

Hands-on 2: Computer vision classification for cattle food grading with a hands-on exercise (Python) on Google Co- lab

Yalong Pi, Associate Research Scientist, Texas A&M Institute of
Data Science

2025 ASAS, Sunday, July 6, 2025, 12:30 PM – 2:20 PM

https://github.com/National-Animal-Nutrition-Program/2025ASAS_Pi

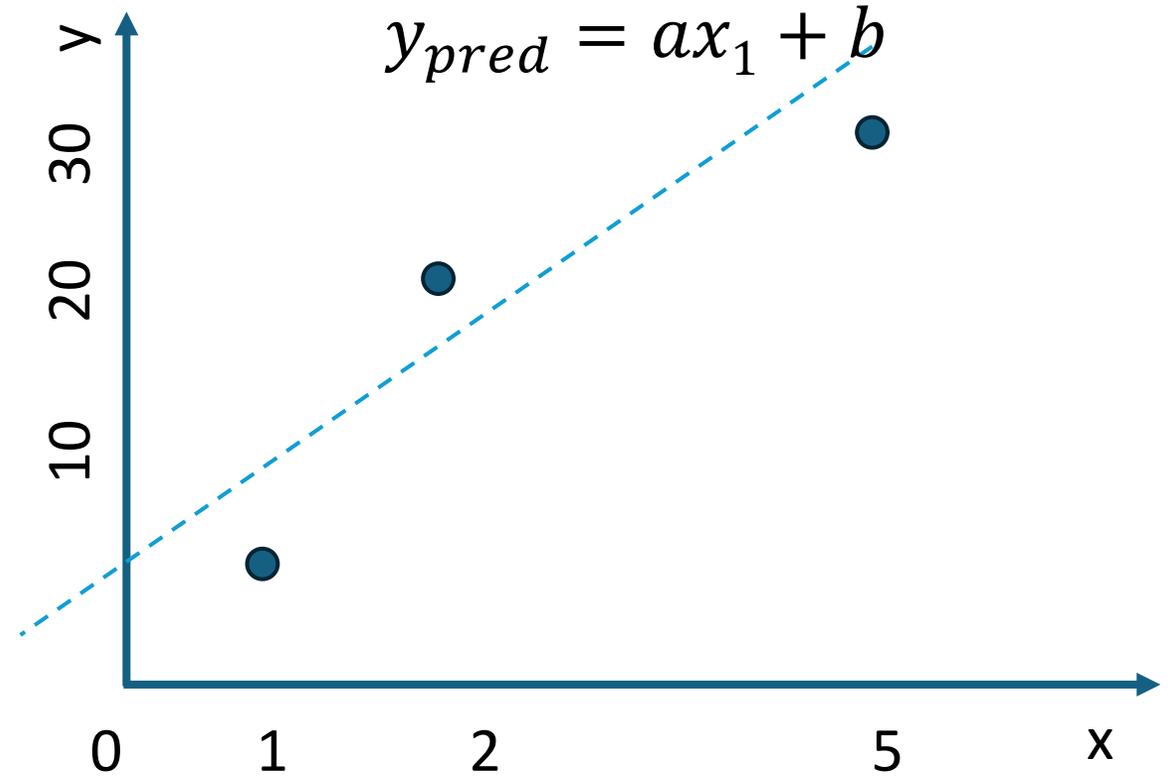
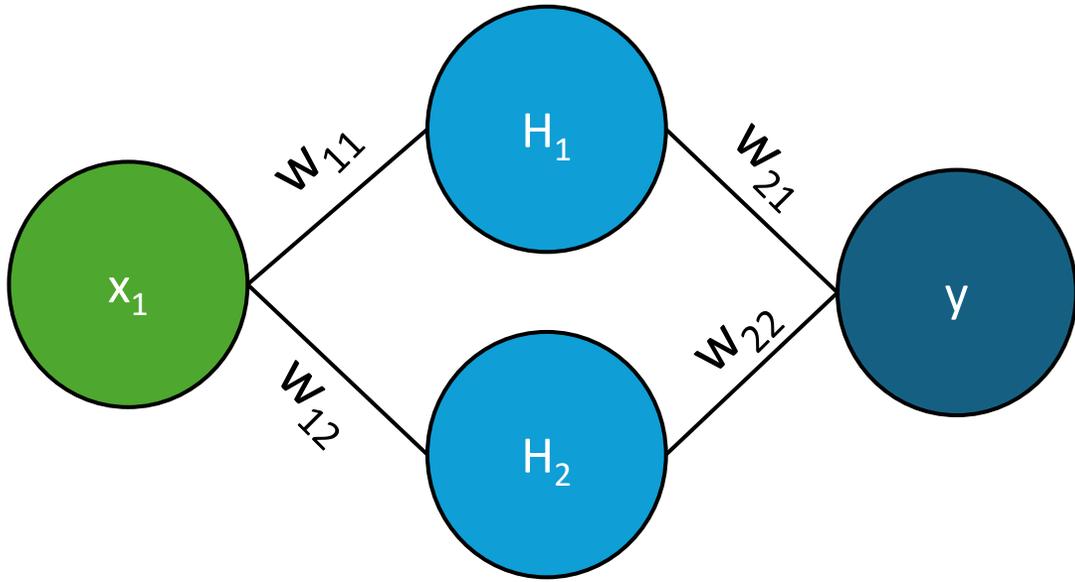
- Objective
- Data
- Simple neural network
- Convolutional network (pre-trained)
- Test on unseen data
- Result visualization

PI



Agenda

- Fundamentals of machine learning (1 hour)
 - Simple example
 - Training and test
 - Understand the results
 - metric
- Feed grading with computer vision part 1 (20 minutes)
- Feed grading with computer vision part 2 (20 minutes)



$$H_1 = x_1 * w_{11}$$

$$H_2 = x_1 * w_{12}$$

$$y_{pred} = x_1 * w_{11} * w_{21} + x_1 * w_{12} * w_{22}$$

x	y
1	10
2	20
5	30

$$y_{pred} = w_{21} * x_1 * w_{11} + w_{22} * x_1 * w_{12}$$

$$loss = (y_{pred} - y_{true})^2$$

$$loss = (y_{true} - x_1 * w_{11} * w_{21} + x_1 * w_{12} * w_{22})^2$$

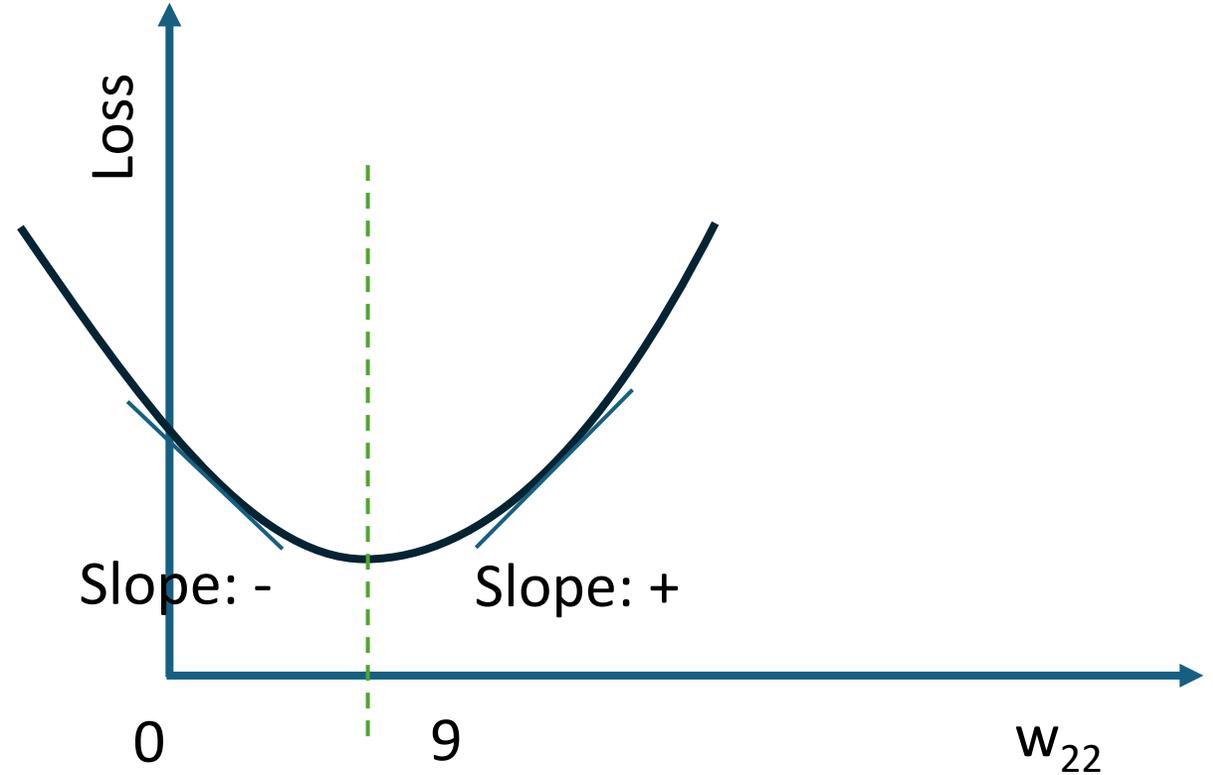
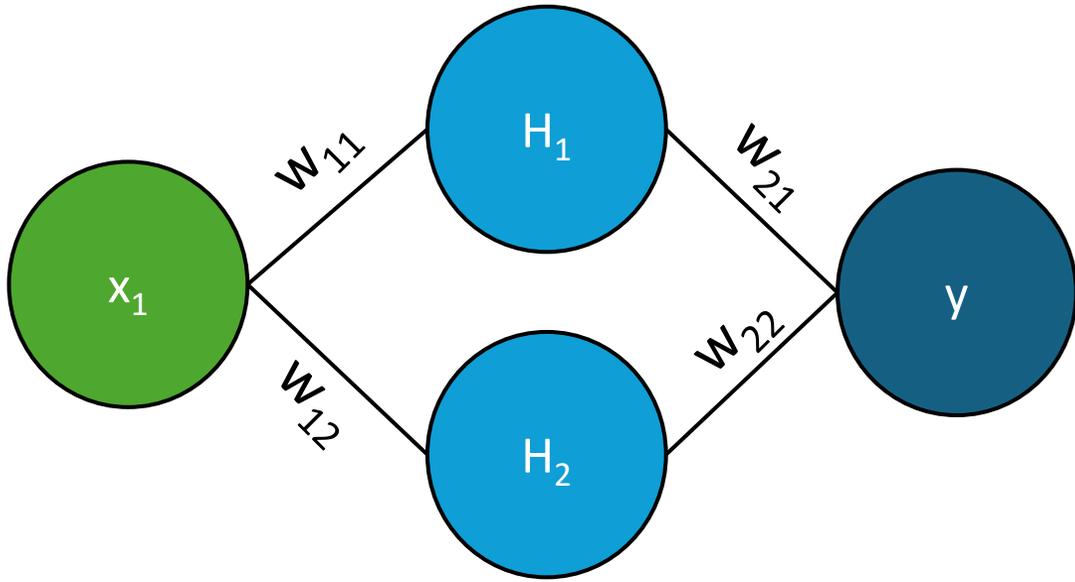
$$= (10 - (1 * 1 * 1 + 1 * 1 * w_{22}))^2$$

$$= (10 - (1 + w_{22}))^2$$

$$= (9 - w_{22})^2$$

$$= 91 - 18w_{22} + w_{22}^2$$

x	y
1	10
2	20
5	30

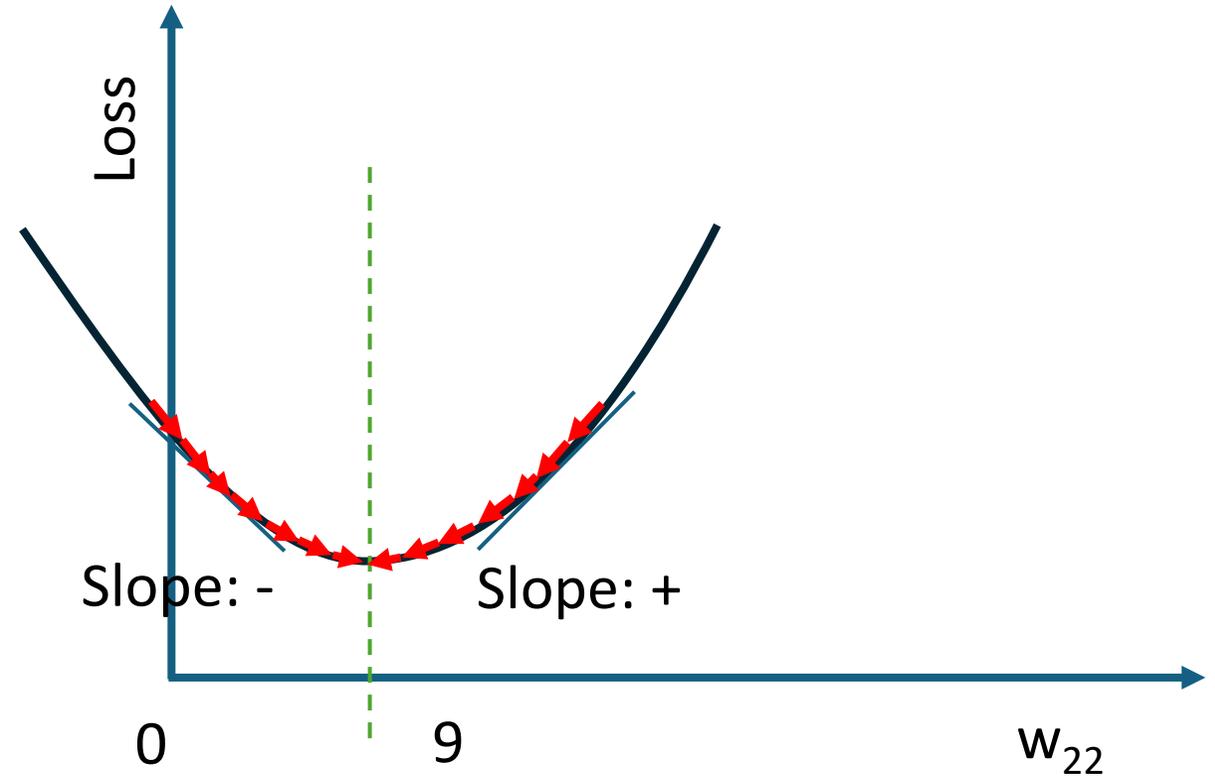
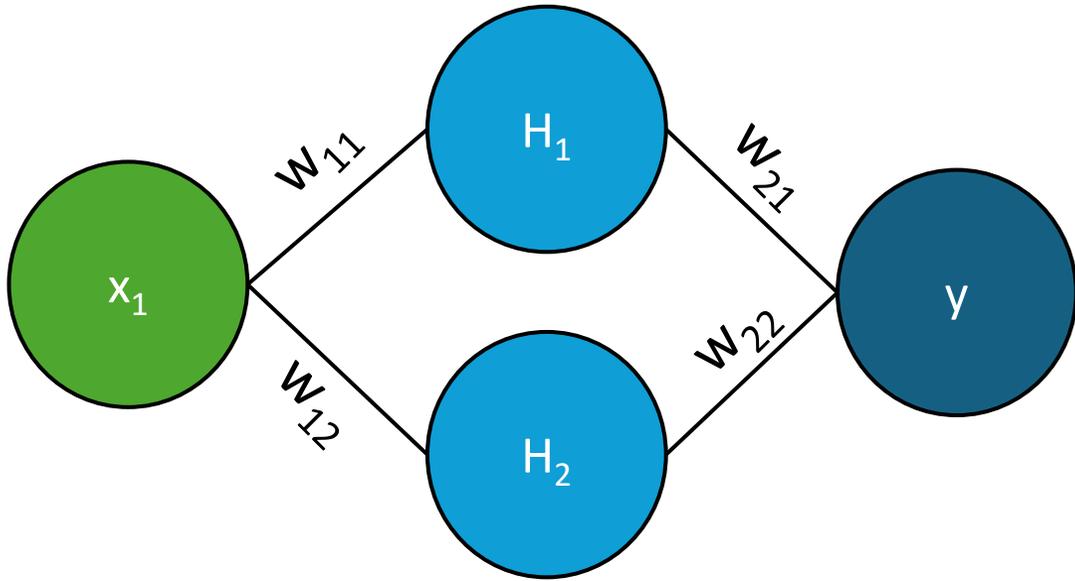


$$loss = 91 - 18w_{22} + w_{22}^2$$

$$slope_{w_{22}} = \frac{loss(w_{22} + \Delta) - loss(w_{22})}{\Delta}$$

$$= 2 * w_{22} - 18 = \frac{\partial loss}{\partial w_{22}}$$

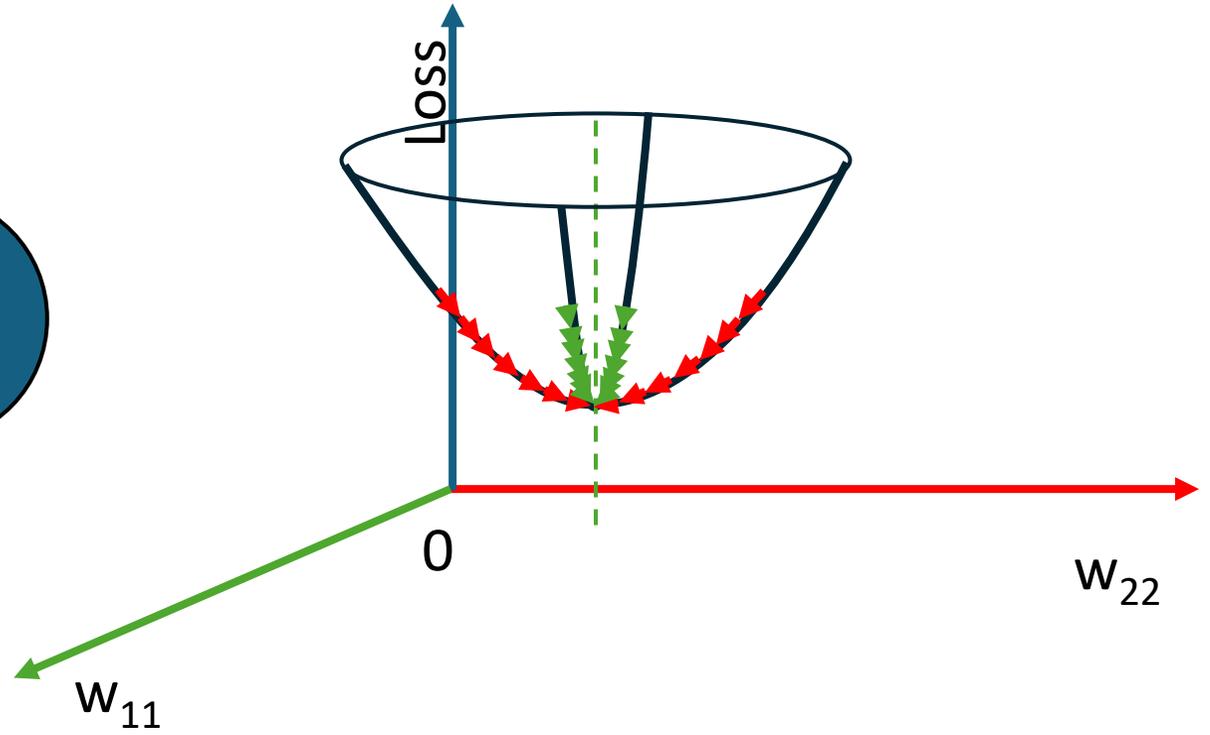
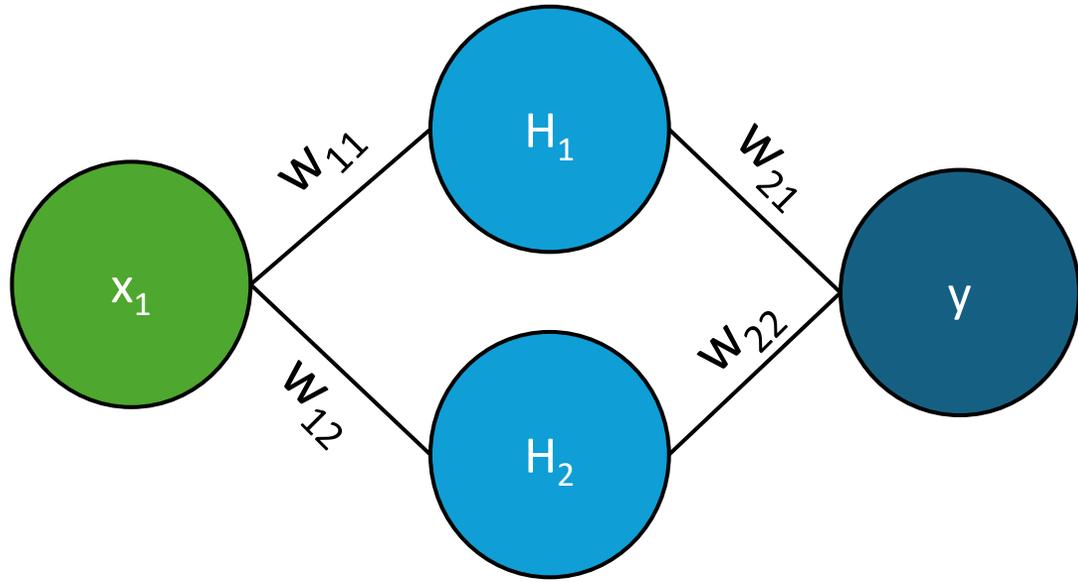
x	y
1	10
2	20
5	30



$$w_{22}(new) = w_{22}(old) - slope_{22} * LR$$

LR is a very small learning rate
 $= 0.0001$

x	y
1	10
2	20
5	30



$$w_{22}(new) = w_{22}(old) - slope_{22} * LR$$

$$w_{11}(new) = w_{11}(old) - slope_{11} * LR$$

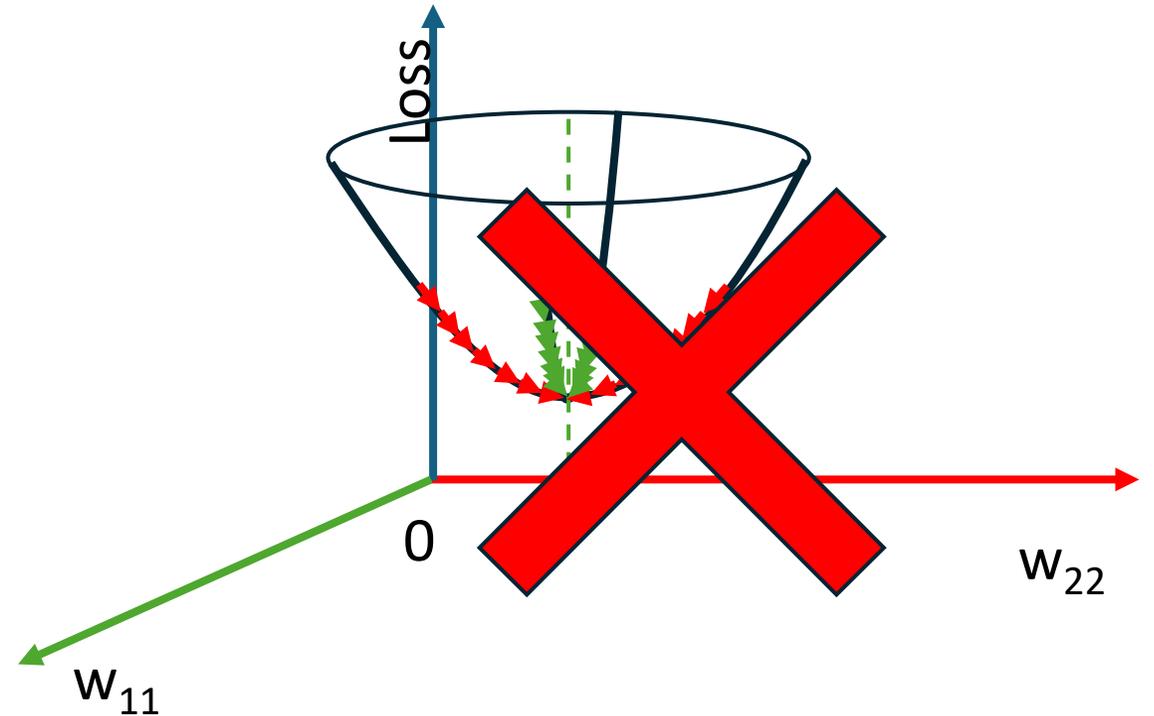
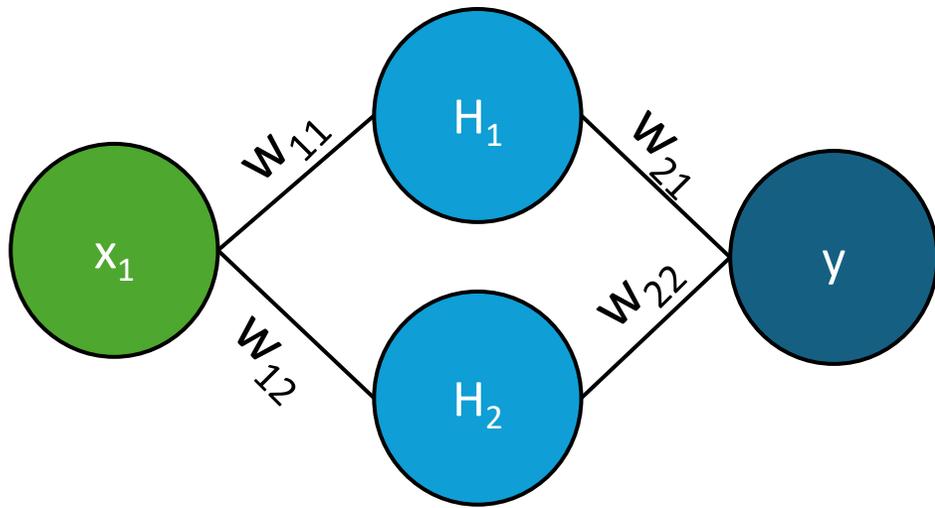
x	y
1	10
2	20
5	30

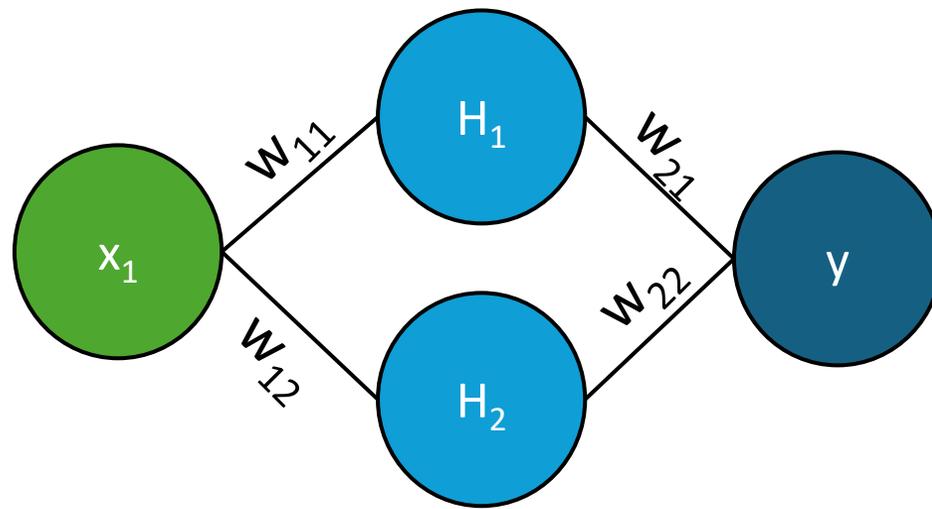
$$w_{22}(new) = w_{22}(old) - slope_{22} * LR$$

$$w_{11}(new) = w_{11}(old) - slope_{11} * LR$$

$$w_{12}(new) = w_{12}(old) - slope_{12} * LR$$

$$w_{21}(new) = w_{21}(old) - slope_{21} * LR$$





$$w_{22}(new) = w_{22}(old) - \frac{\partial loss}{\partial w_{22}} * LR$$

$$w_{11}(new) = w_{11}(old) - \frac{\partial loss}{\partial w_{11}} * LR$$

$$w_{12}(new) = w_{12}(old) - \frac{\partial loss}{\partial w_{12}} * LR$$

$$w_{21}(new) = w_{21}(old) - \frac{\partial loss}{\partial w_{21}} * LR$$

$$\begin{cases}
 w_{22}(new) = w_{22}(old) - \frac{\partial loss}{\partial w_{22}} * LR \\
 w_{11}(new) = w_{11}(old) - \frac{\partial loss}{\partial w_{11}} * LR \\
 w_{12}(new) = w_{12}(old) - \frac{\partial loss}{\partial w_{12}} * LR \\
 w_{21}(new) = w_{21}(old) - \frac{\partial loss}{\partial w_{21}} * LR
 \end{cases}$$



$$\vec{w}(new) = \vec{w}(old) - \nabla * LR \quad \text{↻}$$

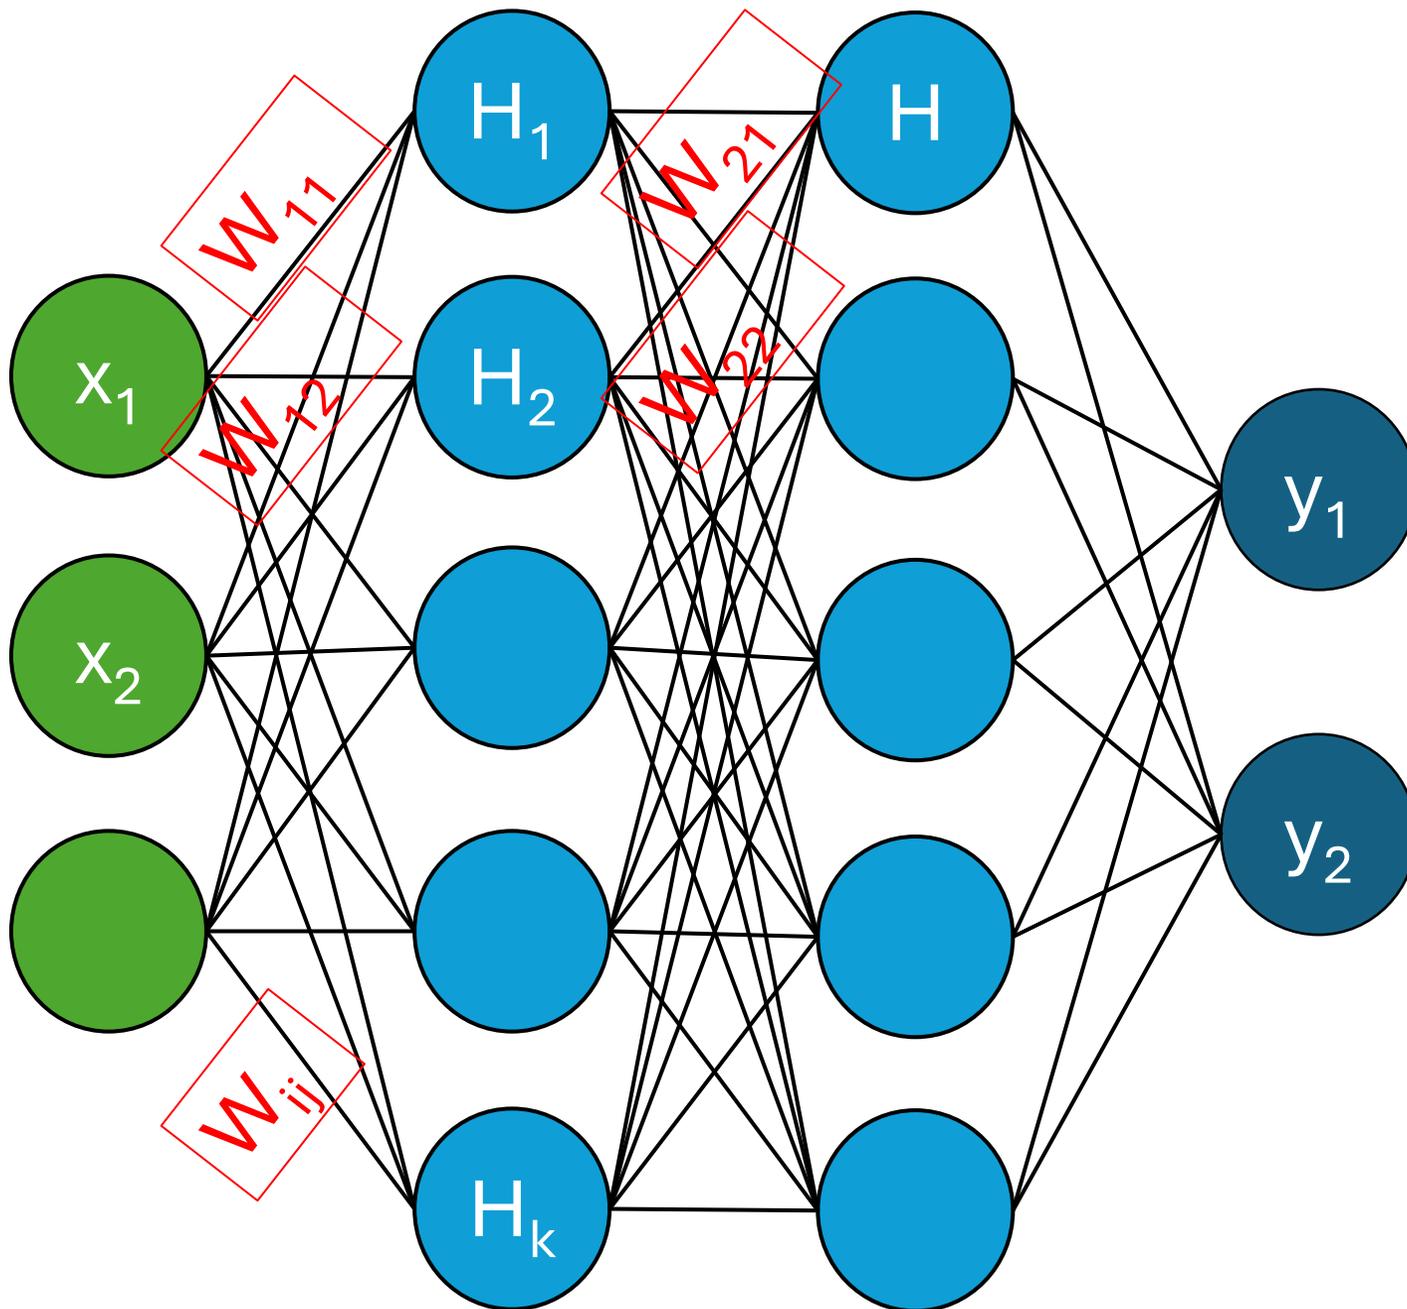
For each neuron H:

$$H = \sum x_k w_{ij} + b$$

□ w_{ij} is the weight

□ b is the bias

□ A DNN has millions of weights and biases



For each neuron H:

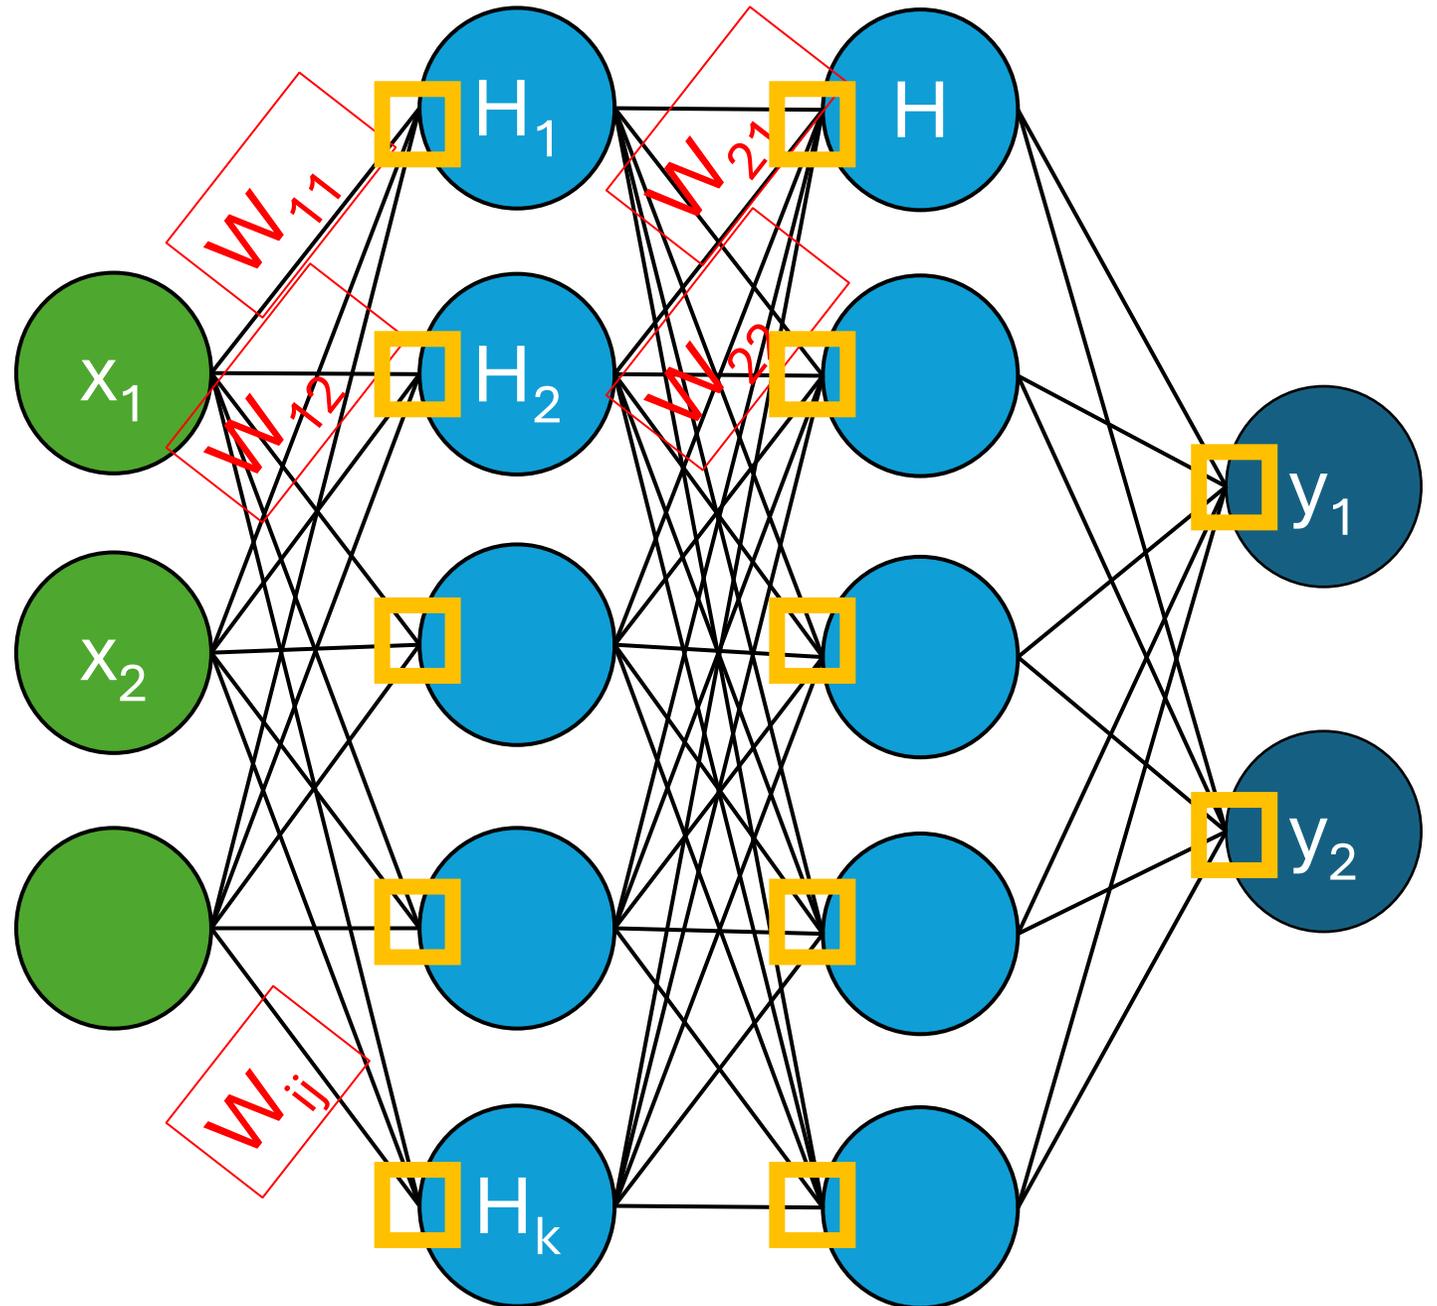
$$H = f\left(\sum x_k w_{ij} + b\right)$$

□ f is the activation function

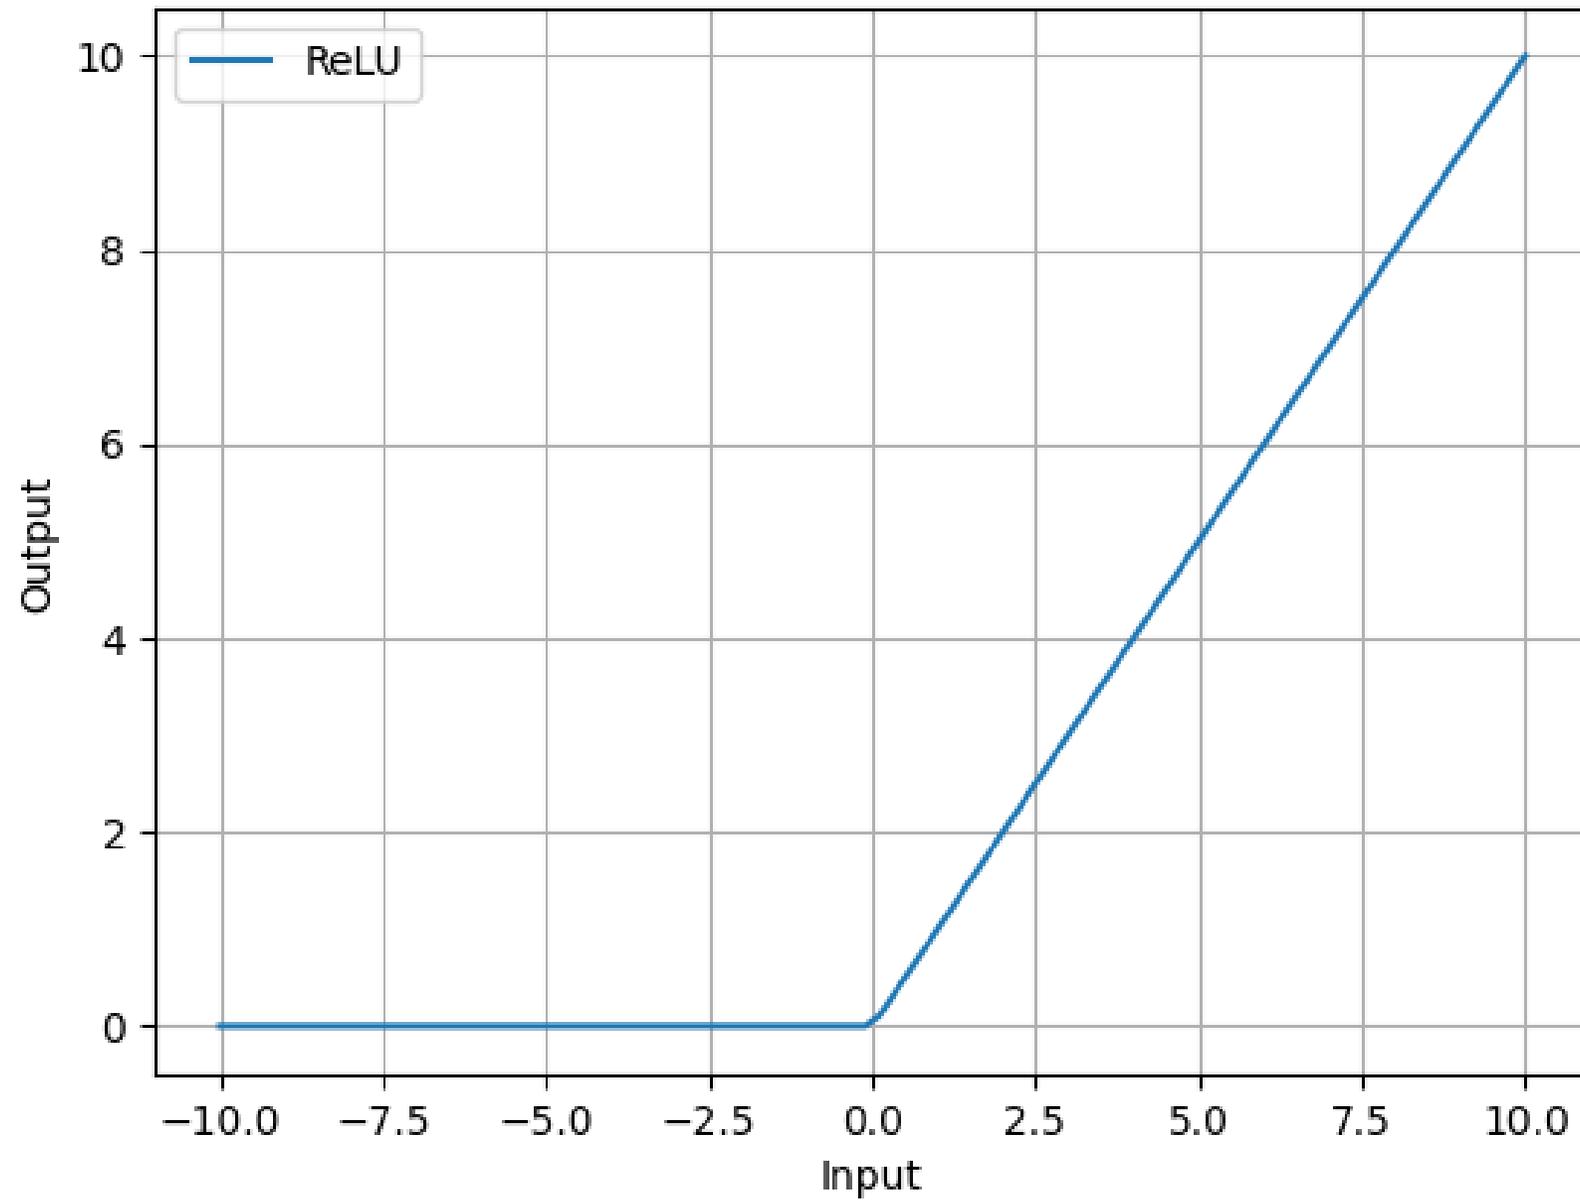
□ w_{ij} is the weight

□ b is the bias

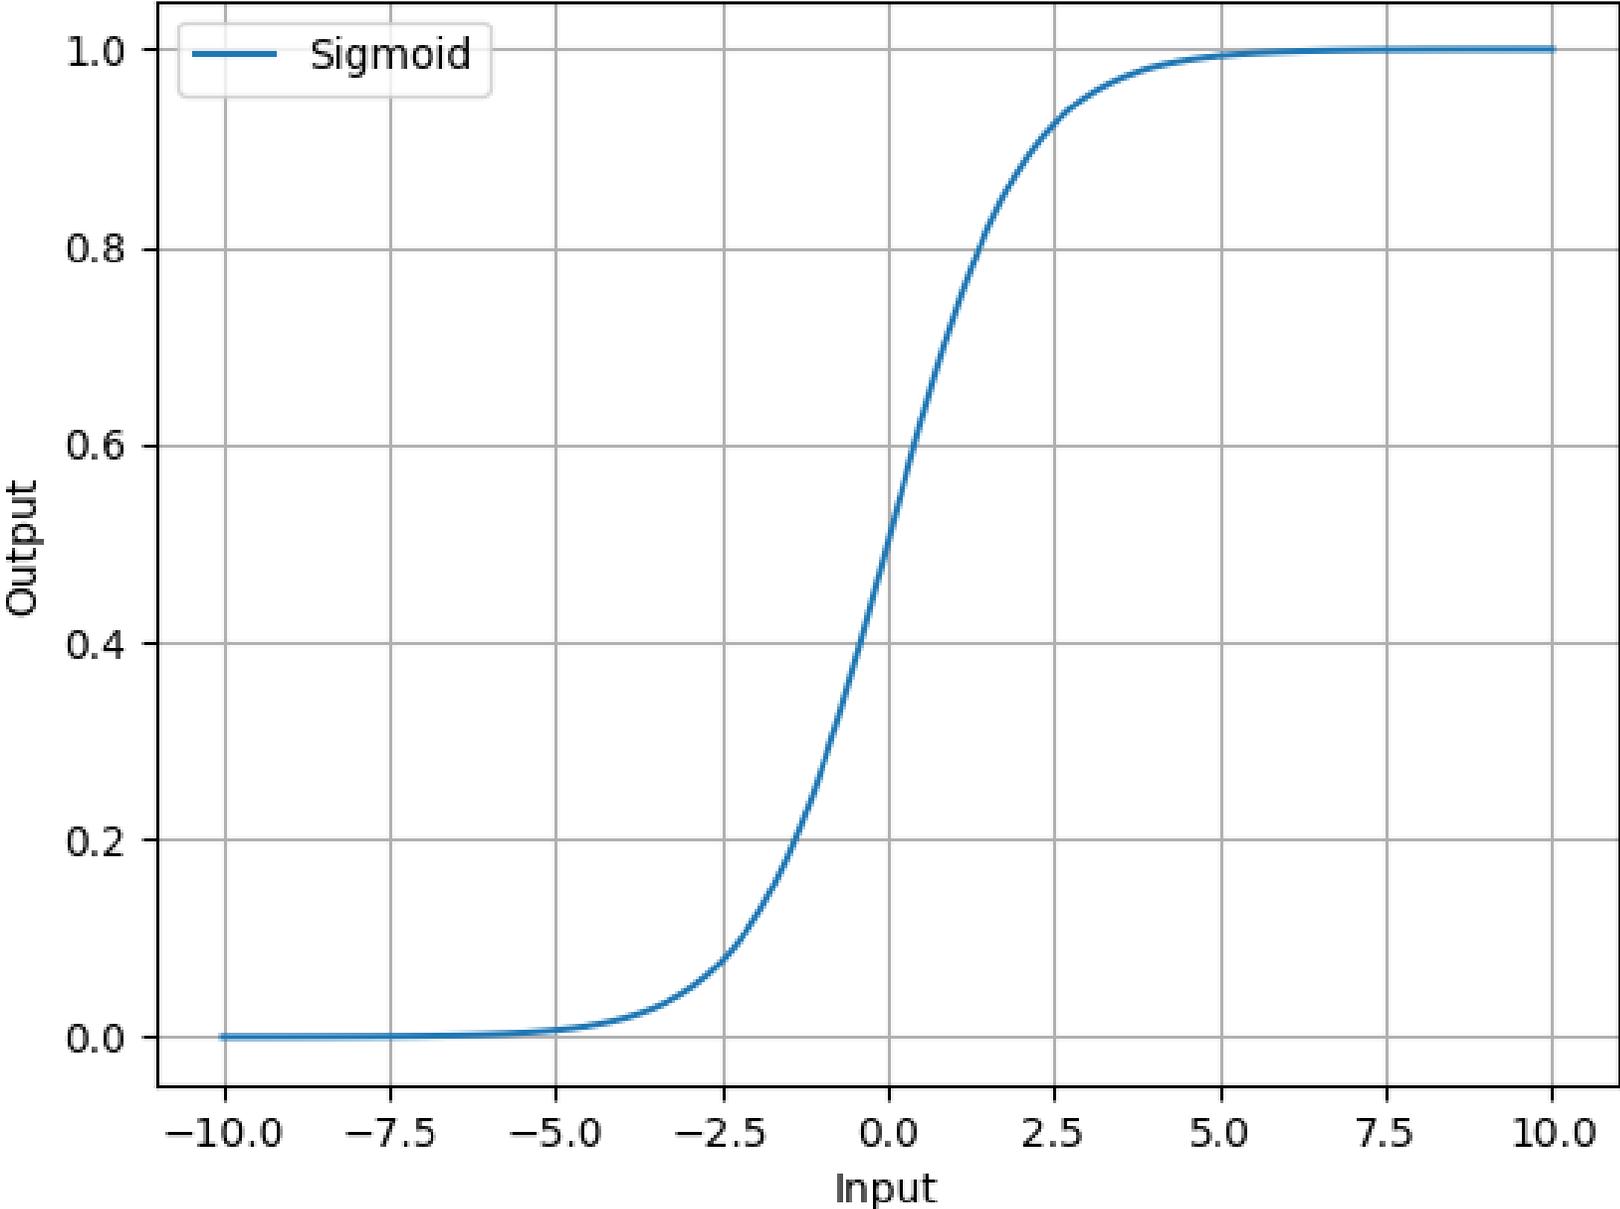
□ A DNN has millions of weights and biases



ReLU Activation Function

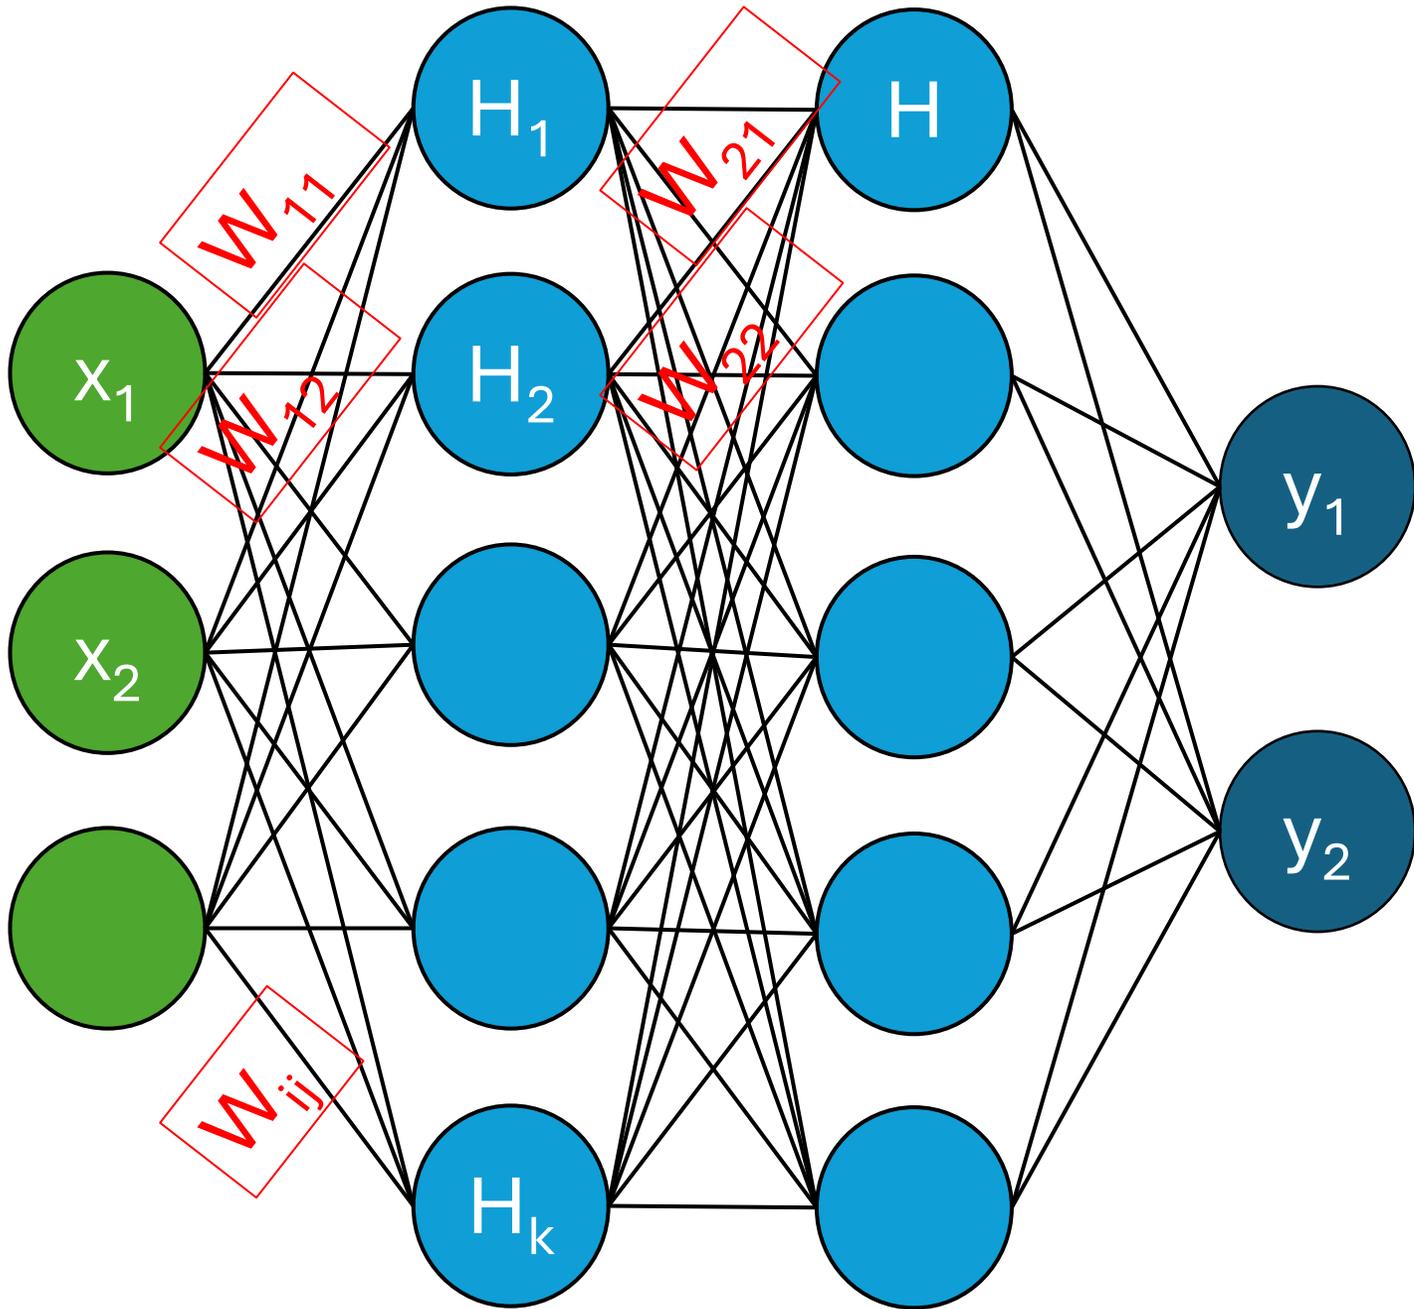


Sigmoid Activation Function



Output types

- Regression
- Classification
- Feature extraction
- Combination



Weight (lbs)

y_{pred}

y_{true}

$\begin{bmatrix} 110 \\ 5.1 \end{bmatrix}$

$\begin{bmatrix} 150 \\ 7.1 \end{bmatrix}$

Height (feet)

Regression loss

- Mean absolute error

$$MAE = \frac{1}{n} |y_{pred} - y_{true}|$$

- Mean square error

$$MSE = \frac{1}{n} (y_{pred} - y_{true})^2$$

One hot encoding for y_{true}

Chicken
Banana

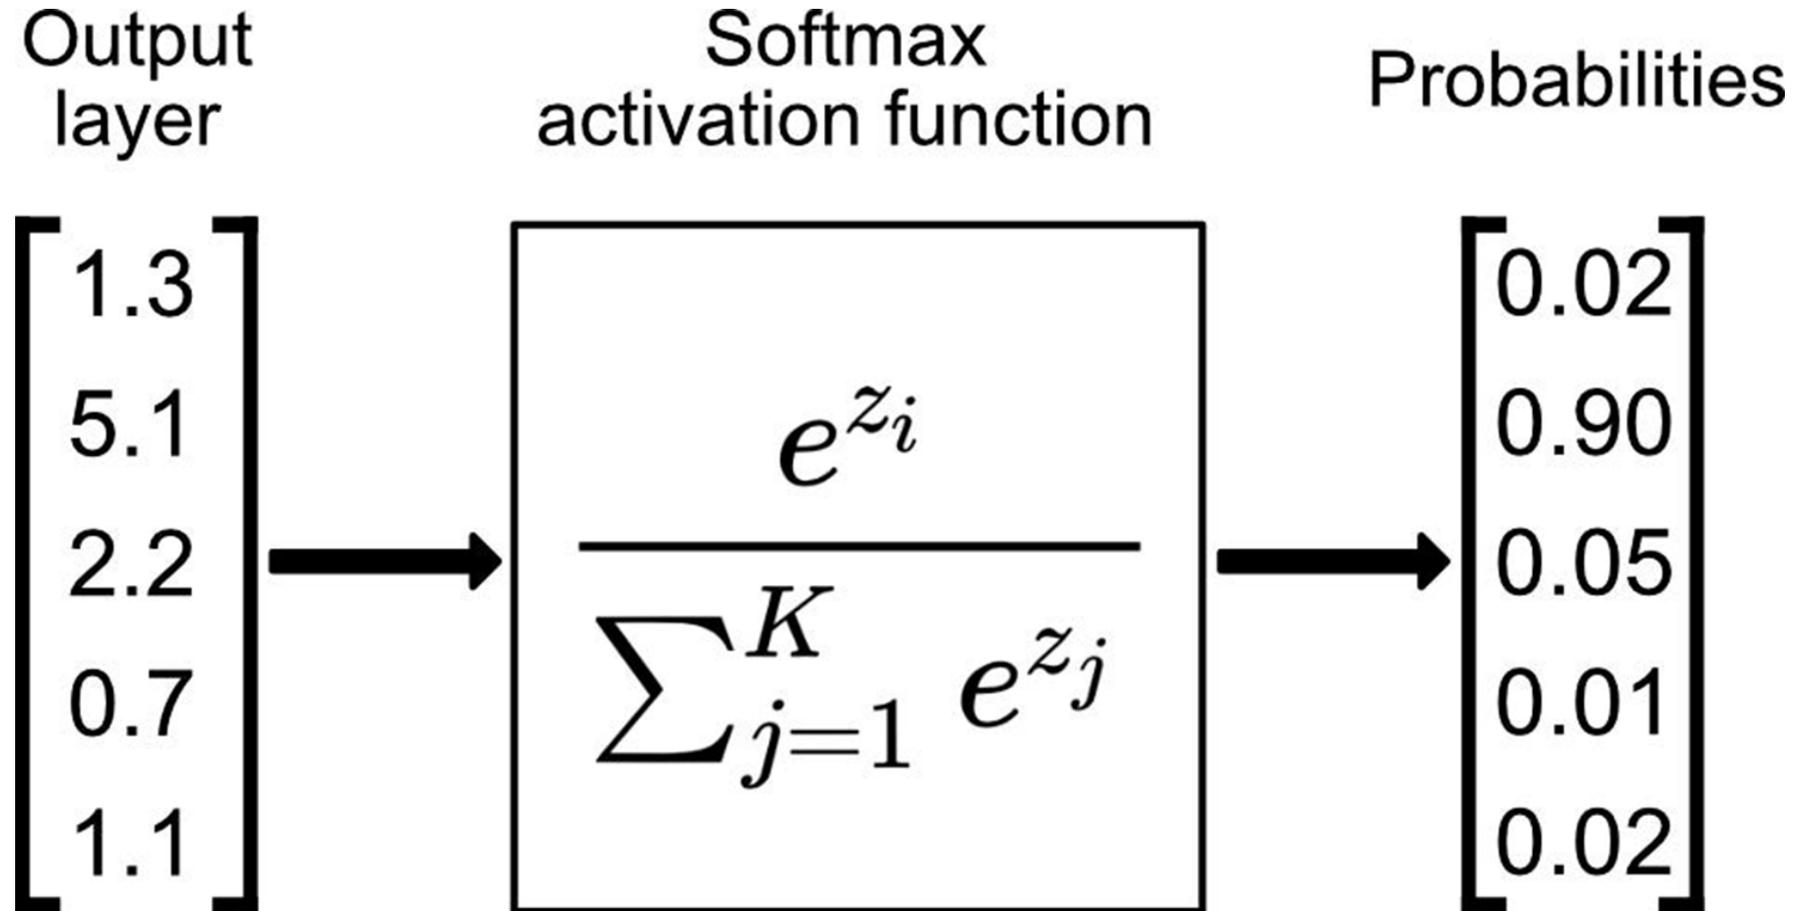


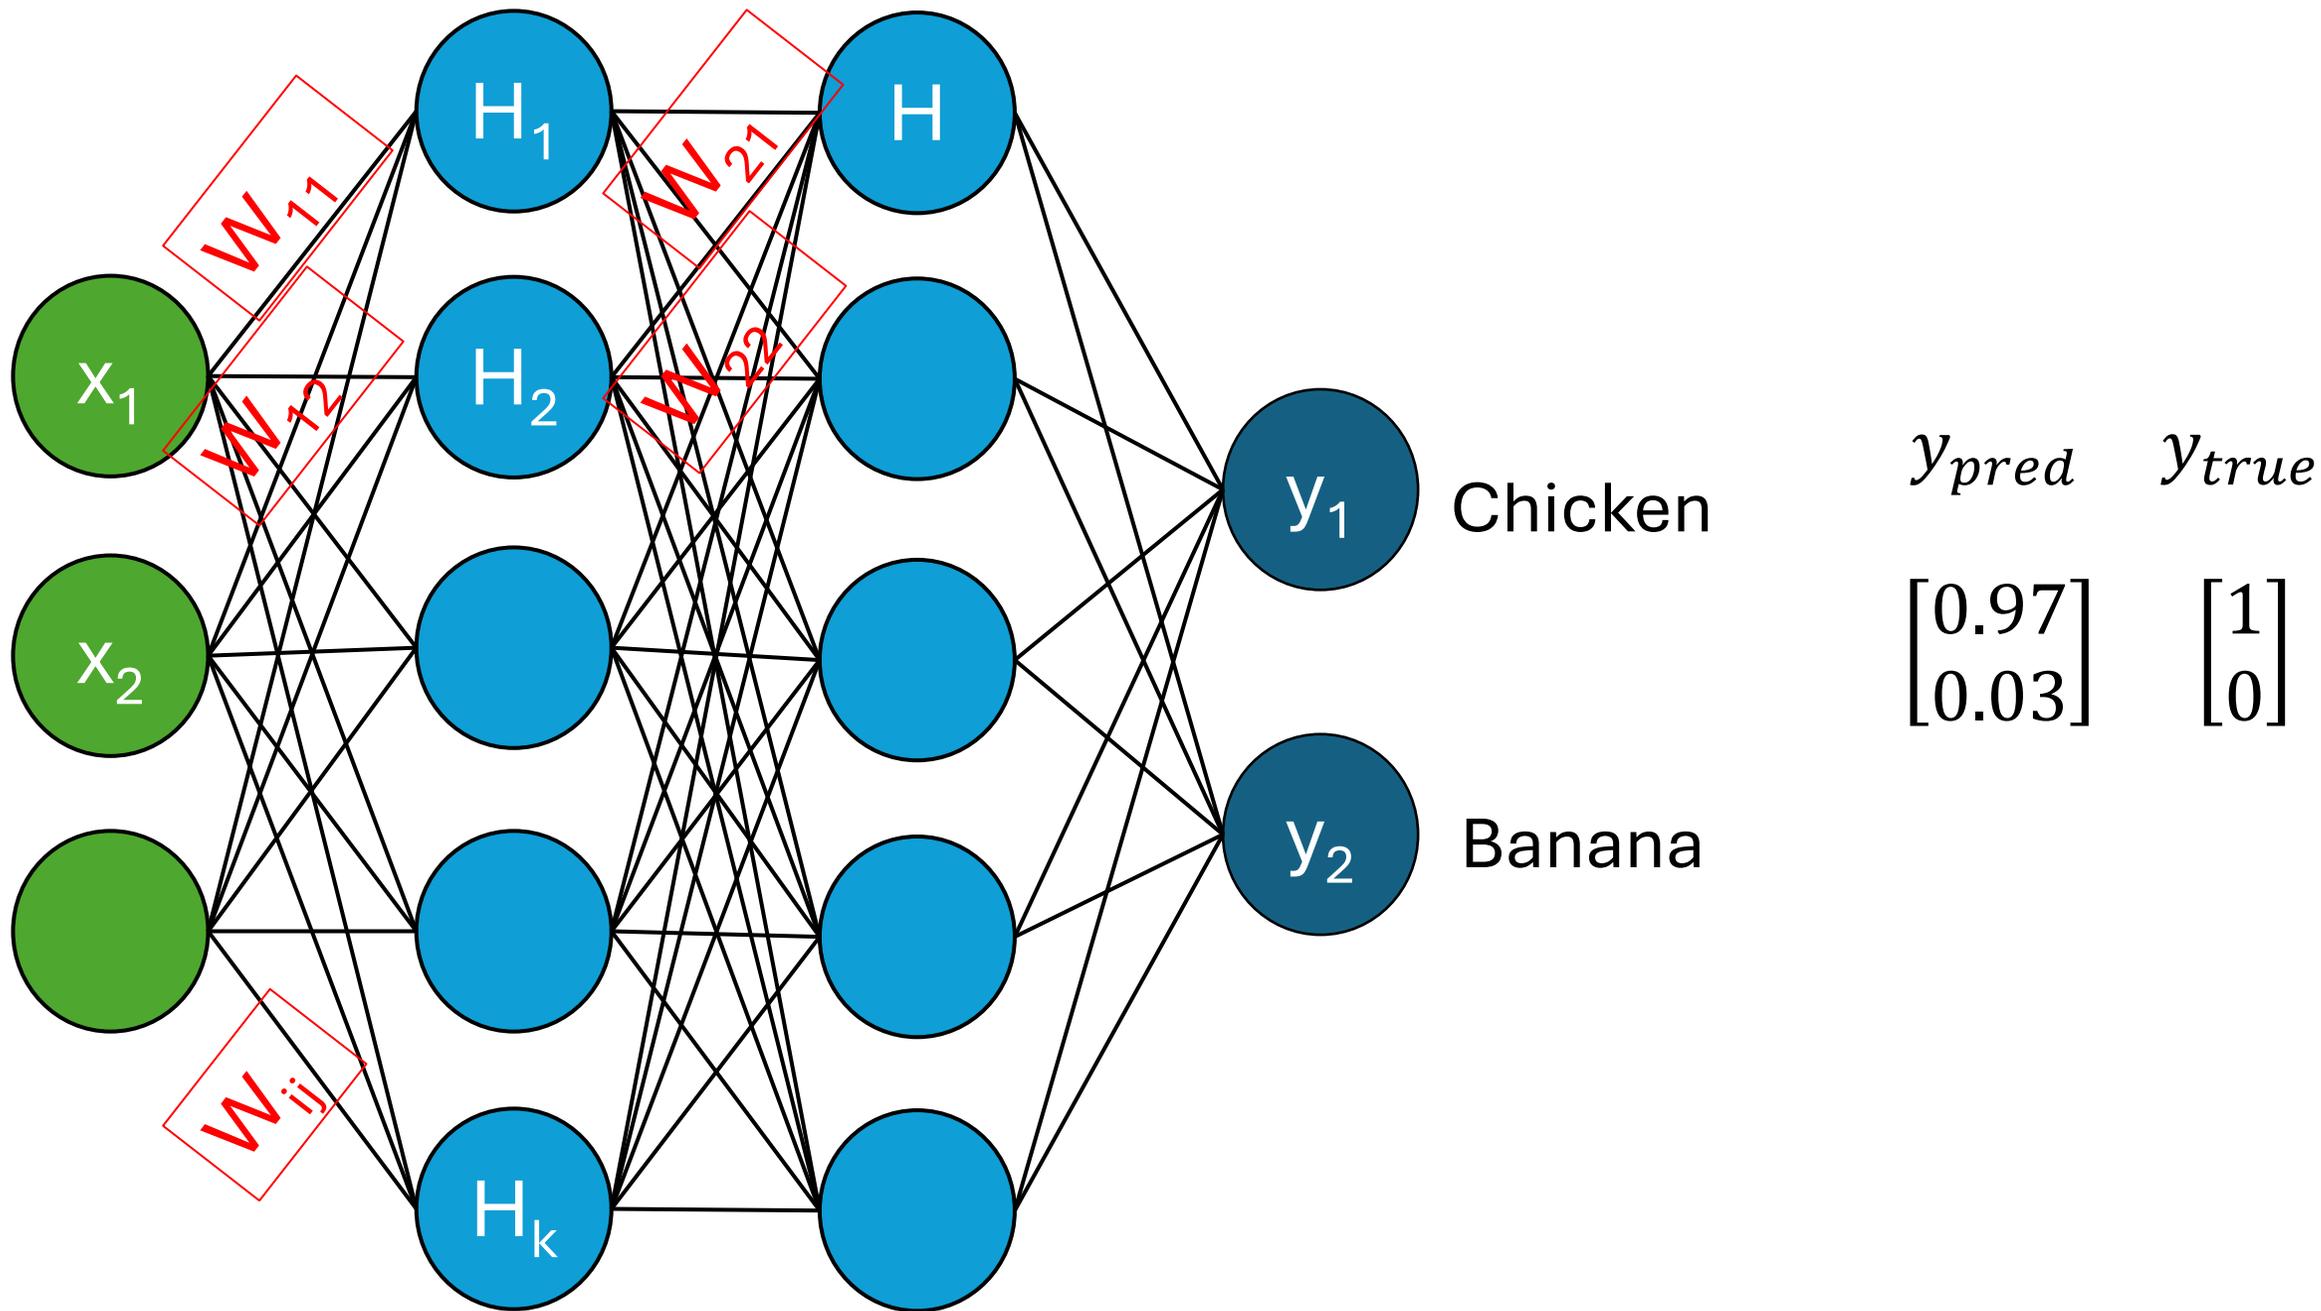
$$\begin{bmatrix} 1 \\ 0 \end{bmatrix}$$



$$\begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

Softmax for y_pred





Cross entropy loss for classification

- Binary cross entropy

- After softmax

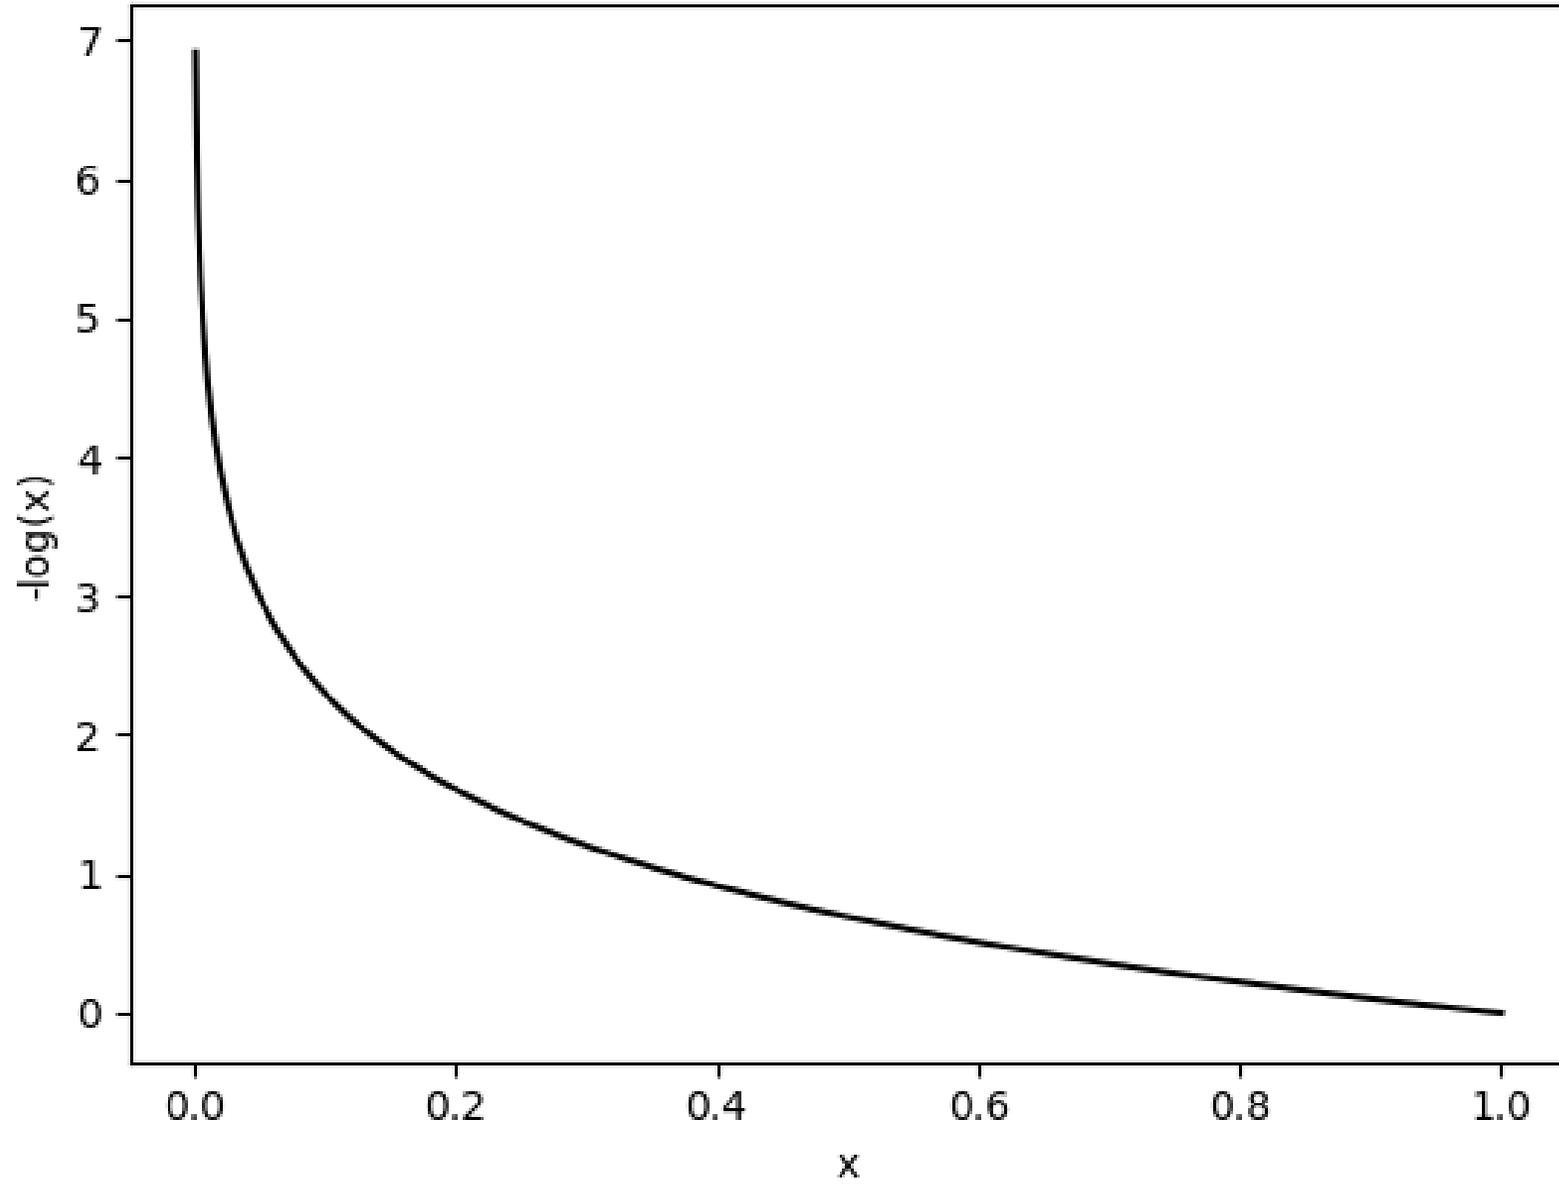
$$BCE = - \sum y_{true} * \log(y_{pred}) + (1 - y_{true}) * \log(1 - y_{pred})$$

- Multi class entropy

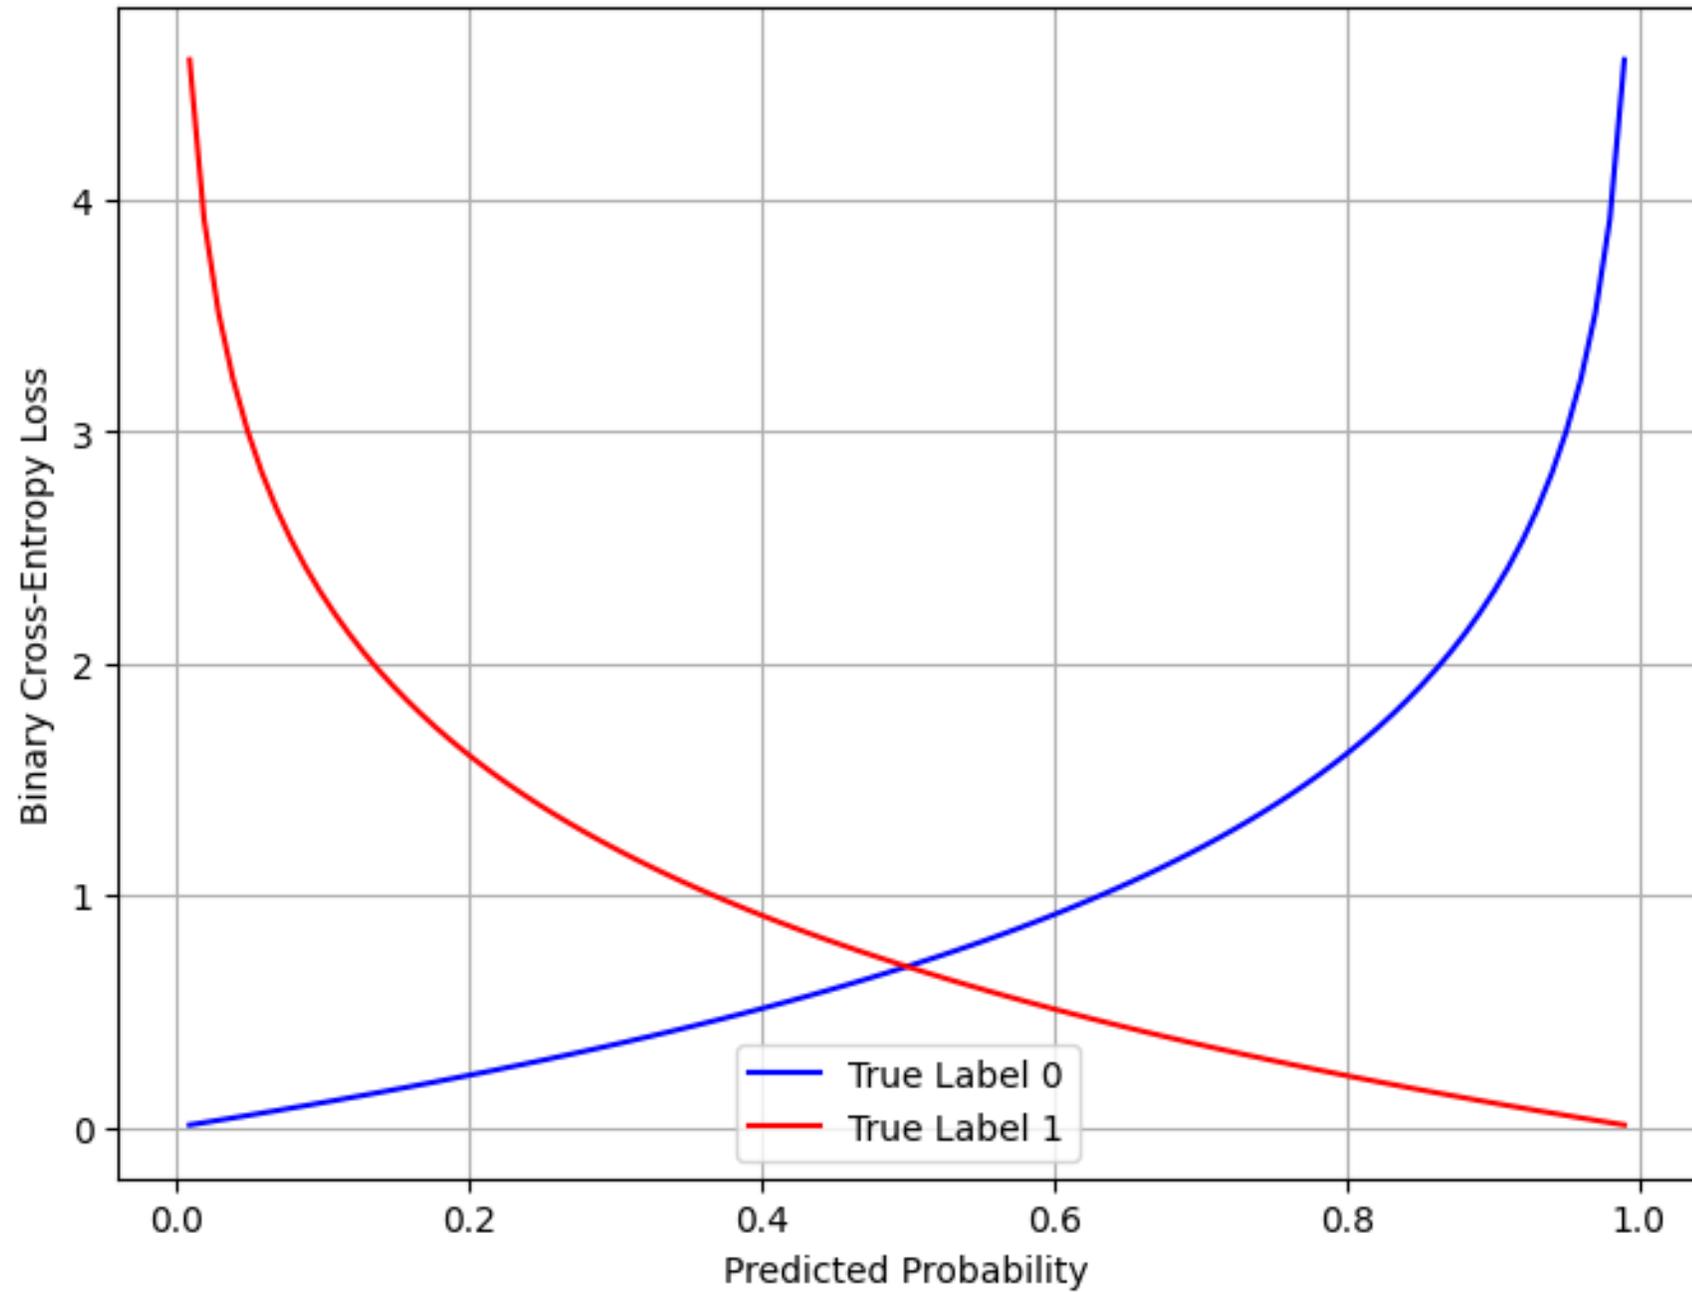
- After softmax

$$CE = - \sum y_{true} * \log(y_{pred})$$

Range of negative log-likelihood

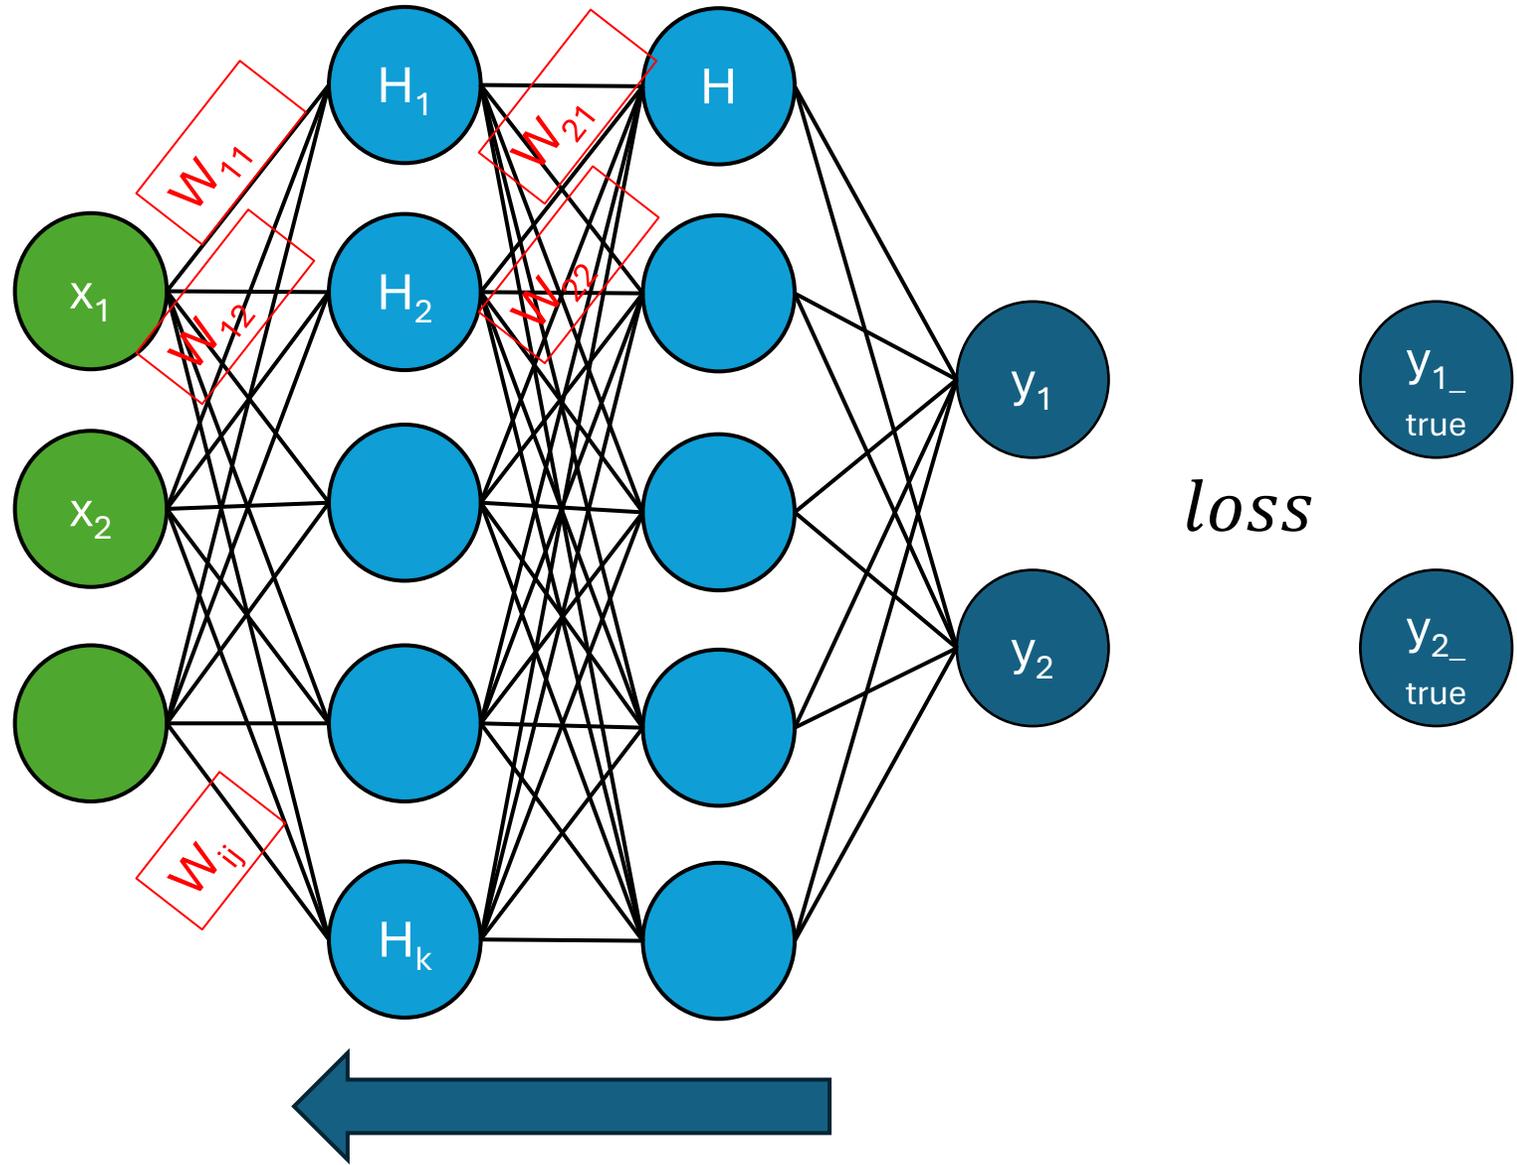


Binary Cross-Entropy Loss for True Labels 0 and 1



Weight update

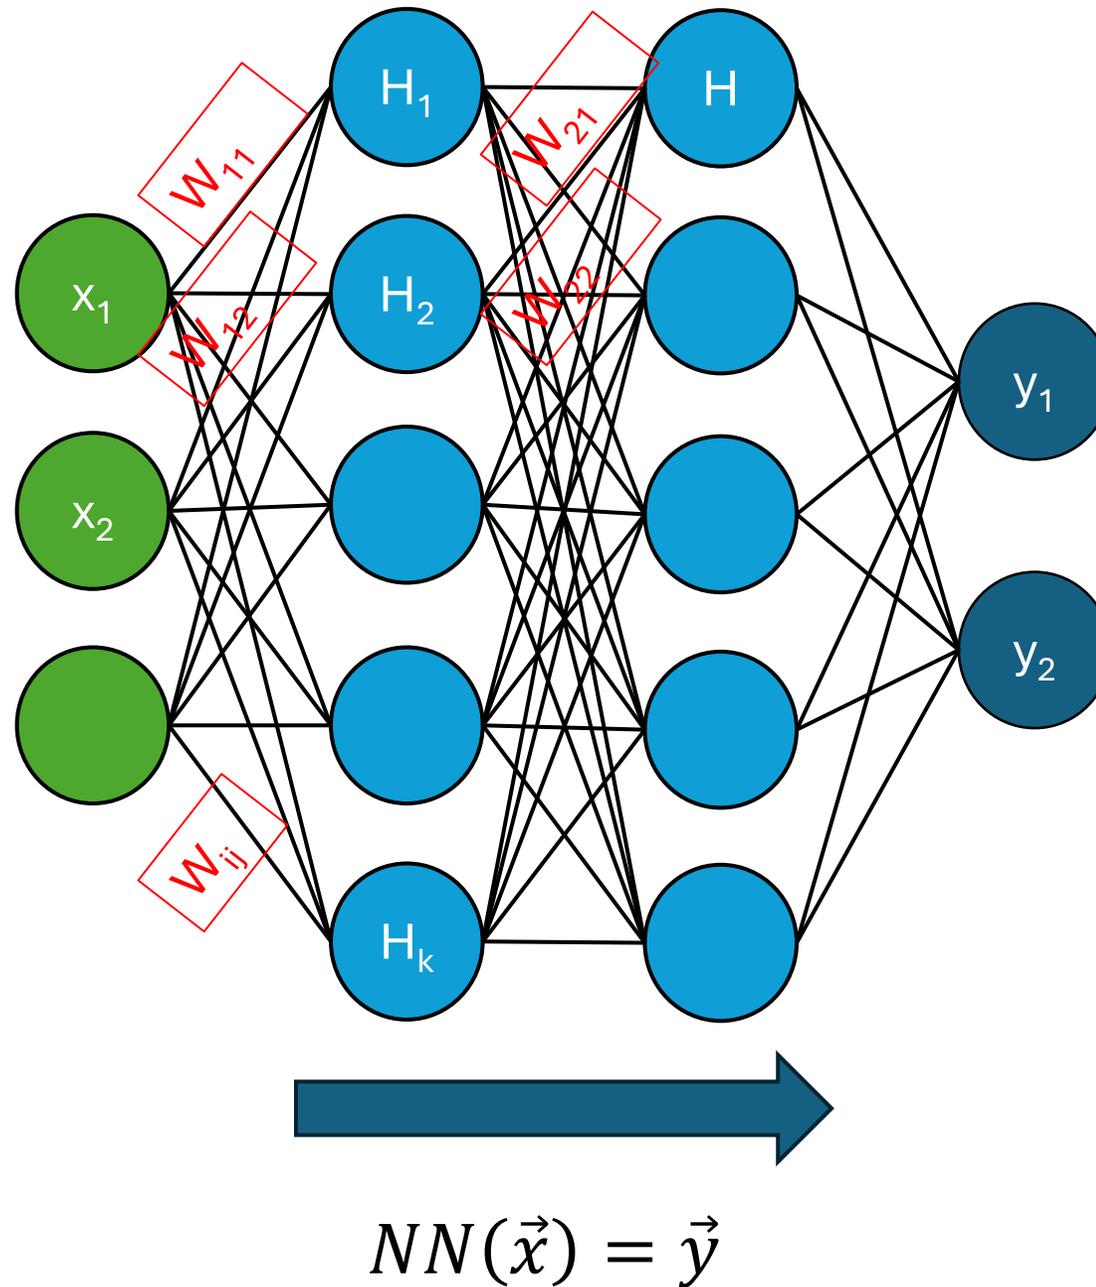
- Learn
- Fit
- Train
- Back propagation



$$new \vec{w} = old \vec{w} - \nabla * LR$$

Use the model

- Forward pass
- Inference
- Test
- Validation
- Deployment



Actual Values

1

0

Predicted Values



Actual Values

1

0

Predicted Values

1	TP	FP
0	FN	TN

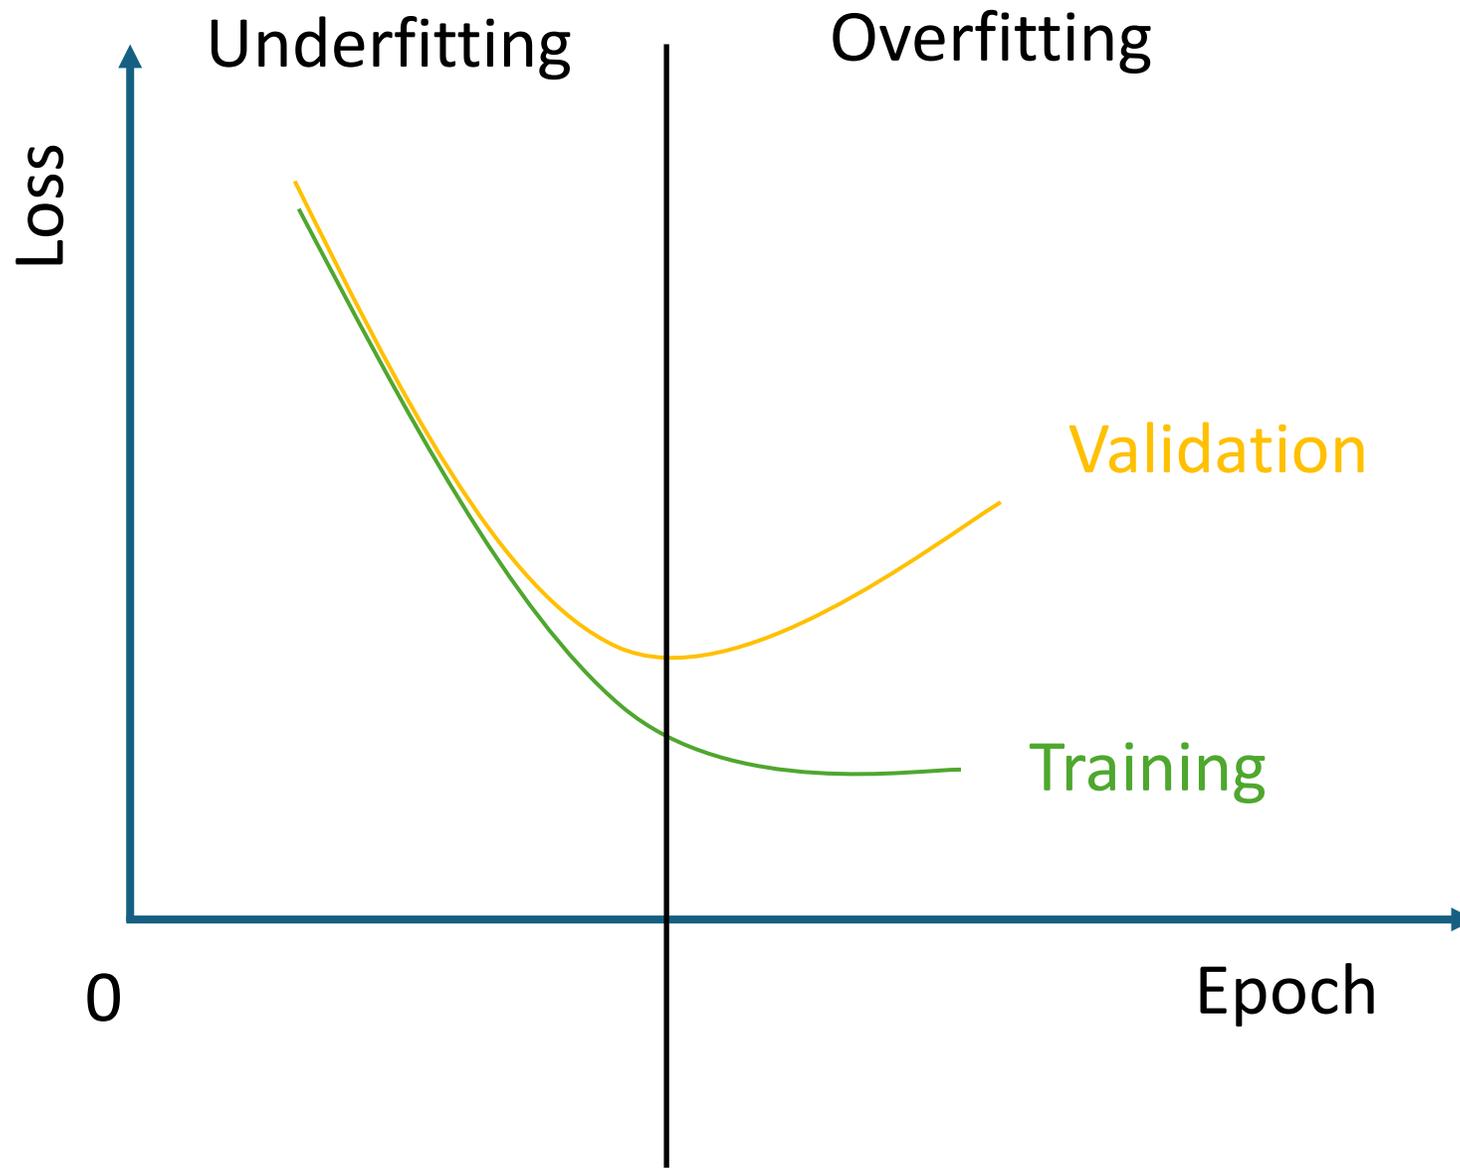
$$precision = \frac{TP}{TP + FP}$$

$$recall = \frac{TP}{TP + FN}$$

$$F1 = \frac{2 \times precision \times recall}{precision + recall}$$

$$accuracy = \frac{TP + TN}{TP + FN + TN + FP}$$

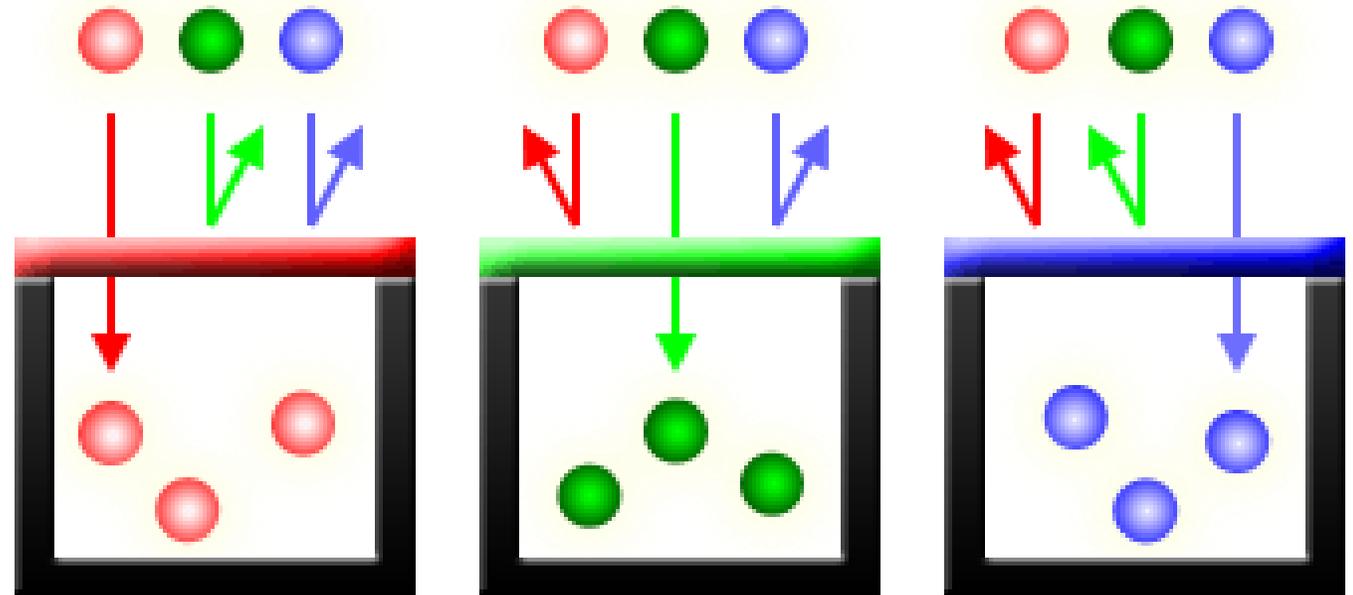
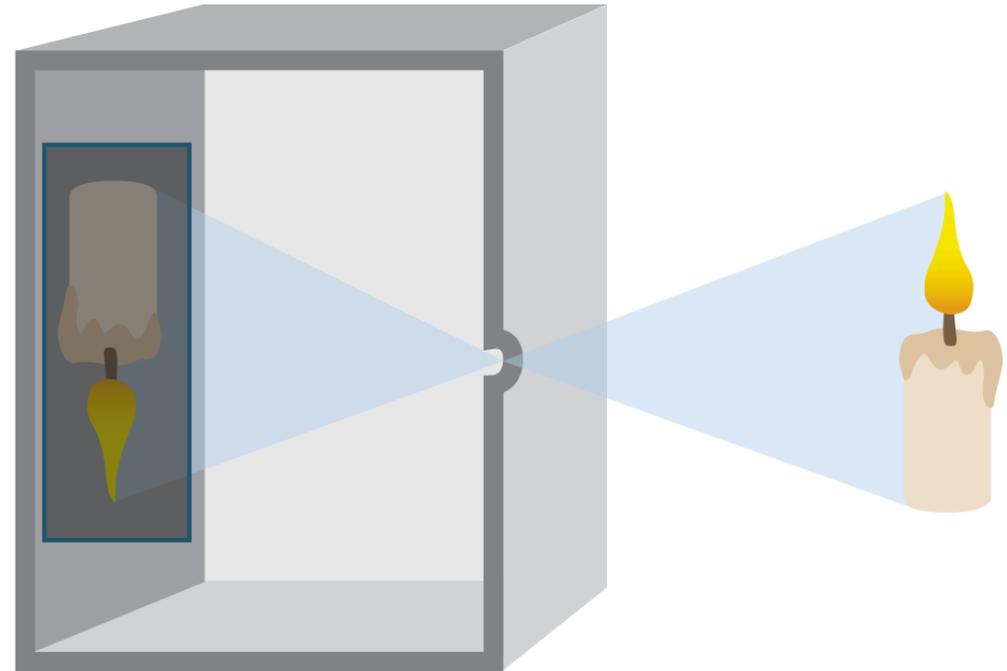
Image Credit: LinkedIn.com



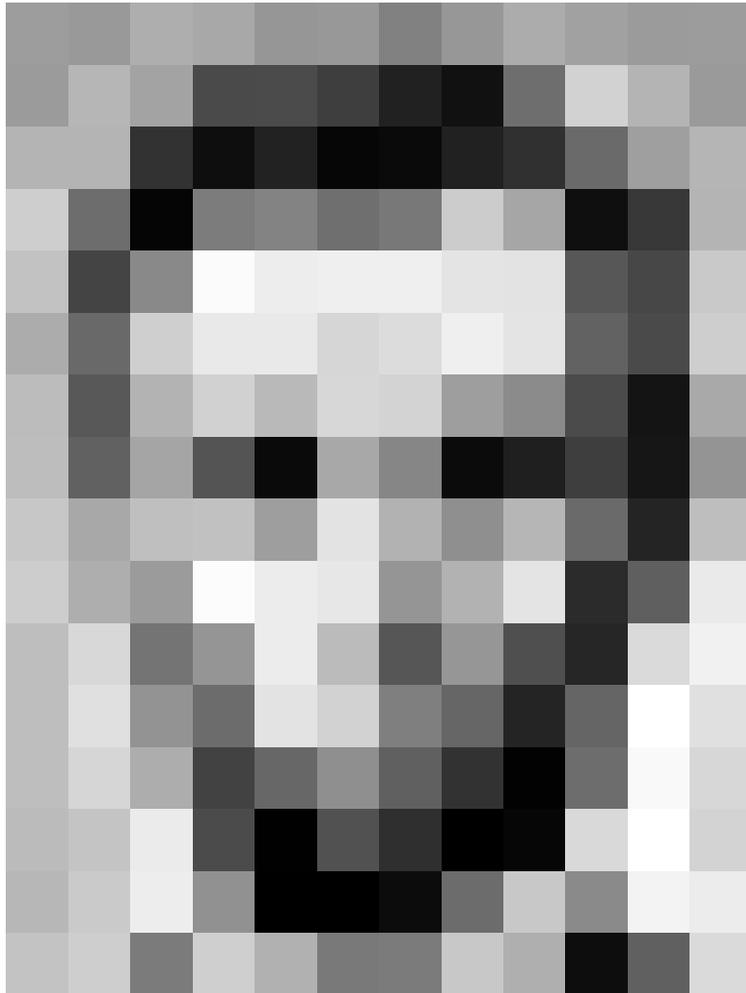
Computer vision



- Pinhole principle
- Traditional film
- Digital sensors (CCD and CMOS)
- Red Green Blue (RGB) channels



Grey Scale

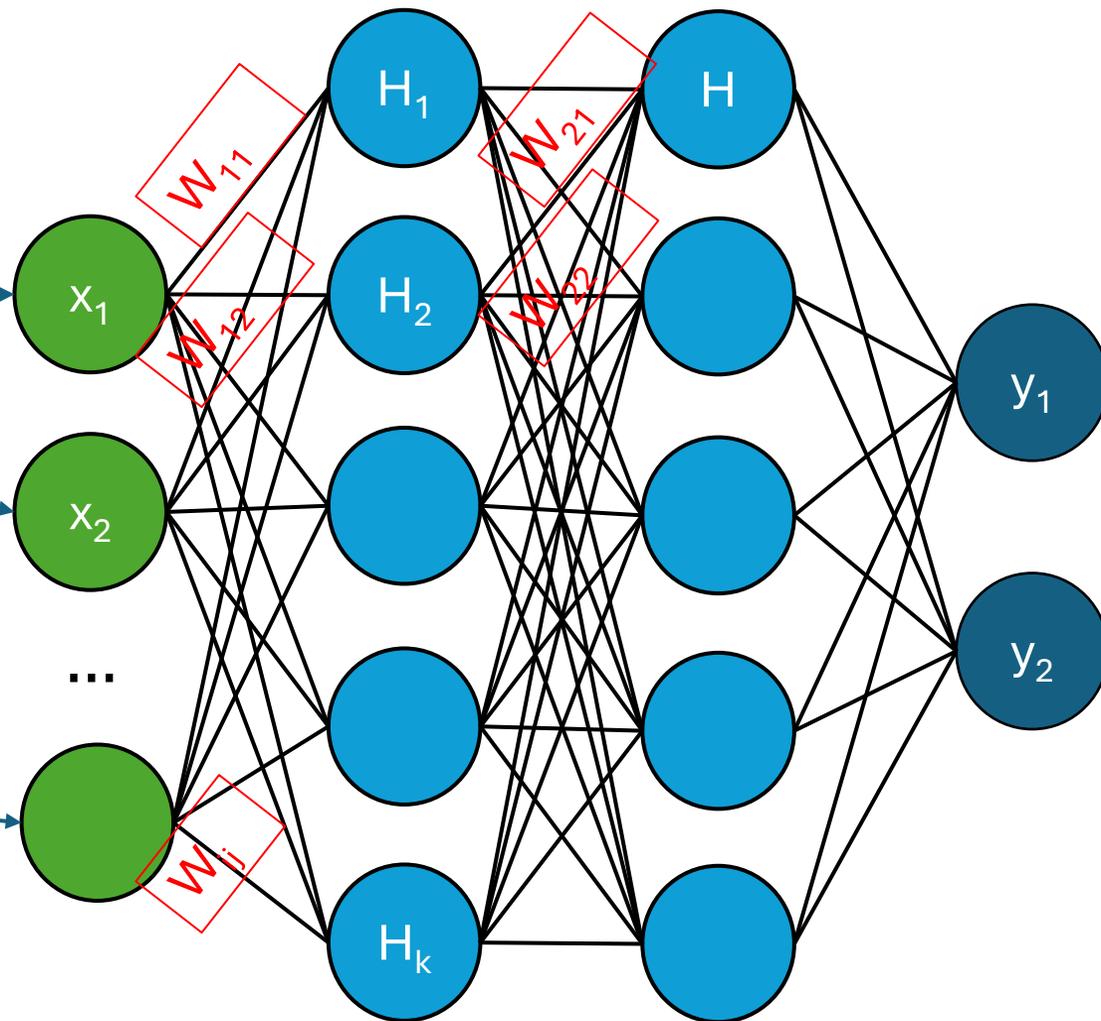


157	153	174	168	150	152	129	151	172	161	155	156
155	182	163	74	75	62	33	17	110	210	180	154
180	180	50	14	34	6	10	33	48	105	159	181
206	109	5	124	131	111	120	204	166	15	56	180
194	68	137	251	237	239	239	228	227	87	71	201
172	105	207	233	233	214	220	239	228	98	74	206
188	88	179	209	185	215	211	158	139	75	20	169
189	97	165	84	10	168	134	11	31	62	22	148
199	168	191	193	158	227	178	143	182	105	36	190
205	174	155	252	236	231	149	178	228	43	95	234
190	216	116	149	236	187	85	150	79	38	218	241
190	224	147	108	227	210	127	102	36	101	255	224
190	214	173	66	103	143	96	50	2	109	249	215
187	196	235	75	1	81	47	0	6	217	255	211
183	202	237	145	0	0	12	108	200	138	243	236
195	206	123	207	177	121	123	200	175	13	96	218

157	153	174	168	150	152	129	151	172	161	155	156
155	182	163	74	75	62	33	17	110	210	180	154
180	180	50	14	34	6	10	33	48	105	159	181
206	109	5	124	131	111	120	204	166	15	56	180
194	68	137	251	237	239	239	228	227	87	71	201
172	105	207	233	233	214	220	239	228	98	74	206
188	88	179	209	185	215	211	158	139	75	20	169
189	97	165	84	10	168	134	11	31	62	22	148
199	168	191	193	158	227	178	143	182	105	36	190
205	174	155	252	236	231	149	178	228	43	95	234
190	216	116	149	236	187	85	150	79	38	218	241
190	224	147	108	227	210	127	102	36	101	255	224
190	214	173	66	103	143	96	50	2	109	249	215
187	196	235	75	1	81	47	0	6	217	255	211
183	202	237	145	0	0	12	108	200	138	243	236
195	206	123	207	177	121	123	200	175	13	96	218

Image of Abraham Lincoln as a matrix of pixel values. Image Credit: Thomas Smits

x_1	x_2	x_3	x_4	x_5
x_6	x_7	x_8	x_9	x_{10}
x_{11}	x_{12}	x_{13}	x_{14}	x_{15}
x_{16}	x_{17}	x_{18}	x_{19}	x_{20}
x_{21}	x_{22}	x_{23}	x_{24}	x_{25}



Files

Analyze your files with code written by Gemini Upload

- ..
- .config
- RANDOM_PICTURES_TO_VERIF...
- mixture
 - S00
 - S05
 - S10
 - S20
 - S30
 - S40
- sample_data
 - Confusing Matrix.pdf
 - RANDOM_PICTURES_TO_VERIF...
 - best_model.pth
 - mixture.zip
 - precision_recall_table.csv

Load Images

torchvision.datasets.ImageFolder is a PyTorch class that simplifies loading image datasets by organizing them into a directory structure where each subdirectory represents a class or category of images. In practice, it is more complicated but doable with care.

```
[4] 1 images = torchvision.datasets.ImageFolder('mixture/')
```

```
[32] 1 images
```

```
Dataset ImageFolder
  Number of datapoints: 183
  Root location: mixture/
  StandardTransform
  Transform: Compose(
    Resize(size=(128, 128), interpolation=bilinear, max_size=None, antialias=True)
    ToTensor()
    Grayscale(num_output_channels=1)
  )
```

```
1 images[4][1]
```

0

```
1 print(np.array(images[4][0]).shape)
```

```
(450, 500, 3)
```

Image Transformation

PyTorch transformations are crucial for preparing datasets before training machine learning models. They ensure standardization and normalization, which scale data points uniformly to prevent features with large ranges from dominating the training process. Transformations also enable data augmentation through techniques like rotation, flipping, and cropping, helping prevent overfitting and improve model

Files

Analyze your files with code written by Gemini Upload

- ..
- .config
- RANDOM_PICTURES_TO_VERI...
- mixture
 - S00
 - S05
 - S10
 - S20
 - S30
 - S40
- sample_data
 - Confusing Matrix.pdf
 - RANDOM_PICTURES_TO_VERIF...
 - best_model.pth
 - mixture.zip
 - precision_recall_table.csv

Simple Neural Network

```
1 import torch.nn as nn
2
3 class SimpleModel(nn.Module):
4     def __init__(self):
5         super(SimpleModel, self).__init__()
6         self.fc1 = nn.Linear(16384, 1000)
7         self.fc2 = nn.Linear(1000, 100)
8         self.fc3 = nn.Linear(100, 6) # Output layer is the same dimension as the classes
9
10    def forward(self, x):
11        x = x.view(-1, 16384) # Flatten the input
12        x = torch.relu(self.fc1(x)) # Activation function for hidden layer
13        x = torch.relu(self.fc2(x))
14        x = self.fc3(x)
15        return x
16
17 model = SimpleModel().to(device)
18
19 # Define loss function and optimizer
20 criterion = nn.CrossEntropyLoss()
21 optimizer = torch.optim.SGD(model.parameters(), lr=0.0001, momentum=0.9)
22
23
```

Loading...

Training

Files

Analyze your files with code written by Gemini Upload

- ..
- .config
- RANDOM_PICTURES_TO_VERI...
- mixture
- sample_data
 - Confusing Matrix.pdf
 - RANDOM_PICTURES_TO_VERIF...
 - best_model.pth
 - mixture.zip
 - precision_recall_table.csv

Training

```
1 since = time.time()
2
3 best_acc = 0.0
4 num_epochs=50
5
6 for epoch in range(num_epochs):
7     print('-' * 10)
8
9     model.train() # Set model to training mode
10    train_running_loss = 0.0
11    train_running_corrects = 0
12
13    for inputs, labels in train_dataloader:
14        inputs = inputs.to(device)
15        labels = labels.to(device)
16
17        # zero the parameter gradients
18        optimizer.zero_grad()
19
20        # forward
21        outputs = model(inputs)
22        _, preds = torch.max(outputs, 1)
23        loss = criterion(outputs, labels)
24
25        # backward + optimize only if in training phase
26        loss.backward()
27        optimizer.step()
28
```

Navigation icons: up, down, search, share, settings, print, delete, menu

- Table of contents
- Problem
- Objectives
- Data
 - Load Images
 - Image Transformation
 - Data Split for Training, Validation, and Testing
 - Dataloader
 - visualize the input examples
 - Result Visualization
 - Simple Neural Network
 - Training
 - Train a Model with Pretrained Weights**
 - Saving the Best Model and Load
 - Test on unseen data
 - Confusion Matrix Result Analysis
- + Section

Train a Model with Pretrained Weights

We need to modify the original pre-defined model because our input channel is only one. Also the output classes of the original resnet is 1000 whereas we only want 6 classes.

```
1
2 model = torchvision.models.resnet18(pretrained=True)
3
4 num_ftrs = model.fc.in_features
5 # The last layer should have len(class_names) neurons
6 model.fc = nn.Linear(num_ftrs, len(class_names))
7
8 criterion = nn.CrossEntropyLoss()
9
10 optimizer = torch.optim.Adam(model.parameters(), lr=0.0001)
11
12 model.conv1 = nn.Conv2d(1, 64, kernel_size=7, stride=2, padding=3, bias=False)
13
14 model = model.to(device)
15
```

/usr/local/lib/python3.11/dist-packages/torchvision/models/_utils.py:208: UserWarning: The parameter 'pretrained' is deprecated since 0.13 and may be removed in a future version. Please use 'weights_only=True' to silence this warning. If you are not using weights, you can safely ignore this warning. (Use warnings.warn(msg) to suppress this warning.)

/usr/local/lib/python3.11/dist-packages/torchvision/models/_utils.py:223: UserWarning: Arguments other than a weight enum or 'None' for 'weights' are deprecated. Please use 'weights_only=True' to silence this warning. If you are not using weights, you can safely ignore this warning. (Use warnings.warn(msg) to suppress this warning.)

```
[18] 1 Start coding or generate with AI.
[19] 1 since = time.time()
```

Convolution Computation

$$K = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \end{bmatrix}$$



Kernel at position 1



Kernel at position 2



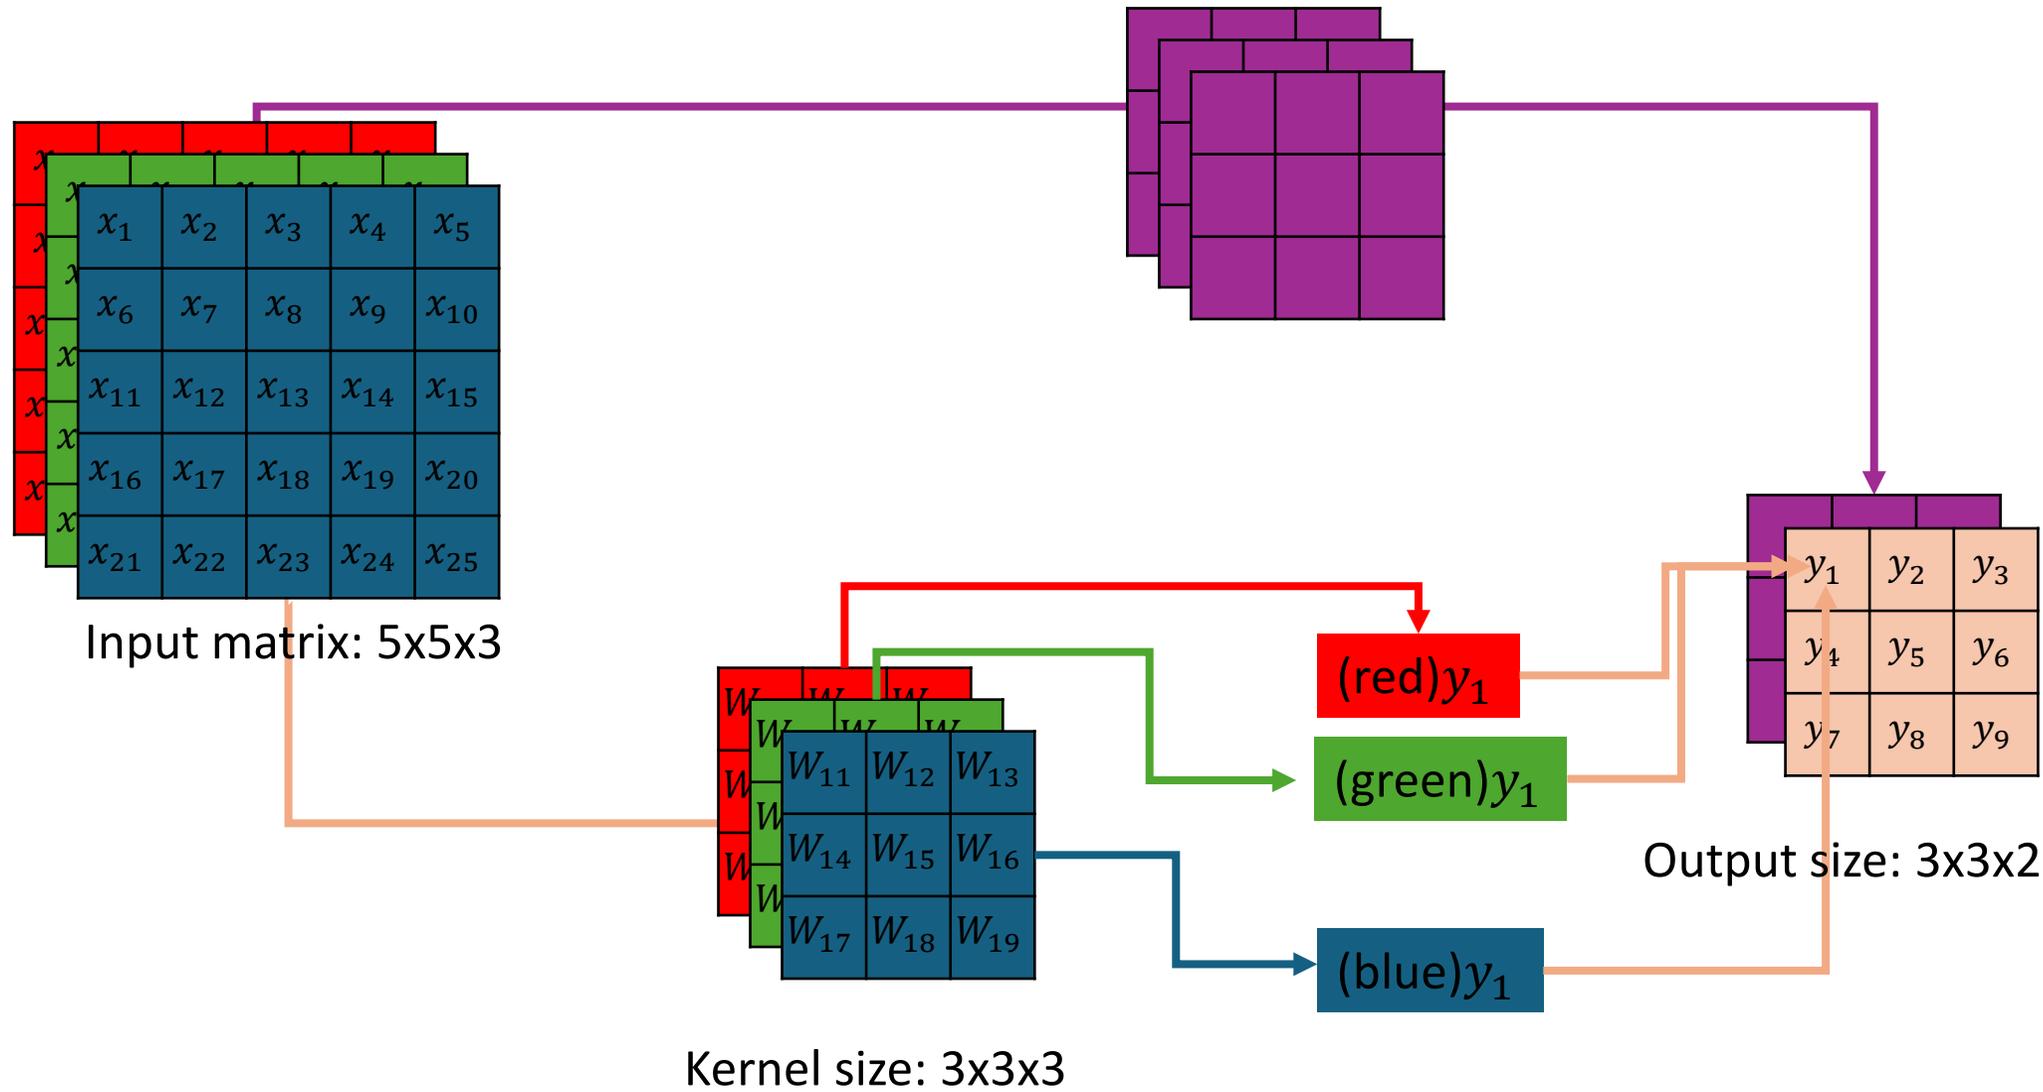
Kernel at position n

7	7	7	7	5
7	7	7	5	5
7	7	5	5	5
7	5	5	5	5
5	5	5	5	5

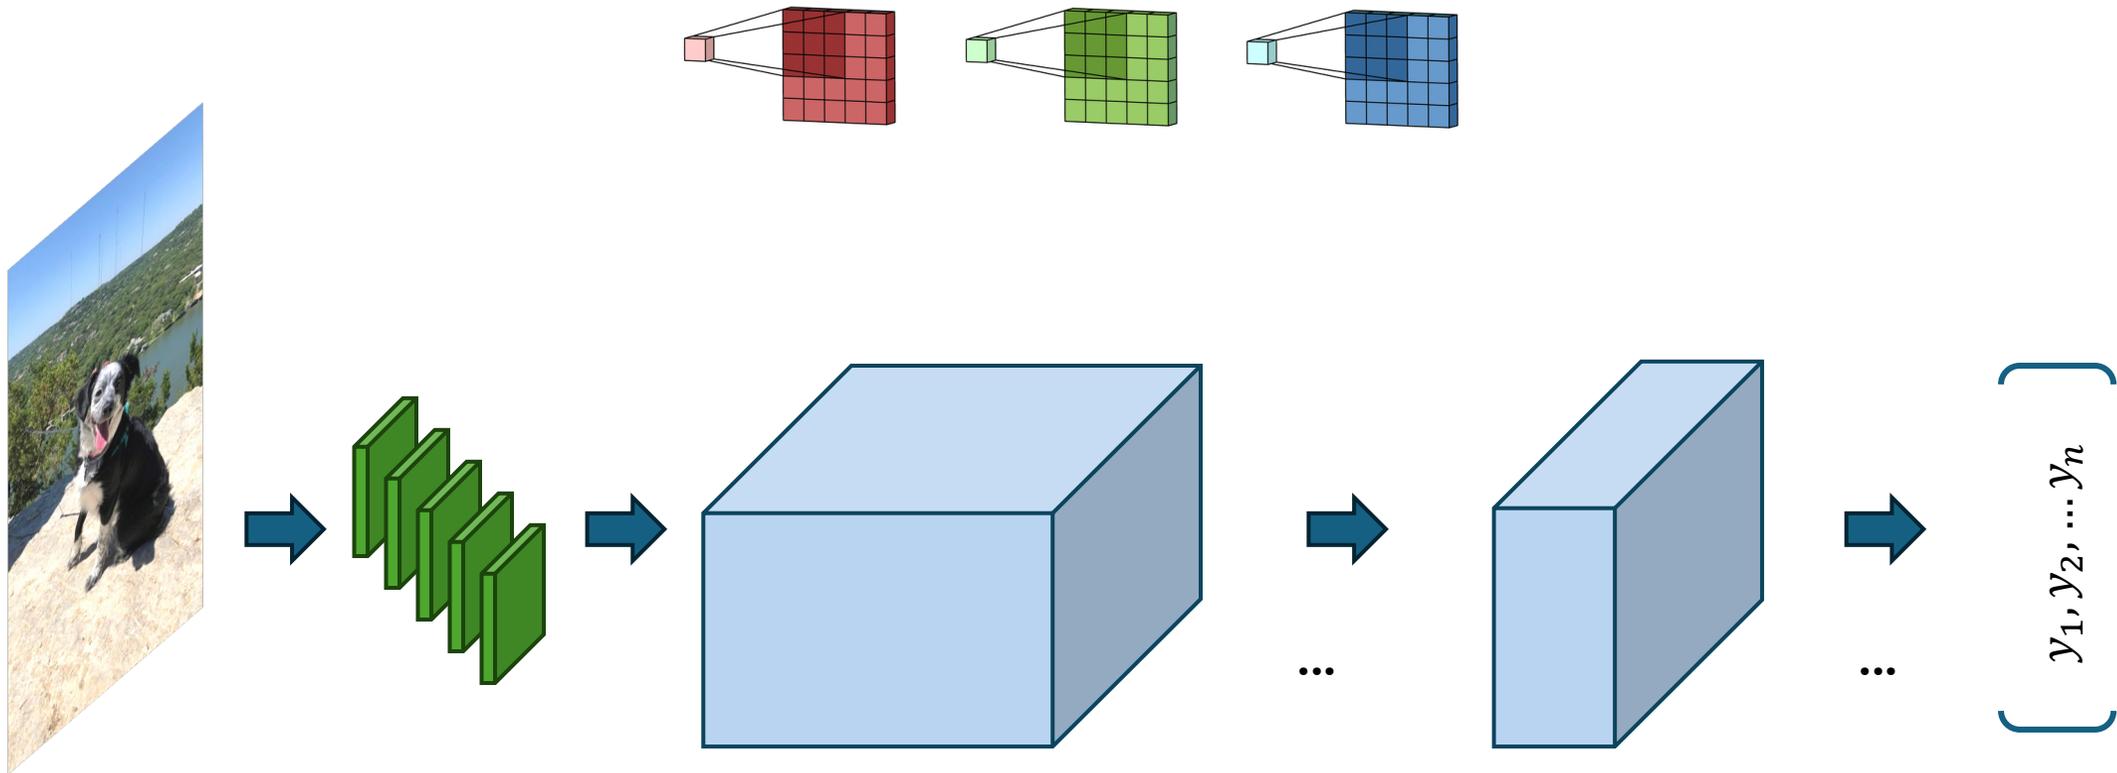
7	7	7	7	5
7	7	7	5	5
7	7	5	5	5
7	5	5	5	5
5	5	5	5	5

7	7	7	7	5
7	7	7	5	5
7	7	5	5	5
7	5	5	5	5
5	5	5	5	5

0	0	0	0	0
0	21	19	17	0
0	19	17	15	0
0	17	15	15	0
0	0	0	0	0



Convolutional Neural Network



RGB Input
(128x128x3)

50 filters
(3x3x3)

Tensor
(128x128x50)

k filters
(3x3x50)

...

Output
(1000 classes)

Large dataset

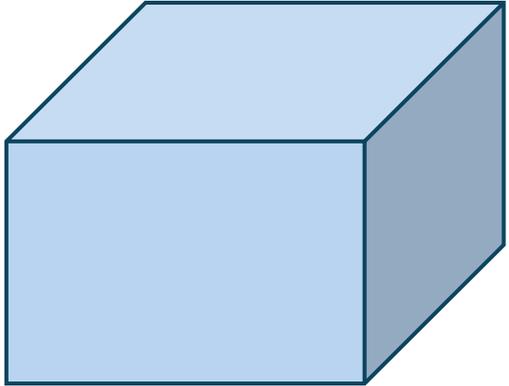
- COCO: <https://cocodataset.org/#home> (80 classes)
- ImageNet (1000 classes with 1.3M labeled images)
- Common Crawl (~100TB+)



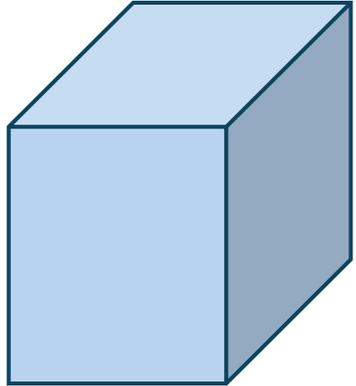
Transfer learning



...



...



Input

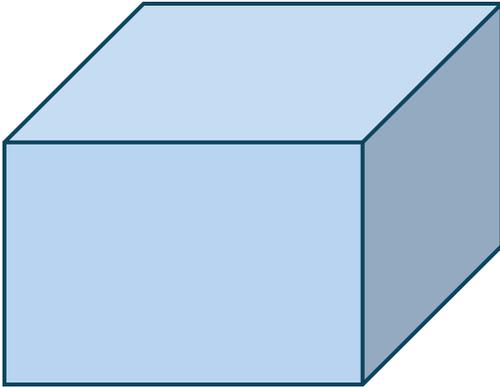
Convolution layer

Output (80)

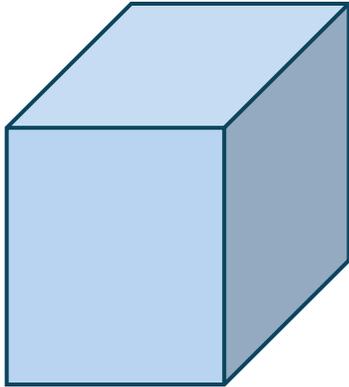
Transfer learning



...



...



...



Input

Convolution layer

Output (6)

Transformer

- GPT: Generative Pre-trained Transformer
- Tokenization: <https://tiktokenizer.vercel.app/>
- Vectorization: <https://projector.tensorflow.org/>



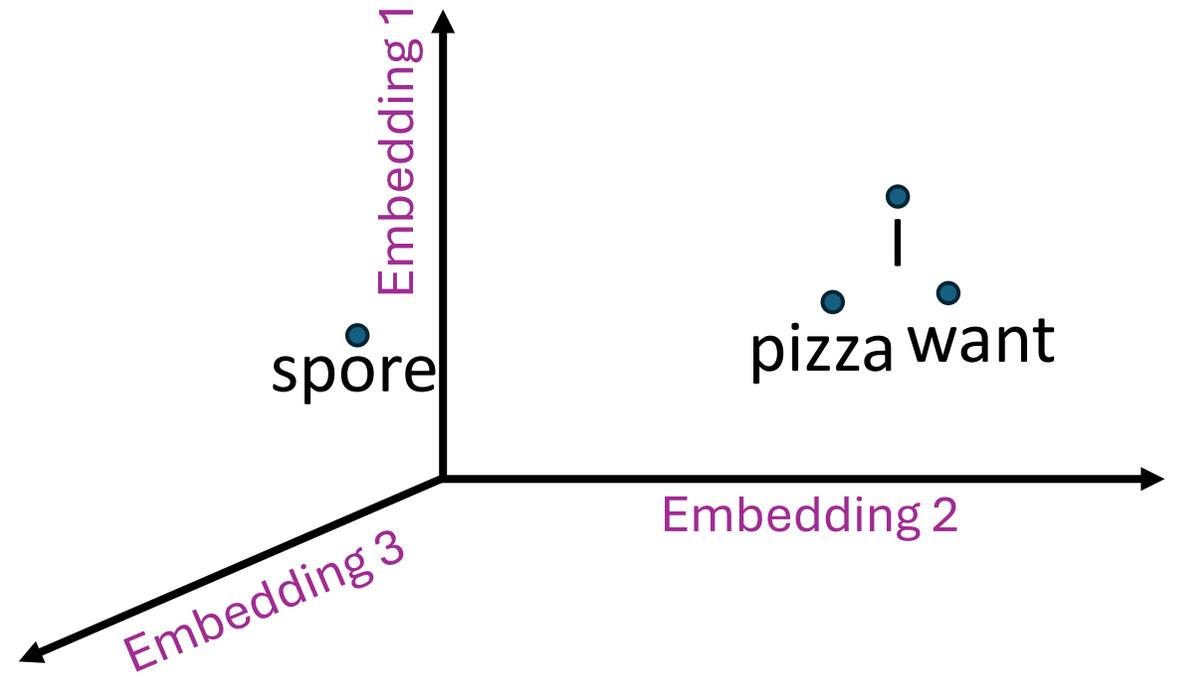
shutterstock.com - 2311573567

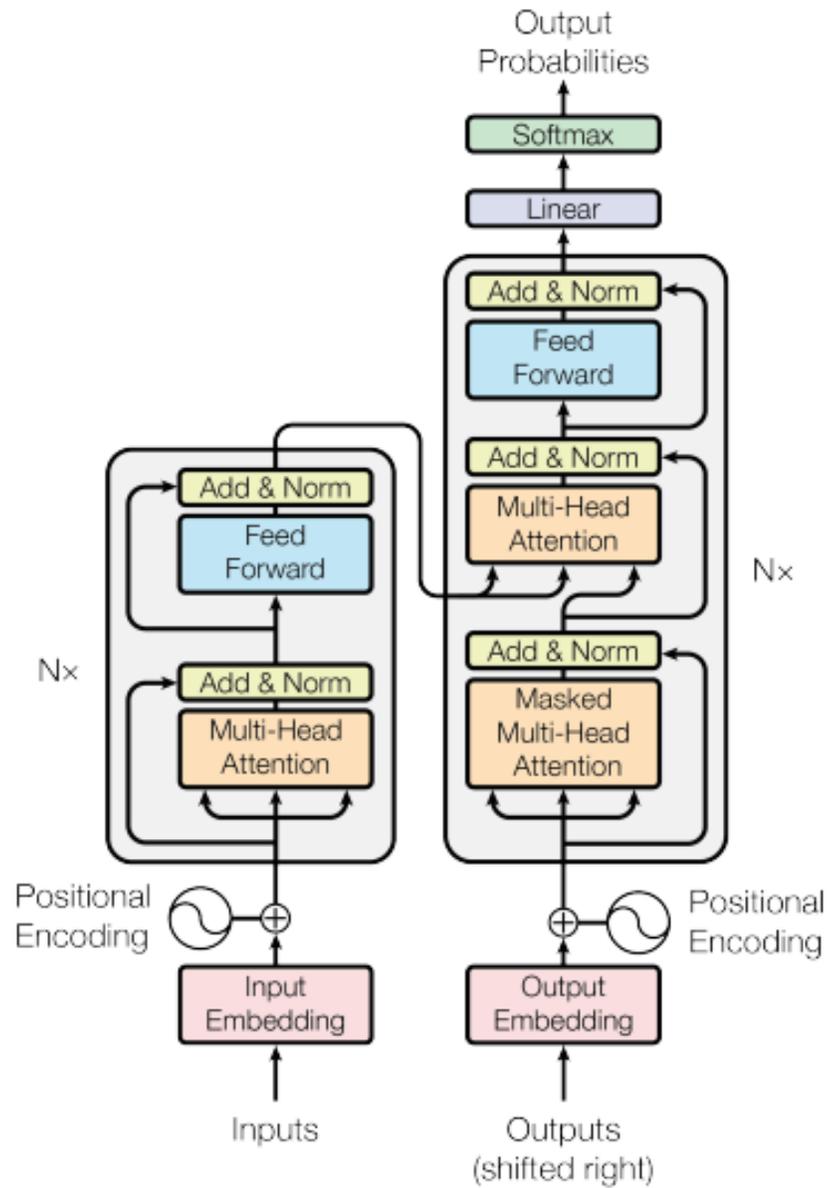


... I want pizza ... spore ...

Total Text Corpus

Vocabulary	Embedding 1	Embedding 2	Embedding 3	...	Embedding n
...
I	e_{i1}	e_{i2}	e_{i3}	...	e_{in}
want	e_{w1}	e_{w2}	e_{w3}	...	e_{wn}
pizza	e_{p1}	e_{p2}	e_{p3}	...	e_{pn}
...
spore	e_{s1}	e_{s2}	e_{s3}	...	e_{sn}
...





Vaswani, A. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*.

GPT 3

- Parameters 175 B
- Dataset 45T
- 96 attention heads
- 2048 token size
- Learn from their chief scientist:
<https://www.youtube.com/watch?v=kCc8FmEb1nY>