Announcements

project: groups of 2-3

2 options

Progress During Learning

plot loss (CE, MSE)

- Track progress during learning by plotting training curves.
- Chose the training criterion (e.g. squared error, cross-entropy) partly to be easy to optimize.
- May wish to track other metrics to measure performance (even if we can't directly optimize them).

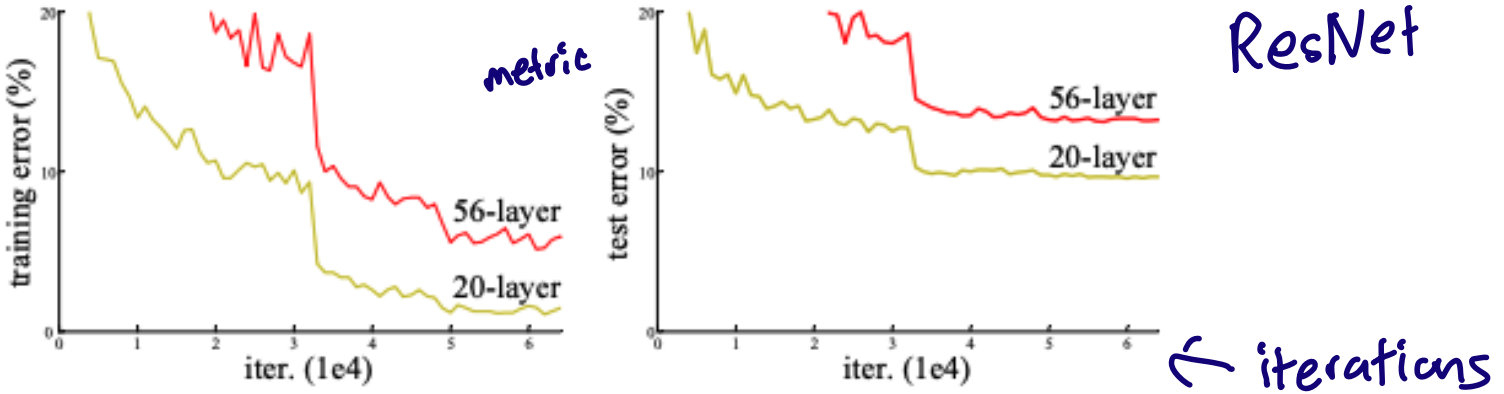

Figure 1. Training error (left) and test error (right) on CIFAR-10 with 20-layer and 56-layer "plain" networks. The deeper network has higher training error, and thus test error. Similar phenomena on ImageNet is presented in Fig. 4.

Tracking Accuracy for Binary Classification

We can track accuracy, or fraction correctly classified.

- Equivalent to the average 0–1 loss, the error rate, or fraction incorrectly classified.
- Useful metric to track even if we couldn't optimize it.

Another way to break down the accuracy:

$$Acc = \frac{TP + TN}{P + N} = \frac{TP + TN}{(TP + FN) + (TN + FP)}$$

- P=num positive; N=num negative;
- TP=true positives; TN=true negatives
- FP=false positive or a type I error
- FN=false negative or a type II error

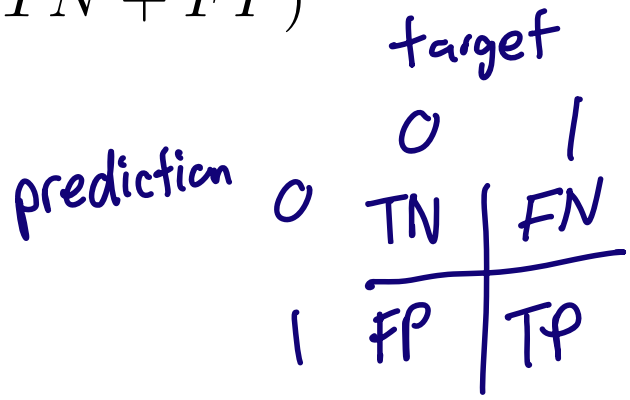

Accuracy is Highly Sensitive to Class Imbalance

Suppose you are screening patients for a particular disease.
It's known that 1% of patients have that disease.

- What is the simplest model that can achieve 99% accuracy?

predict everyone does not have disease

Sensitivity and Specificity

Useful metrics even under class imbalance!

$$\text{Sensitivity} = \frac{TP}{TP+FN} \text{ [True positive rate]}$$

$$\text{Specificity} = \frac{TN}{TN+FP} \text{ [True negative rate]}$$

What happens if our problem is not linearly separable?
How do we pick a threshold for $y = \sigma(x)$?

$$\hat{y} = \begin{cases} 1 & \text{if logit} \geq 0.5 \\ 0 & \text{if logit} < 0.5 \end{cases}$$

prediction

0.5 ↖ threshold

Designing Diagnostic Tests

- A binary model to predict whether someone has a disease.
- What happens to sensitivity and specificity as you slide the threshold from left to right?

e.g. 0.5 → 0.4
change threshold to predict more positives

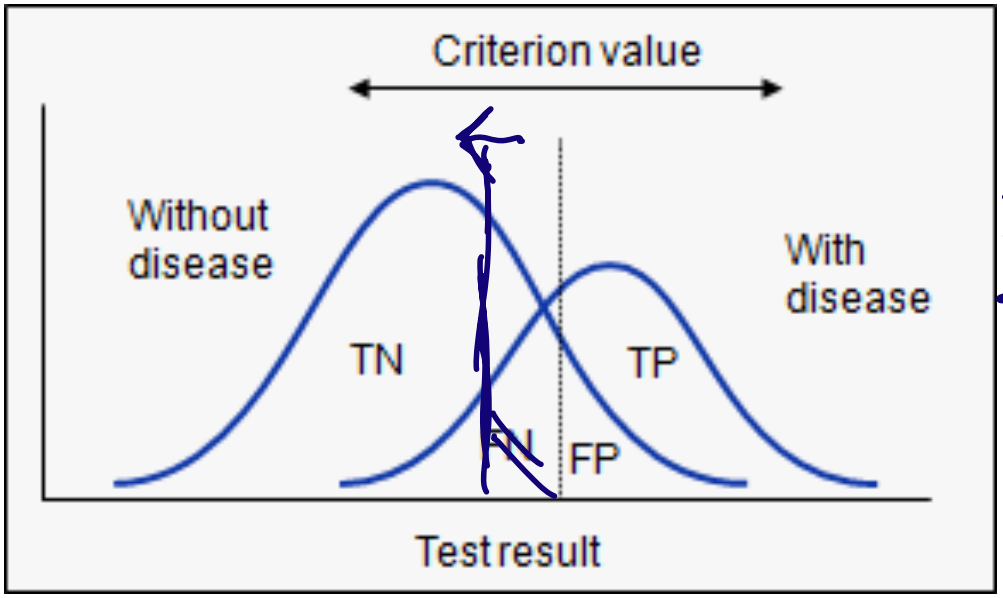

FP TP ↑
TN FN ↓

Tradeoff between Sensitivity and Specificity

As we increase the criterion value (i.e. move from left to right), how do the sensitivity and specificity change?

trade-off

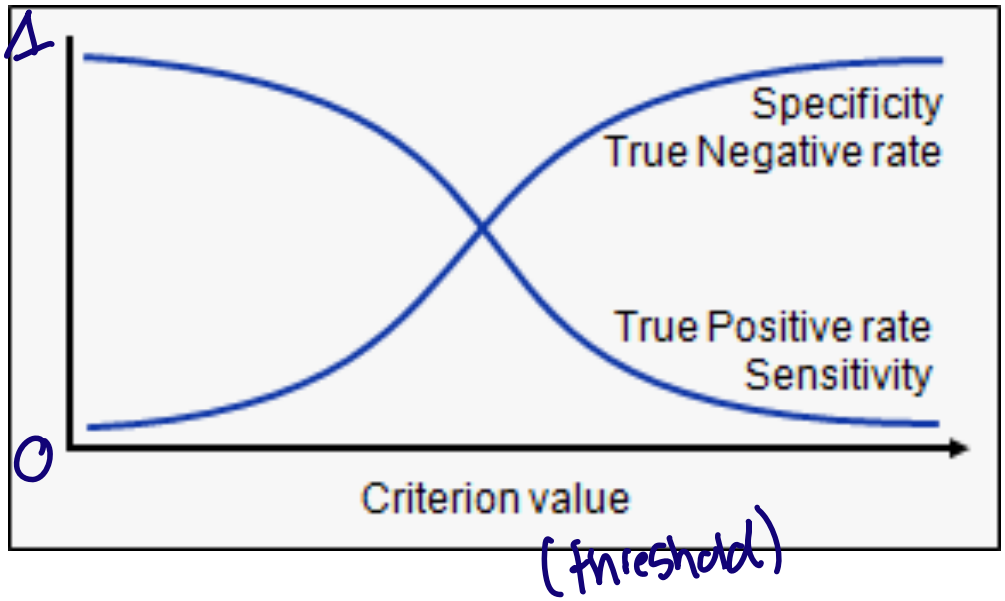

close to 1

Receiver Operating Characteristic (ROC) Curve

Area under the ROC curve (AUC) can quantify if a binary classifier achieves a good tradeoff between sensitivity and specificity.

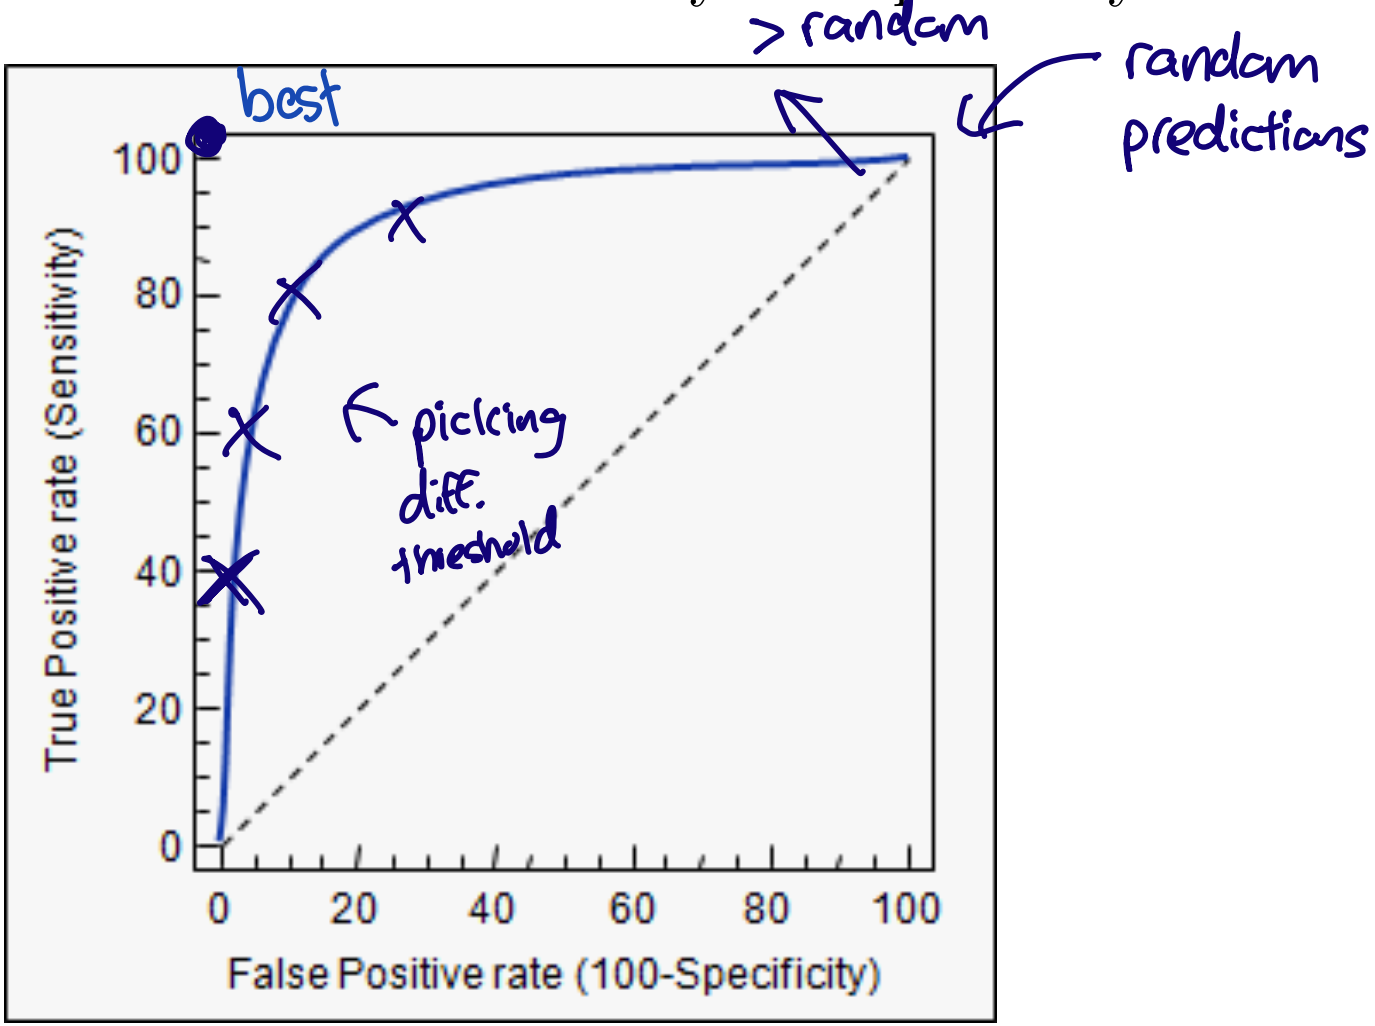

Confusion Matrix for Multi-Class classification

- Visualizes how frequently certain classes are confused.
- $K \times K$ matrix; rows are true labels, columns are predicted labels, entries are frequencies
- What does the confusion matrix for a perfect classifier look like?

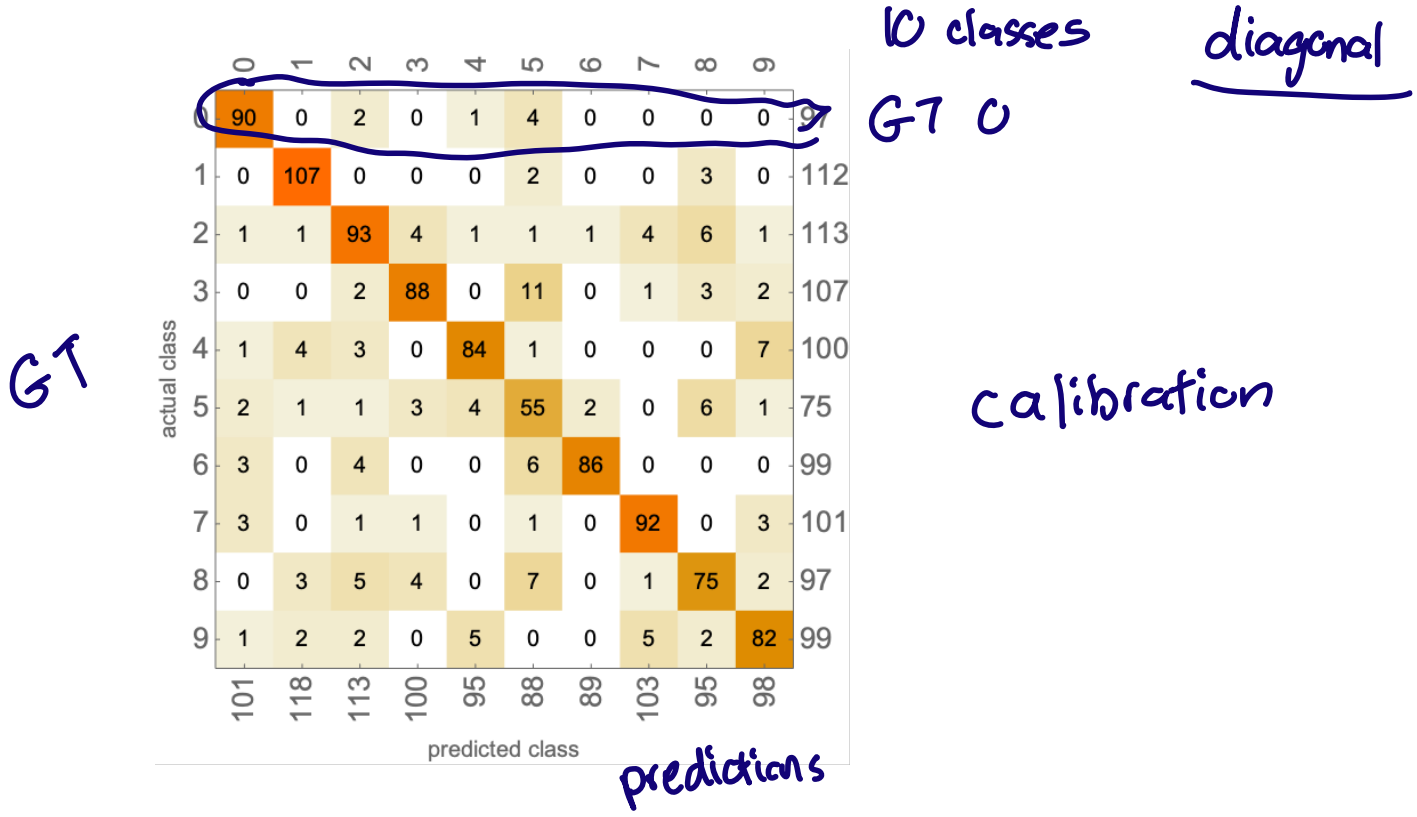

XOR is Not Linearly Separable

Some datasets are not linearly separable, e.g. **XOR**.

exclusive or

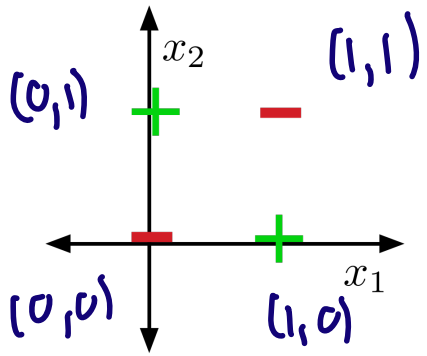

convexity

Visually obvious, but how can we prove this formally?

Proof That XOR is Not Linearly Separable

Proof by Contradiction: $a^T x < b$

- Half-spaces are convex. That is, if two points lie in a half-space, the line segment connecting them also lie in the same half-space.
- Suppose that the problem is feasible.
- If the positive examples are in the positive half-space, then the green line segment must be as well.
- Similarly, the red line segment must lie in the negative half-space.
- But, the intersection can't lie in both half-spaces. Contradiction!

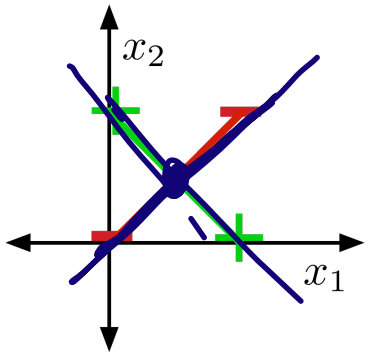

Classifying XOR Using Feature Maps

Sometimes, we can overcome this limitation using feature maps, e.g., for **XOR**.

$$\boldsymbol{\psi}(\mathbf{x}) = \begin{pmatrix} x_1 \\ x_2 \\ x_1 x_2 \end{pmatrix}$$

x_1	x_2	$\psi_1(\mathbf{x})$	$\psi_2(\mathbf{x})$	$\psi_3(\mathbf{x})$	t
0	0	0	0	0	0
0	1	0	1	0	1
1	0	1	0	0	1
1	1	1	1	1	0

- This is linearly separable. (Try it!)
- Designing feature maps can be hard. Can we learn them?

Neurons in the Brain

Neurons receive input signals and accumulate voltage.
After some threshold, they will fire spiking responses.

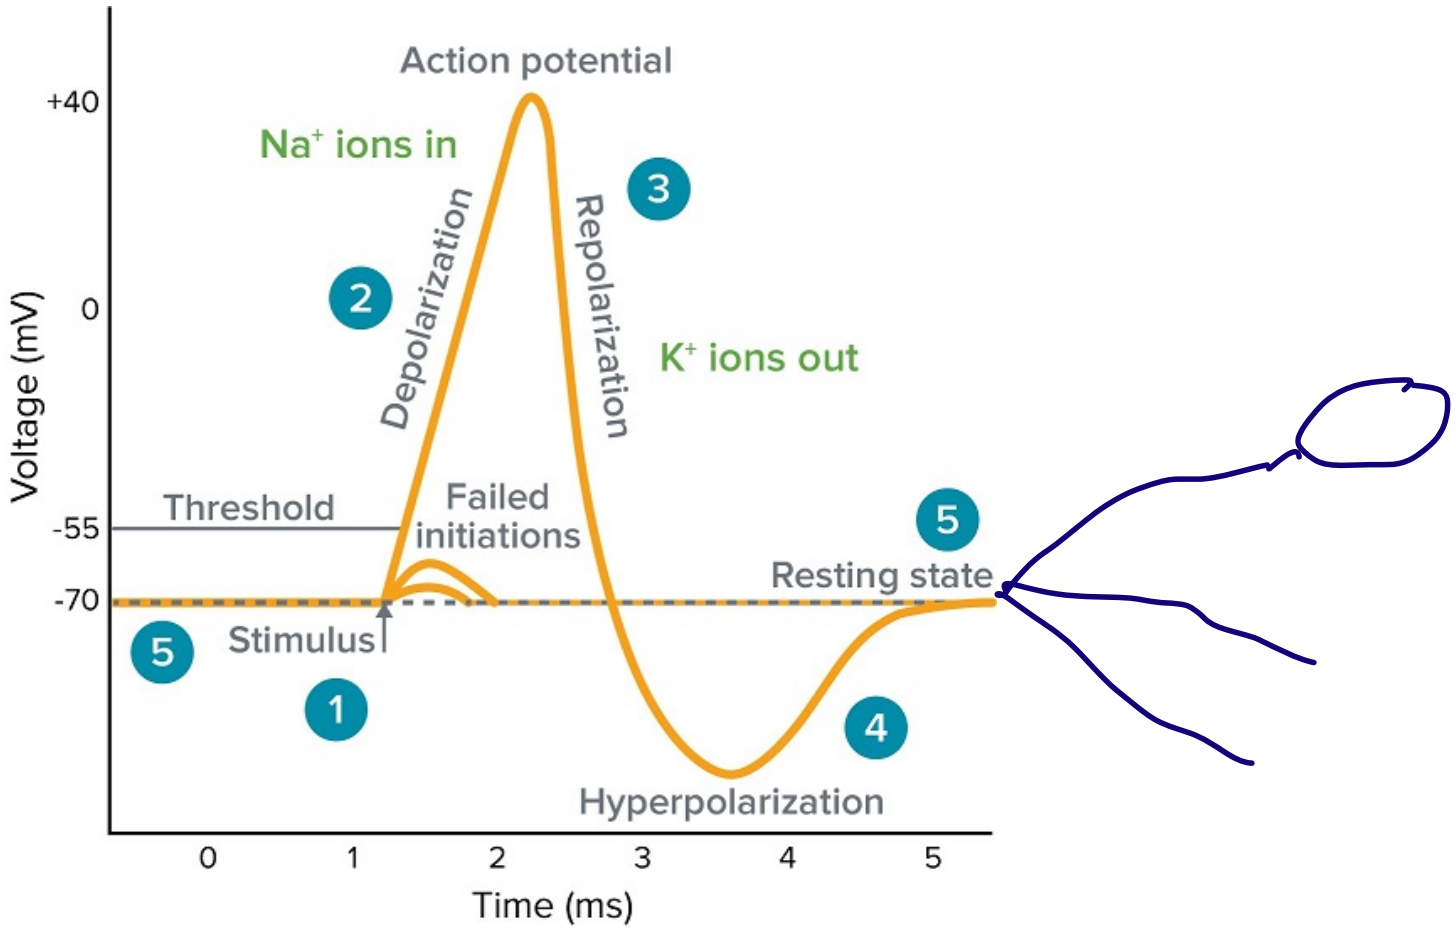

[Pic credit: www.moleculardevices.com]

A Simpler Neuron

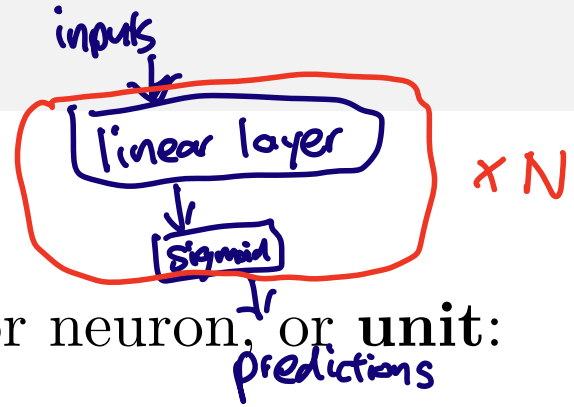

fully connected neural networks

For neural nets, we use a much simpler model for neuron, or **unit**:

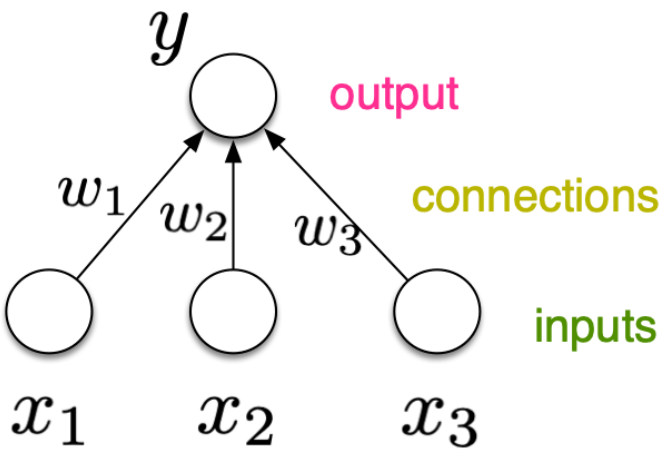

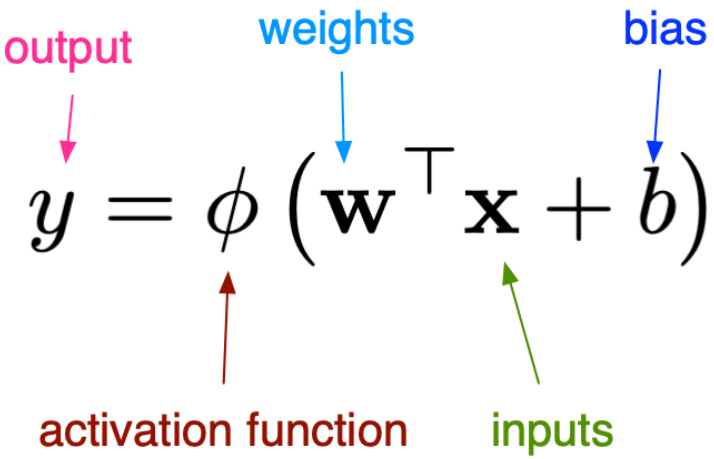

- Similar to logistic regression: $y = \sigma(\mathbf{w}^\top \mathbf{x} + b)$
- By throwing together lots of these simple neuron-like processing units, we can do some powerful computations!

A Feed-Forward Neural Network

model: (linear) (neural nets) → $f(x)$ expressive

$\nabla f(x)$

optim.

no cycles

- A directed acyclic graph (DAG)
- Units are grouped into layers

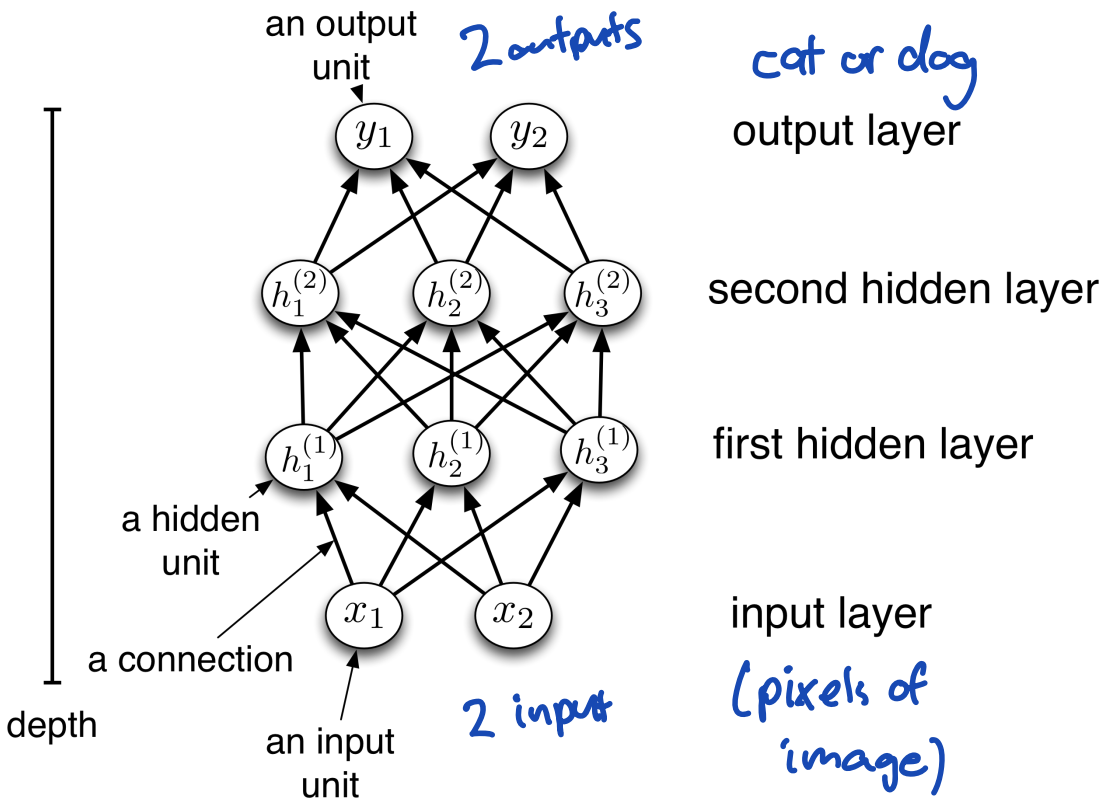

Multilayer Perceptrons

- A multi-layer network consists of fully connected layers.
- In a fully connected layer, all input units are connected to all output units.
- Each hidden layer i connects N_{i-1} input units to N_i output units. Weight matrix is N_i x N_{i-1}.
- The outputs are a function of the input units:

$$\mathbf{y} = f(\mathbf{x}) = \phi(\mathbf{W}\mathbf{x} + \mathbf{b})$$

ϕ is applied component-wise.

modularity

MNIST

28x28 dim

input: 784

200

784x200

200x10

output: 10

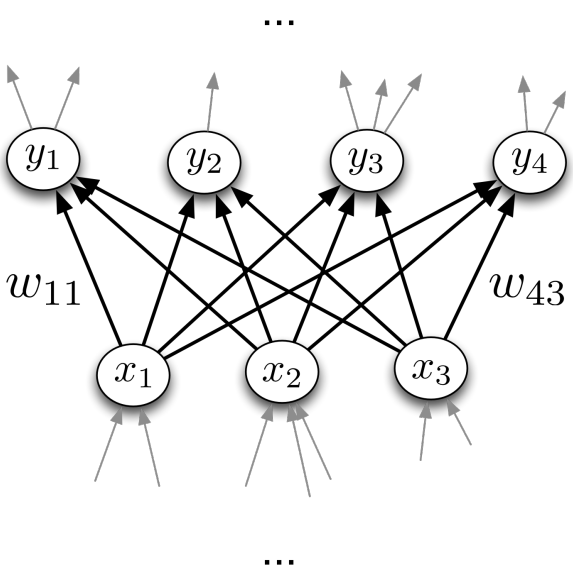

4 units

3 units

12 weights

3x4

Some Activation Functions

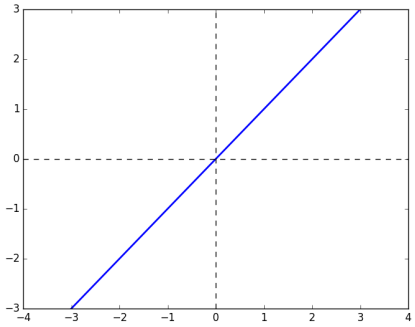

Identity

$$y = z$$

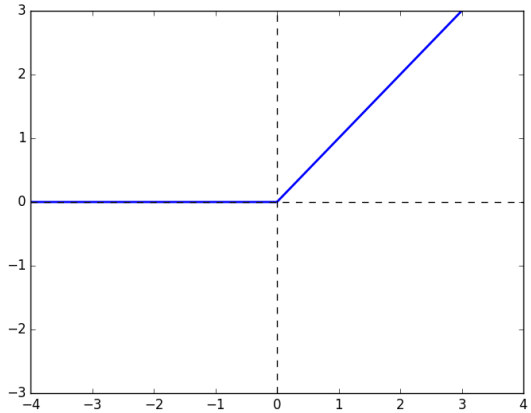

Rectified Linear Unit (ReLU)

$$y = \max(0, z)$$

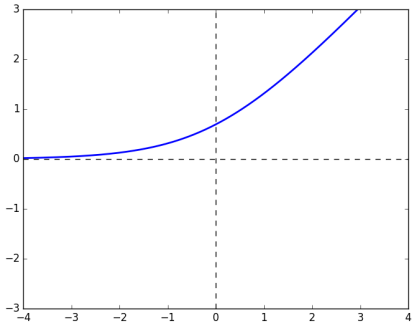

Soft ReLU

$$y = \log 1 + e^z$$

More Activation Functions

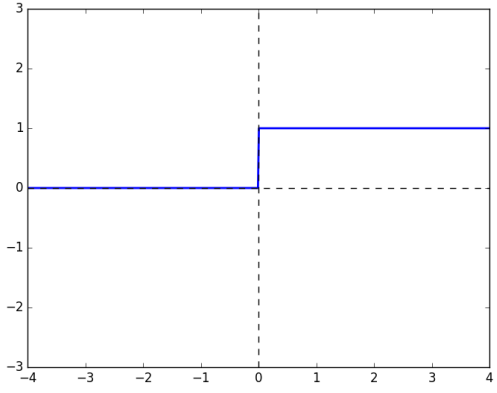

Hard Threshold

$$y = \begin{cases} 1 & \text{if } z > 0 \\ 0 & \text{if } z \leq 0 \end{cases}$$

not used in practice

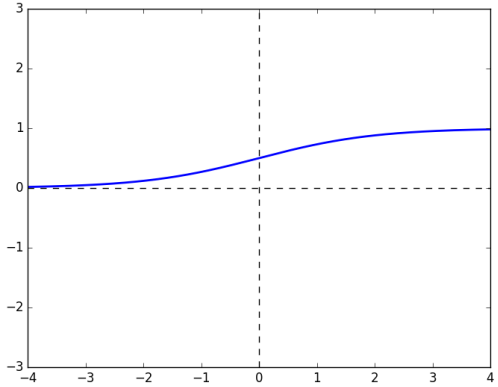

Logistic

$$y = \frac{1}{1 + e^{-z}}$$

$(0, 1)$

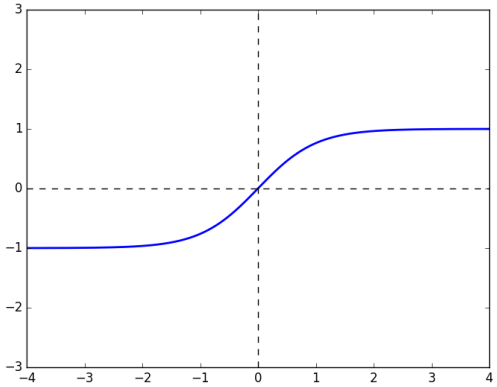

Hyperbolic Tangent (tanh)

$$y = \frac{e^{z} - e^{-z}}{e^{z} + e^{-z}}$$

$(-1, 1)$

Computation in Each Layer

Each layer computes a function.

input

$$\mathbf{h}^{(1)} = f^{(1)}(\mathbf{x}) = \phi(\mathbf{W}^{(1)}\mathbf{x} + \mathbf{b}^{(1)})$$

$$\mathbf{h}^{(2)} = f^{(2)}(\mathbf{h}^{(1)}) = \phi(\mathbf{W}^{(2)}\mathbf{h}^{(1)} + \mathbf{b}^{(2)})$$

$$\vdots$$

$$\mathbf{y} = f^{(L)}(\mathbf{h}^{(L-1)})$$

prediction

$y = W^3 W^2 W^1 x$ (just linear)

$t \in \mathbb{R}$

If task is regression: choose

$$\mathbf{y} = f^{(L)}(\mathbf{h}^{(L-1)}) = (\mathbf{w}^{(L)})^\top \mathbf{h}^{(L-1)} + b^{(L)}$$

If task is binary classification: choose

$$\mathbf{y} = f^{(L)}(\mathbf{h}^{(L-1)}) = \sigma((\mathbf{w}^{(L)})^\top \mathbf{h}^{(L-1)} + b^{(L)})$$

sigmoid

$t \in (0,1)$

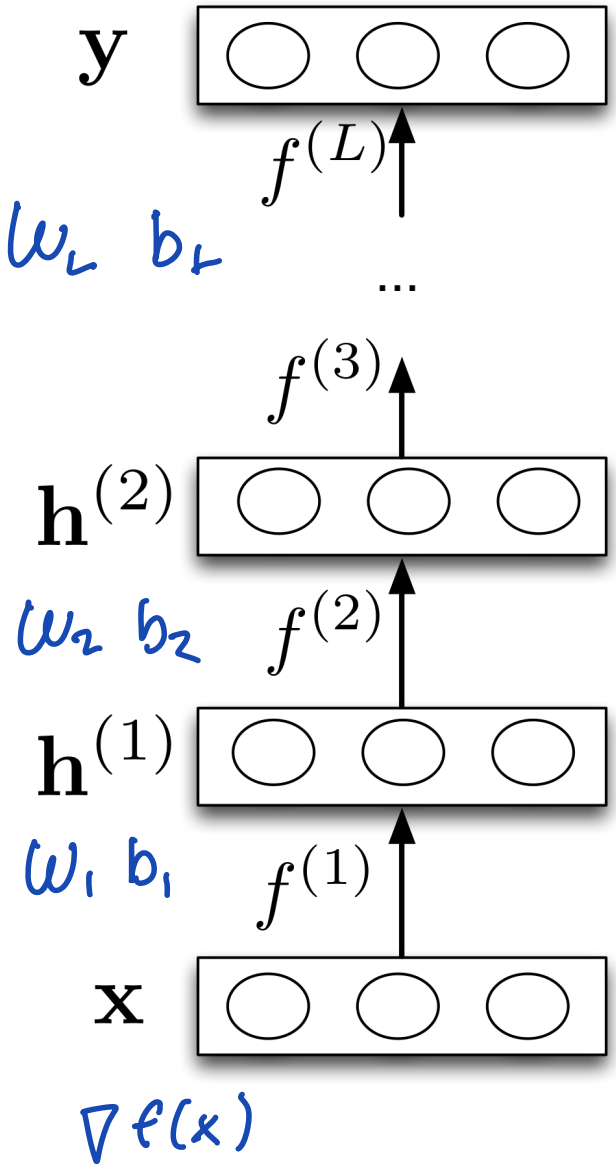

A Composition of Functions

The network computes a composition of functions.

$$\mathbf{y} = f^{(L)} \circ \cdots \circ f^{(1)}(\mathbf{x}).$$

$f^{(L)}(f^{(L-1)} \dots (f^{(1)}(x)))$

Modularity: We can implement each layer's computations as a black box.

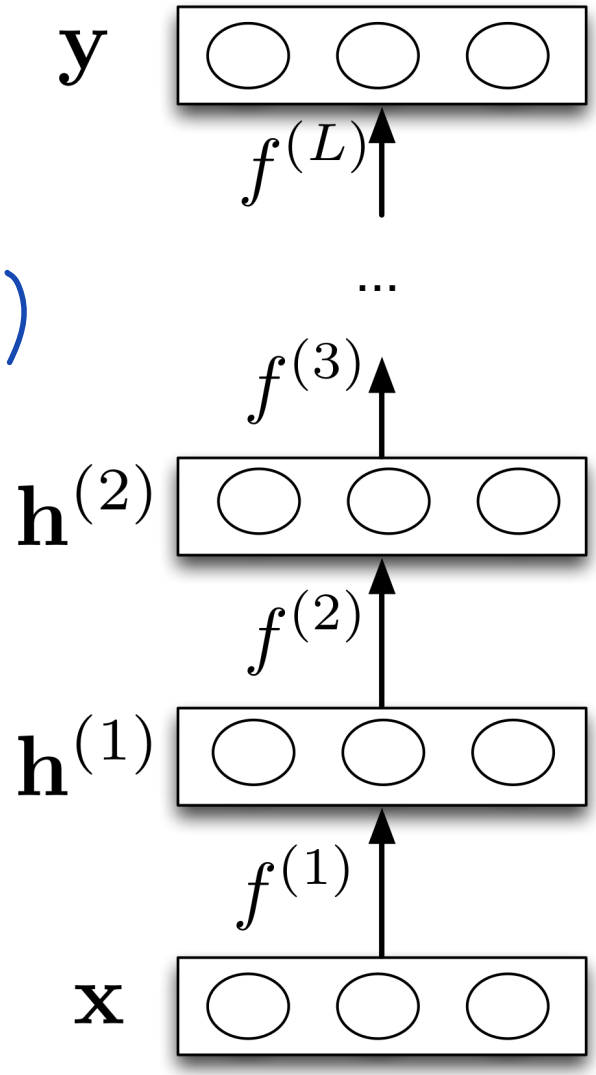

Feature Learning

Neural nets can be viewed as a way of learning features:

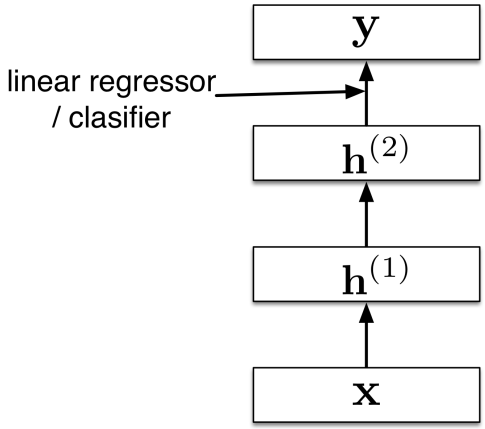

$= \psi(\mathbf{x})$

$w^3 \psi(x)$

useful features for linear classifier

The goal:

poly

$1, x, x^2, \ldots x^n$

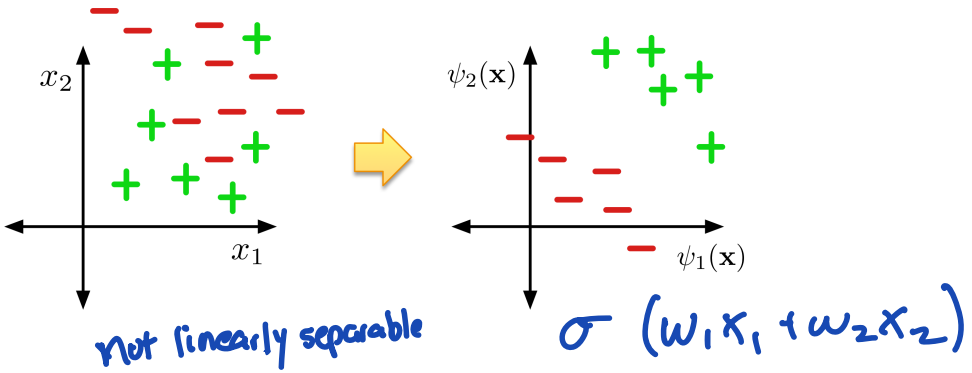

not linearly separable

$\sigma(w_1 x_1 + w_2 x_2)$

Feature Learning

- Suppose we're trying to classify images of handwritten digits.
- Each image is represented as a vector of $28 \times 28 = 784$ pixel values.
- Each hidden unit in the first layer acts as a **feature detector**.
- We can visualize $\mathbf{w}$ by reshaping it into an image.
  Below is an example that responds to a diagonal stroke.

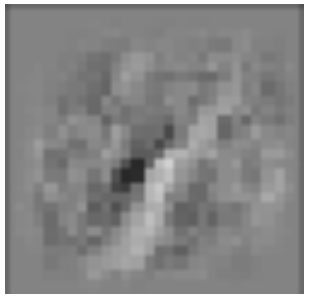

Features for Classifying Handwritten Digits

Features learned by the first hidden layer of a handwritten digit classifier:

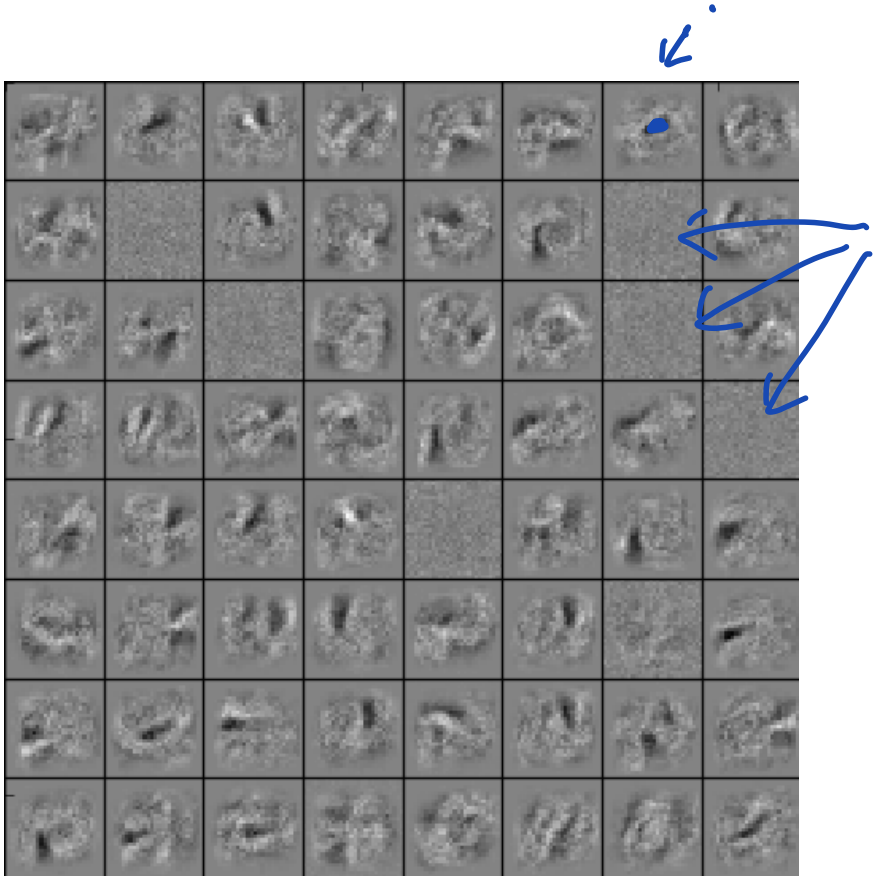

Unlike hard-coded feature maps (e.g., in polynomial regression), features learned by neural networks adapt to patterns in the data.

Expressivity

- A hypothesis space $\mathcal{H}$ is the set of functions that can be represented by some model.
- Consider two models A and B with hypothesis spaces $\mathcal{H}_A, \mathcal{H}_B$.
- If $\mathcal{H}_B \subseteq \mathcal{H}_A$, then A is more expressive than B.
  A can represent any function f in $\mathcal{H}_B$.

linear classifiers $\subseteq$ neural networks

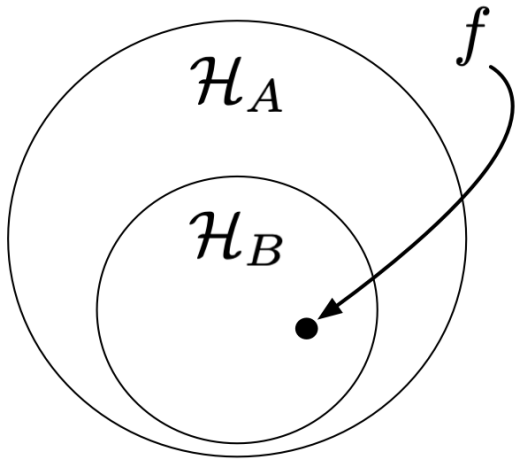

expressivity $\neq$ better

- Some functions (XOR) can't be represented by linear classifiers.
  Are deep networks more expressive?

Expressive Power of Linear Networks

- Consider a linear layer: the activation function was the identity. The layer just computes an affine transformation of the input.
- Any sequence of linear layers is equivalent to a single linear layer.

$$\mathbf{y} = \underbrace{\mathbf{W}^{(3)}\mathbf{W}^{(2)}\mathbf{W}^{(1)}}_{\triangleq \mathbf{W}'}\mathbf{x}$$

$W' = W^3 W^2 W^1$

- Deep linear networks can only represent linear functions — no more expressive than linear regression.

Expressive Power of Non-linear Networks

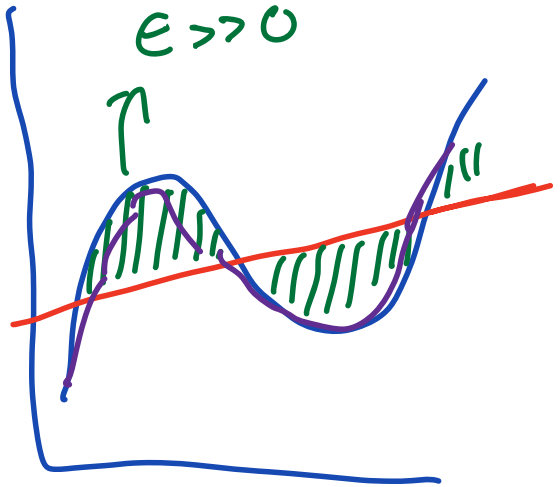

- Multi-layer feed-forward neural networks with non-linear activation functions

- **Universal Function Approximators**:
  They can approximate any function arbitrarily well,
  i.e., for any $f : \mathcal{X} \to \mathcal{T}$ there is a sequence $f_i \in \mathcal{H}$ with $f_i \to f$.

linear ≠ not UFA

- True for various activation functions
  (e.g. thresholds, logistic, ReLU, etc.)

✓ ✓

NNs are UFA

Designing a Network to Classify XOR

Assume a hard threshold activation function.

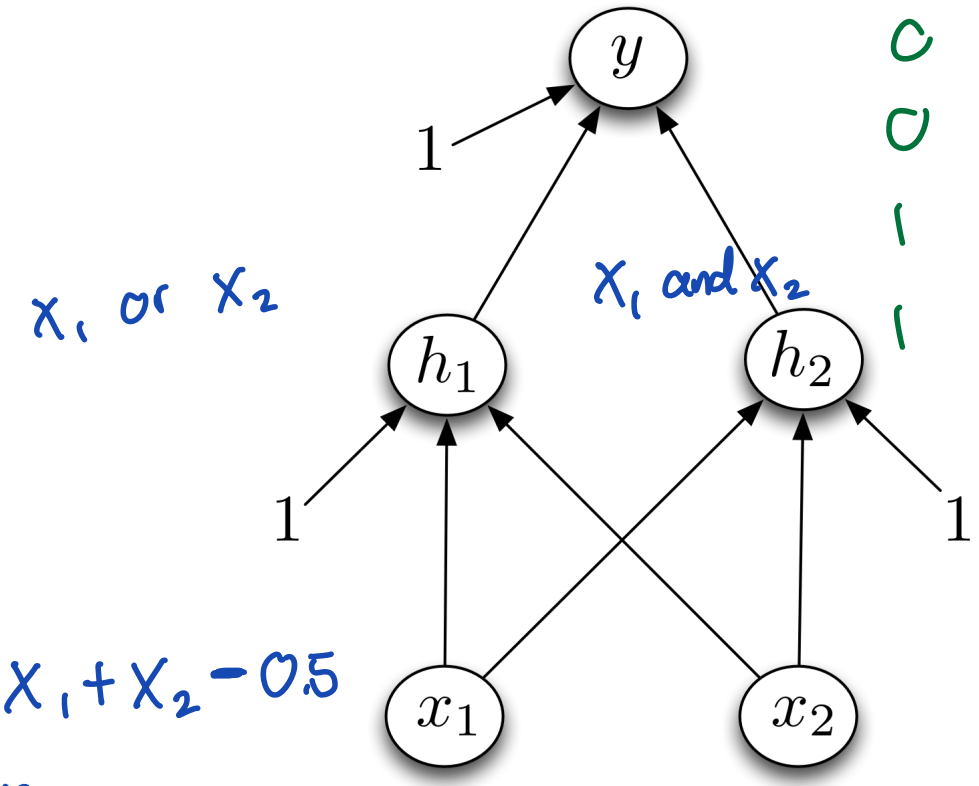

x_1 or x_2

x_1 and x_2

x_1	x_2	y	$\tilde{h}_1$	h_1	h_2
0	0	0	-0.5	0	0
0	1	1	0.5	1	0
1	0	1	0.5	1	0
1	1	0	1.5	1	1

$y = h_1 - h_2$

$\tilde{h}_1 = x_1 + x_2 - 0.5$

$h_1 = (\;\_\!\_\!\_\;)(\tilde{h}_1)$

Designing a Network to Classify XOR

h_1 computes $x_1 \vee x_2$

$$\mathbb{I}[x_1 + x_2 - 0.5 > 0]$$

1 if true
0 if false

h_2 computes $x_1 \wedge x_2$

$$\mathbb{I}[x_1 + x_2 - 1.5 > 0]$$

y computes $h_1 \wedge (\neg h_2) = x_1 \oplus x_2$

$$\mathbb{I}[h_1 - h_2 - 0.5 > 0]$$
$$\equiv \mathbb{I}[h_1 + (1 - h_2) - 1.5 > 0]$$

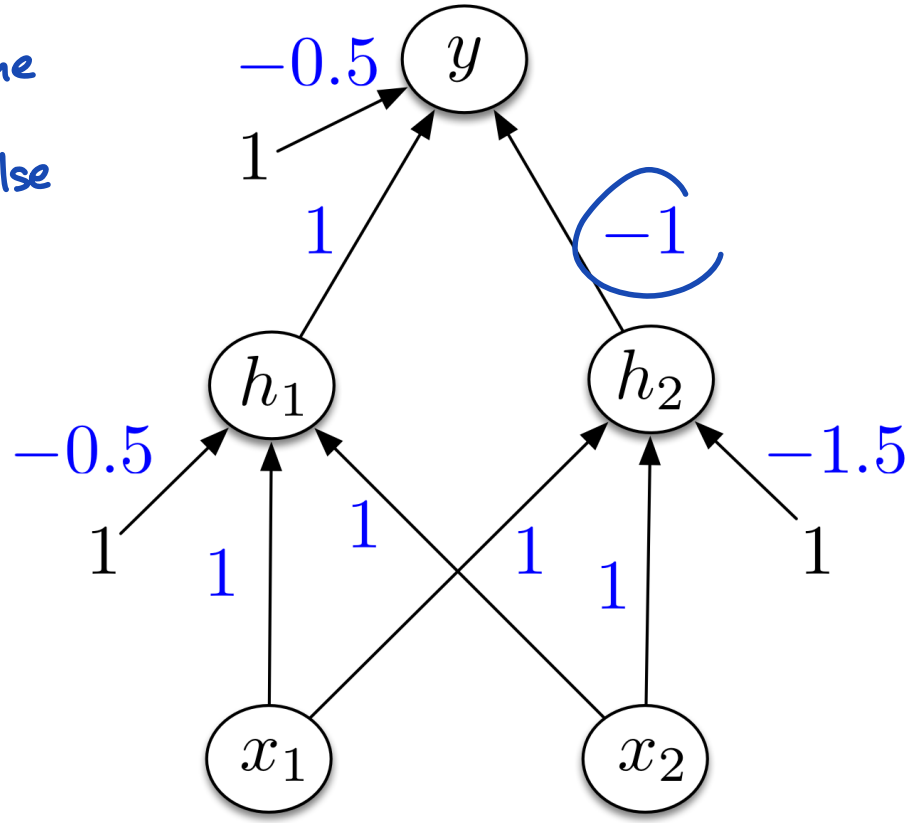

Universality for Binary Inputs and Targets

- Hard threshold hidden units, linear output
- Strategy: 2^D hidden units, each of which responds to one particular input configuration

hidden unit activates for one data point

x_1	x_2	x_3	t
	$\vdots$		$\vdots$
-1	-1	1	-1
-1	1	-1	1
-1	1	1	1
	$\vdots$		$\vdots$

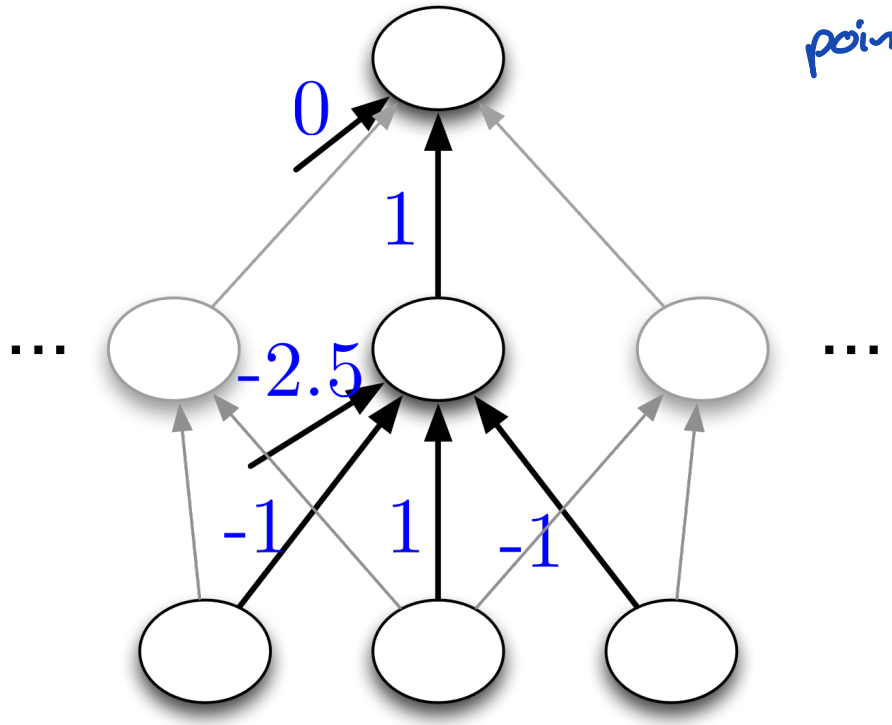

- Only requires one hidden layer, though it is extremely wide.

Expressivity of the Logistic Activation Function

- What about the logistic activation function?
- Approximate a hard threshold by scaling up w and b.

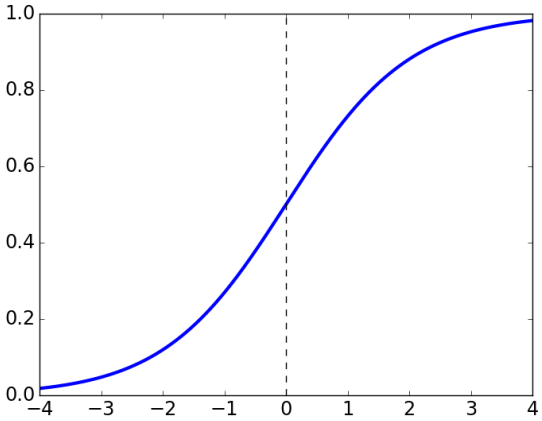

$y = \sigma(x)$

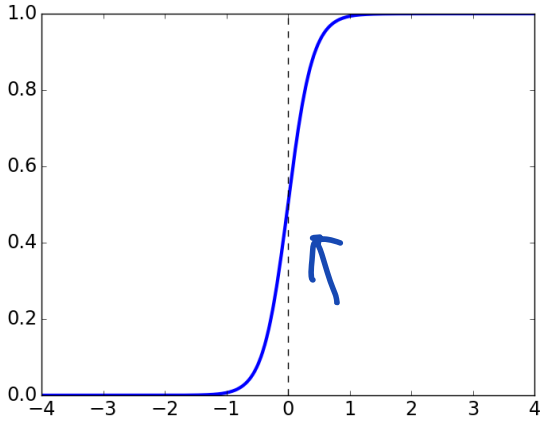

$y = \sigma(5x)$

- Logistic units are differentiable, so we can learn weights with gradient descent.

What is Expressivity Good For?

- May need a very large network to represent a function.
- Non-trivial to learn the weights that represent a function.
- If you can learn any function, over-fitting is potentially a serious concern!

  For the polynomial feature mappings, expressivity increases with the degree M, eventually allowing multiple perfect fits to the training data. This motivated L^2 regularization.

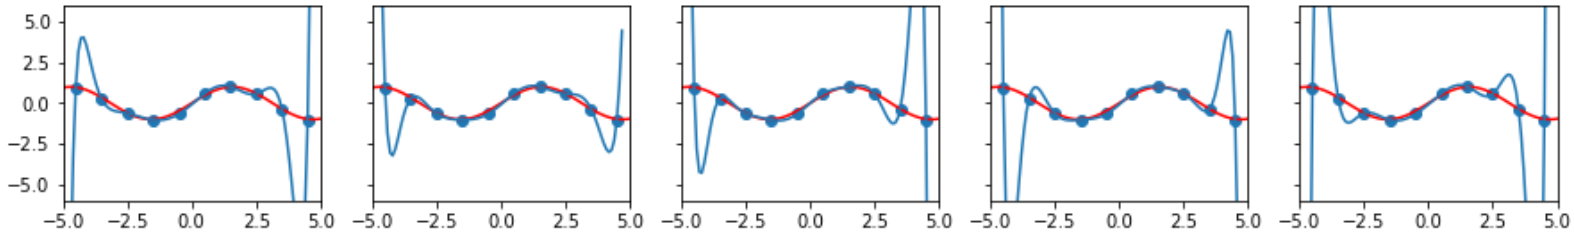

- Do neural networks over-fit and how can we regularize them?

Regularization and Over-fitting for Neural Networks

- The topic of over-fitting (when & how it happens, how to regularize, etc.) for neural networks is not well-understood, even by researchers!
  - In principle, you can always apply L^2 regularization.
  - You will learn more in CSC413.

dropout

batch normalization

- A common approach is early stopping, or stopping training early, because over-fitting typically increases as training progresses.

in practice, curve is usually stochastic

- Don't add an explicit $\mathcal{R}(\boldsymbol{\theta})$ term to our cost.

Conclusion

- Multi-class classification
- Selecting good metrics to track performance in models
- From linear to non-linear models