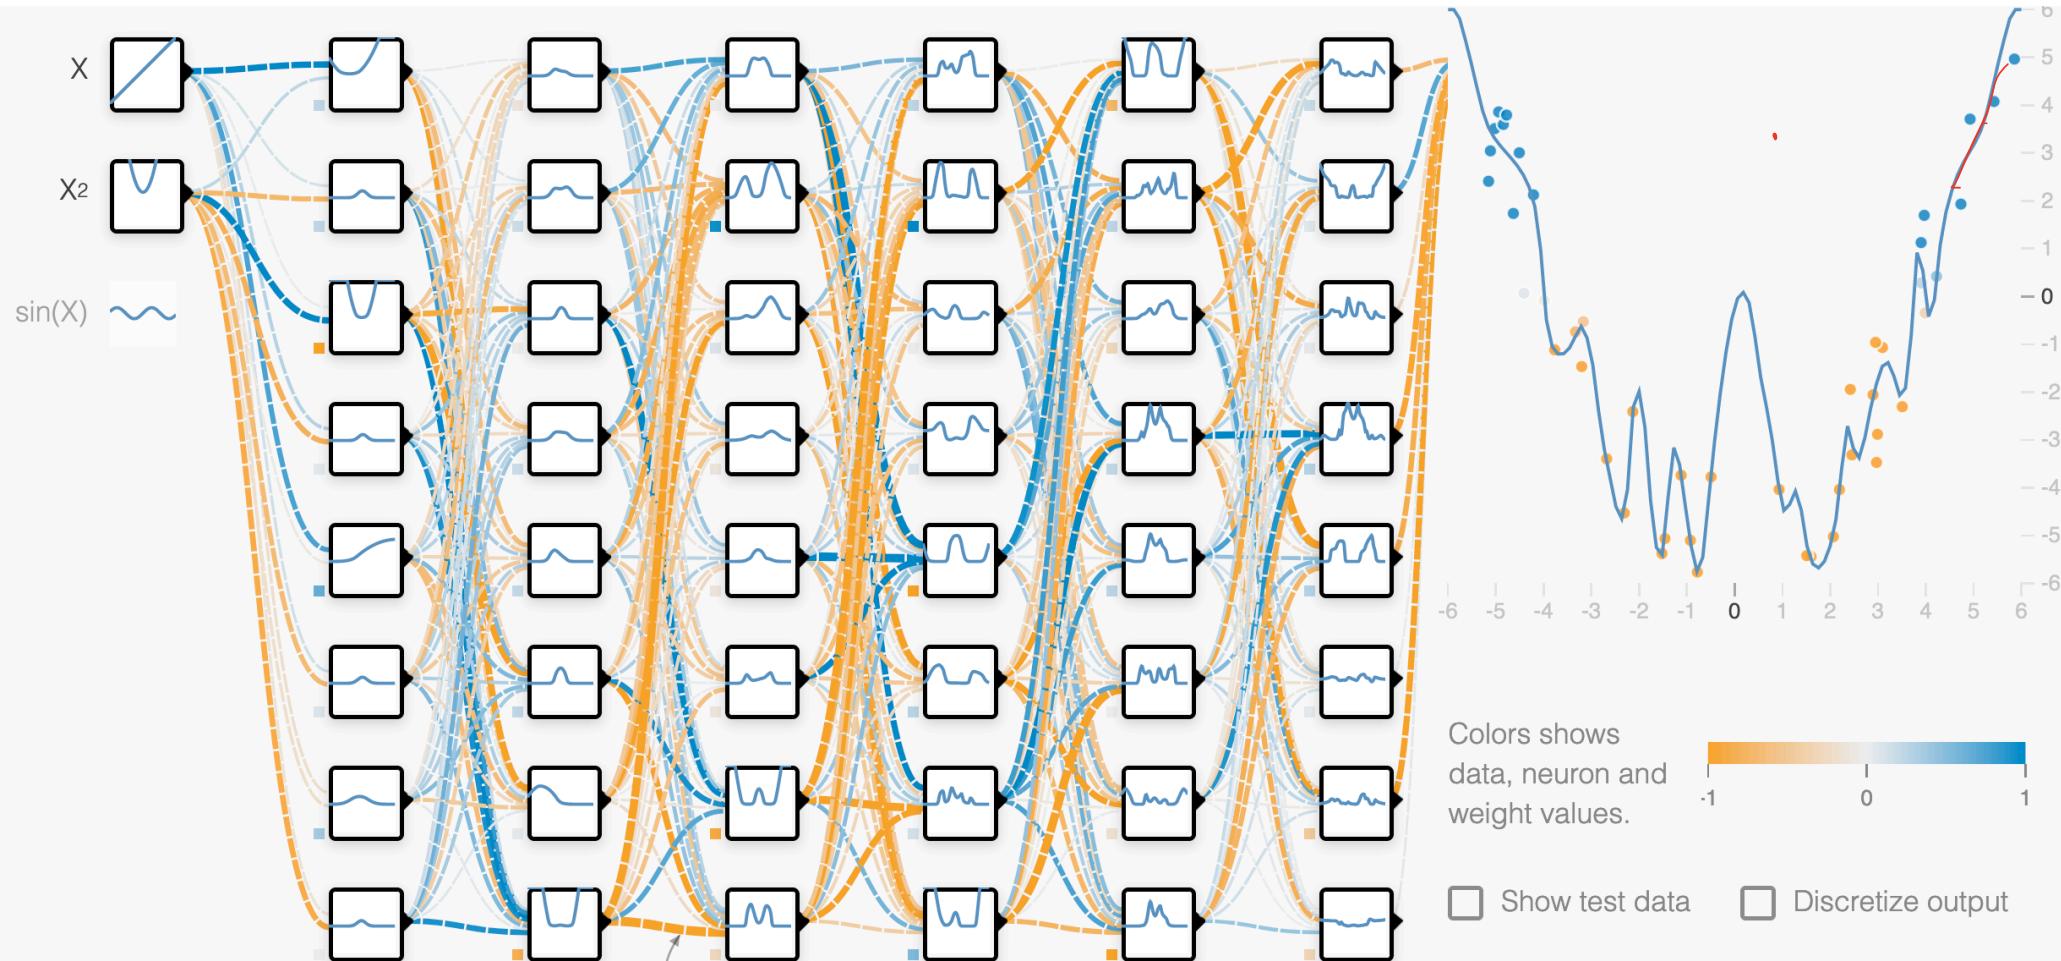
Sampling data



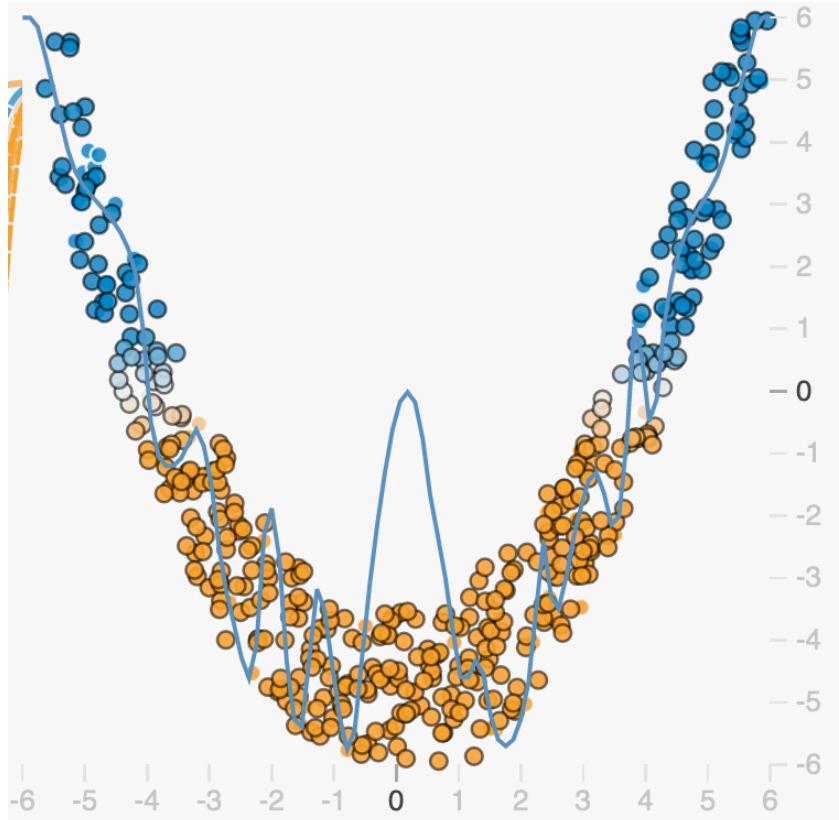
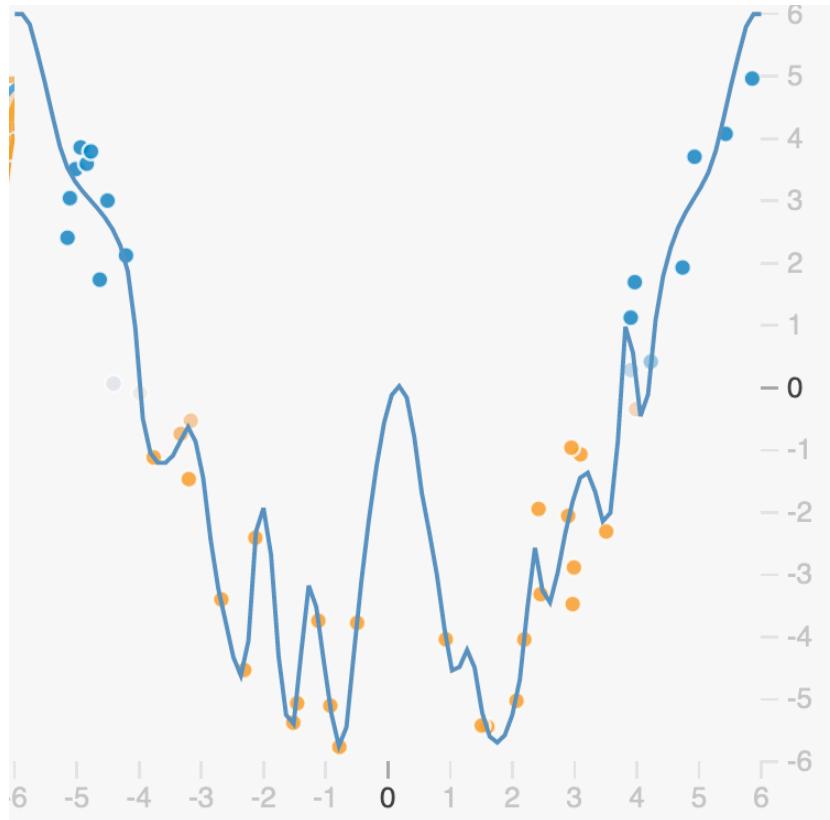
Underfitting



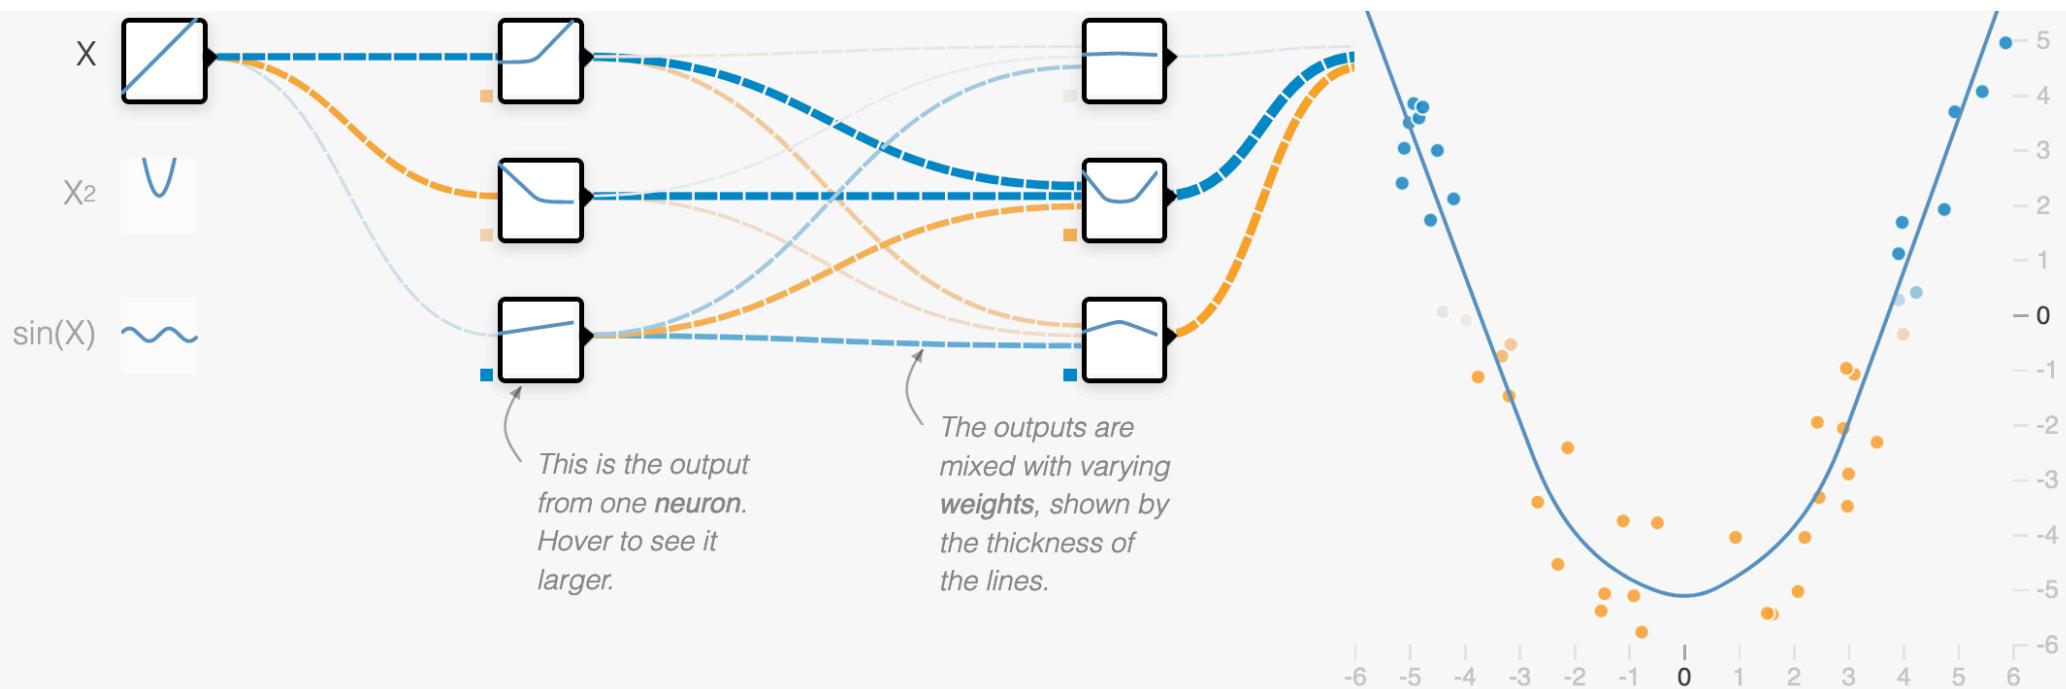
Overfitting



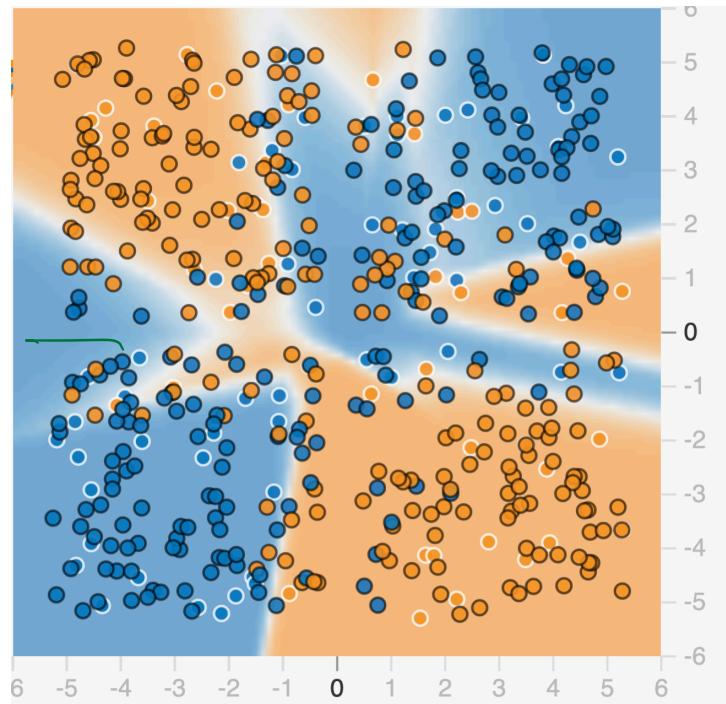
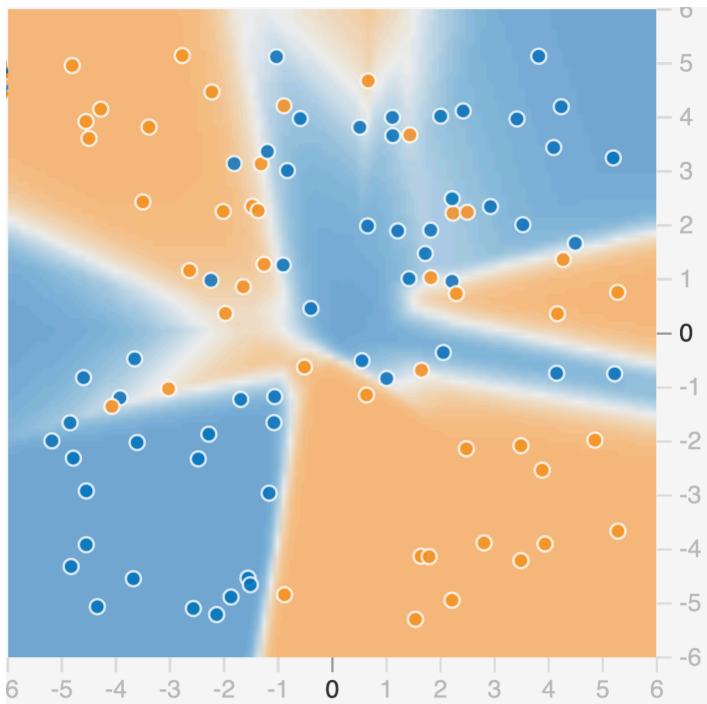
Overfitting



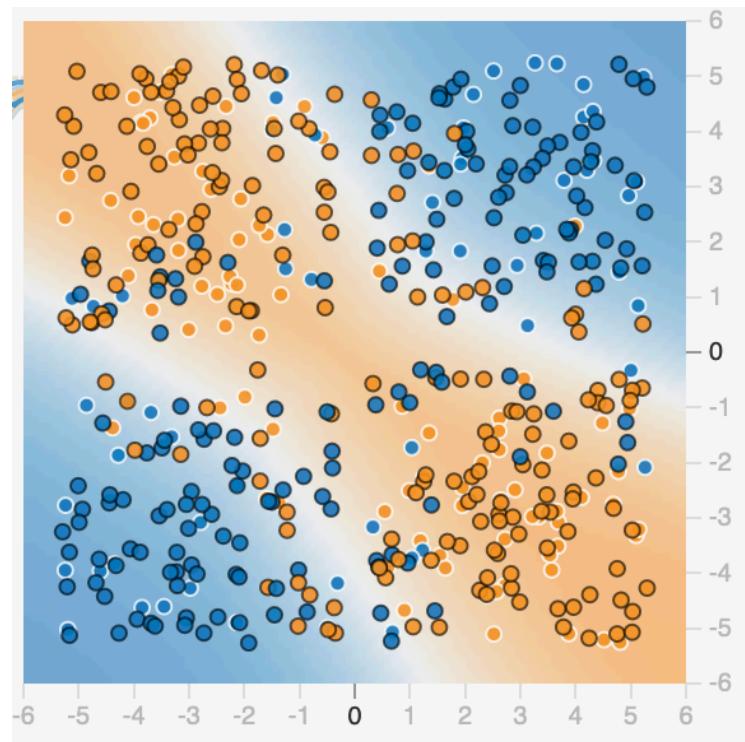
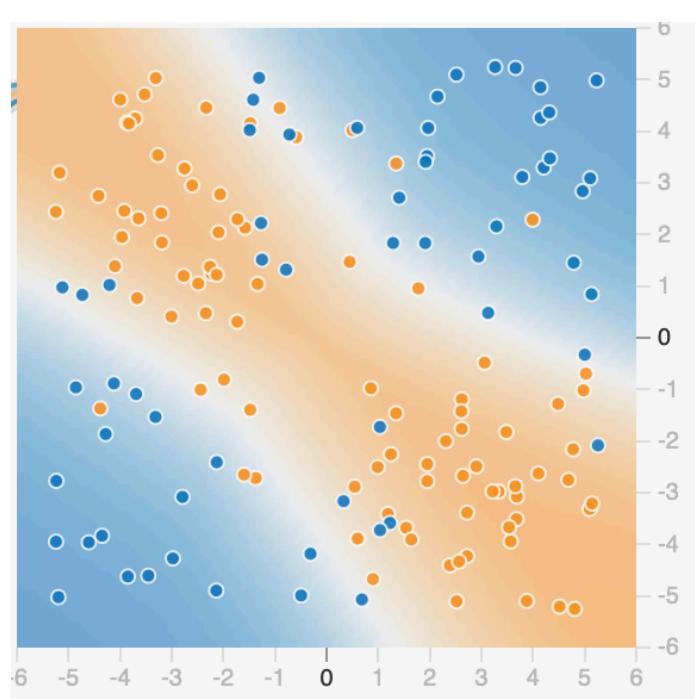
Good fit



Overfitting

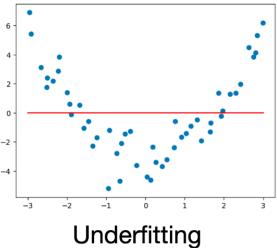


Better fit

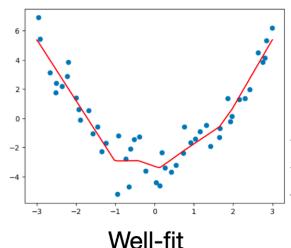


Loss

Model complexity curve



Validation loss



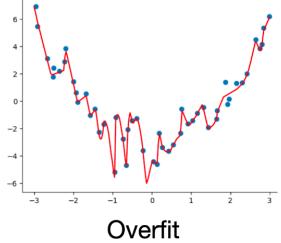
Training loss

Underfitting

Well fitting

Overfitting

Model complexity



Underfitting

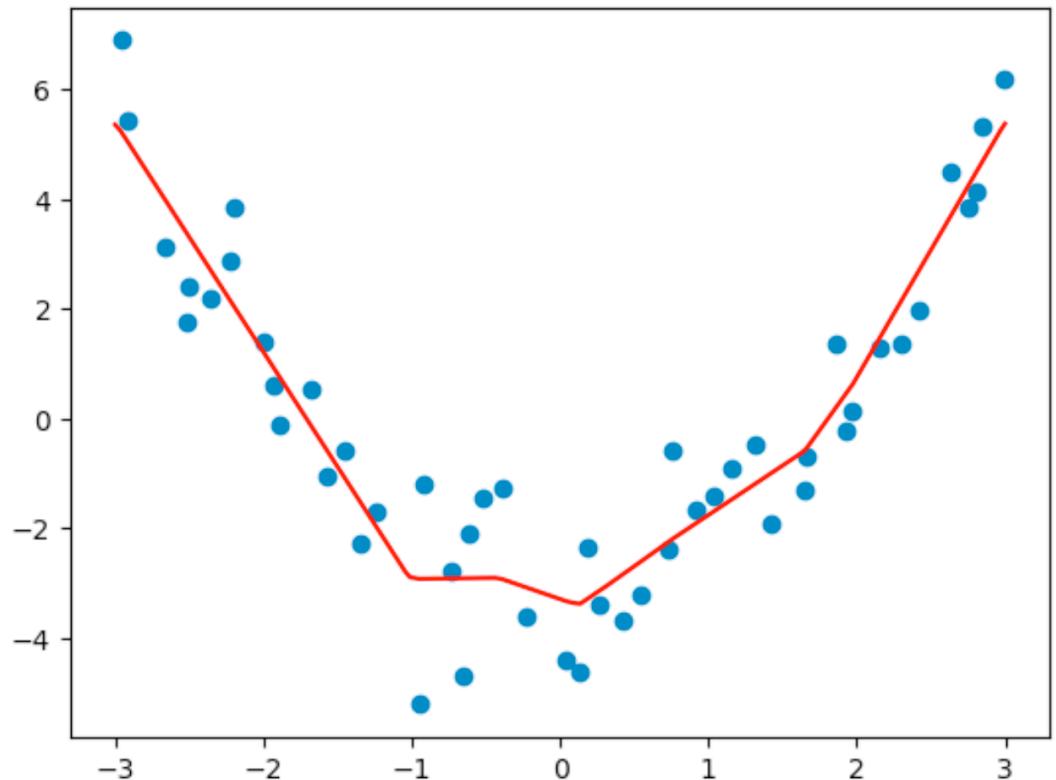
Overfitting

Good fit

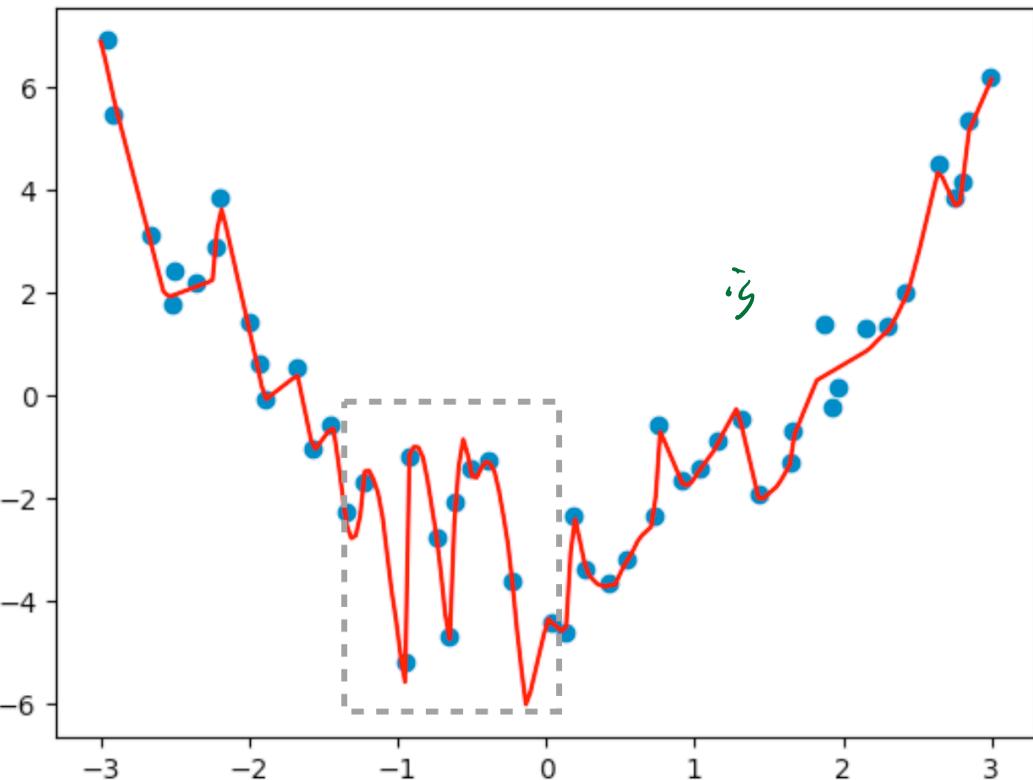
Training loss:

Test loss:

Overfitting with high weights

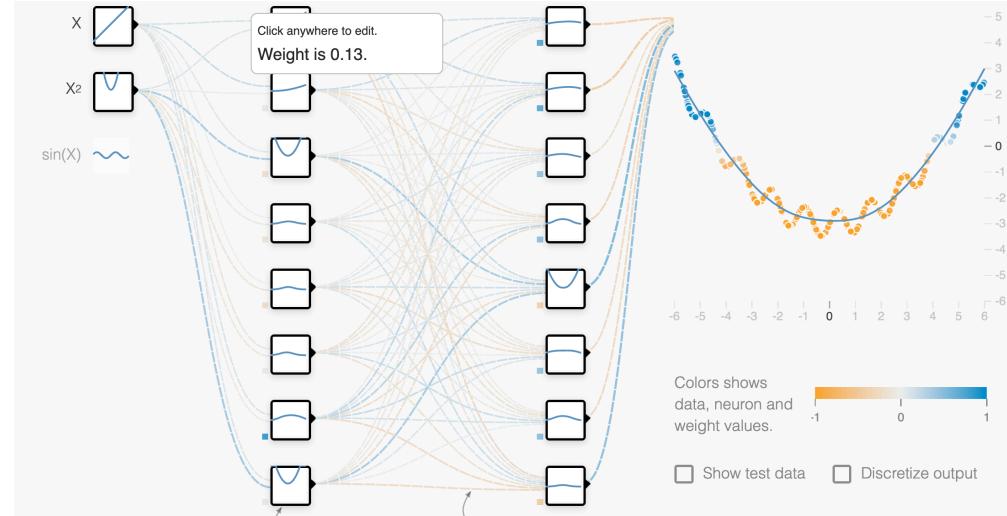
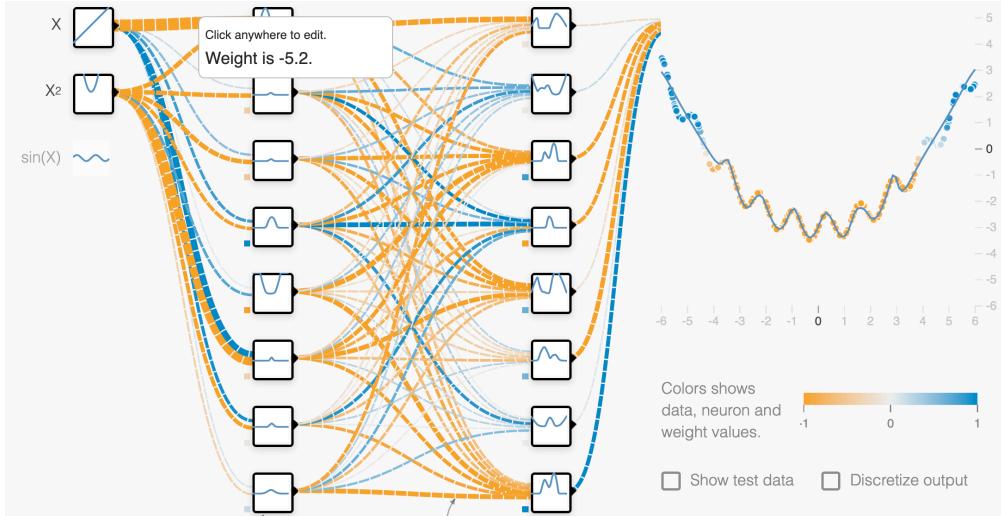


Well-fit



Overfit

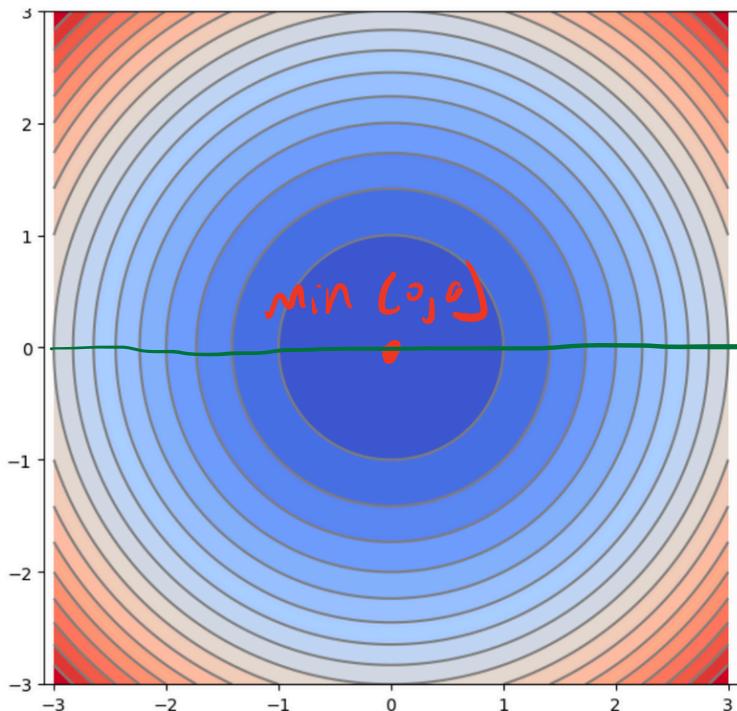
Overfitting with high weights



L2 Regularization

L2-Loss

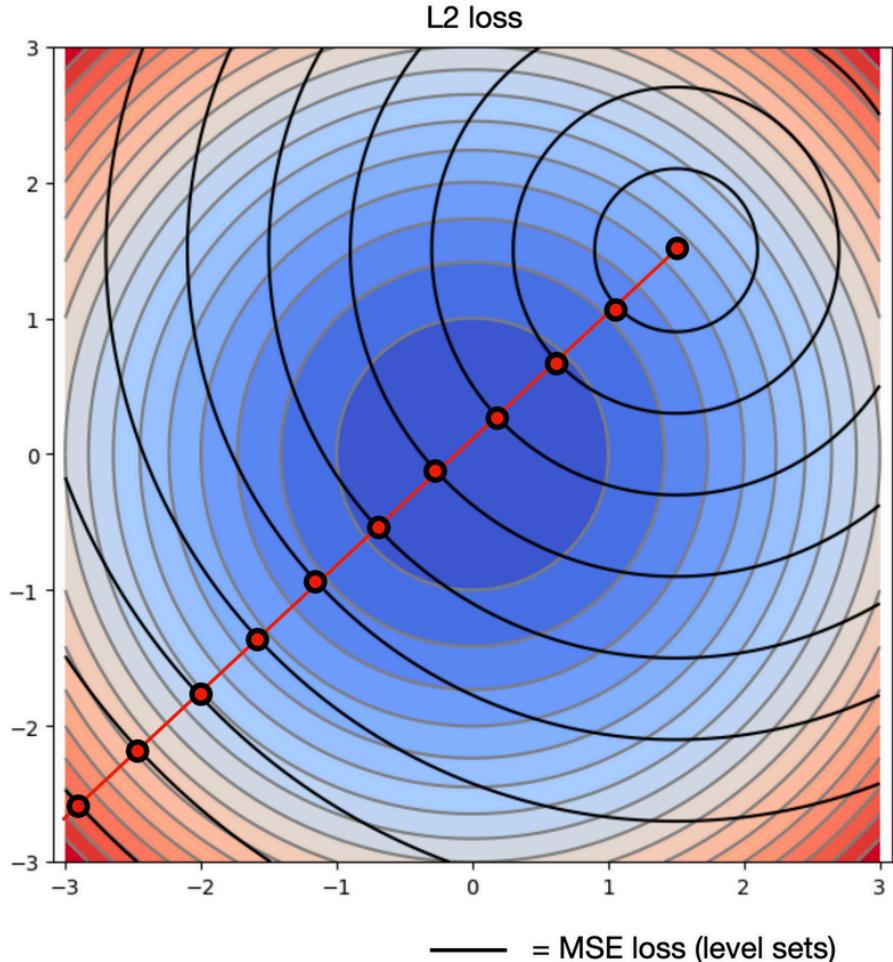
$$\ell^2 = w_1^2 + w_2^2$$



$$\mathbf{L}_2(\mathbf{W}) = \|\mathbf{W}\|_2^2 = \sum_{i=1}^d \sum_{j=1}^e w_{ij}^2$$

i

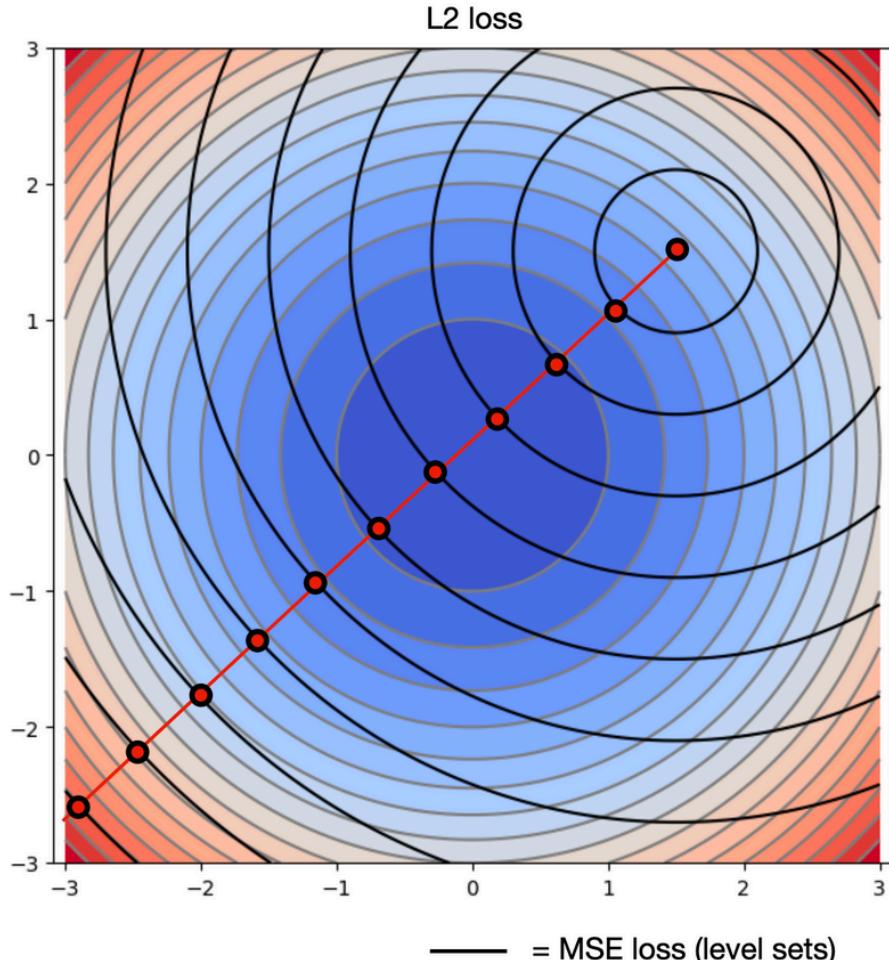
L2 Regularization



$$\mathbf{L}_2(\mathbf{W}) = \|\mathbf{W}\|_2^2 = \sum_{i=1}^d \sum_{j=1}^e w_{ij}^2$$

$$\text{Loss}(\mathbf{X}, \mathbf{y}, \mathbf{w}) = \text{MSE}(\mathbf{X}, \mathbf{y}, \mathbf{w}) + \lambda \mathbf{L}_2(\mathbf{w})$$

L2 Regularization



$$\mathbf{L}_2(\mathbf{W}) = \|\mathbf{W}\|_2^2 = \sum_{i=1}^d \sum_{j=1}^e w_{ij}^2$$

$$\text{Loss}(\mathbf{X}, \mathbf{y}, \mathbf{w}) = \text{MSE}(\mathbf{X}, \mathbf{y}, \mathbf{w}) + \lambda \mathbf{L}_2(\mathbf{w})$$

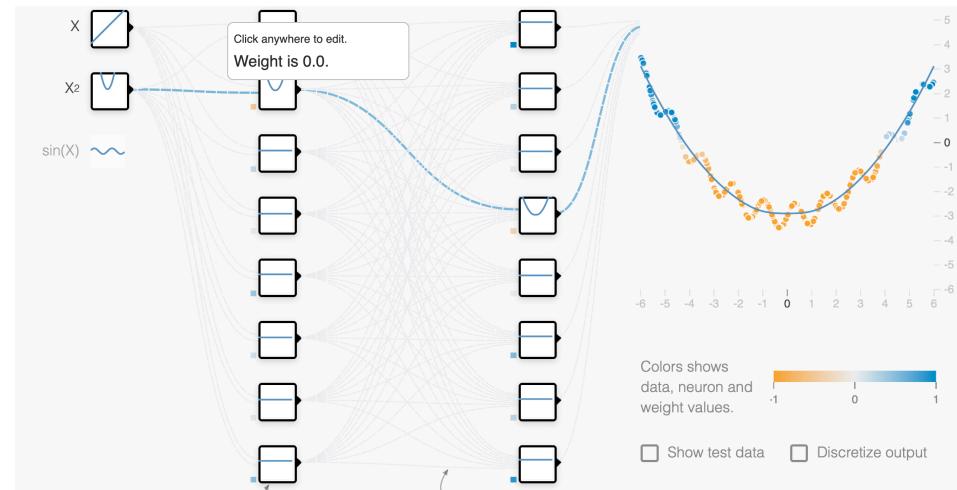
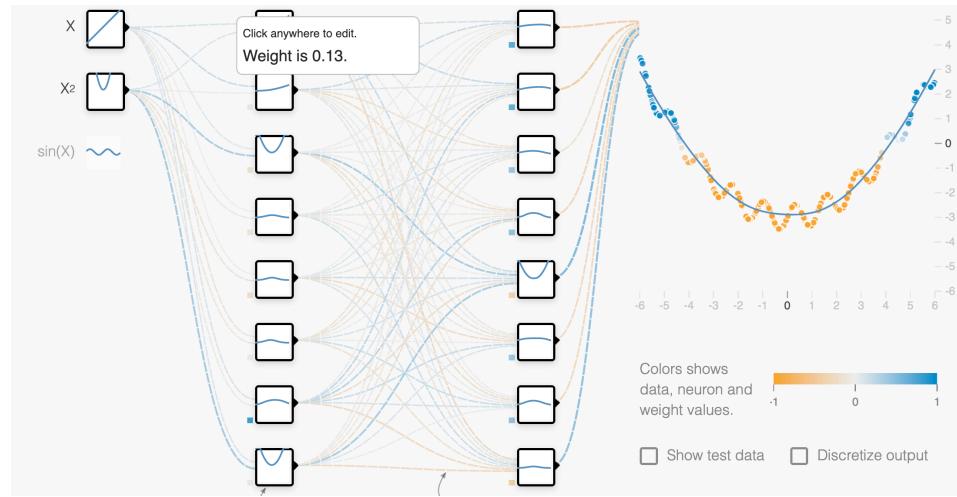
```
from torch import optim
optimizer = optim.SGD(model.parameters(), lr=0.1, weight_decay=0.01)
```

L1 Regularization

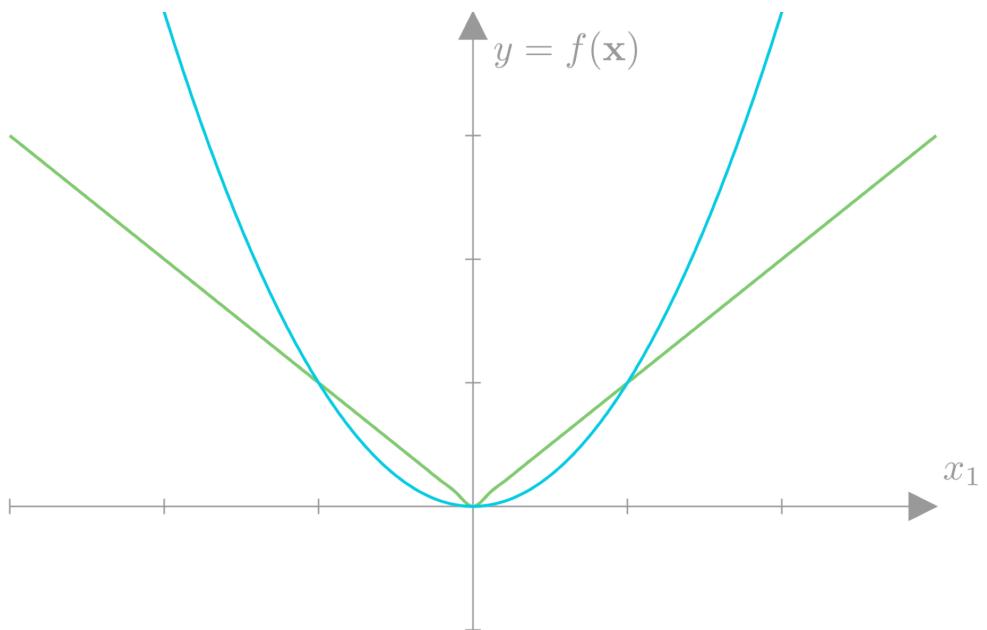
$$\text{Vector: } \mathbf{L}_2(\mathbf{w}) = \|\mathbf{w}\|_2^2 = \sum_{i=1}^d w_i^2 \quad \text{Matrix: } \mathbf{L}_2(\mathbf{W}) = \|\mathbf{W}\|_2^2 = \sum_{i=1}^d \sum_{j=1}^e w_{ij}^2$$

$$\text{Vector: } \mathbf{L}_1(\mathbf{w}) = \|\mathbf{w}\|_1 = \sum_{i=1}^d |w_i|, \quad \text{Matrix: } \mathbf{L}_1(\mathbf{W}) = \|\mathbf{W}\|_1 = \sum_{i=1}^d \sum_{j=1}^e |w_{ij}|$$

L1 Regularization



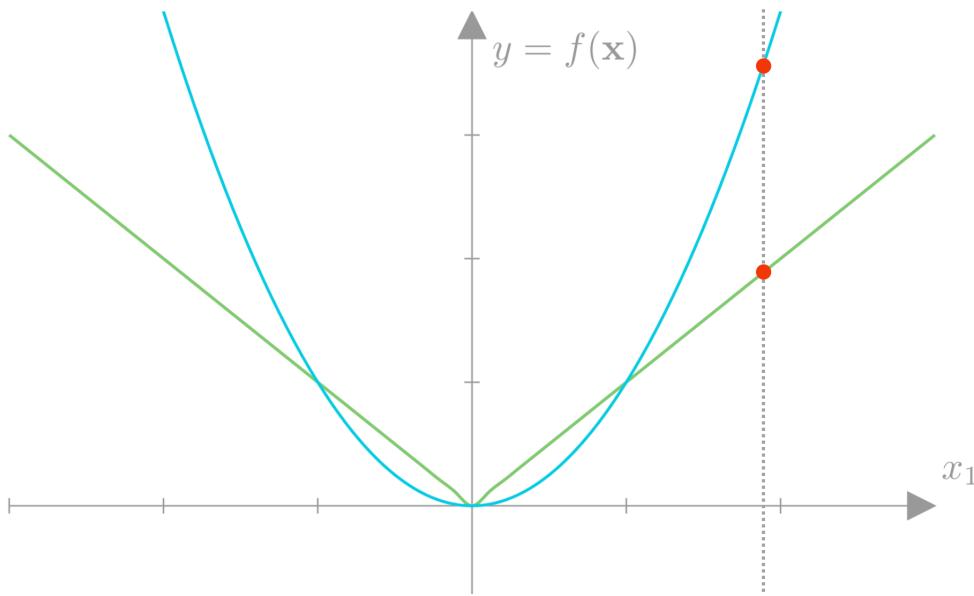
L1 Vs. L2 Regularization



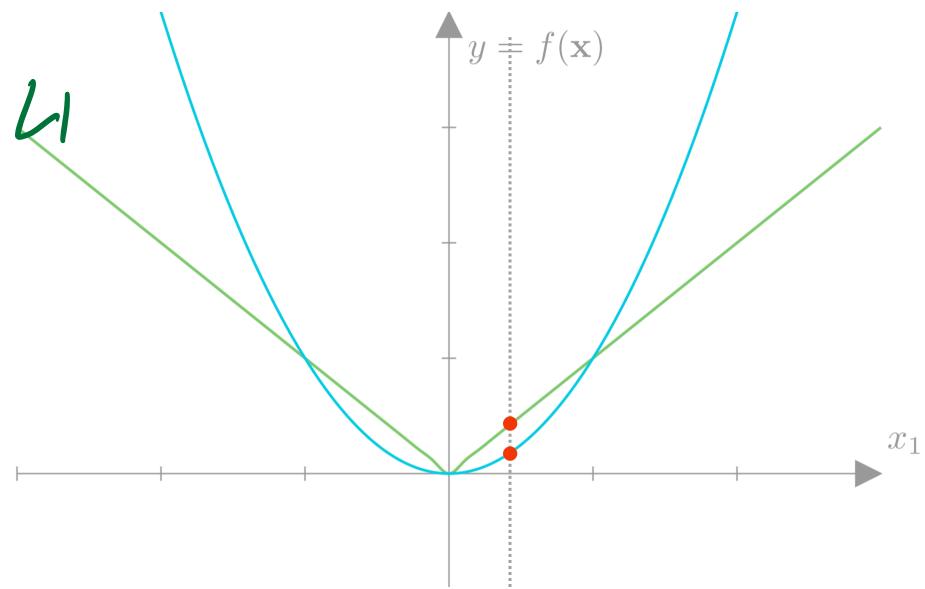
Vector: $\mathbf{L}_2(\mathbf{w}) = \|\mathbf{w}\|_2^2 = \sum_{i=1}^d w_i^2$

Vector: $\mathbf{L}_1(\mathbf{w}) = \|\mathbf{w}\|_1 = \sum_{i=1}^d |w_i|$

L1 Vs. L2 Regularization



L_2



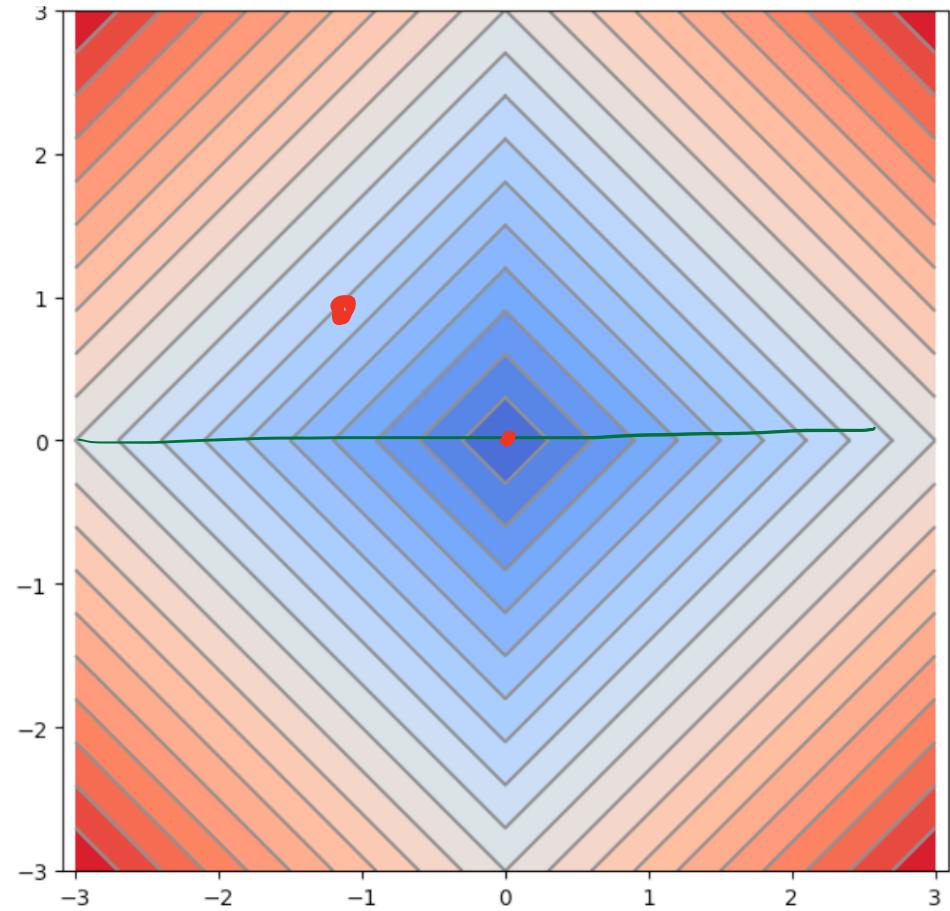
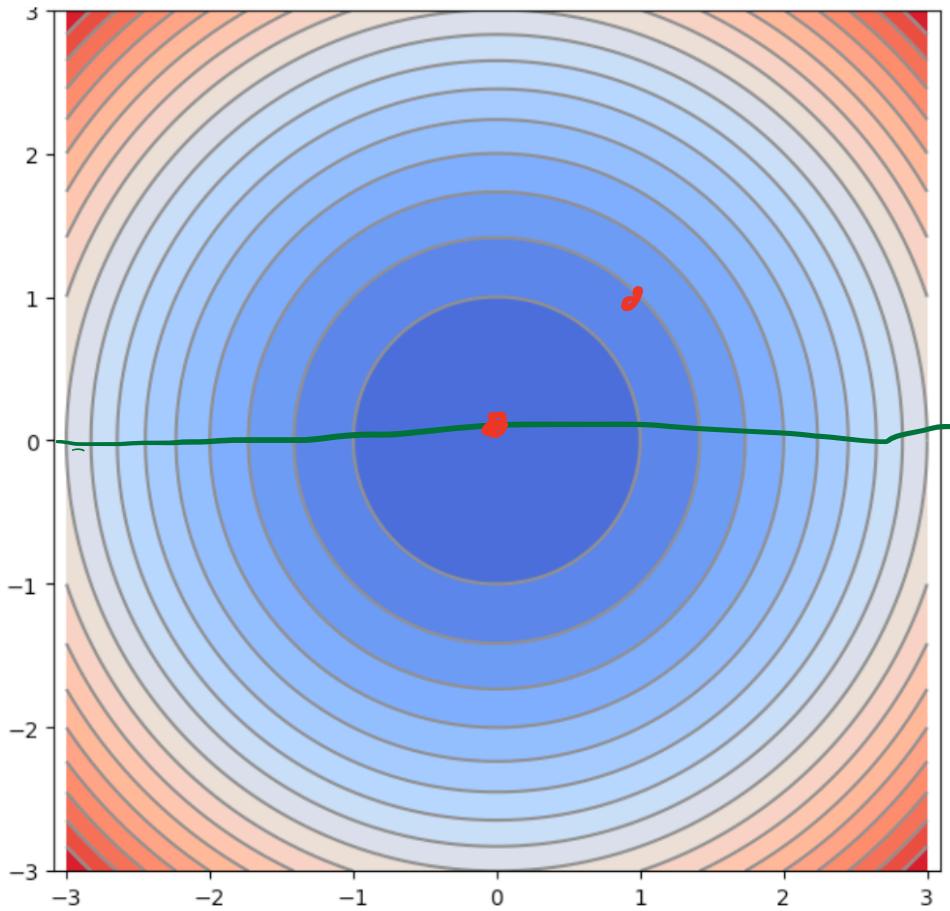
L2-Loss

$$\ell^2 = w_1^2 + w_2^2$$

L1 Vs. L2 Regularization

L1-Loss

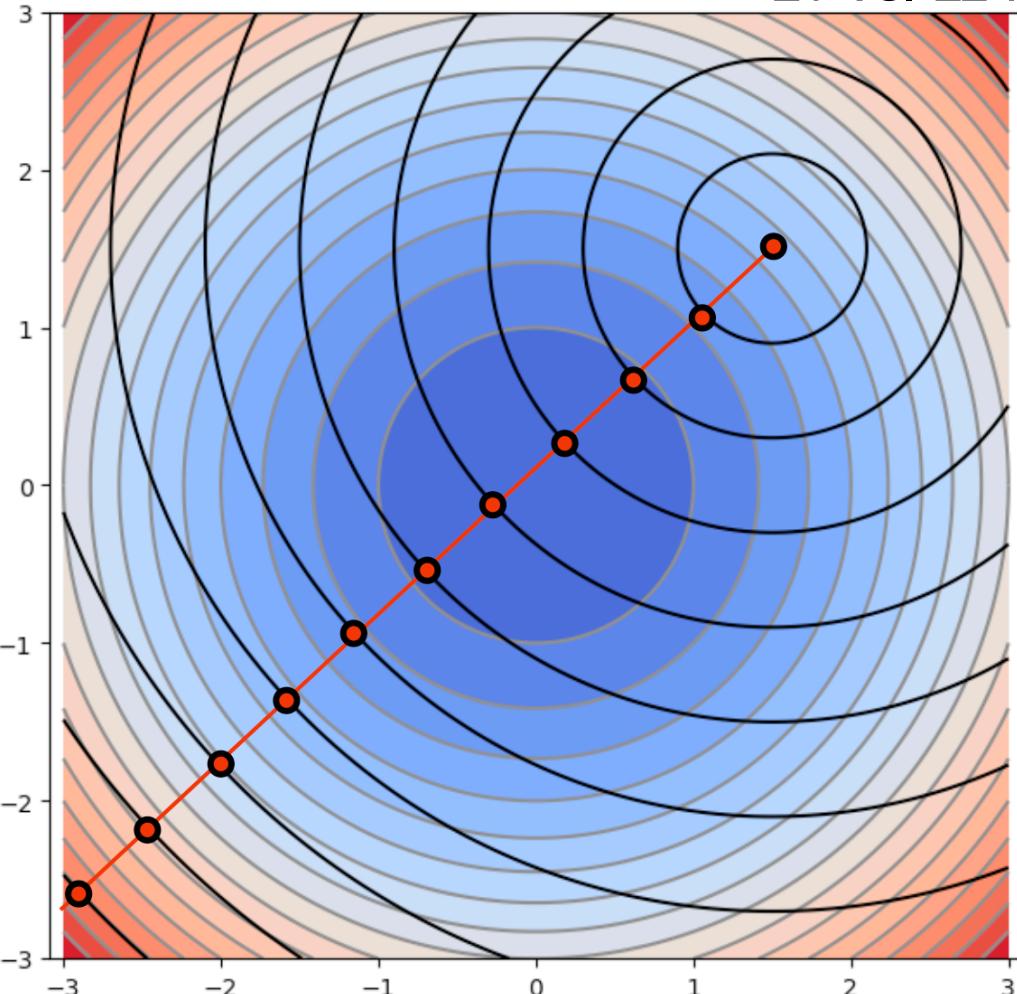
$$\ell^1 = |w_1| + |w_2|$$



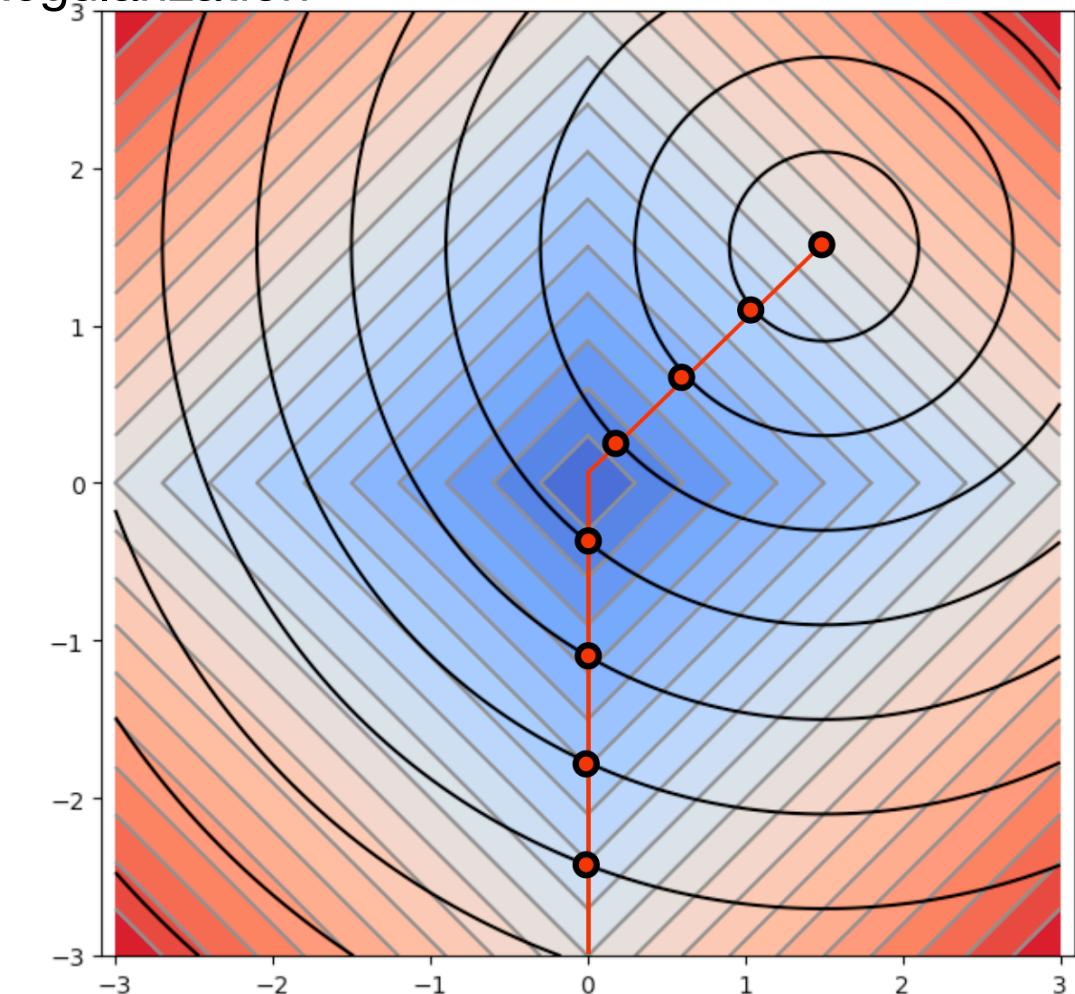
L2 loss

L1 Vs. L2 Regularization

L1 loss

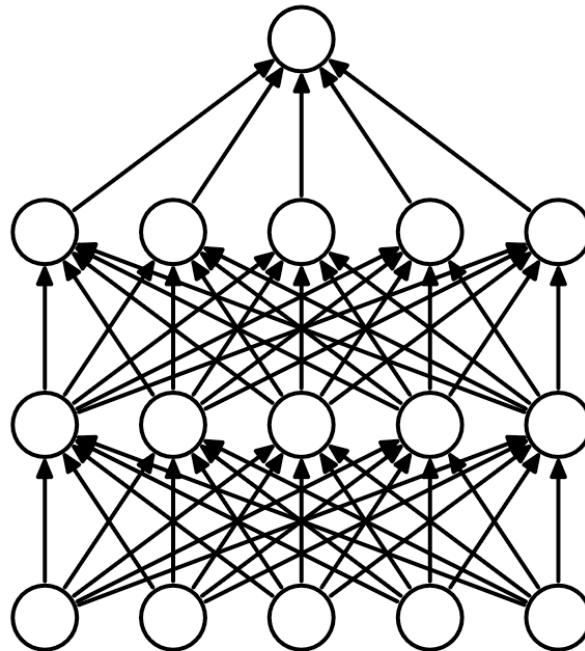


— = MSE loss (level sets)

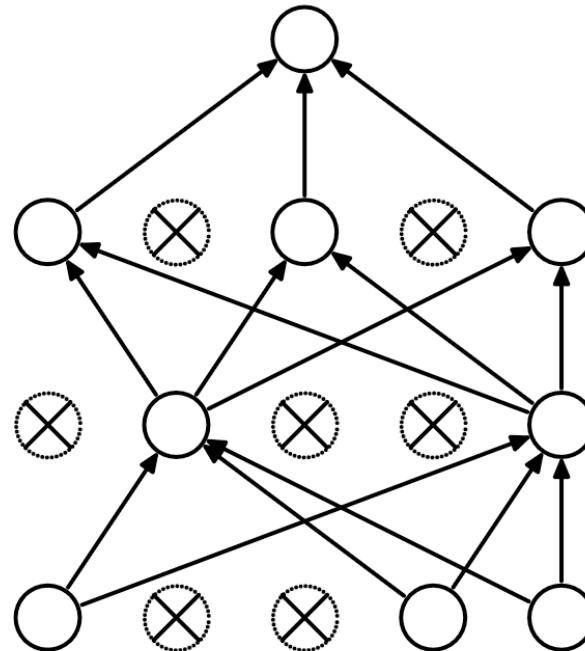


● = Minimum L2 or L1 loss for constant MSE

Dropout



(a) Standard Neural Net



(b) After applying dropout.

Figure 1: Dropout Neural Net Model. **Left:** A standard neural net with 2 hidden layers. **Right:** An example of a thinned net produced by applying dropout to the network on the left. Crossed units have been dropped.

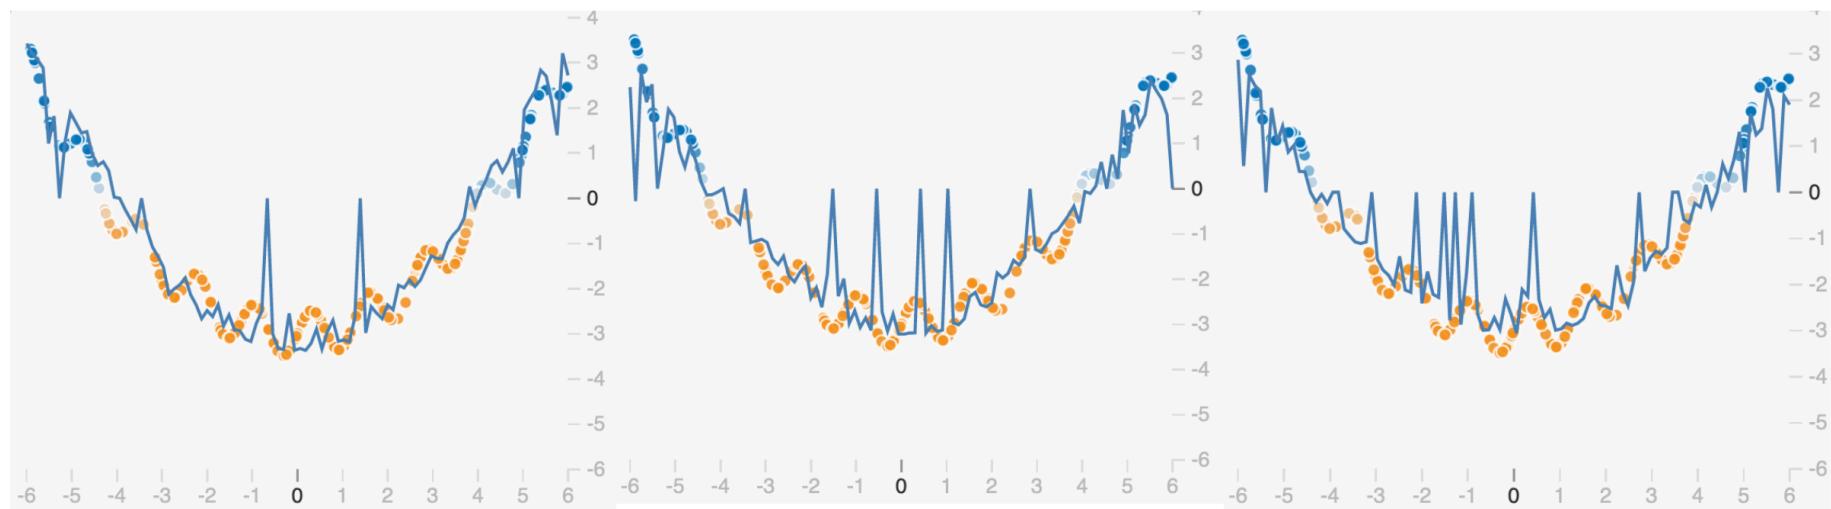
Dropout

$$\text{Dropout}(\mathbf{X}, r) = \mathbf{D} \odot \mathbf{X}, \quad \mathbf{D} = \begin{bmatrix} d_{11} & d_{12} & \dots & d_{1n} \\ d_{21} & d_{22} & \dots & d_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ d_{m1} & d_{m2} & \dots & d_{mn} \end{bmatrix}, \quad d_{ij} \sim \text{Bernoulli}(1 - r)$$

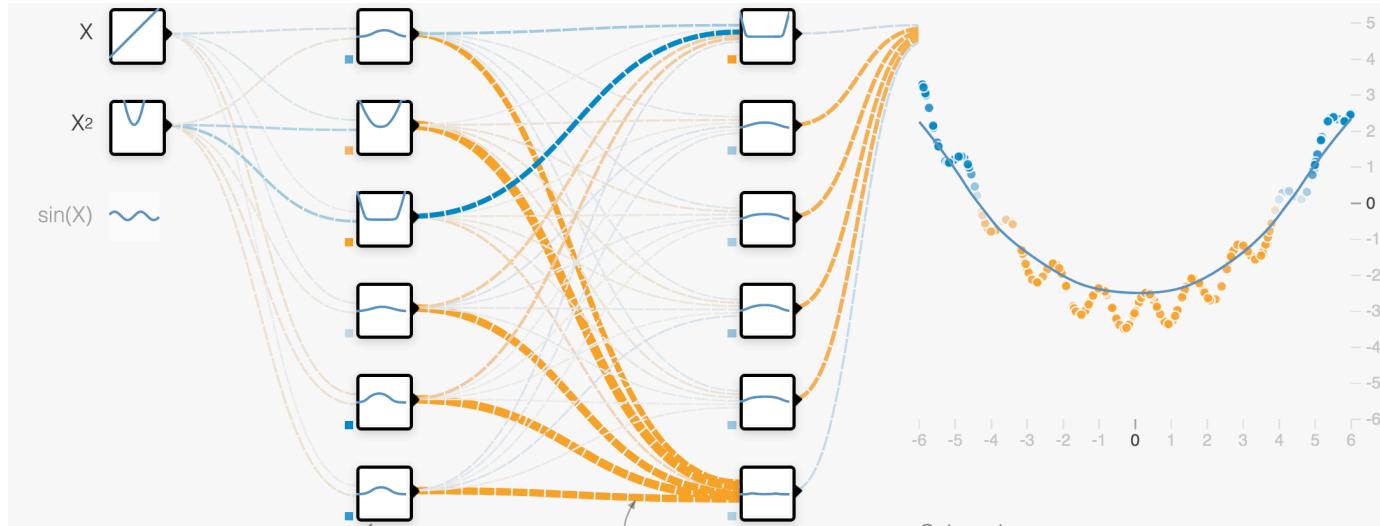
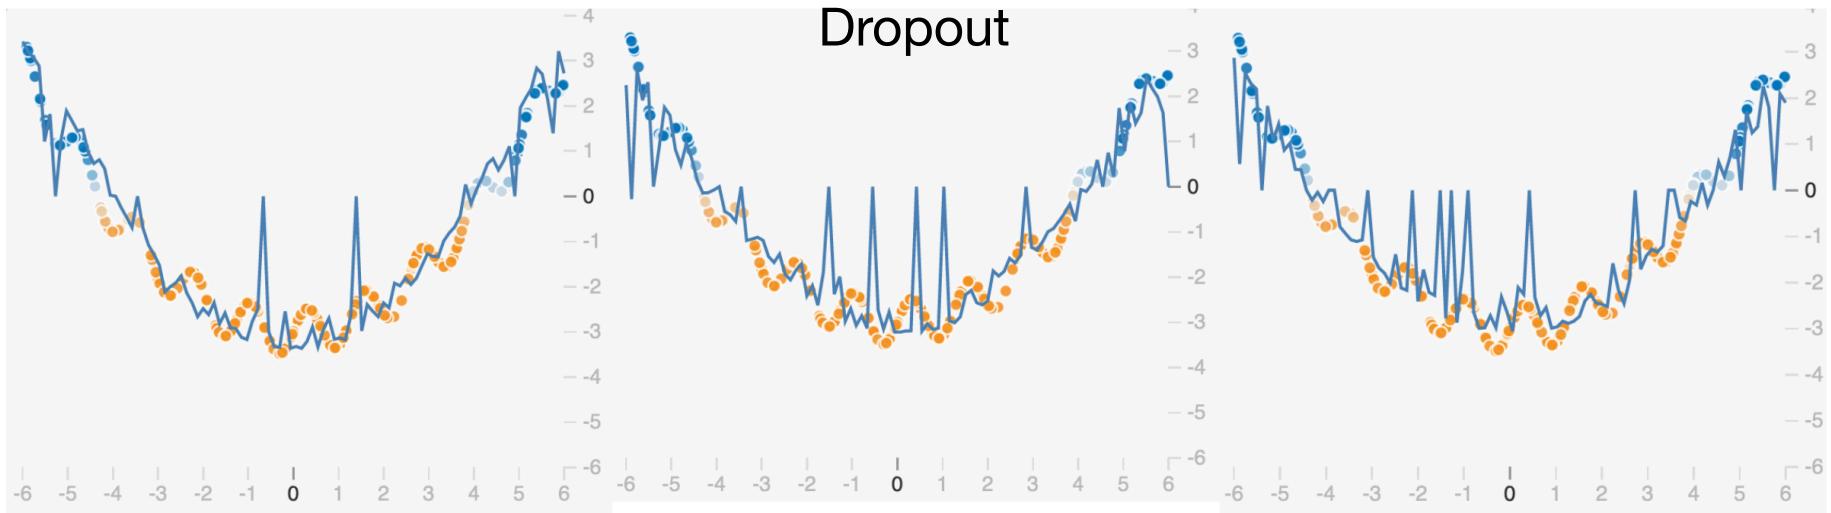
$$\phi(\mathbf{x}) = \sigma(\text{DO}_r(\mathbf{x})^T \mathbf{W} + \mathbf{b})$$

Dropout

$$\text{Dropout}(\mathbf{X}, r) = \mathbf{D} \odot \mathbf{X}, \quad \mathbf{D} = \begin{bmatrix} d_{11} & d_{12} & \dots & d_{1n} \\ d_{21} & d_{22} & \dots & d_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ d_{m1} & d_{m2} & \dots & d_{mn} \end{bmatrix}, \quad d_{ij} \sim \text{Bernoulli}(1 - r)$$

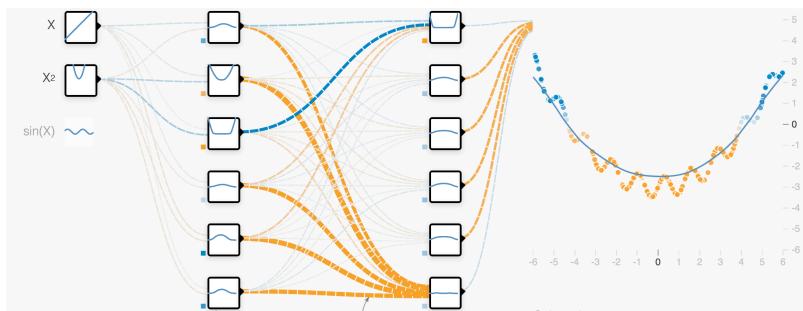
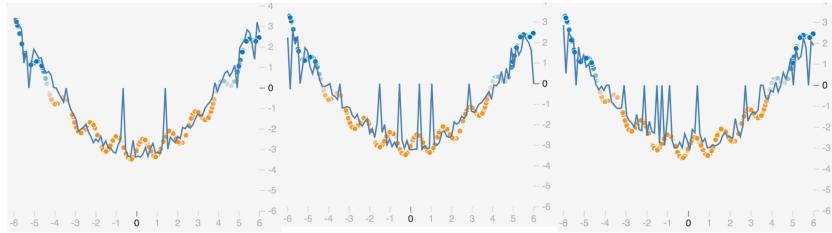


Dropout



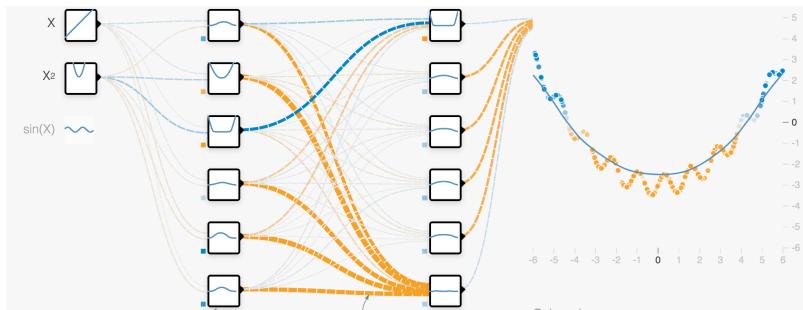
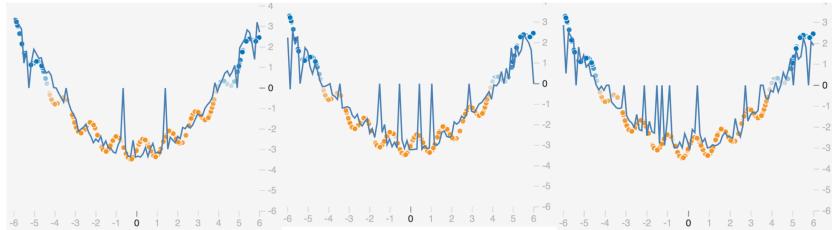
Dropout

$$\phi(\mathbf{x})_{train} = \sigma(\text{DO}_r(\mathbf{x})^T \mathbf{W} + \mathbf{b}) \quad \rightarrow \quad \phi(\mathbf{x})_{eval} = \sigma(\mathbf{x}^T \mathbf{W} + \mathbf{b})$$



Dropout

$$\phi(\mathbf{x})_{train} = \sigma(\text{DO}_r(\mathbf{x})^T \mathbf{W} + \mathbf{b}) \quad \rightarrow \quad \phi(\mathbf{x})_{eval} = \sigma(\mathbf{x}^T \mathbf{W} + \mathbf{b})$$



$$\mathbb{E}[\text{DO}_r(\mathbf{x})^T \mathbf{w}] = \sum_i d_i x_i w_i, \quad d_i \sim \text{Bernoulli}(1 - r)$$

$$= \sum_i p(d_i = 1) x_i w_i = (1 - r) \sum_i x_i w_i < \sum_i x_i w_i$$

Dropout