

Unstructured Data Analysis

Lecture 12: Intro to neural nets & deep learning

George Chen

What is deep learning?



Classification units



PIT/AIT




V4/PIT




V2/V4



V1/V2




Basic Idea



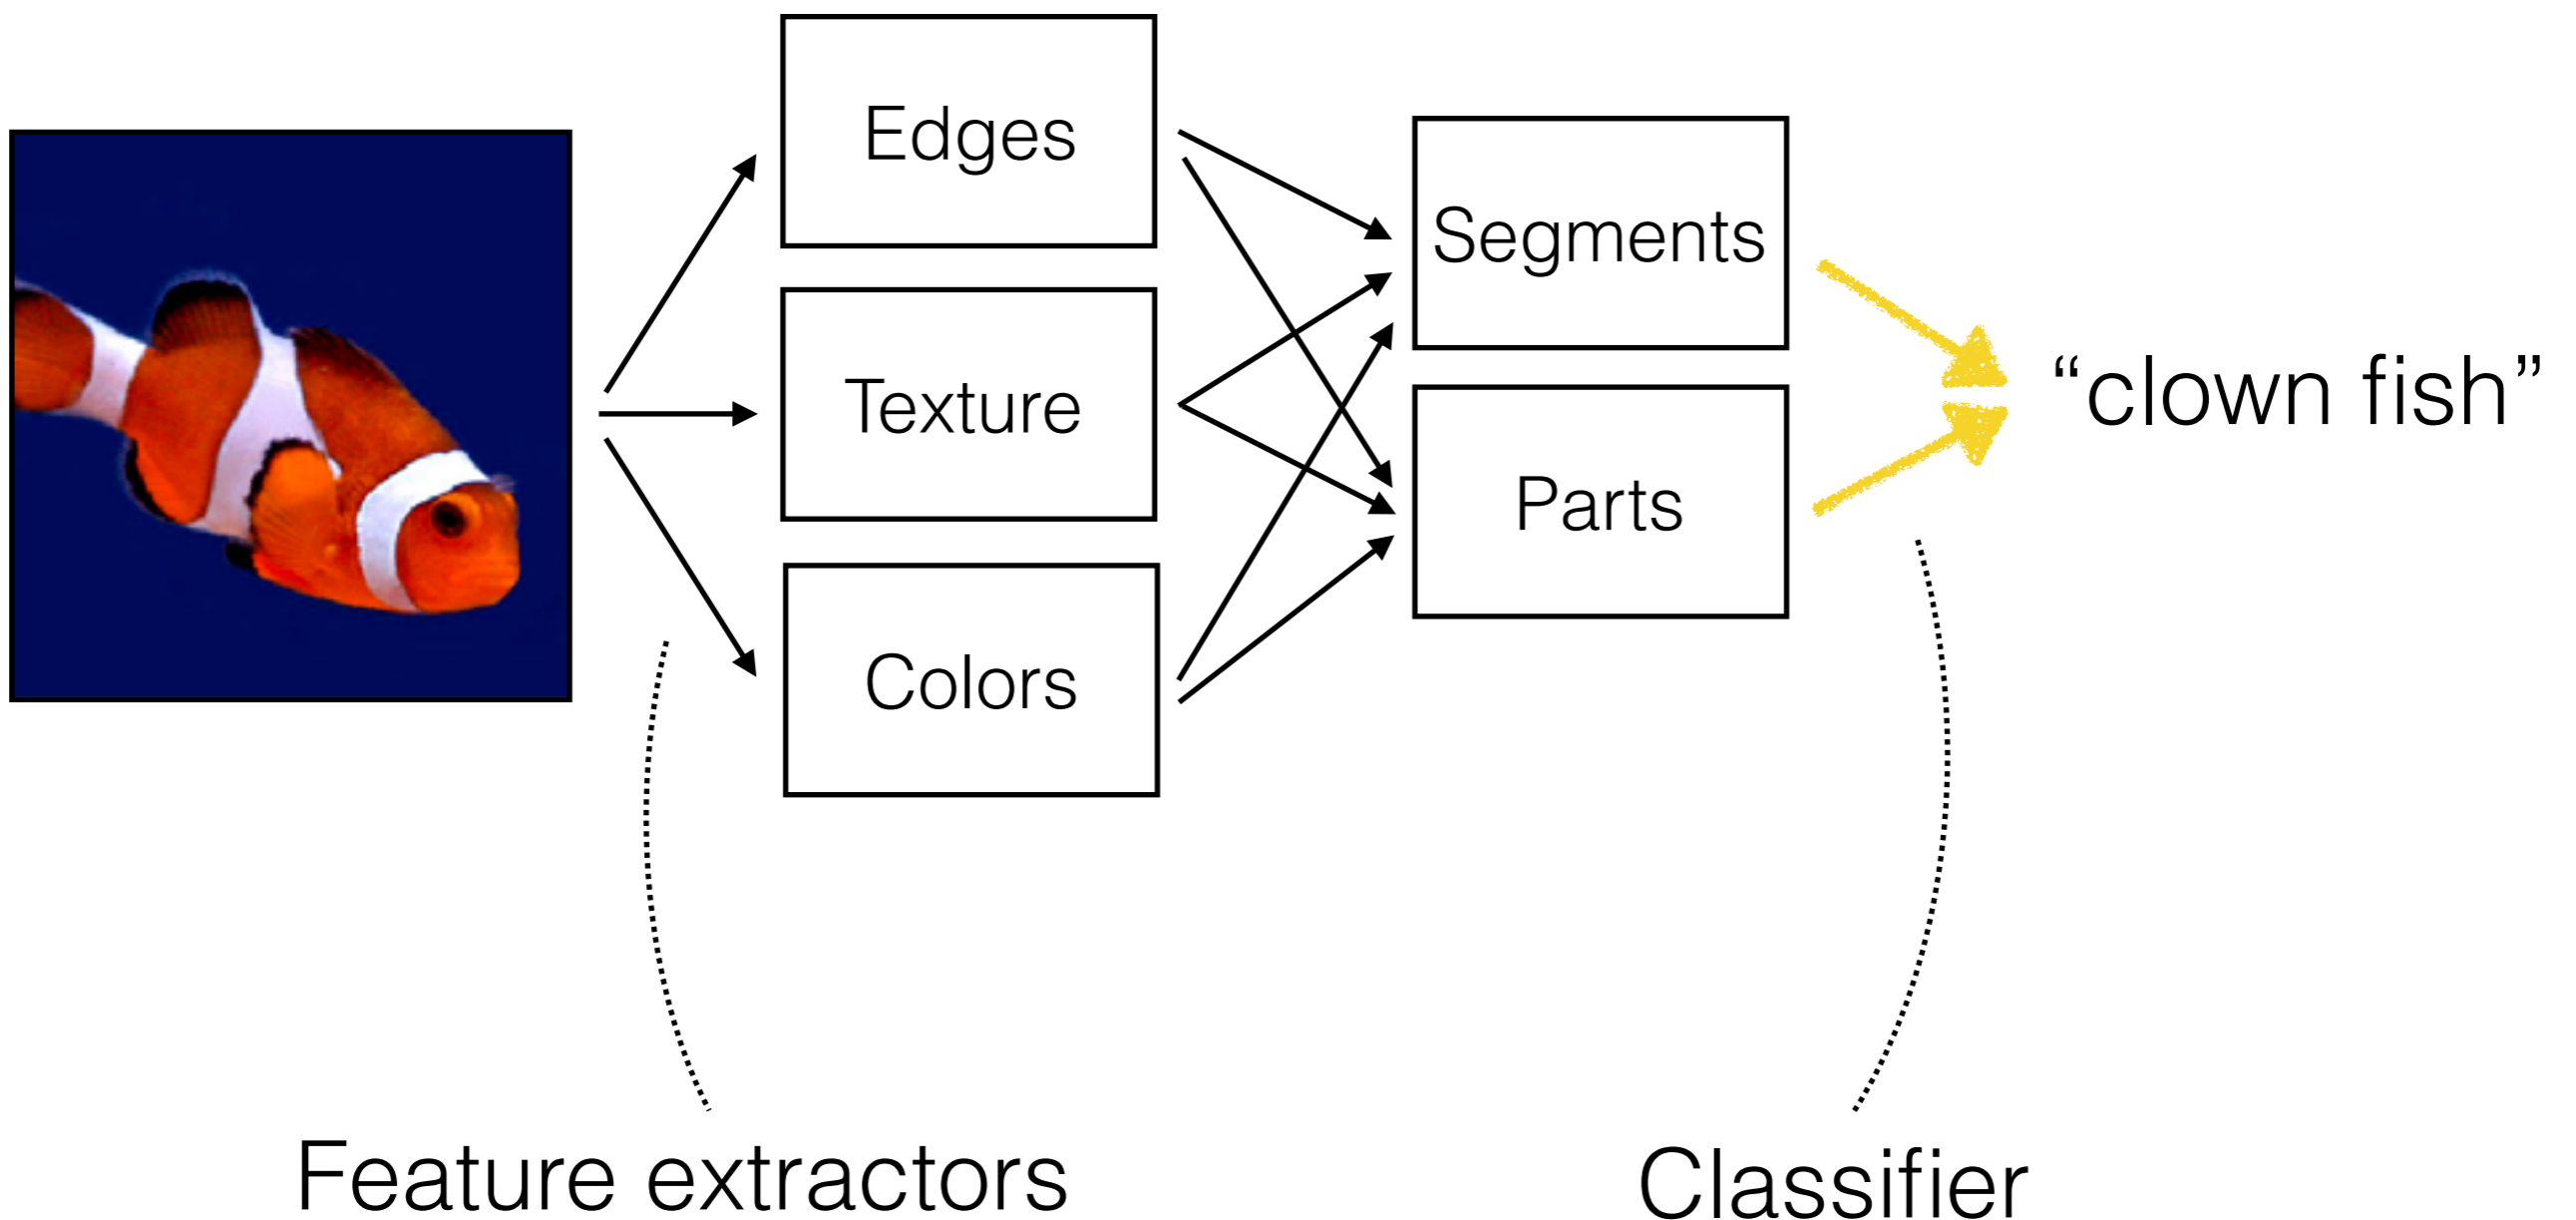
“clown fish”

Object Recognition





Object Recognition

Learned



Neural Network


Learned



"clown fish"



Neural Network

Learned




Deep Neural Network

Learned



Crumpled Paper Analogy




binary classification: 2 crumpled sheets of paper corresponding to the different classes

deep learning: series (“layers”) of simple unfolding operations to try to disentangle the 2 sheets

Analogy: Francois Chollet, photo: George Chen

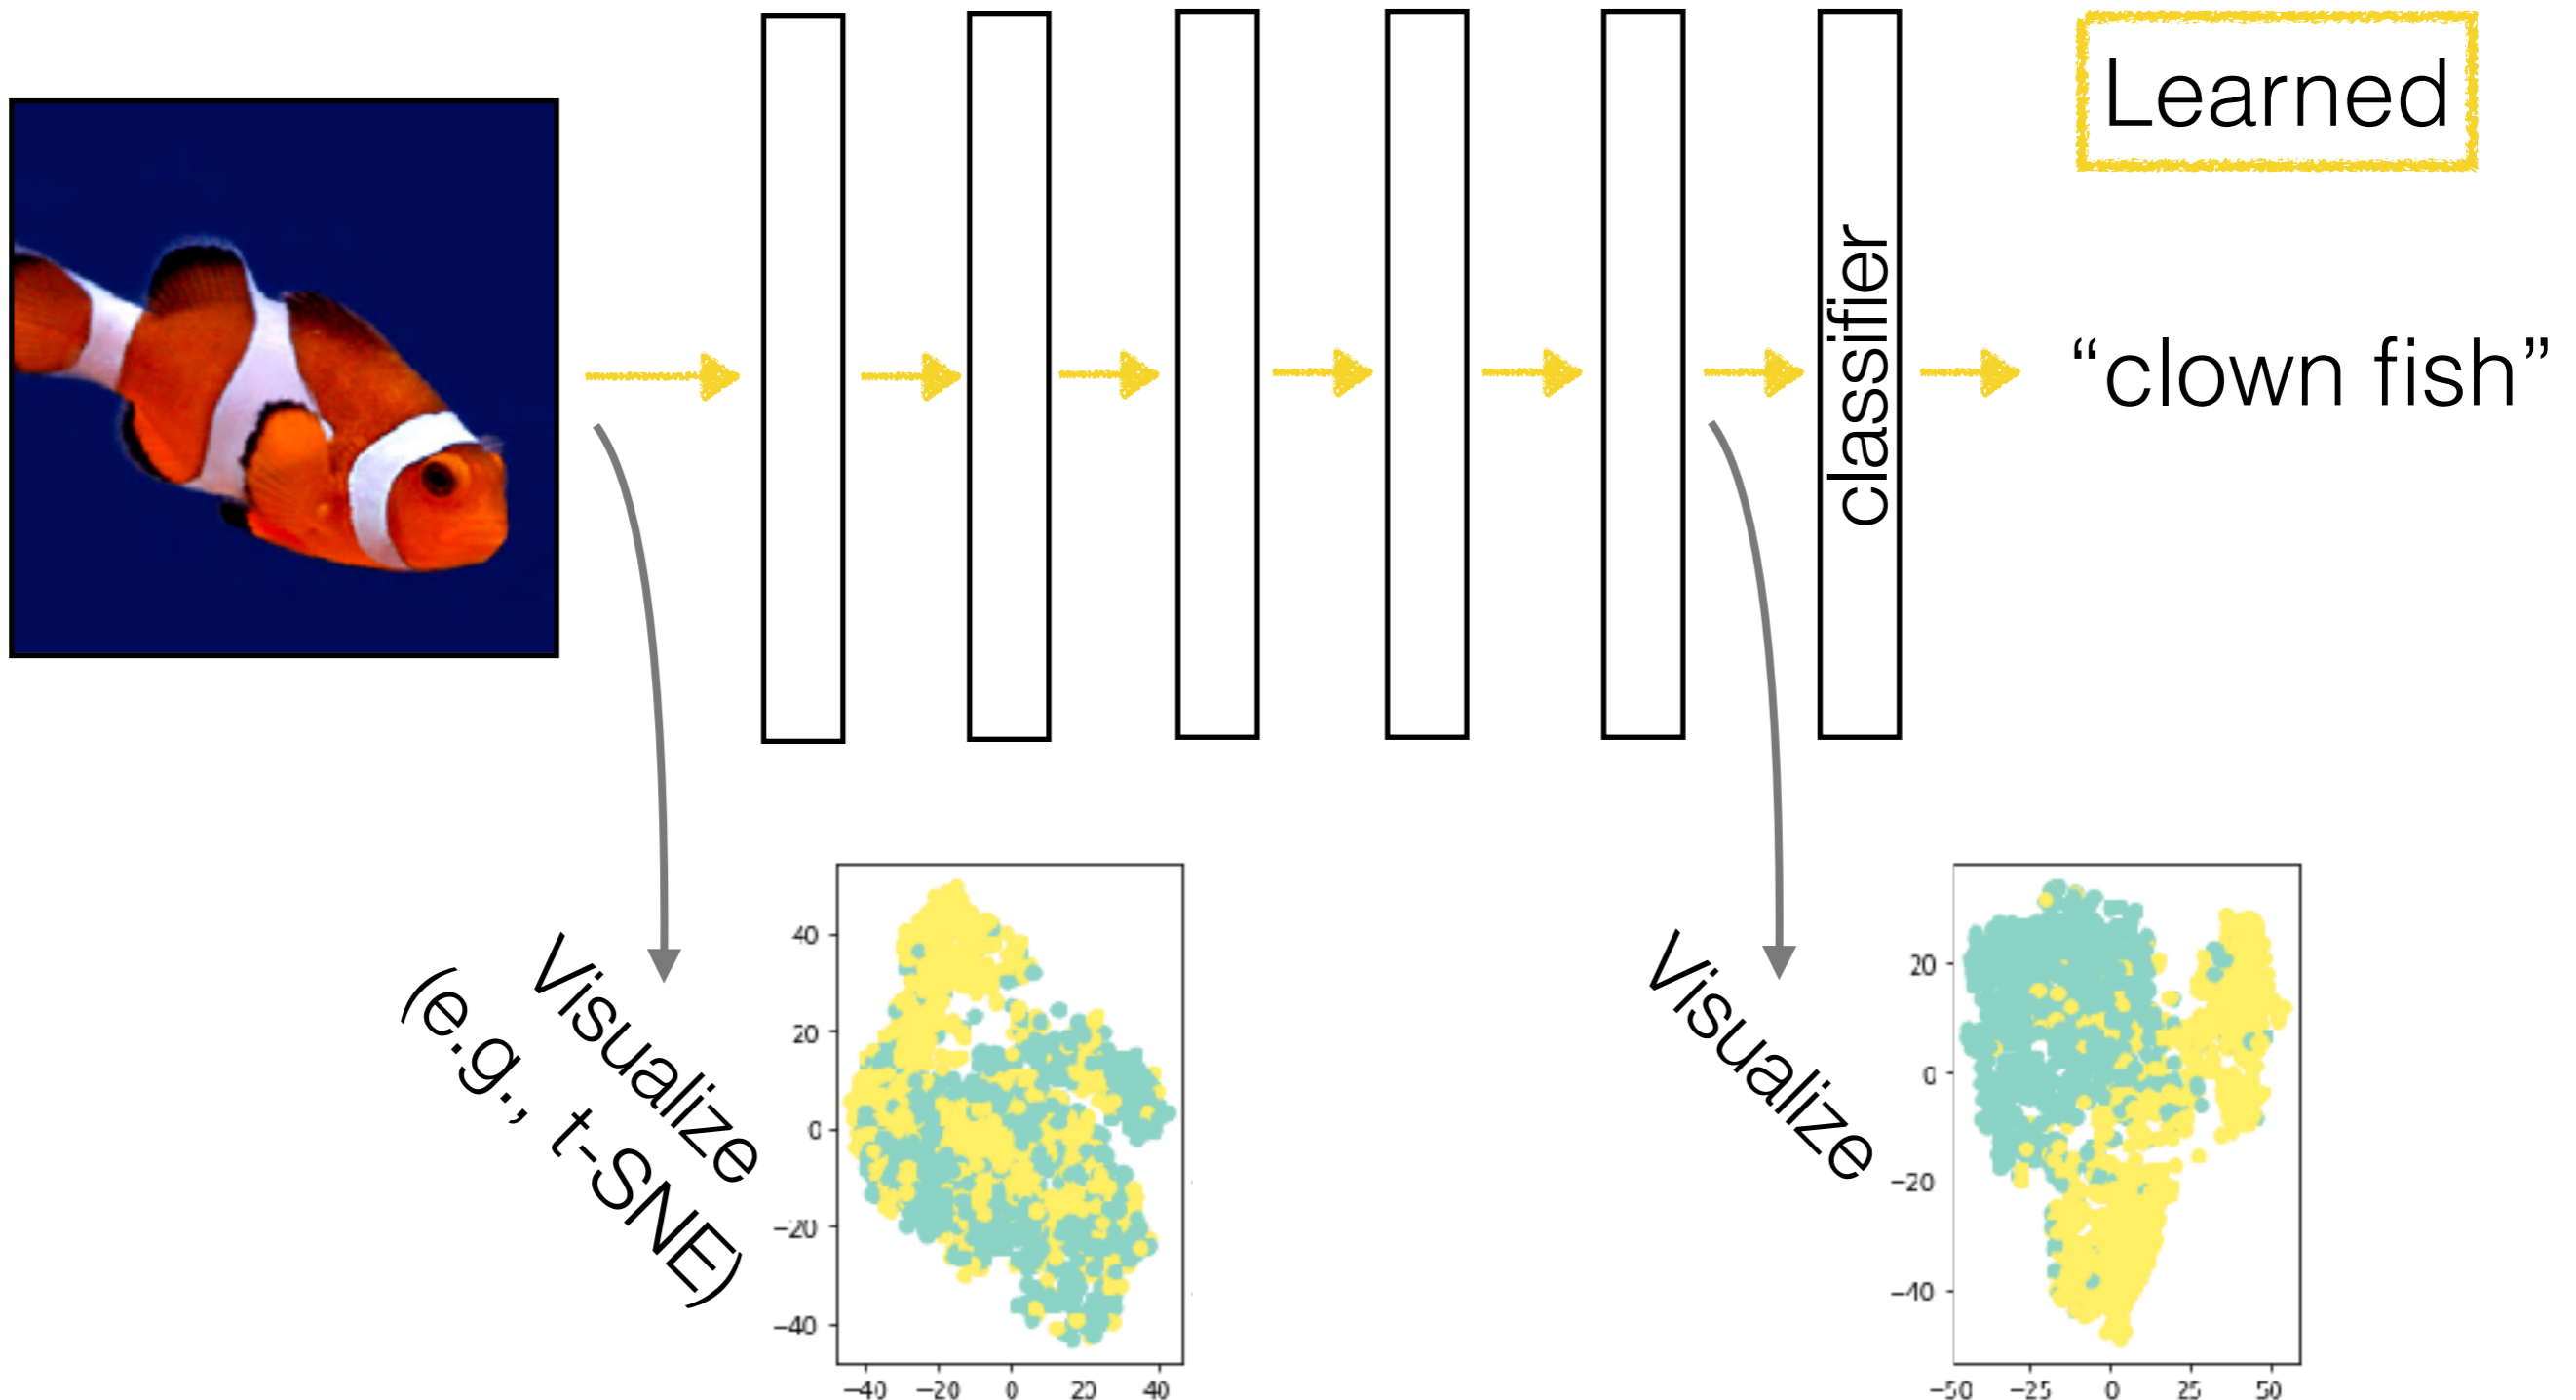
Representation Learning

Each layer's output is *another way we could represent the input data*



Representation Learning

Each layer's output is *another way we could represent the input data*



Why Does Deep Learning Work?

Actually the ideas behind deep learning are old (~1980's)

- Big data




- Better hardware



CPU's
& Moore's law



GPU's



TPU's

- Better algorithms

Structure Present in Data Matters


Neural nets aren't doing black magic

- **Image analysis:** convolutional neural networks (convnets) neatly incorporates basic image processing structure
- **Time series analysis:** recurrent neural networks (RNNs) incorporates ability to remember and forget things over time
 - Note: text is a time series
 - Note: video is a time series

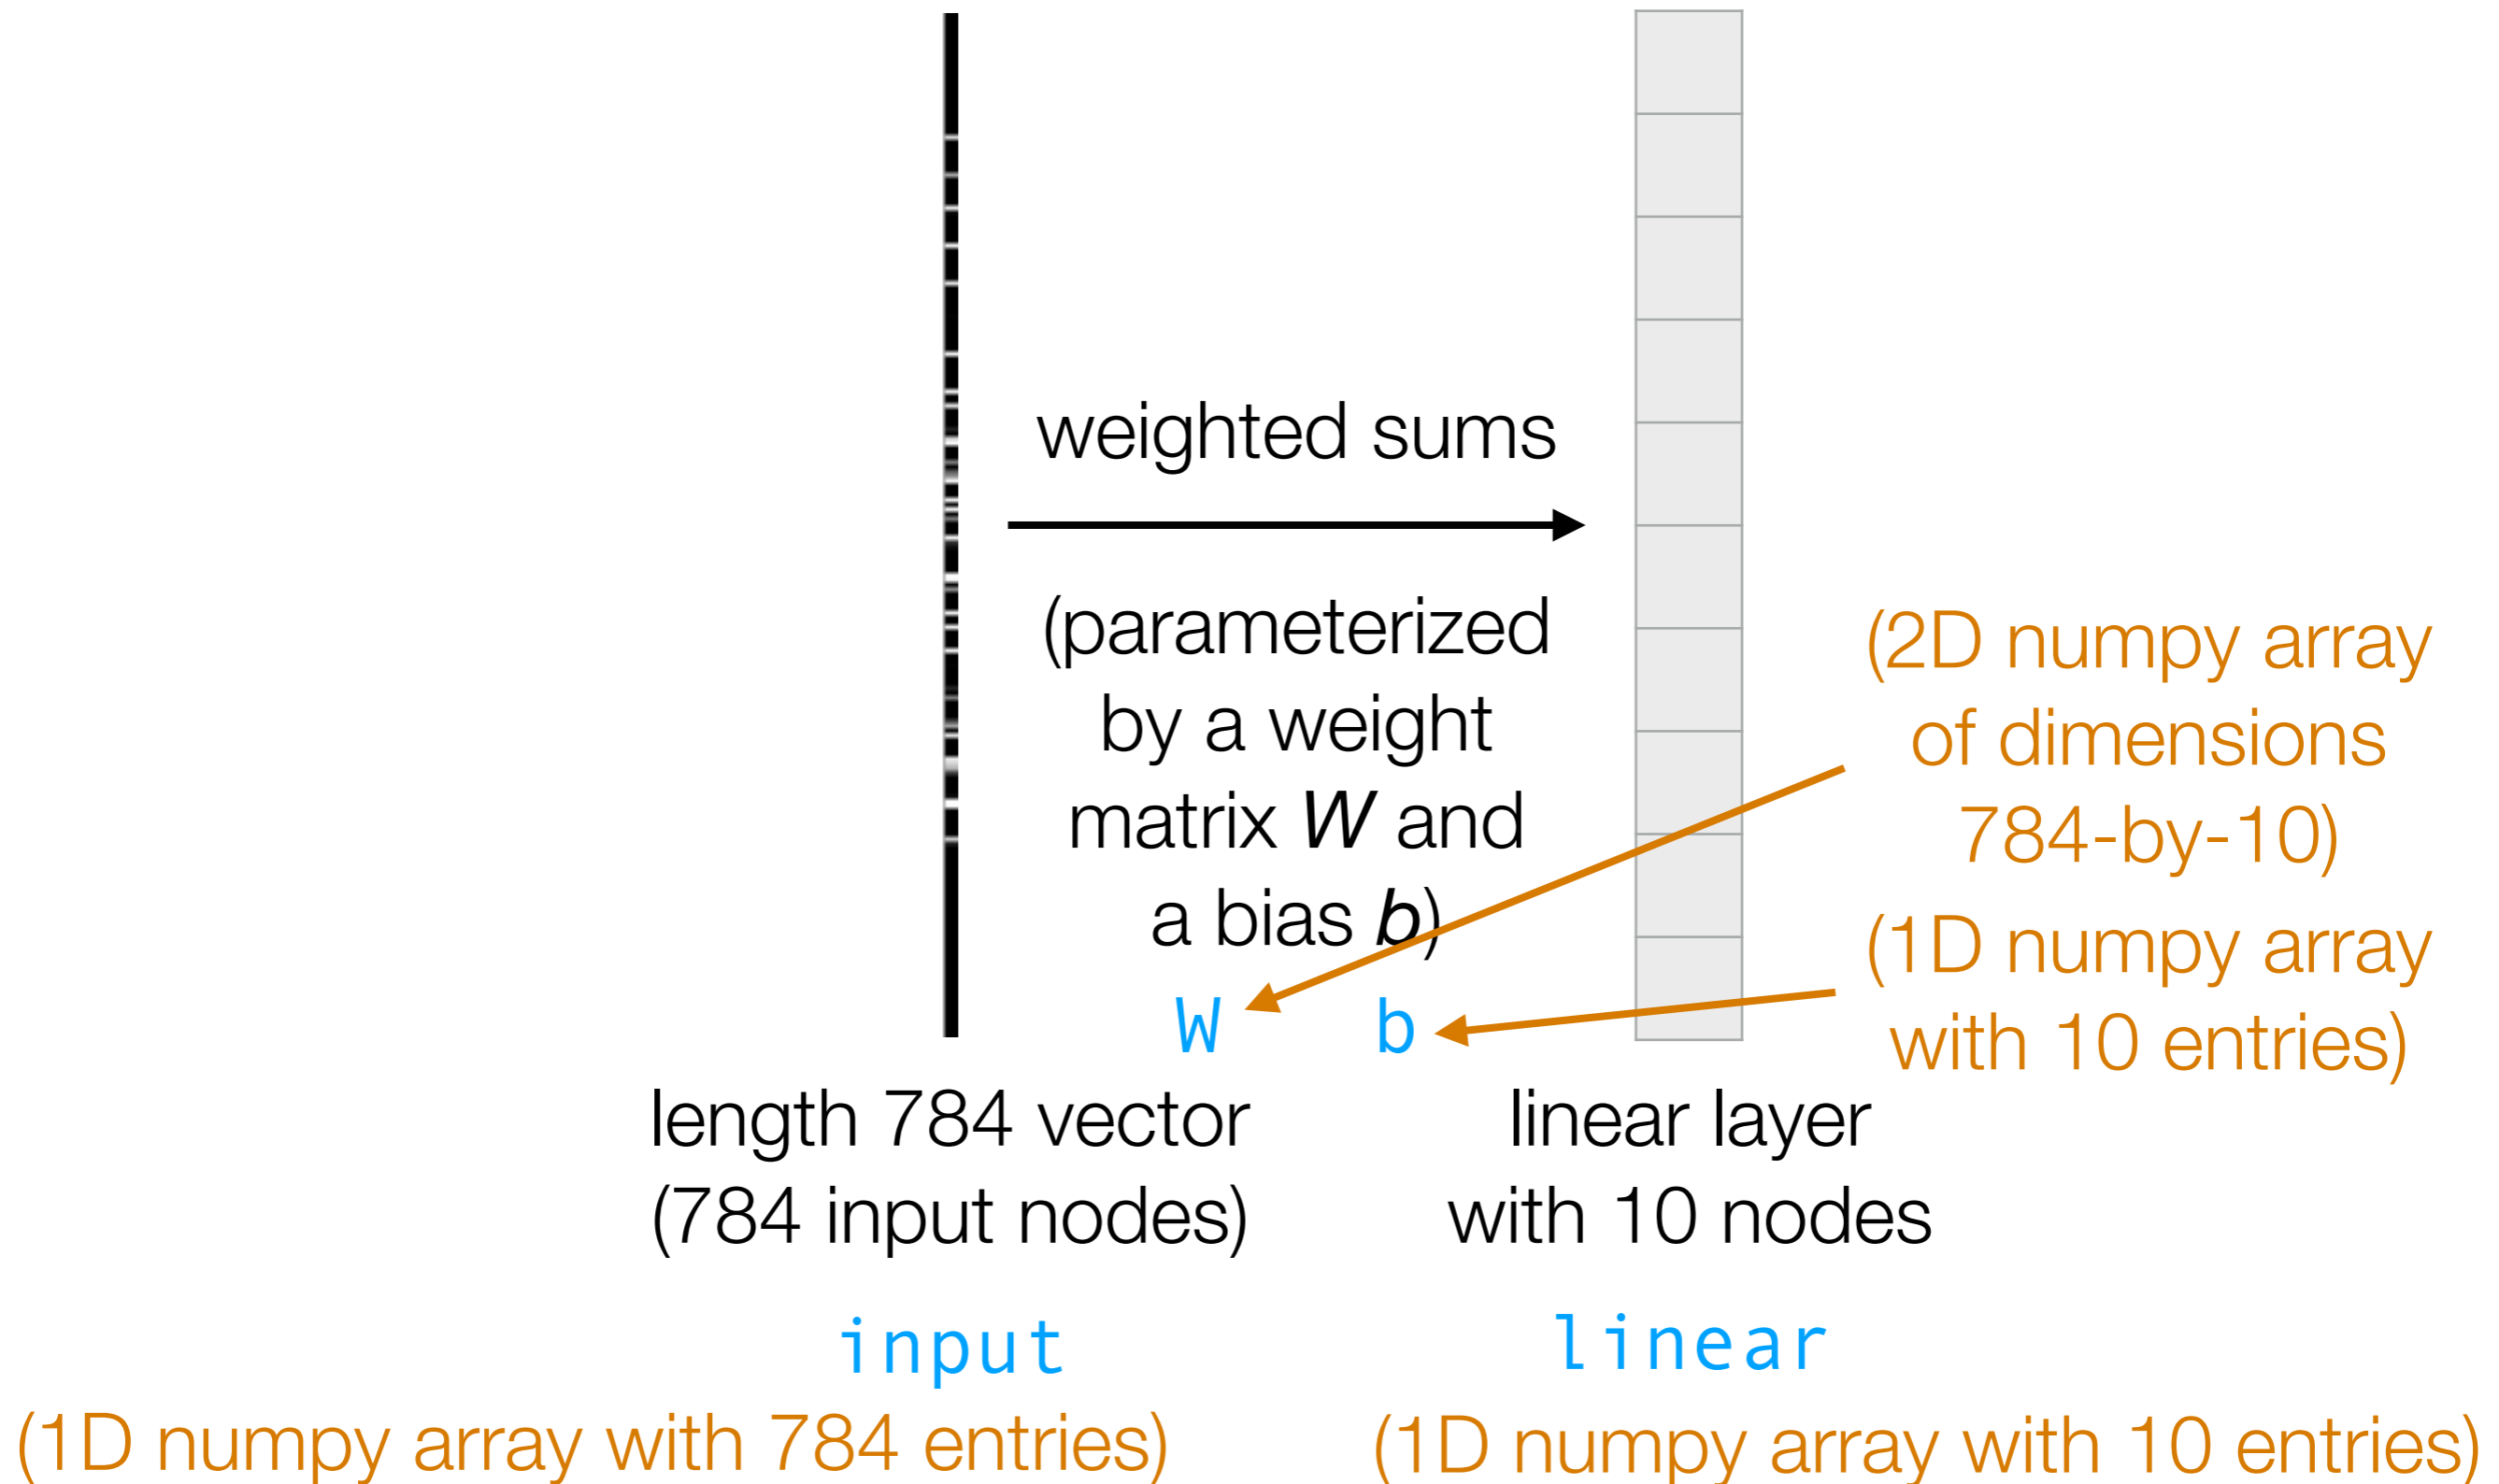
Handwritten Digit Recognition Example

Walkthrough of 2 extremely simple neural nets

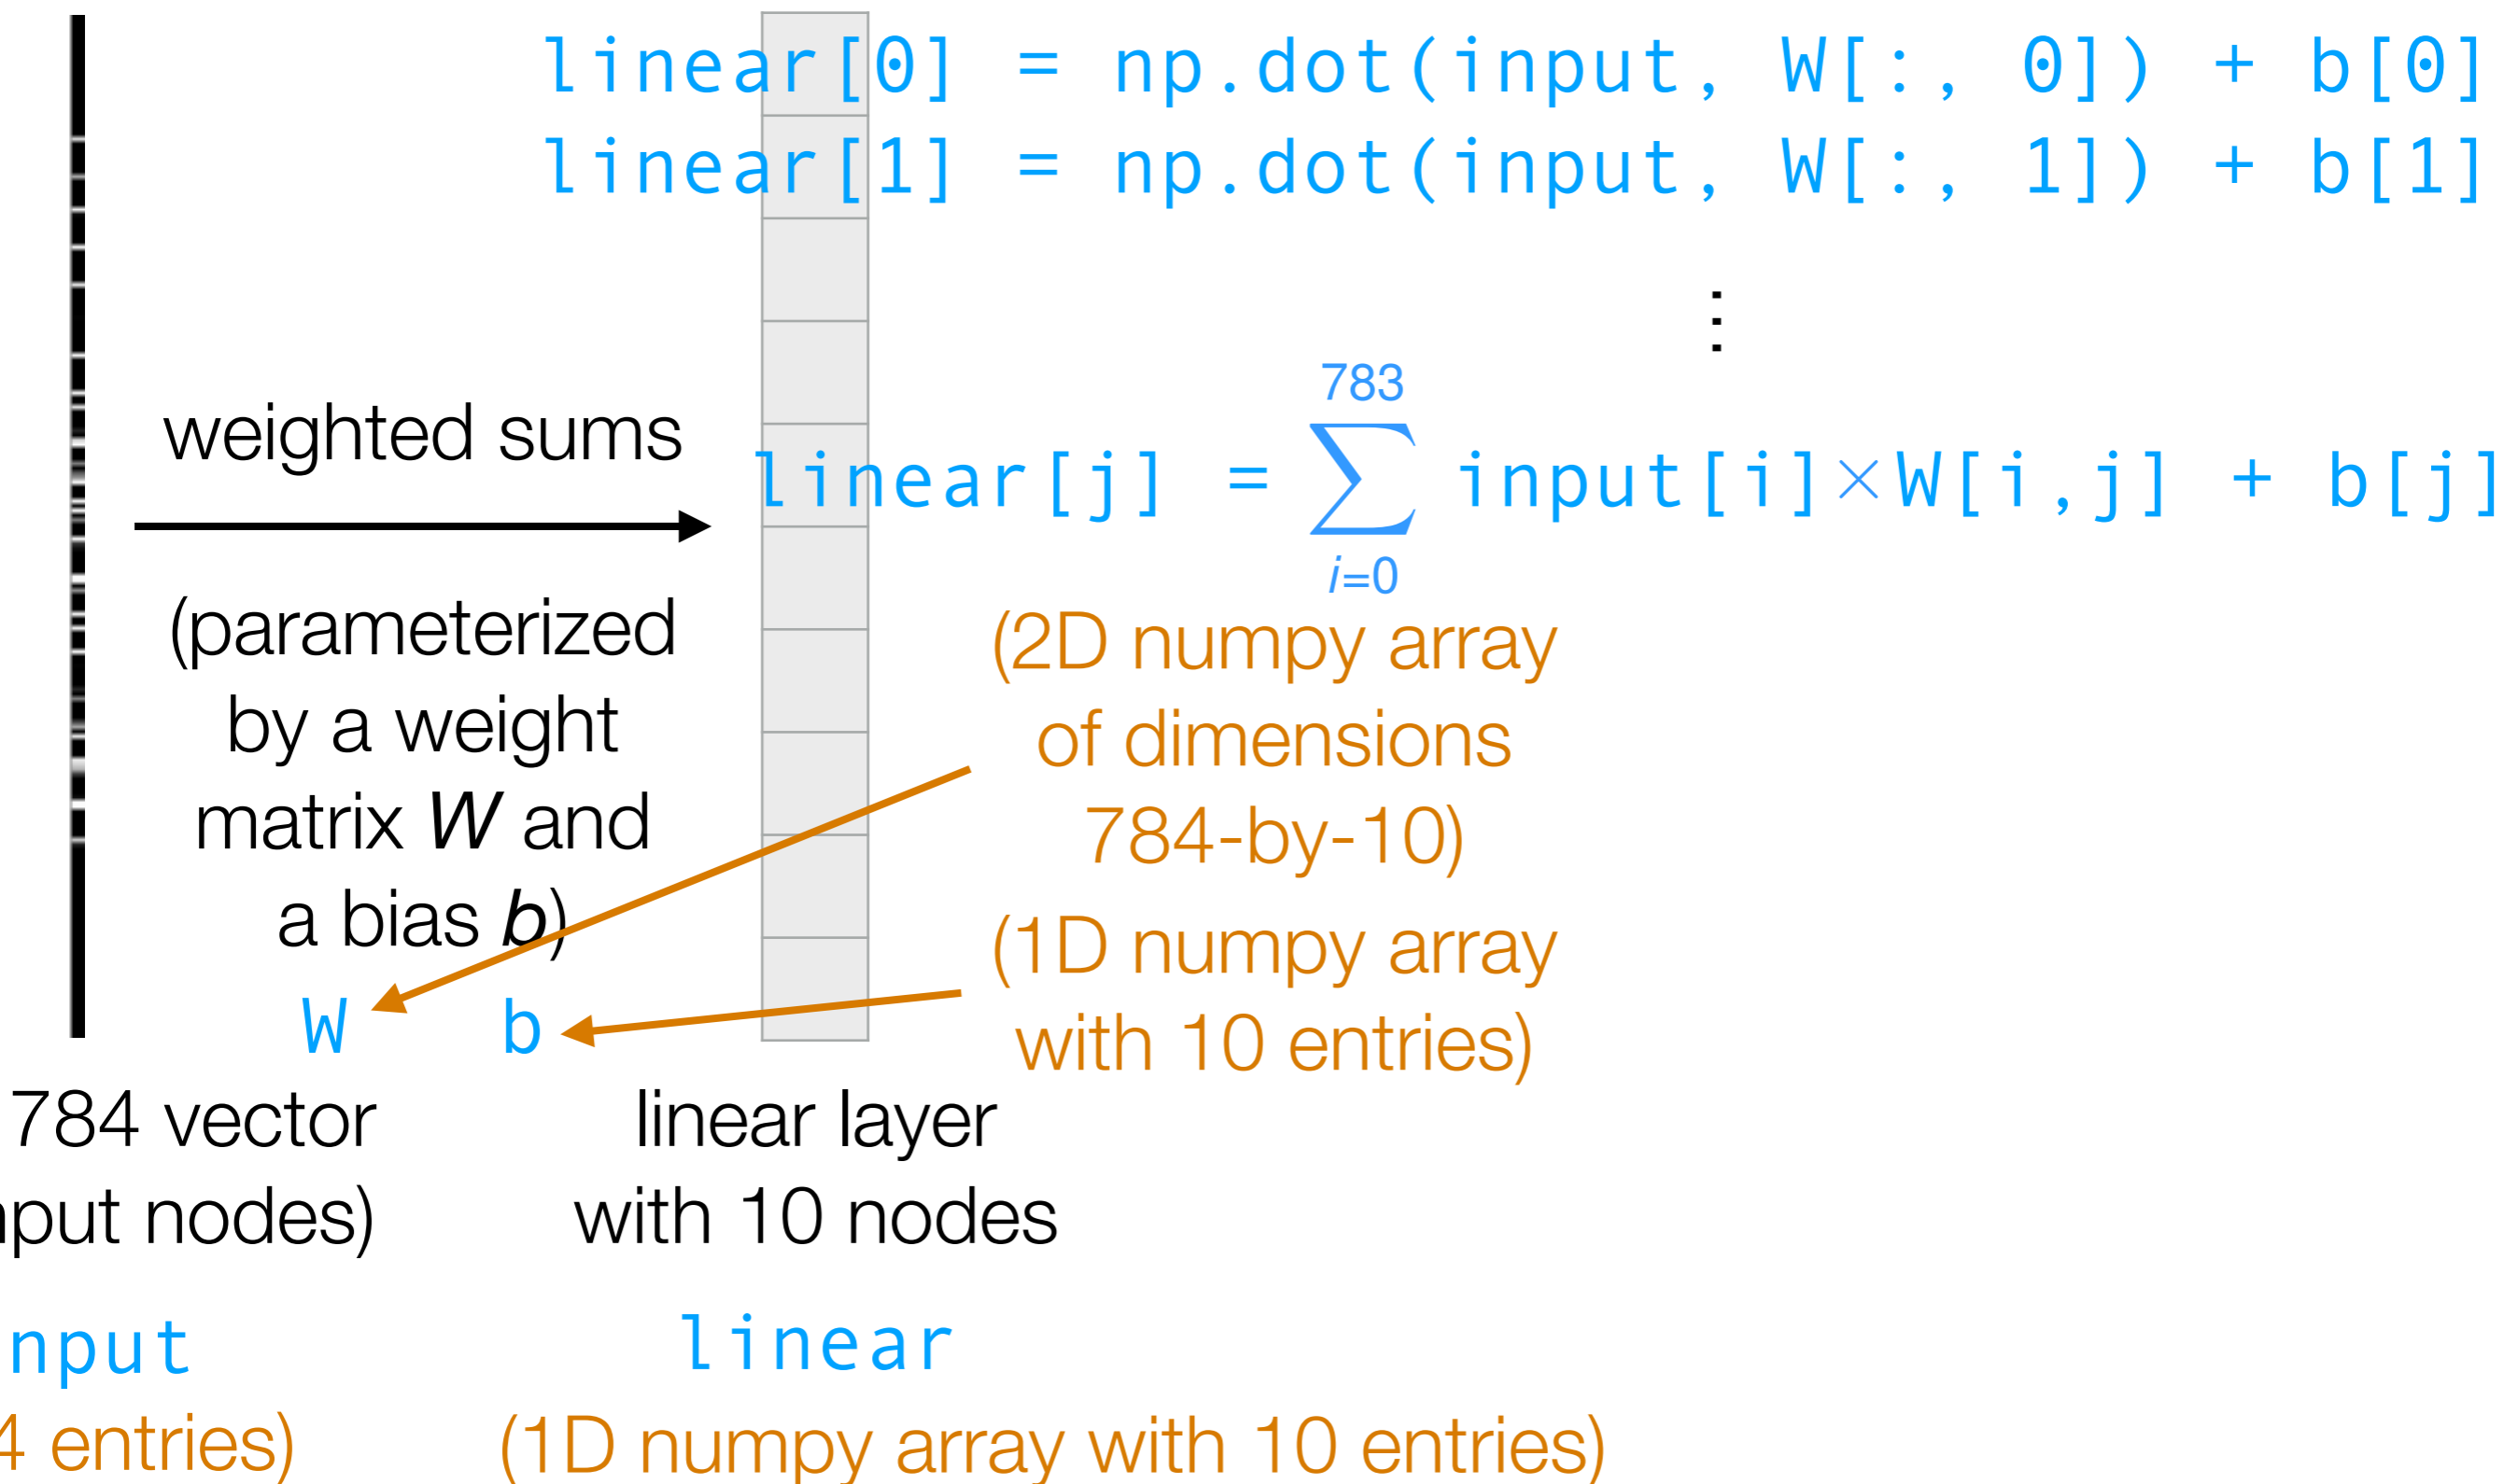
Handwritten Digit Recognition



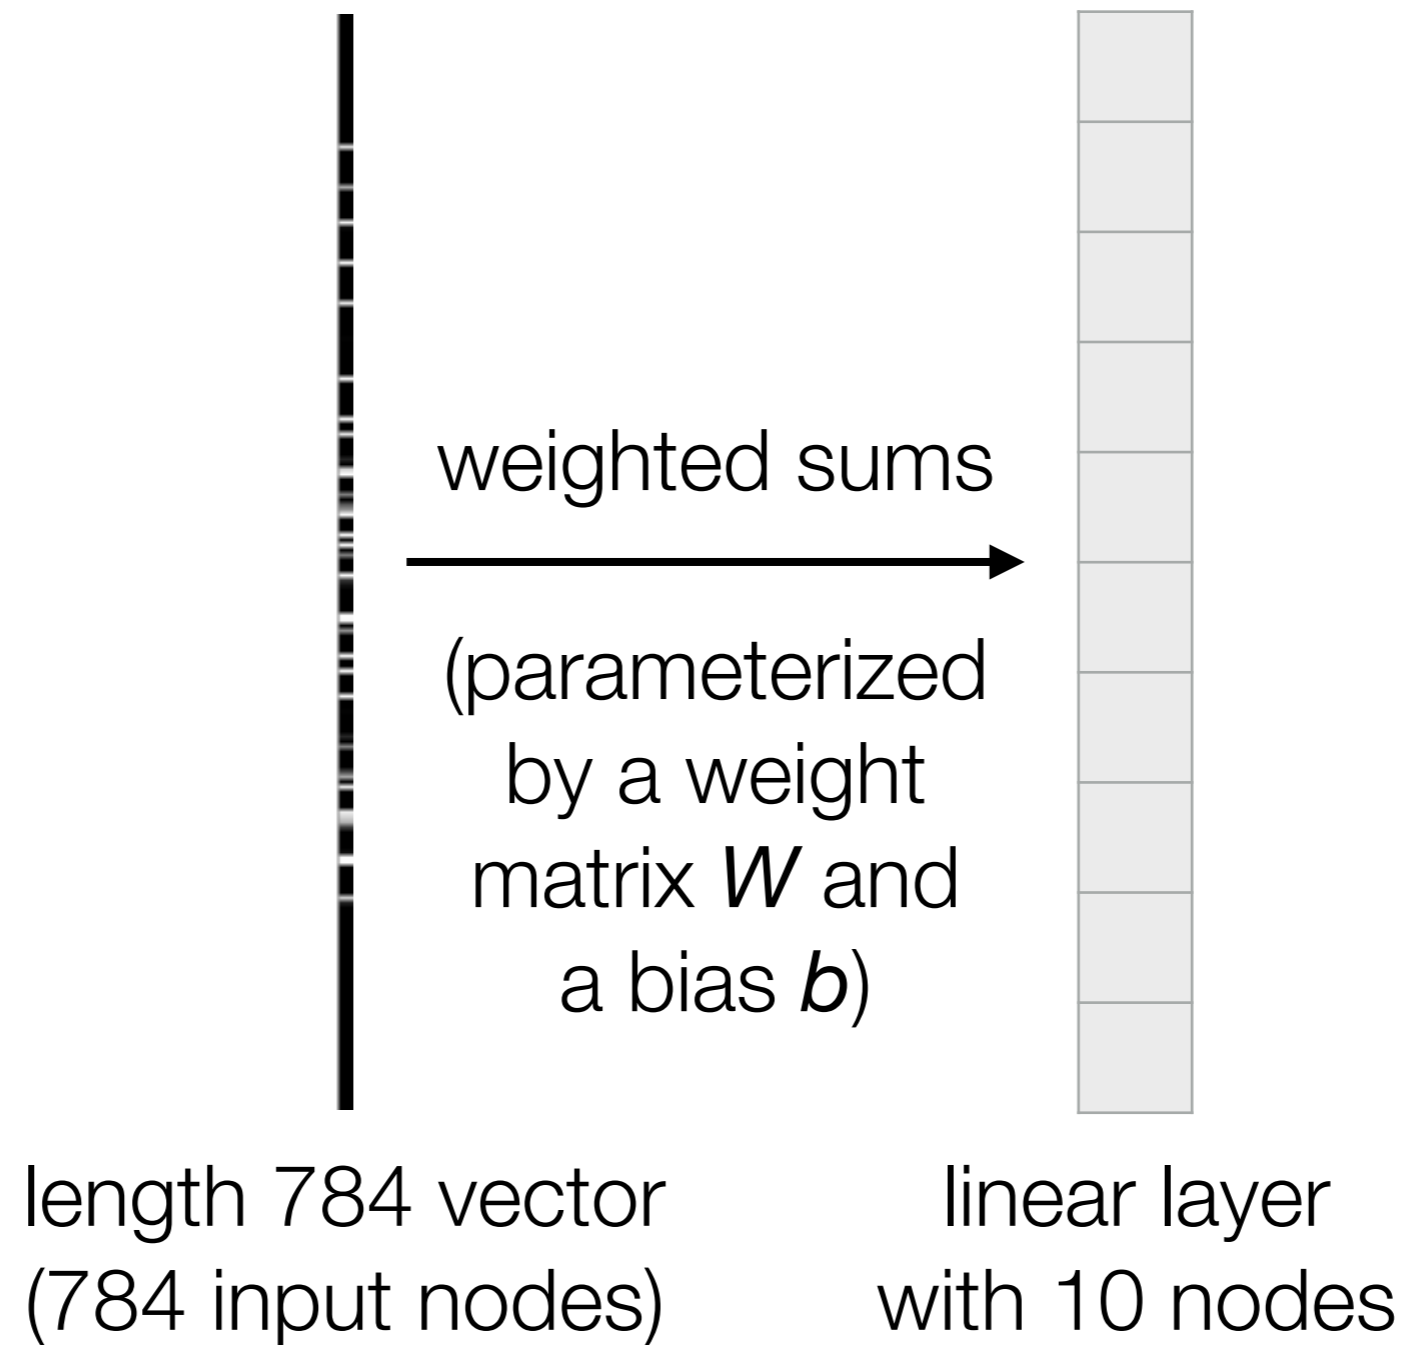
Handwritten Digit Recognition




Handwritten Digit Recognition



Handwritten Digit Recognition



Handwritten Digit Recognition



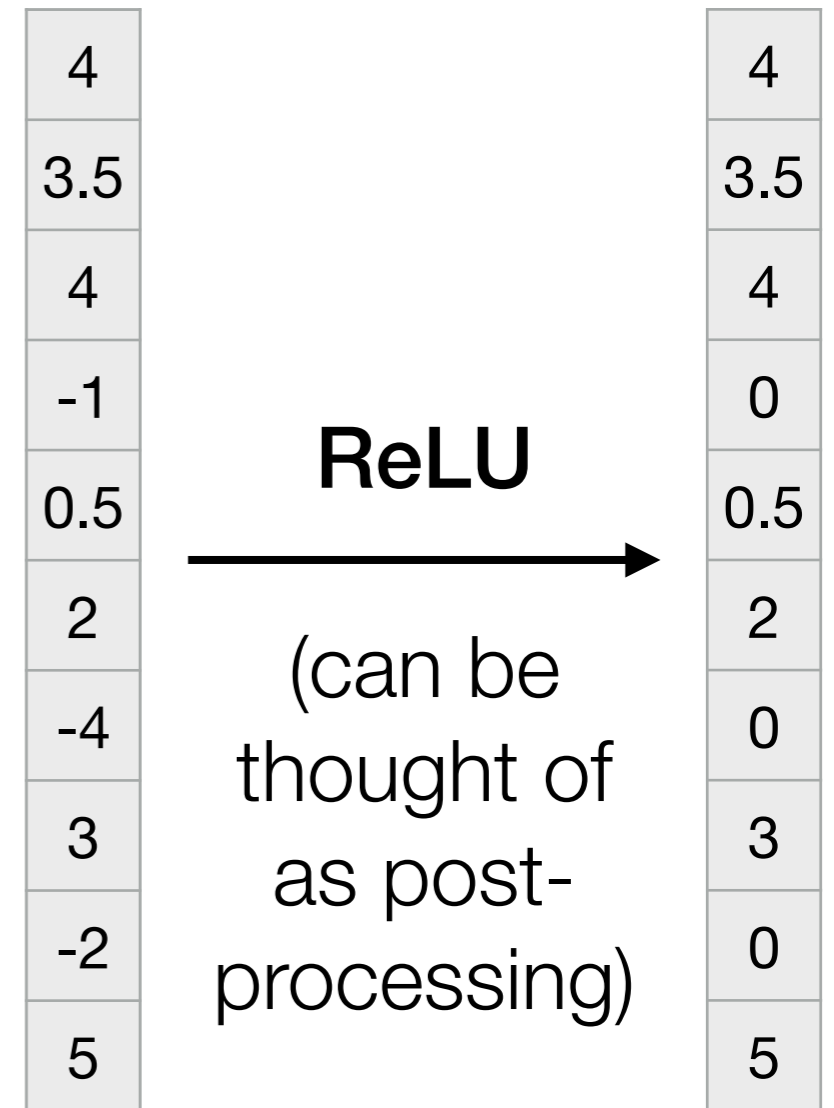
Handwritten Digit Recognition

Many different activation functions possible

Example: **Rectified linear unit (ReLU)**

zeros out entries that are negative

```
final = np.maximum(0, linear)
```



linear layer
with 10 nodes

final
output

`linear`

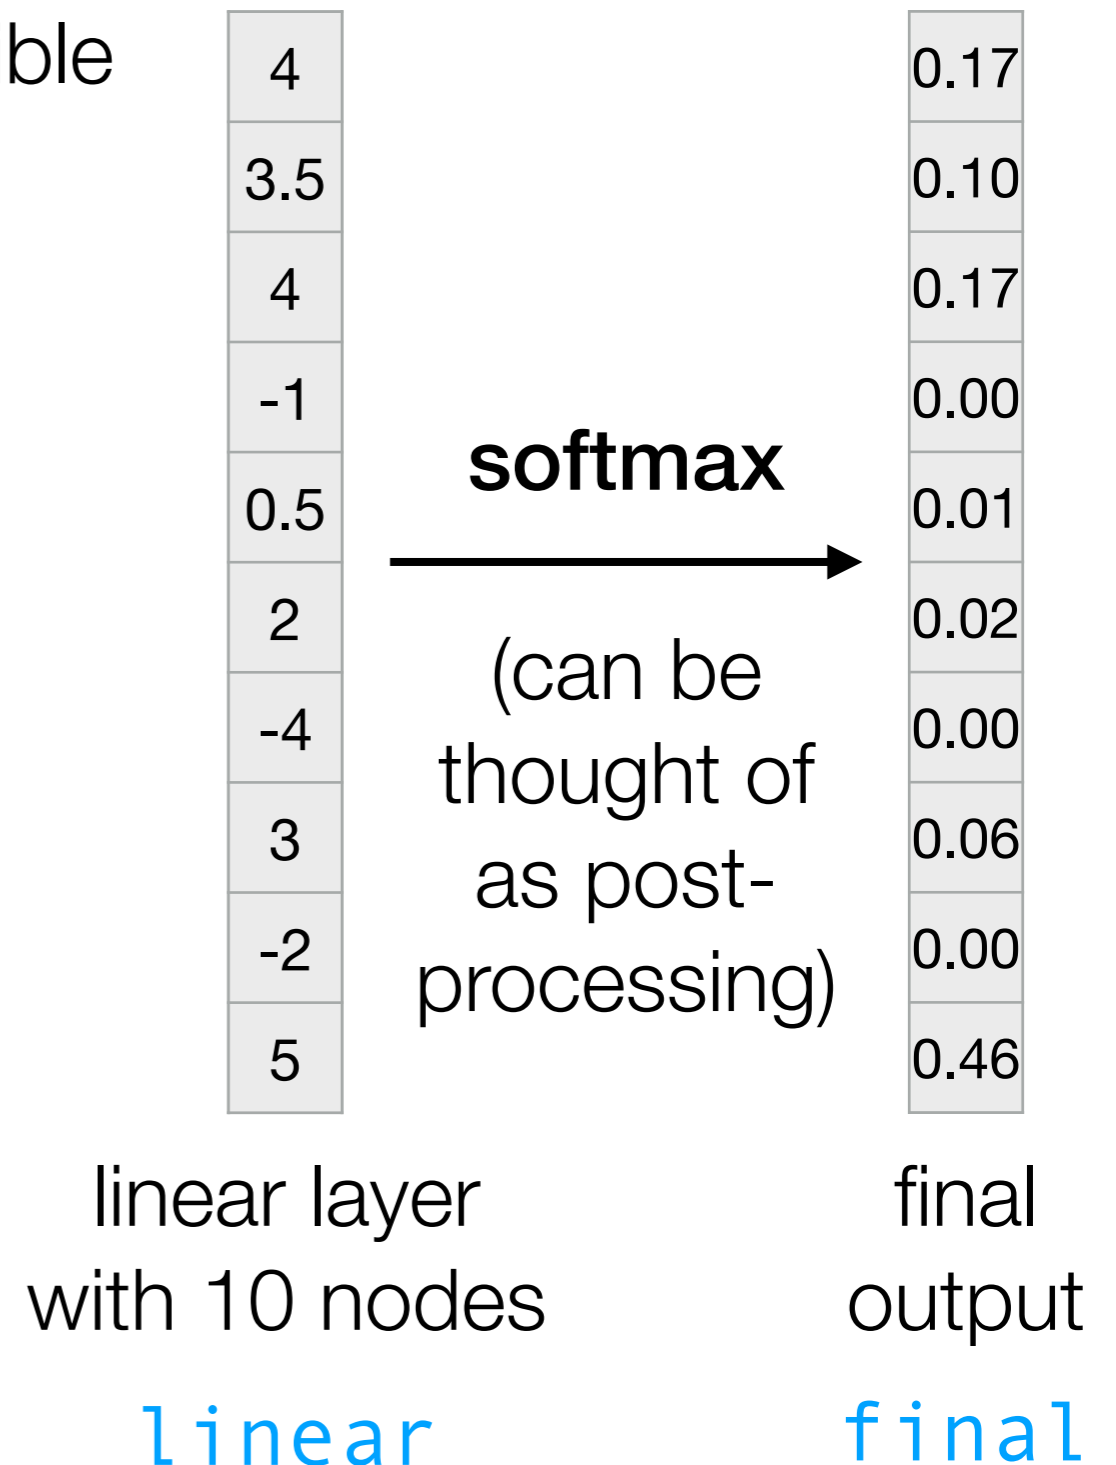
`final`

Handwritten Digit Recognition

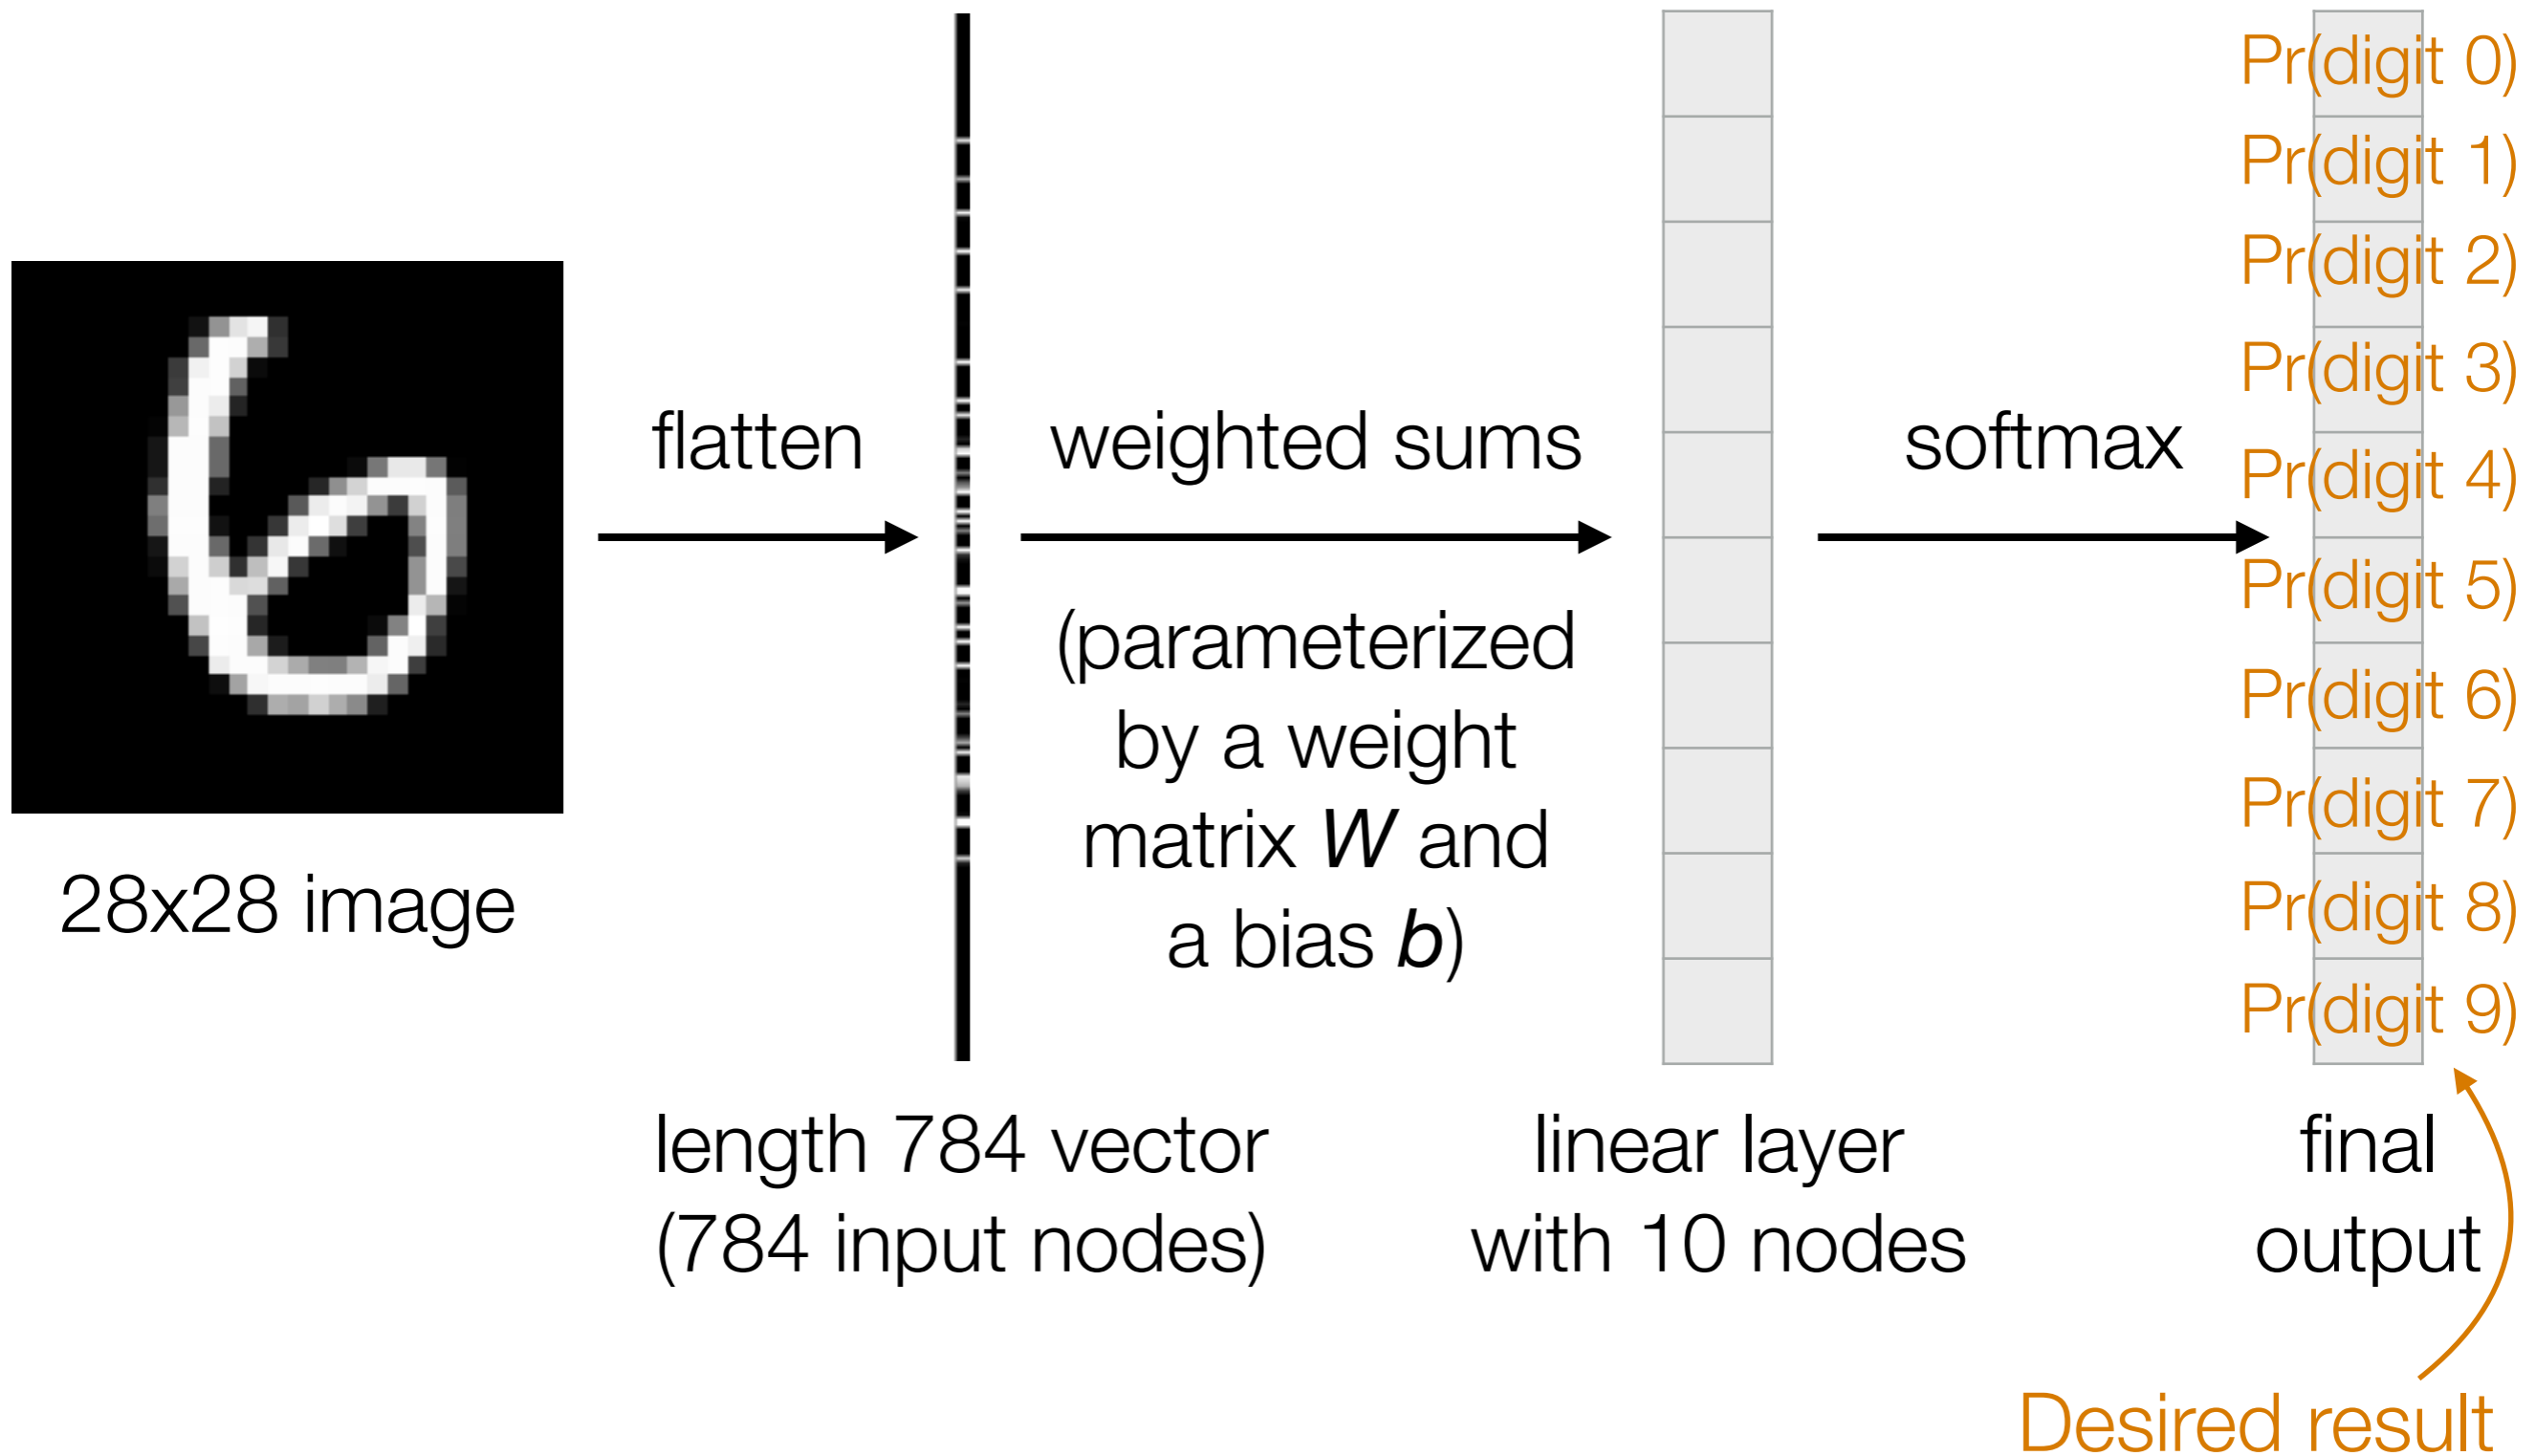
Many different activation functions possible

Example: **softmax** converts a table of numbers into a probability distribution

```
exp = np.exp(linear)
final = exp / exp.sum()
```




Handwritten Digit Recognition

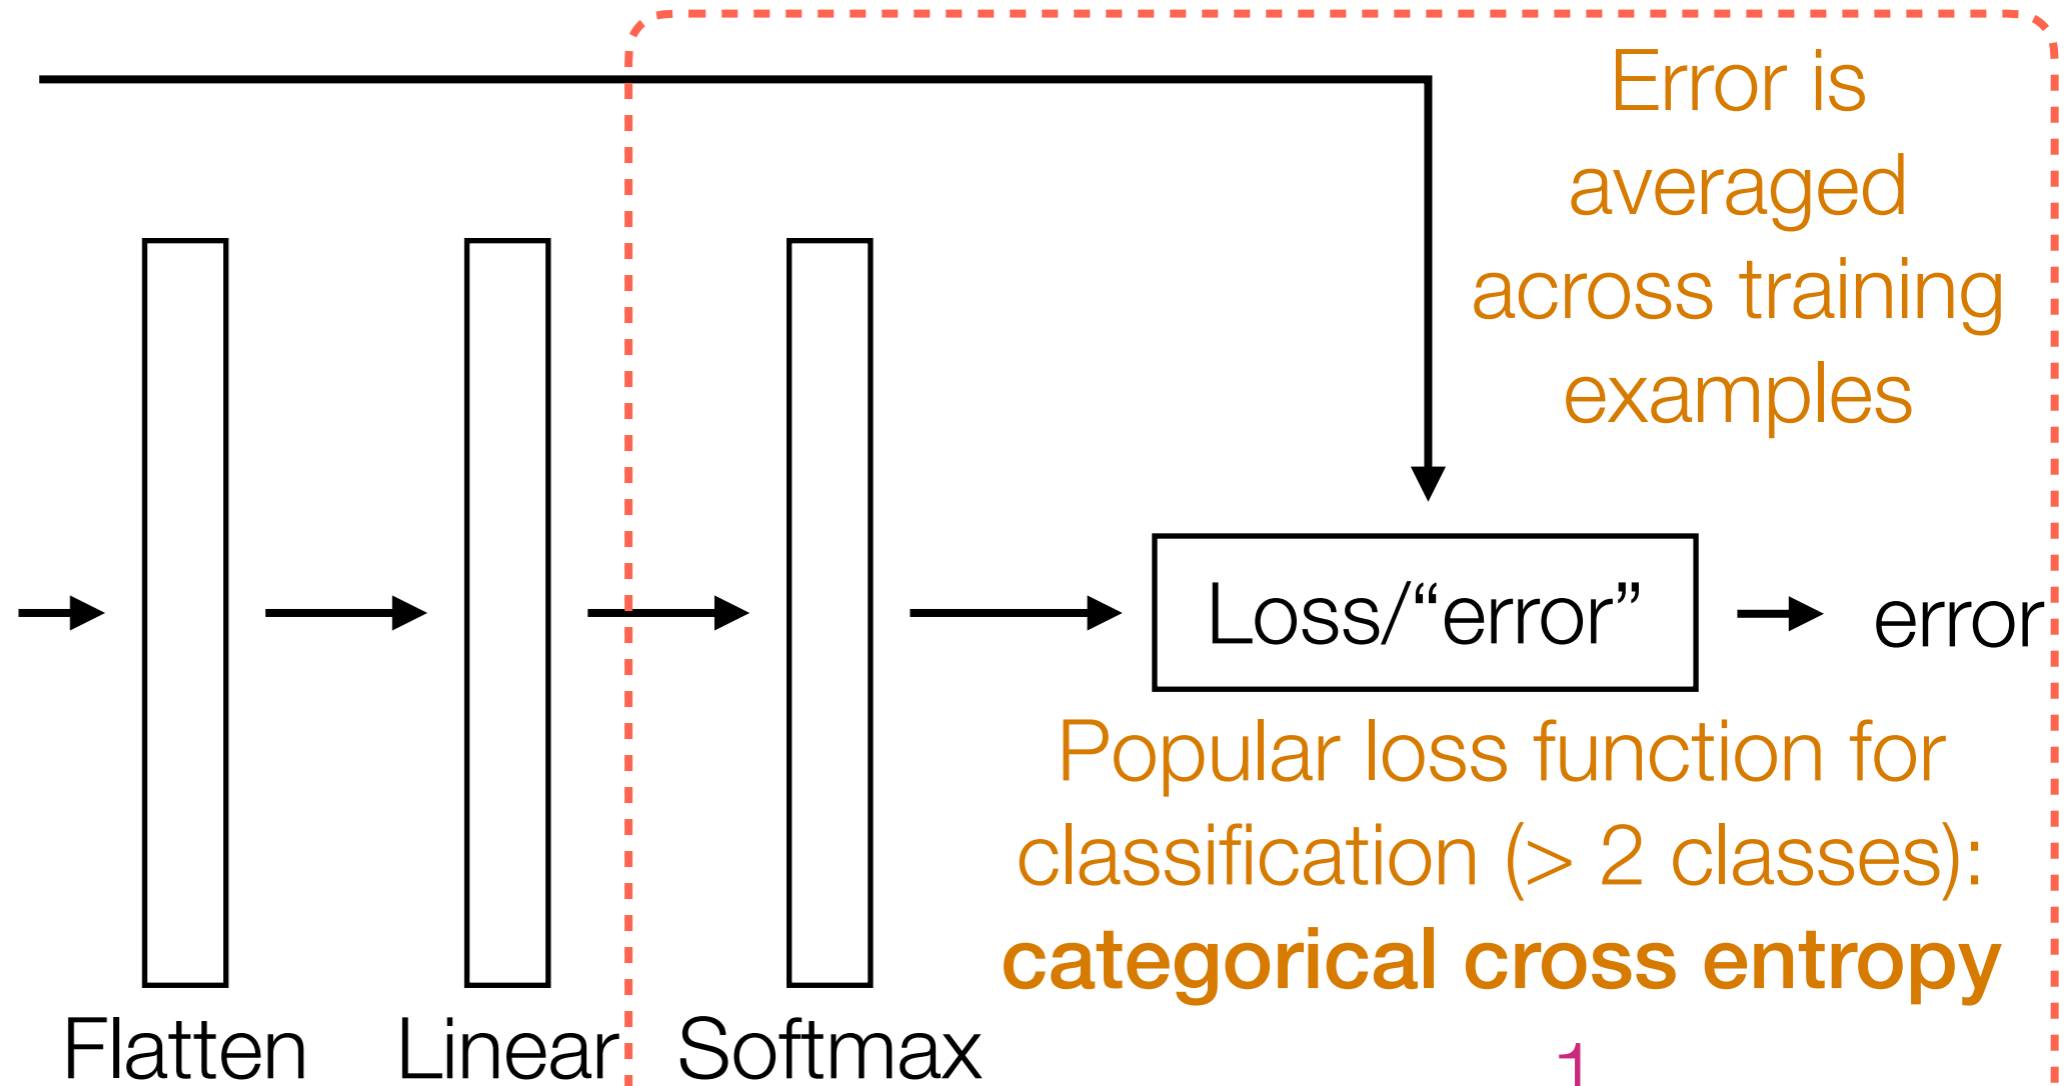


Handwritten Digit Recognition

Training label: 6



Input



Learning this neural net means learning W and b

Also called *fully-connected* or *dense* layer

⚠ In PyTorch, softmax is included as part of the cross entropy loss


$$\log \frac{1}{\text{estimated Pr(digit 6)}}$$

Handwritten Digit Recognition

Demo part 1

Handwritten Digit Recognition


Training label: 6



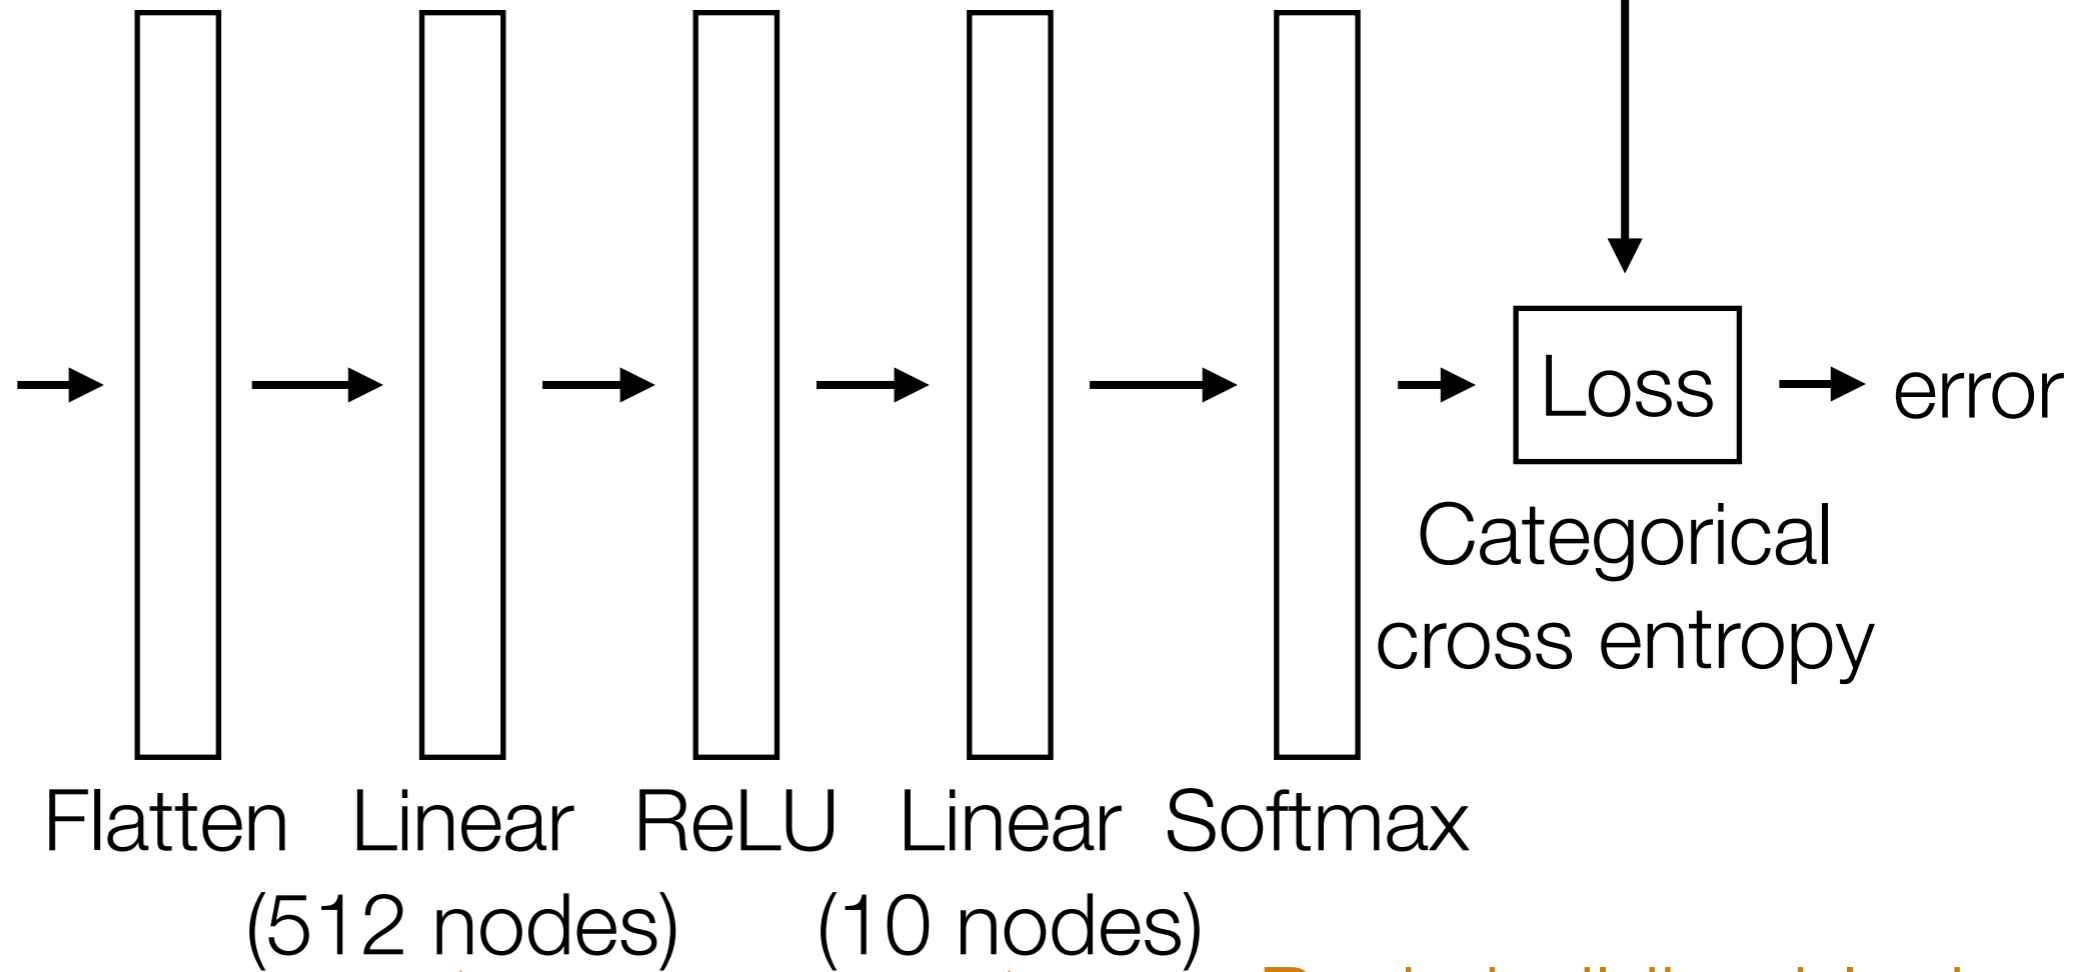
This neural net has a name: **multinomial logistic regression** (when there are only 2 classes, it's called **logistic regression**)

Handwritten Digit Recognition

Training label: 6



Input




Learning this neural net means learning parameters of both linear layers!


Basic building block of neural nets: *linear layer with nonlinear activation*

Handwritten Digit Recognition

Training label: 6



Input



This neural net is called a **multilayer perceptron** (# nodes need not be 512 & 10; activations need not be ReLU and softmax)

Important: in lecture, I will some times use this notation instead

Handwritten Digit Recognition

Demo part 2

Architecting Neural Nets

- Basic building block that is often repeated: linear layer followed by nonlinear activation
 - Without nonlinear activation, two consecutive linear layers is mathematically equivalent to having a single linear layer!
- How to select # of nodes in a layer, or # of layers?
 - These are hyperparameters! Infinite possibilities!
 - Can choose between different options using hyperparameter selection strategy from earlier lectures
 - Very expensive in practice!
(Active area of research: neural architecture search)
- Much more common in practice: modify existing architectures that are known to work well (e.g., ResNet for image classification/object recognition)

PyTorch GitHub Has Lots of Examples

github.com/pytorch/examples

PyTorch Examples

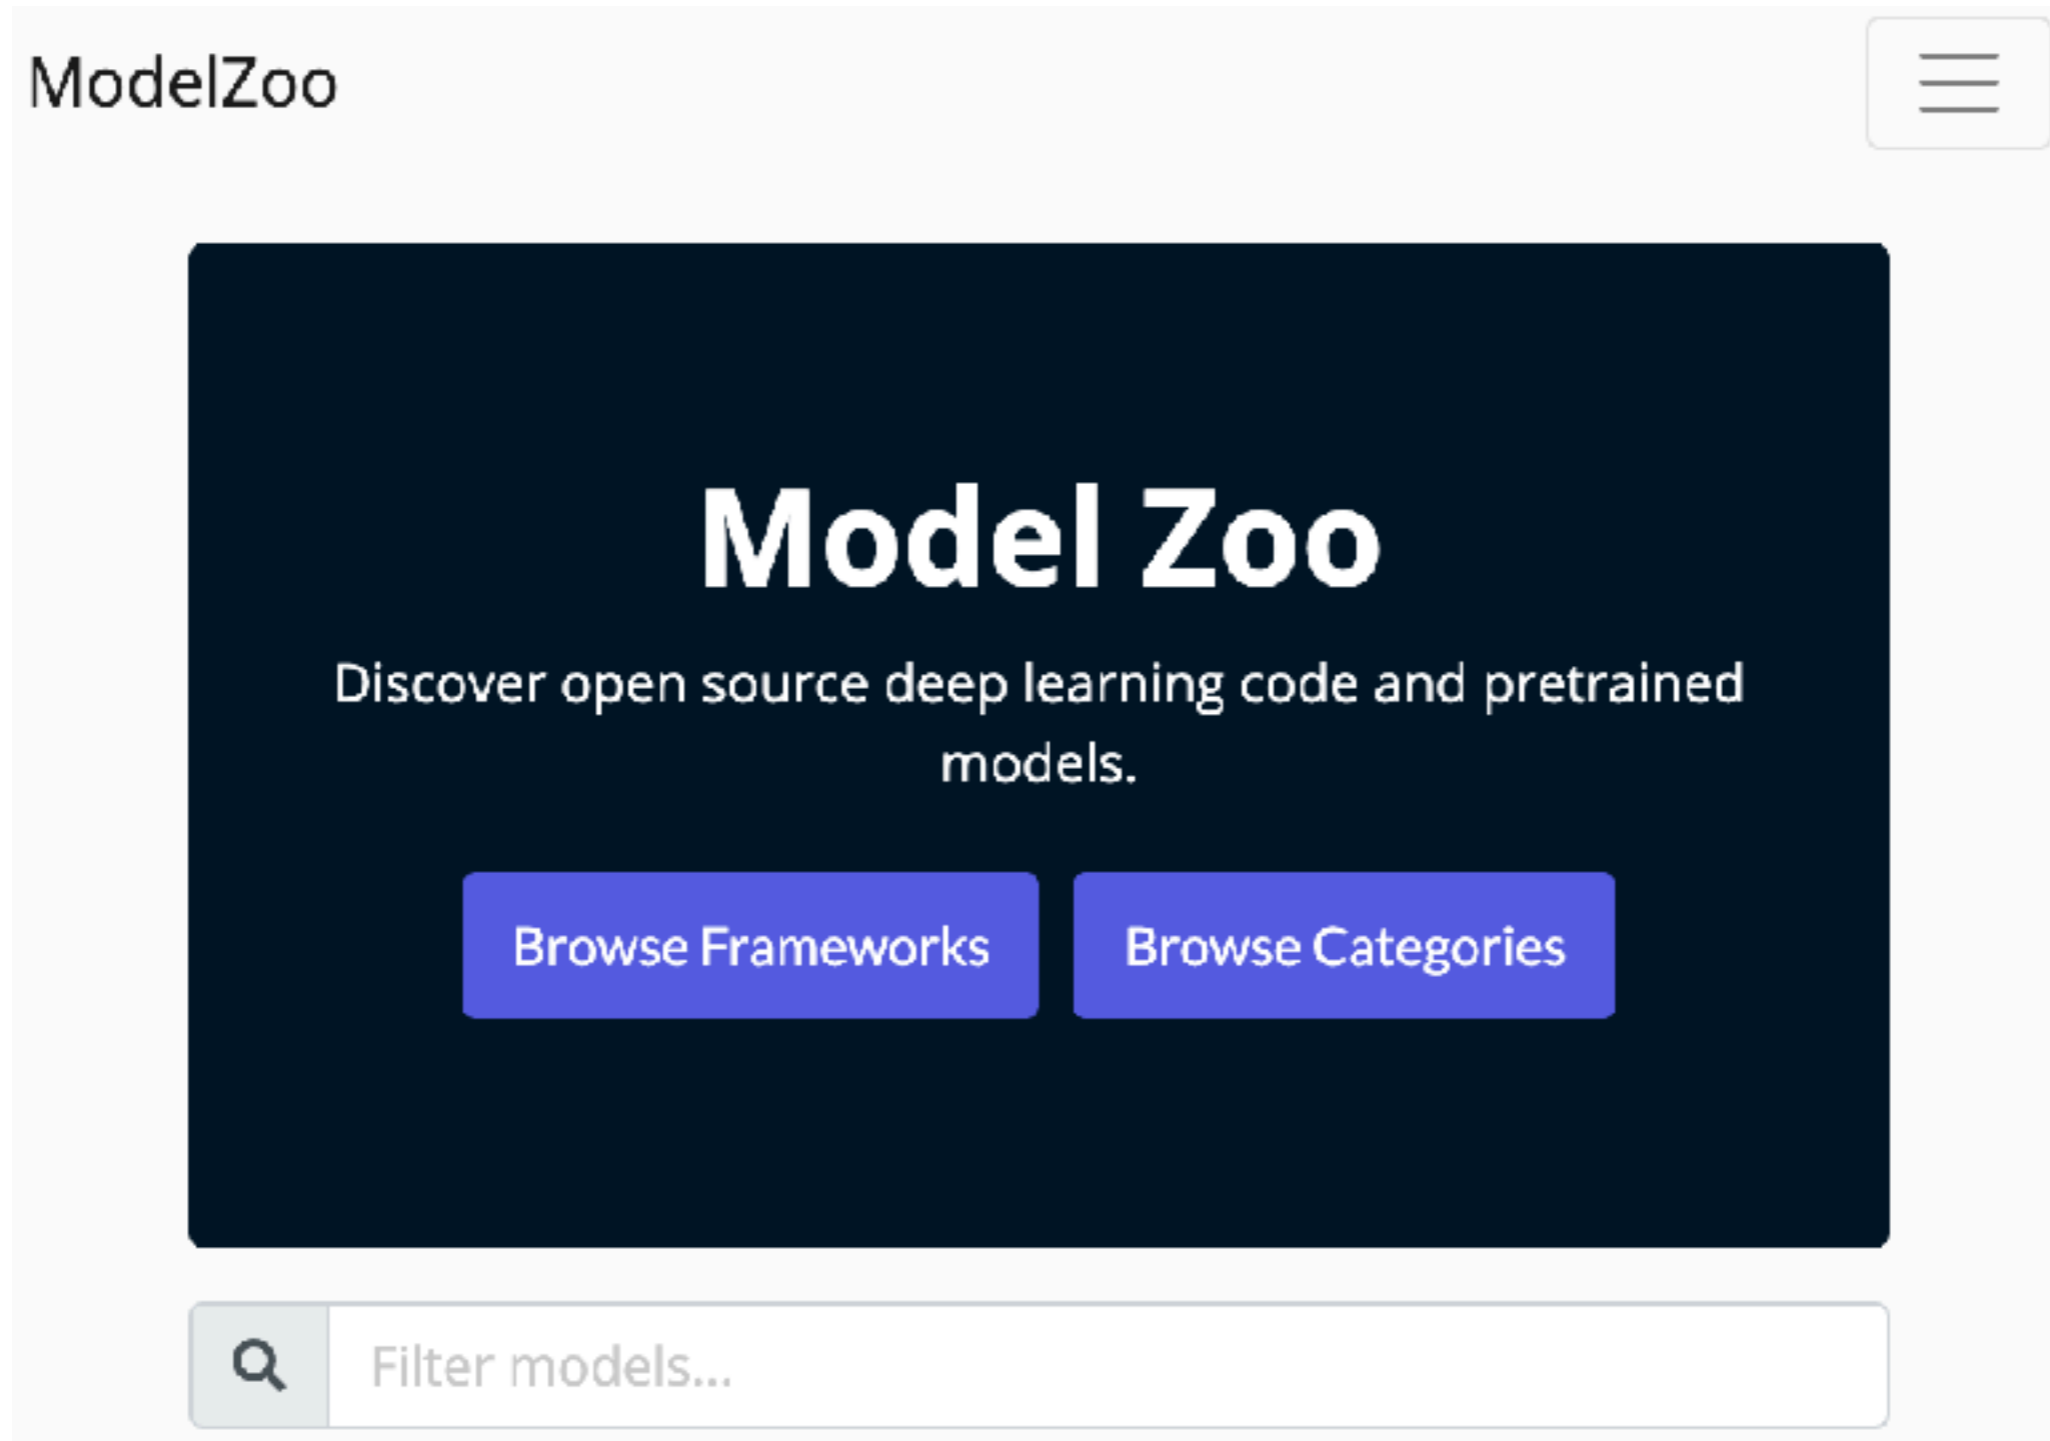
A repository showcasing examples of using [PyTorch](#)

- [Image classification \(MNIST\) using Convnets](#)
- [Word level Language Modeling using LSTM RNNs](#)
- [Training Imagenet Classifiers with Residual Networks](#)
- [Generative Adversarial Networks \(DCGAN\)](#)
- [Variational Auto-Encoders](#)
- [Superresolution using an efficient sub-pixel convolutional neural network](#)
- [Hogwild training of shared ConvNets across multiple processes on MNIST](#)
- [Training a CartPole to balance in OpenAI Gym with actor-critic](#)
- [Natural Language Inference \(SNLI\) with GloVe vectors, LSTMs, and torchtext](#)
- [Time sequence prediction - use an LSTM to learn Sine waves](#)
- [Implement the Neural Style Transfer algorithm on images](#)
- [Several examples illustrating the C++ Frontend](#)

Additionally, a list of good examples hosted in their own repositories:

- [Neural Machine Translation using sequence-to-sequence RNN with attention \(OpenNMT\)](#)

Find a Massive Collection of Models at the Model Zoo



Learning a neural net amounts to curve fitting

We're just estimating a function

Neural Net as Function Approximation

Given `input`, learn a computer program that computes `output`

this is a **function**

Multinomial logistic regression:

```
def f(input):
```

```
    output = softmax(np.dot(input,  $W$ ) +  $b$ )
```

```
    return output
```

the only things that we are learning
(we fix their dimensions in advance)

We are fixing what the function `f` looks like in code
and are only adjusting `W` and `b`!!!

Neural Net as Function Approximation

Given `input`, learn a computer program that computes `output`

Multinomial logistic regression:

```
output = softmax(np.dot(input, W) + b)
```

Multilayer perceptron:

```
intermediate = relu(np.dot(input, W1) + b1)
```

```
output = softmax(np.dot(intermediate, W2) + b2)
```

Learning a neural net: learning a simple computer program that maps inputs (raw feature vectors) to outputs (predictions)

Complexity of a Neural Net?

- Increasing number of layers (depth) makes neural net more “complex”
 - Learn computer program that has more lines of code
 - Some times, more parameters may be needed
 - If so, more training data may be needed

Earlier: multinomial logistic regression had fewer parameters than multilayer perceptron example

Upcoming: we'll see examples of deep nets with *fewer* parameters than “shallower” nets