

27: Intro to Deep Learning + Wrap-up

Lisa Yan

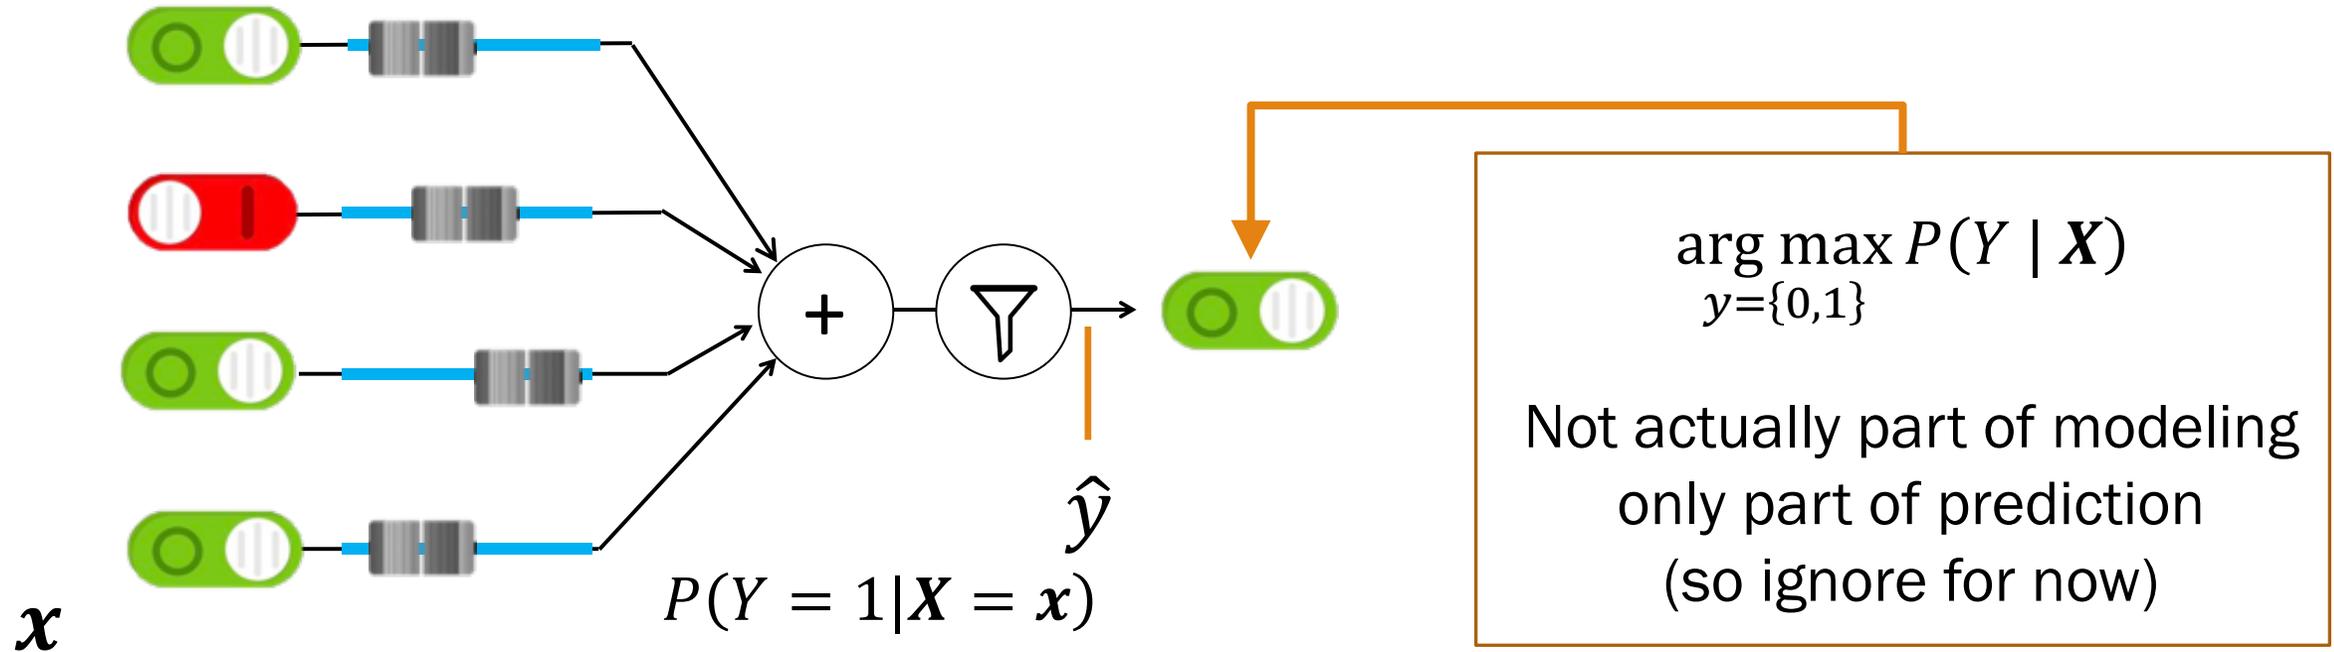
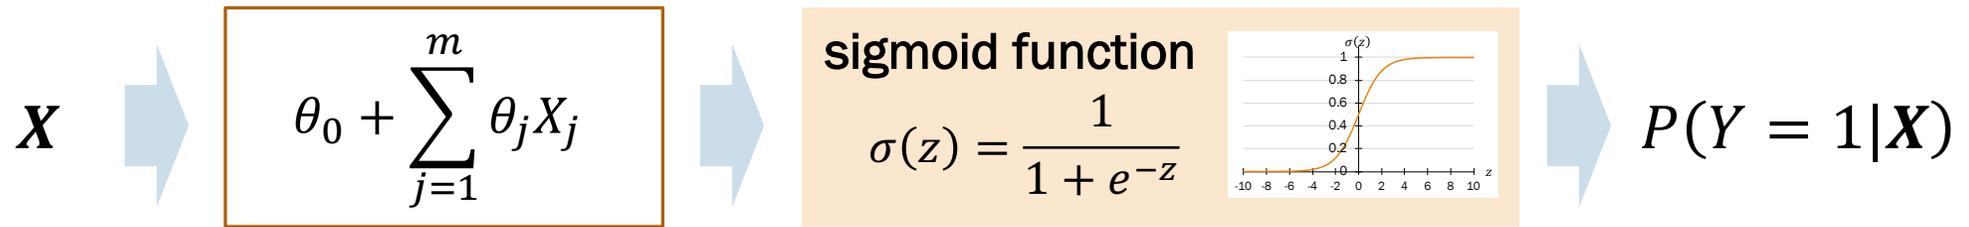
June 8, 2020

Quick slide reference

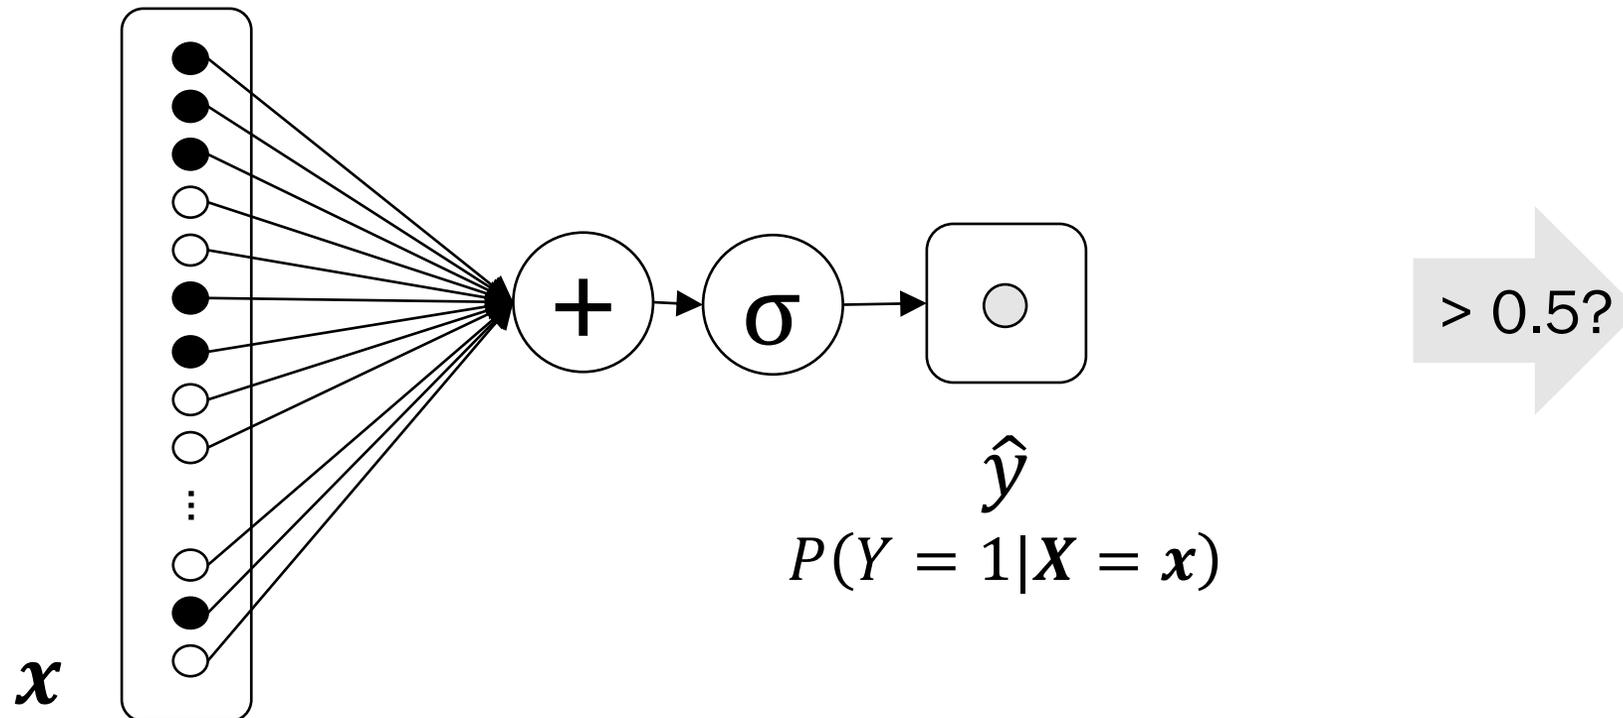
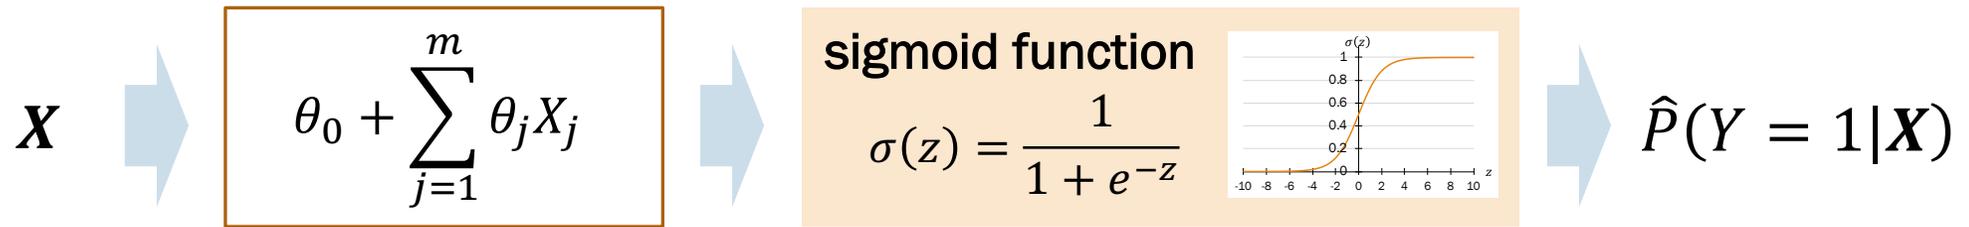
3	0.1% of Deep Learning	LIVE
39	Beyond the Basics	LIVE
63	CS109 Wrap-Up	LIVE

Deep Learning

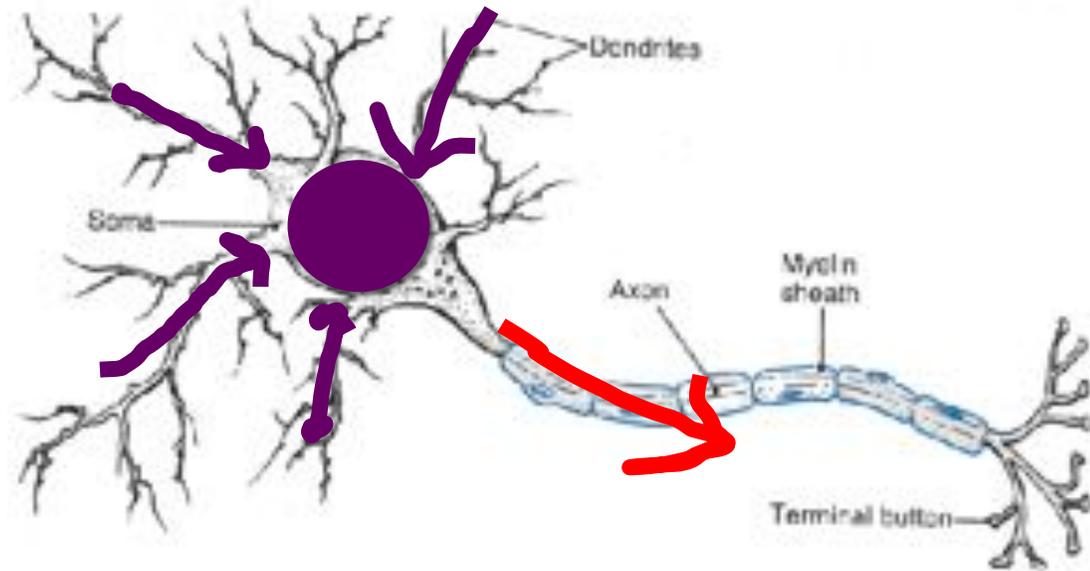
Logistic Regression Model



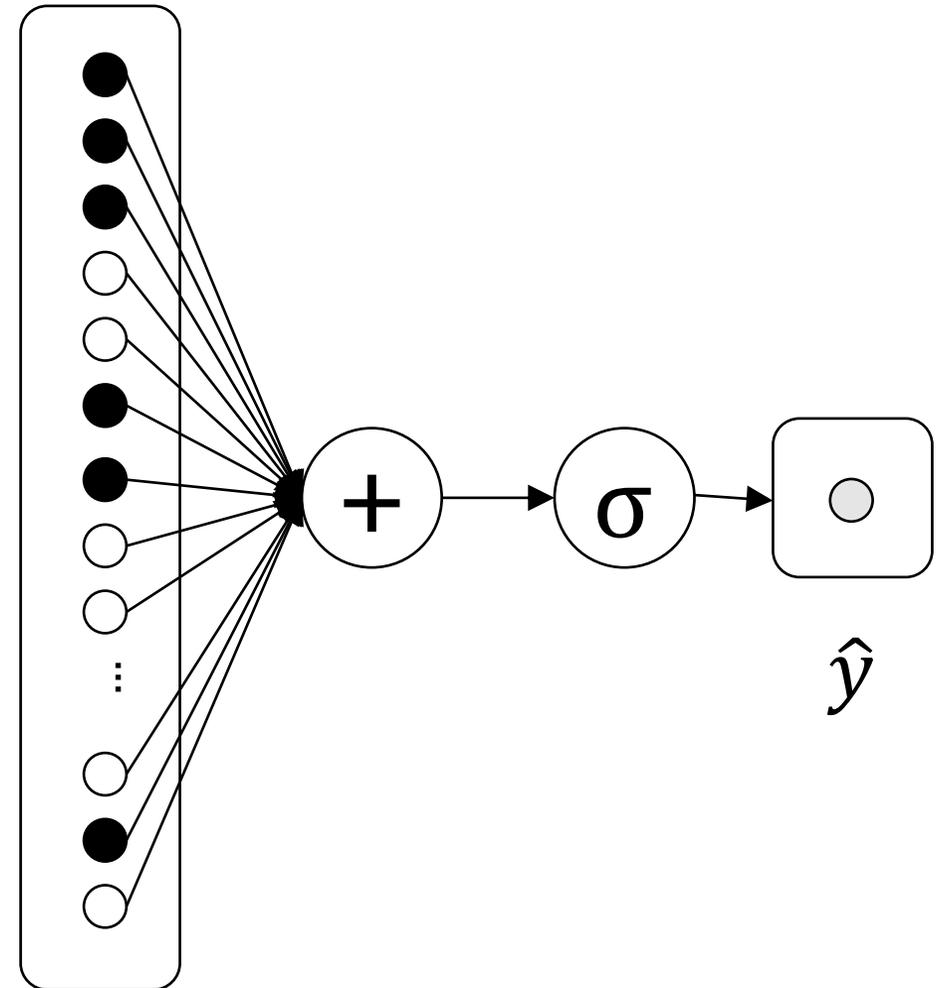
Big ideas from Logistic Regression



One neuron = One logistic regression

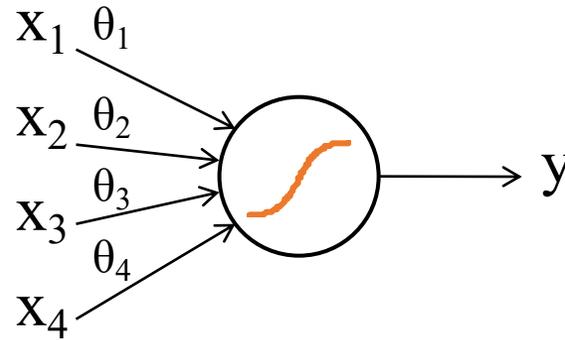
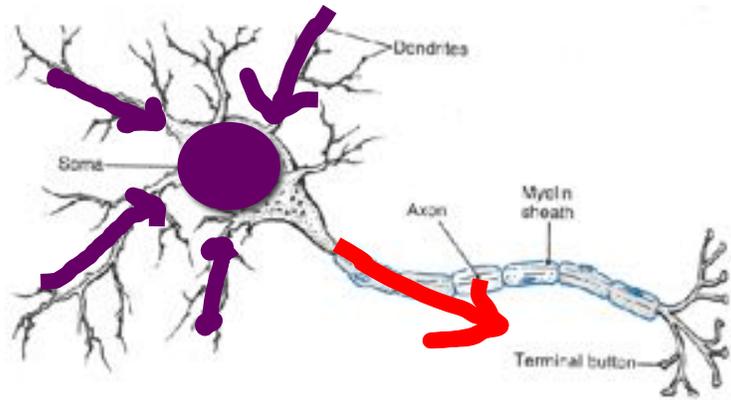


=



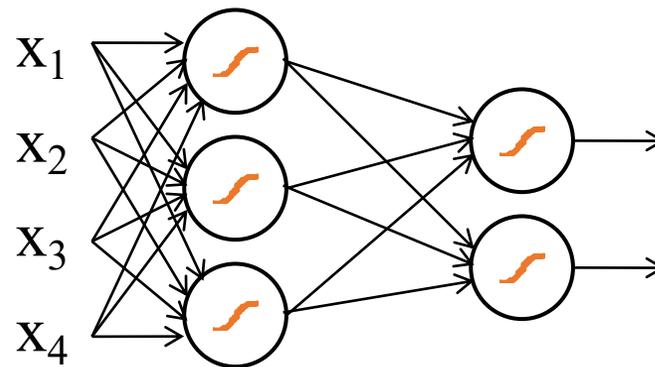
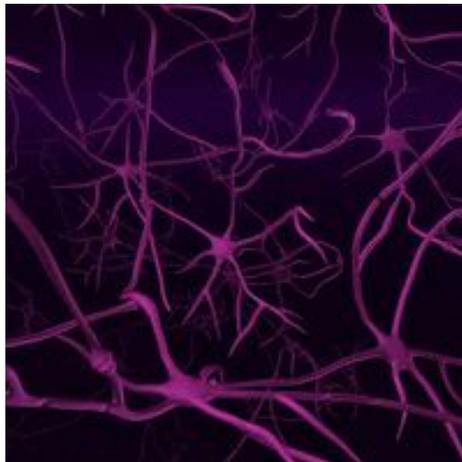
Biological basis for neural networks

A neuron



One neuron =
one logistic
regression

Your brain



Neural network =
many logistic
regressions

(actually, probably someone else's brain)

Innovations in deep learning



AlphaGO (2016)

Deep learning (neural networks) is the core idea behind the current revolution in AI.

Computers making art

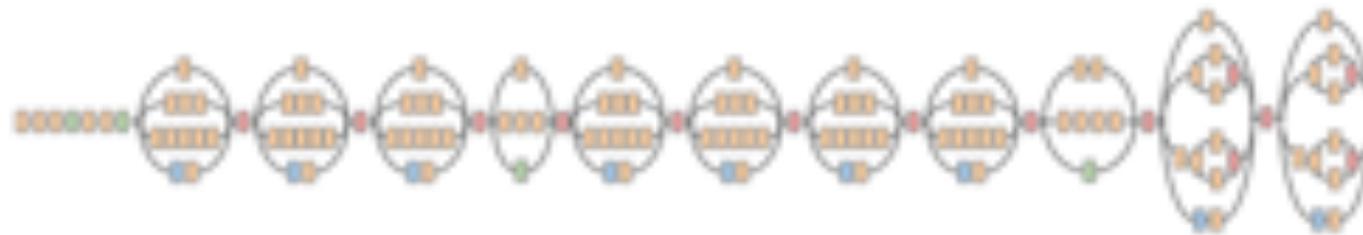


Detecting skin cancer

Skin Lesion Image



Deep Convolutional Neural Network (Inception-v3)



Training Classes (757)

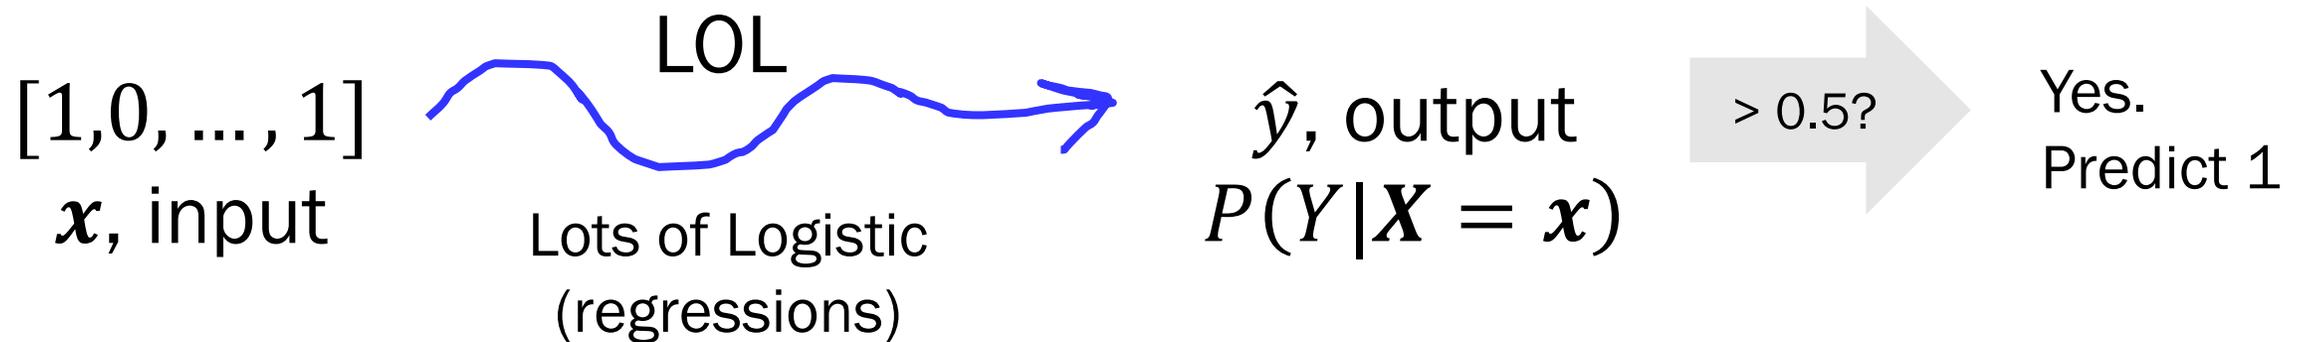


Esteva, Andre, et al. "Dermatologist-level classification of skin cancer with deep neural networks." *Nature* 542.7639 (2017): 115-118.

Deep learning

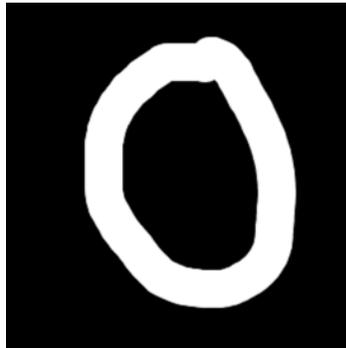
def **Deep learning** is
maximum likelihood estimation
with neural networks.

def A **neural network** is
(at its core) many logistic
regression pieces stacked on
top of each other.



Digit recognition example

Input image



Input feature vector

$$\mathbf{x}^{(i)} = [0,0,0,0, \dots, 1,0,0,1, \dots, 0,0,1,0]$$

Output label

$$y^{(i)} = 0$$

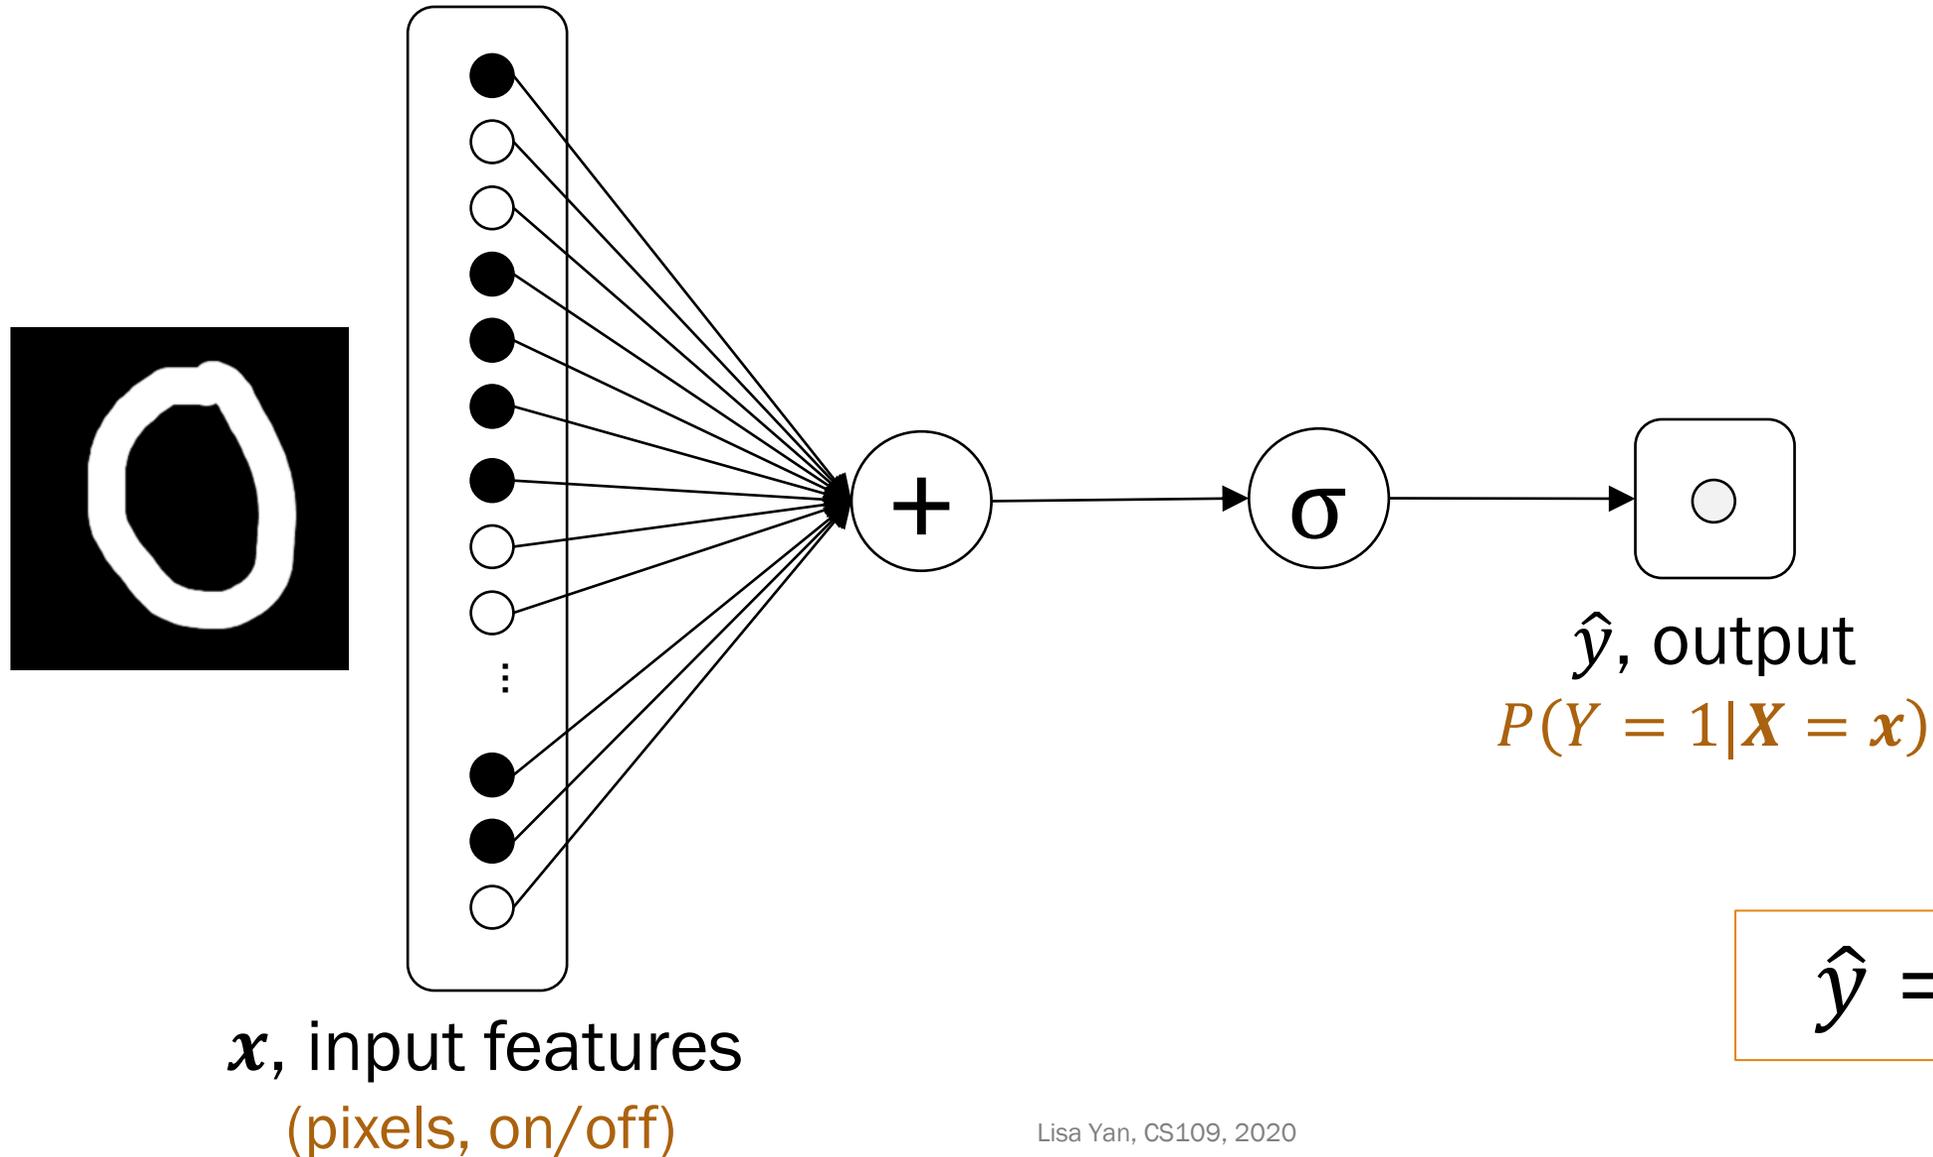


$$\mathbf{x}^{(i)} = [0,0,1,1, \dots, 0,1,1,0, \dots, 0,1,0,0]$$

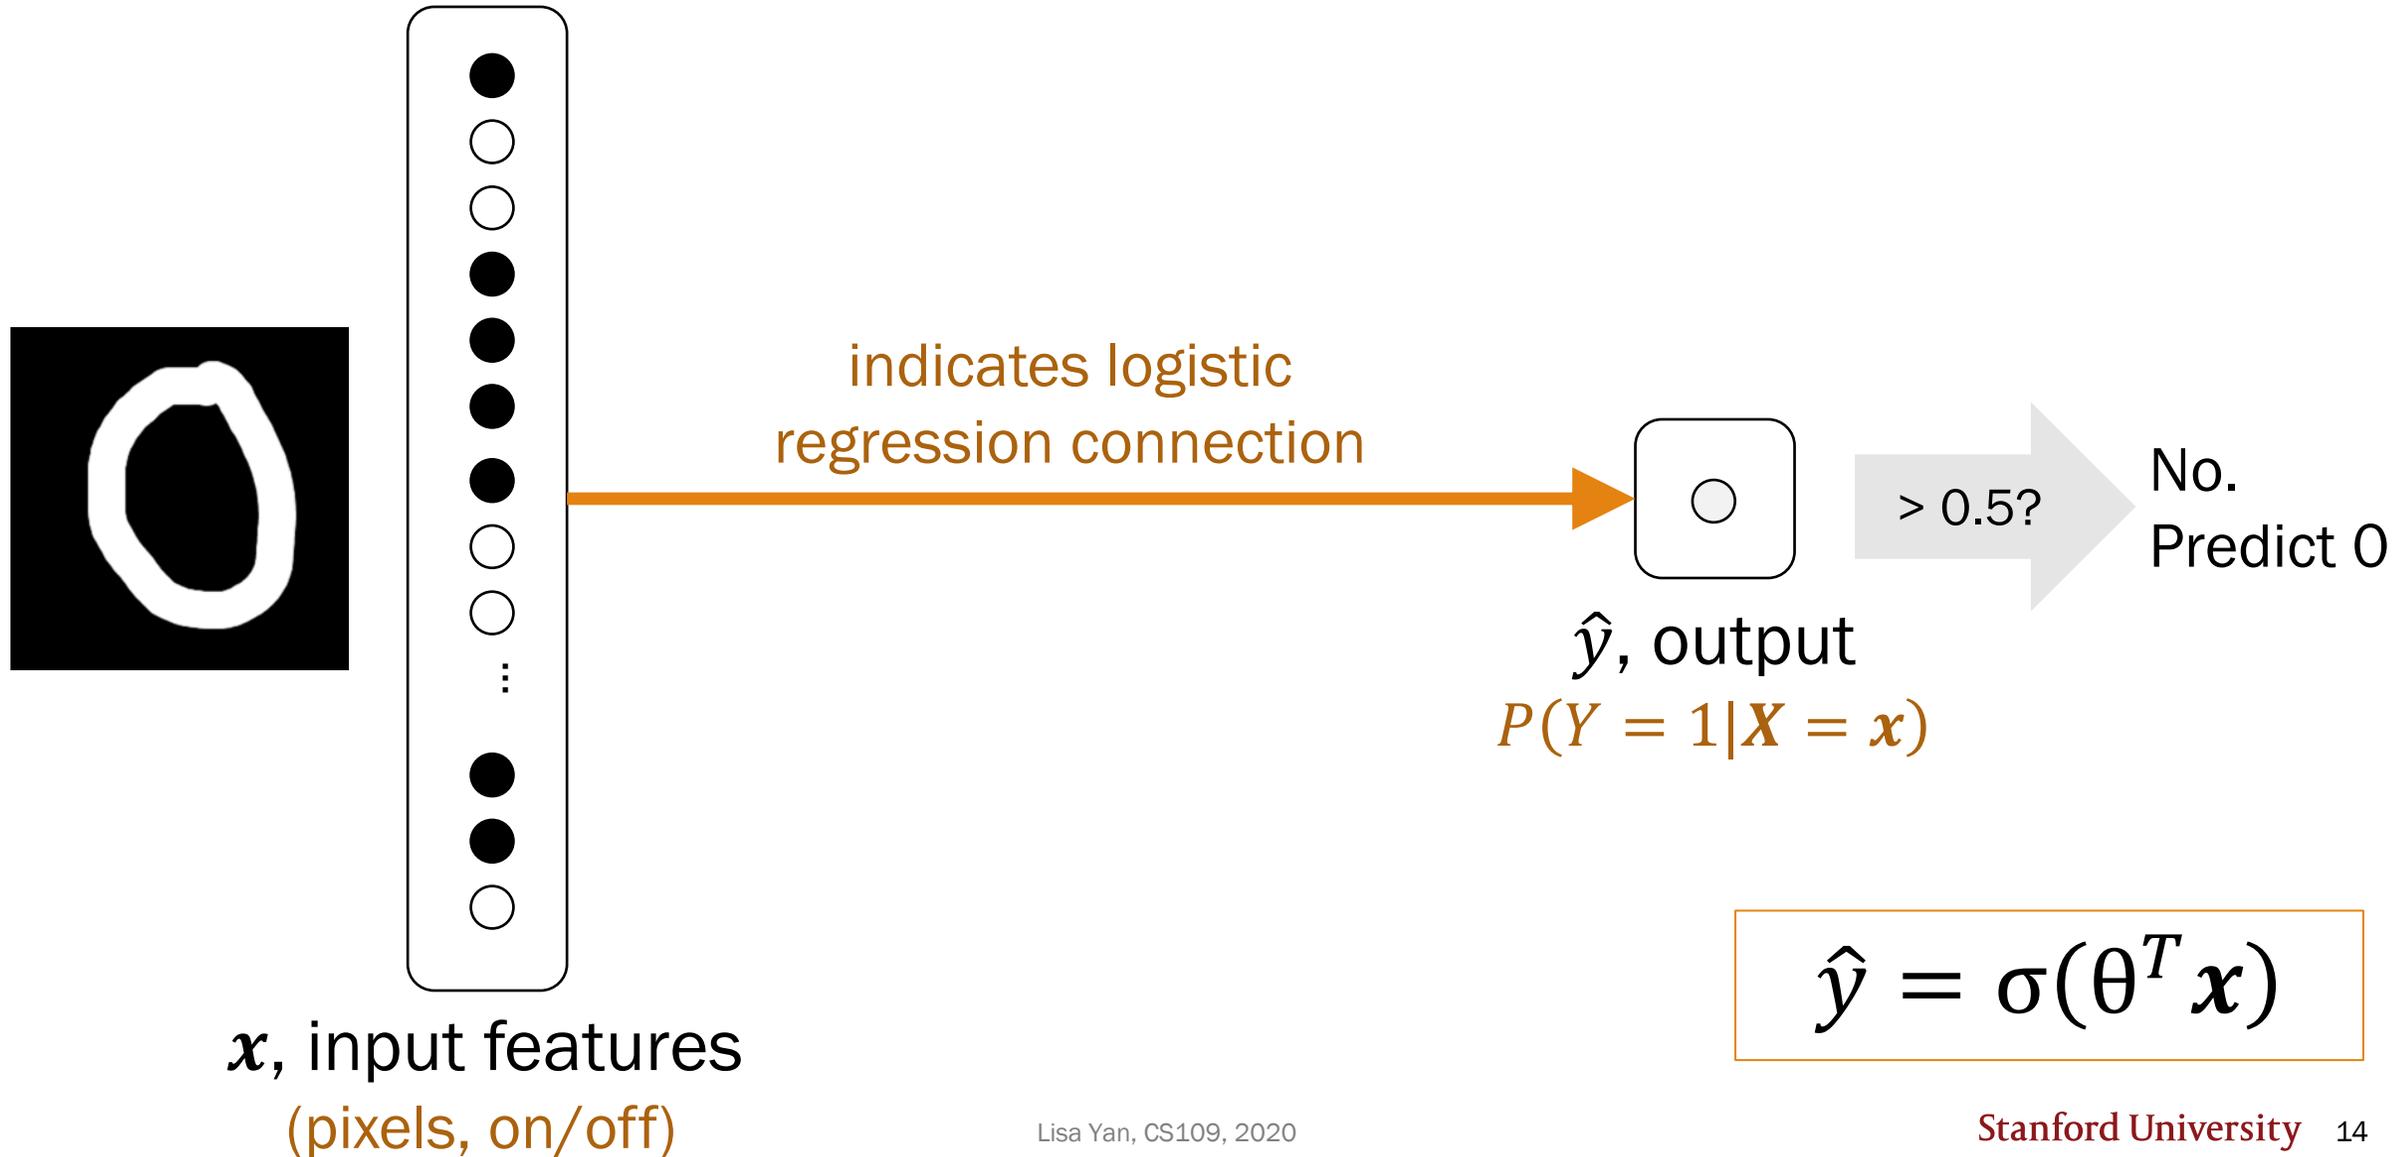
$$y^{(i)} = 1$$

We make feature vectors from (digitized) pictures of numbers.

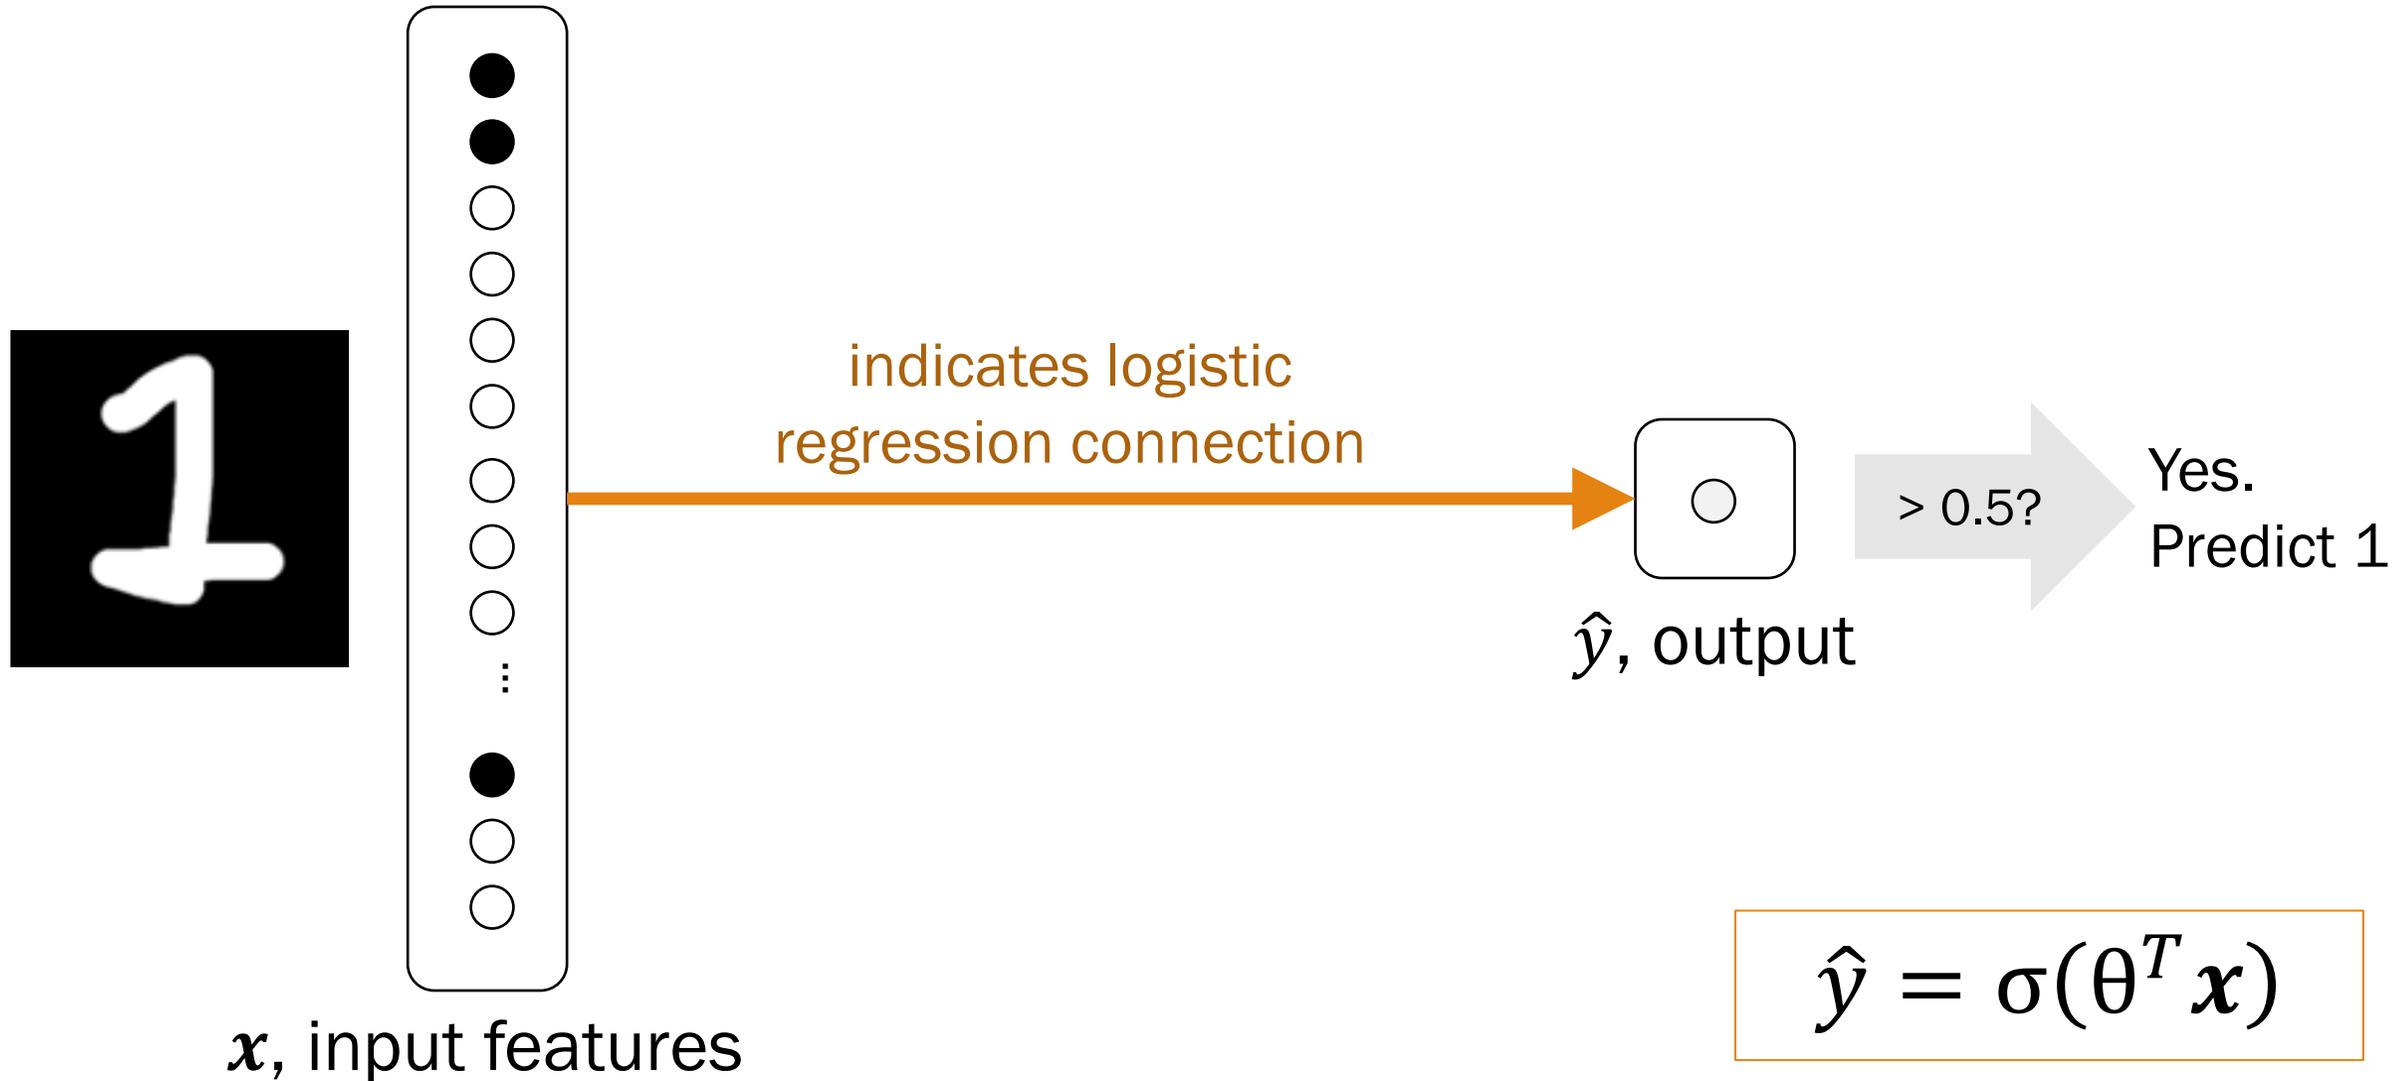
Logistic Regression



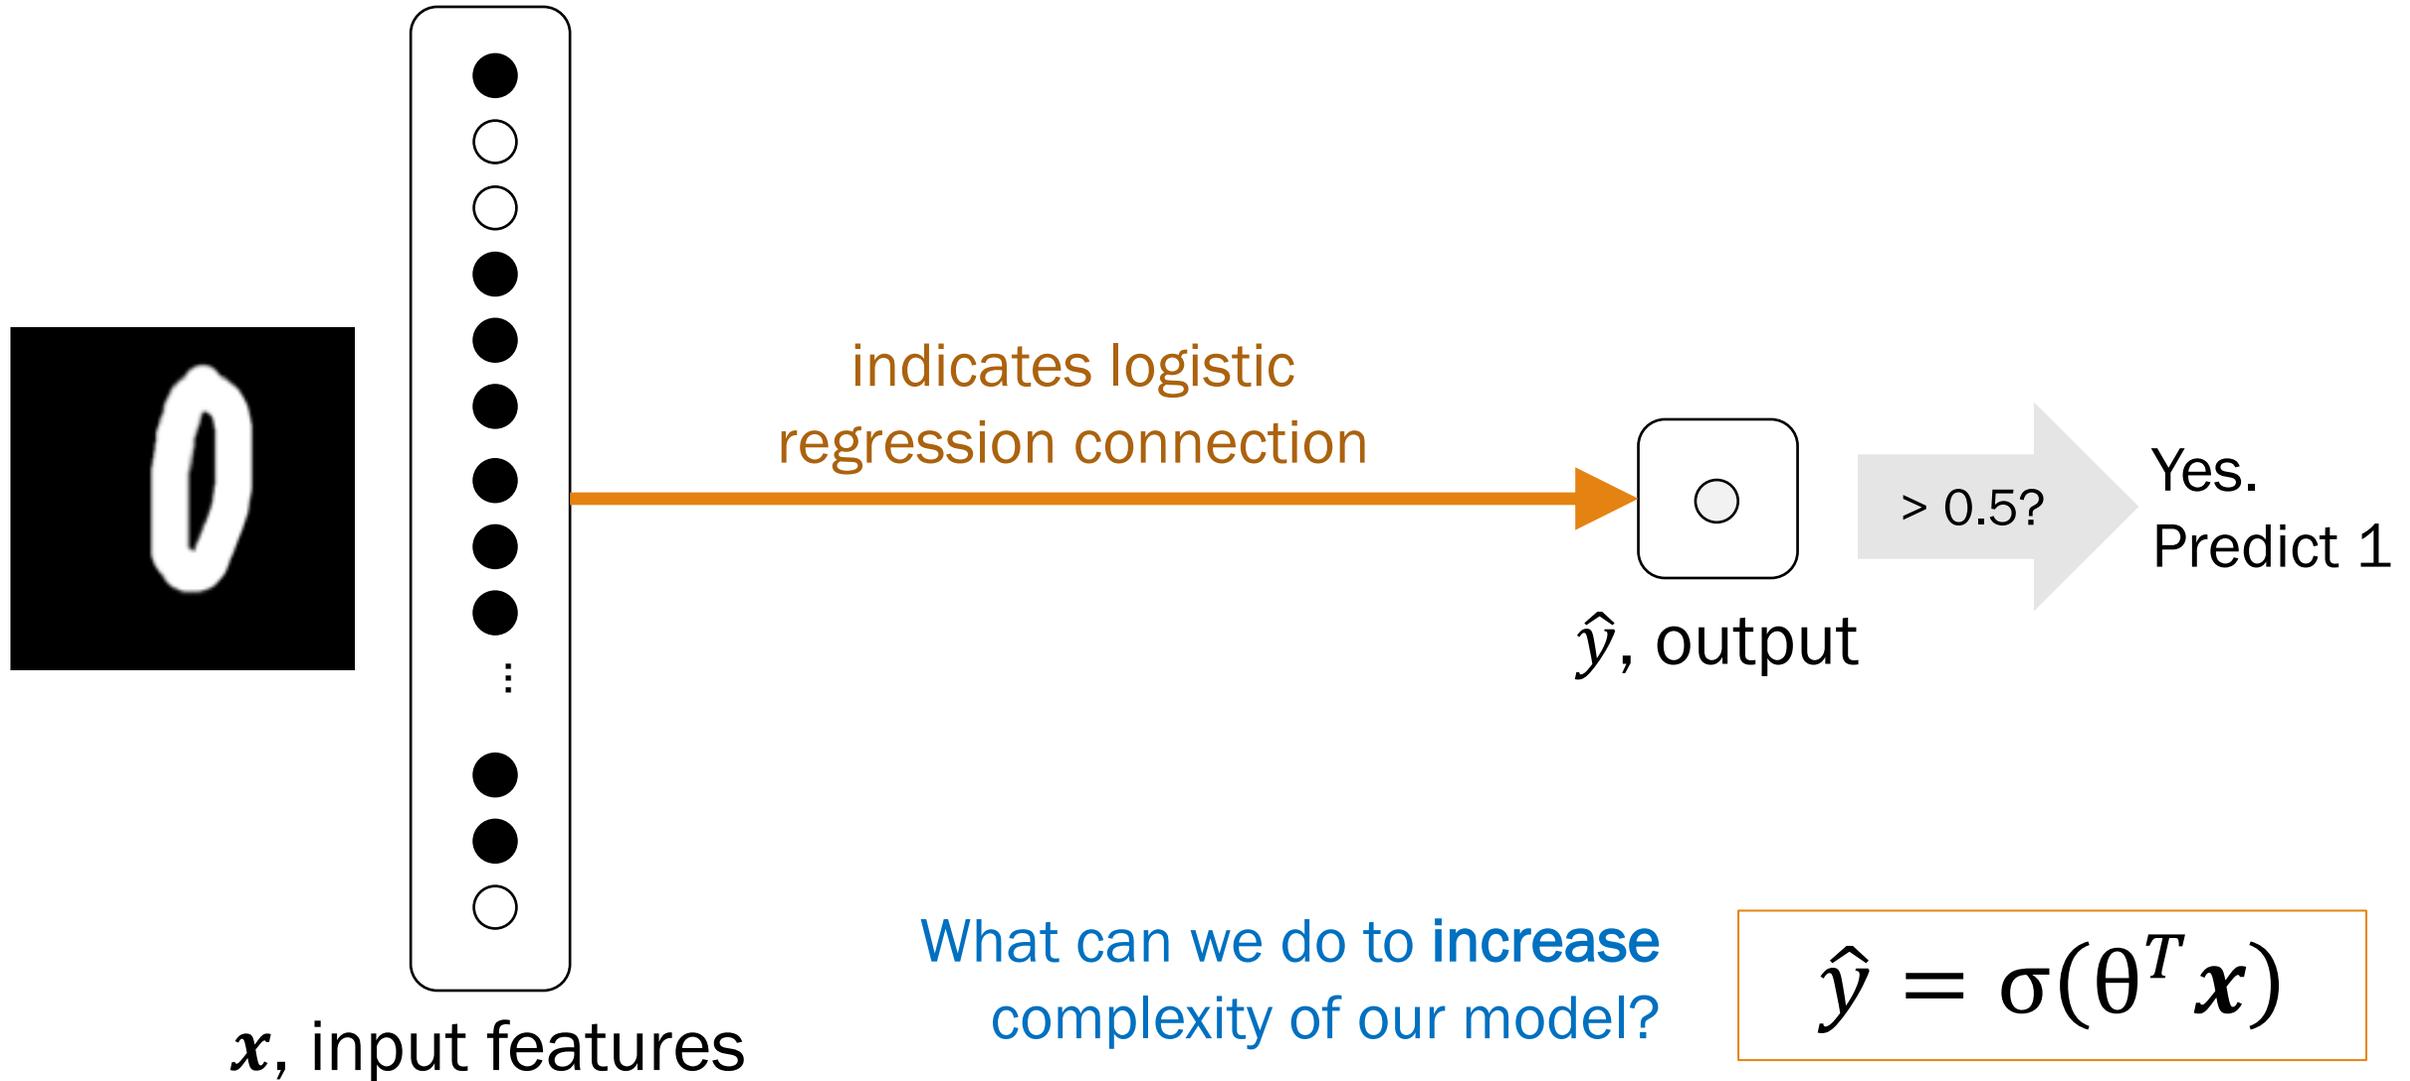
Logistic Regression



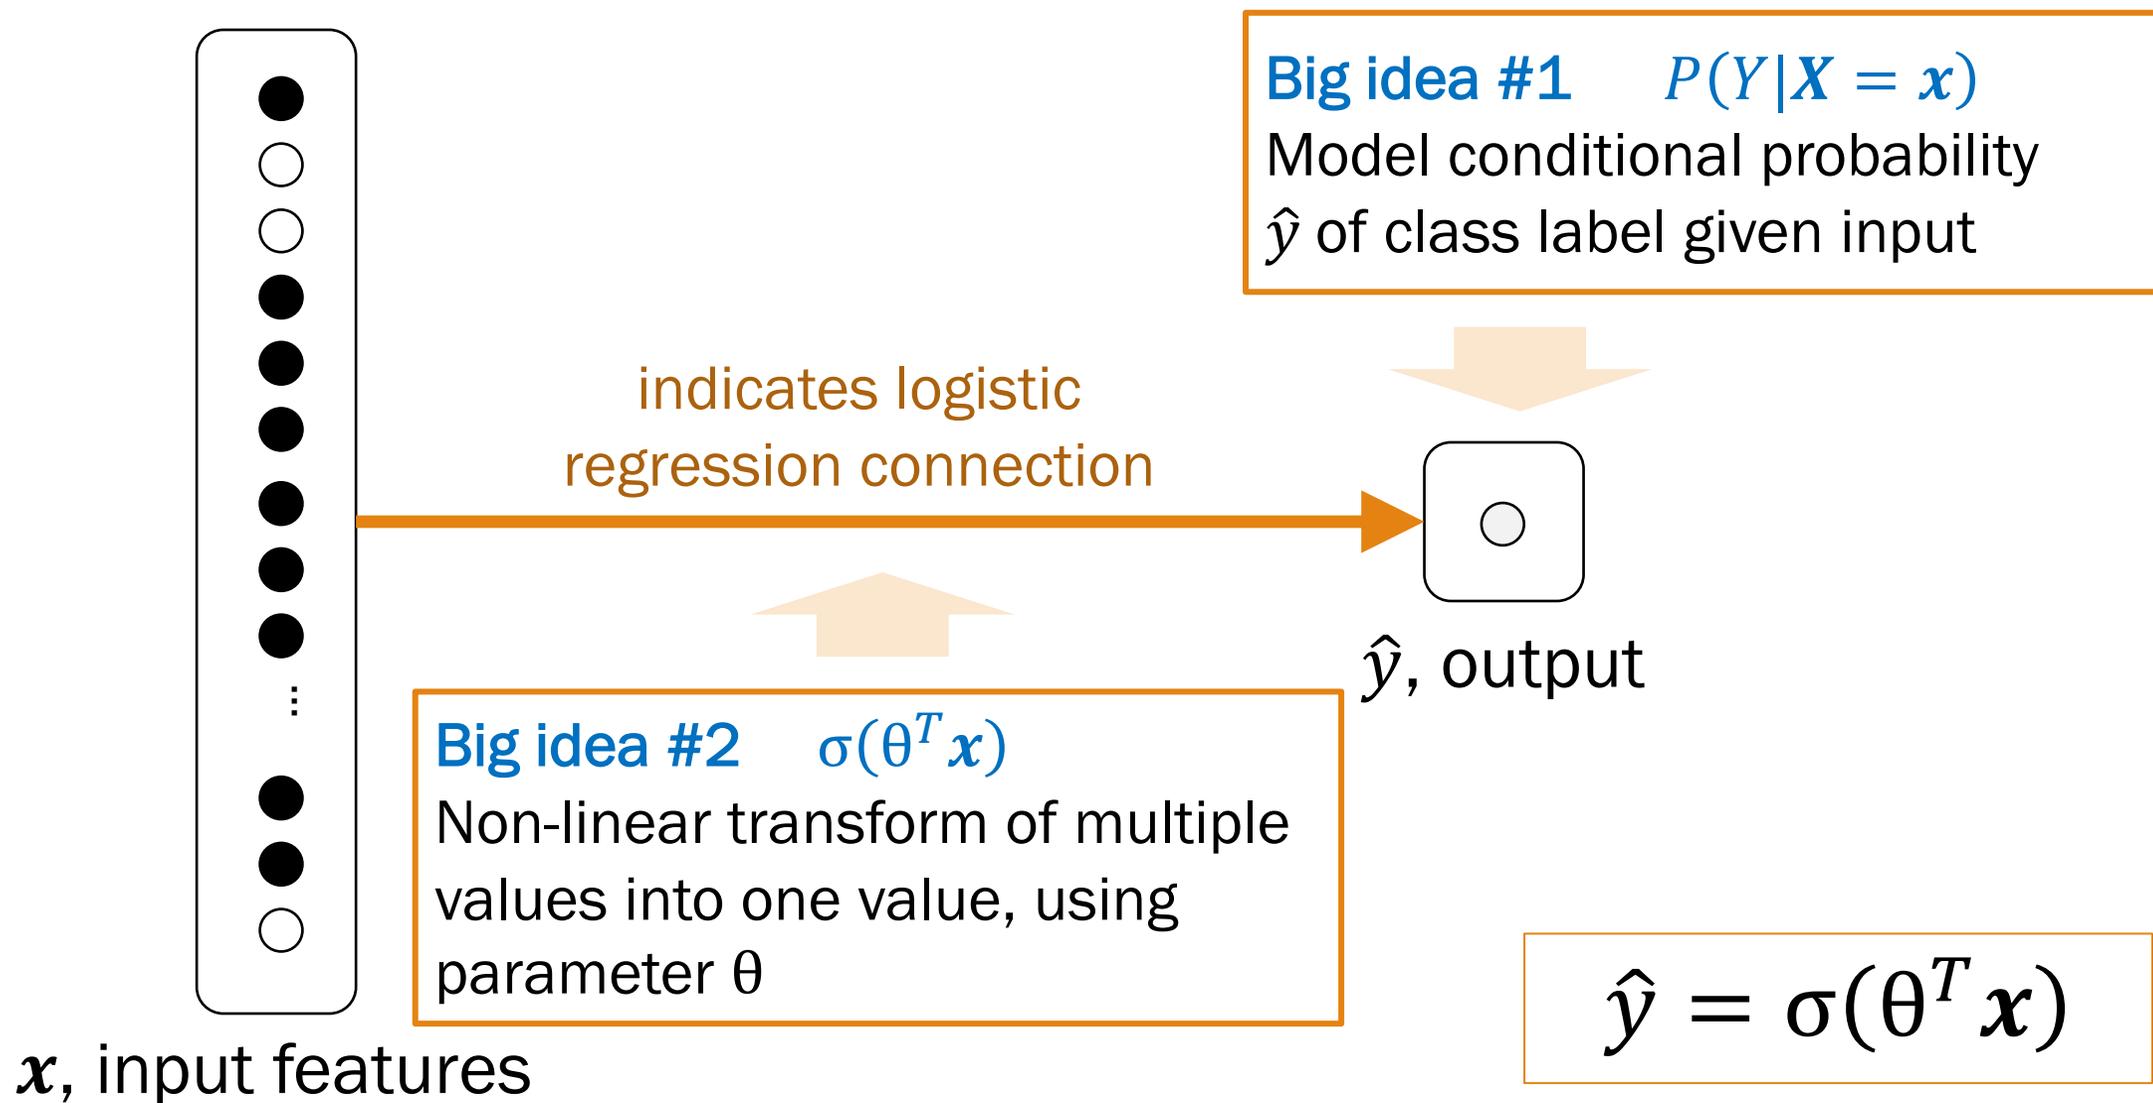
Logistic Regression



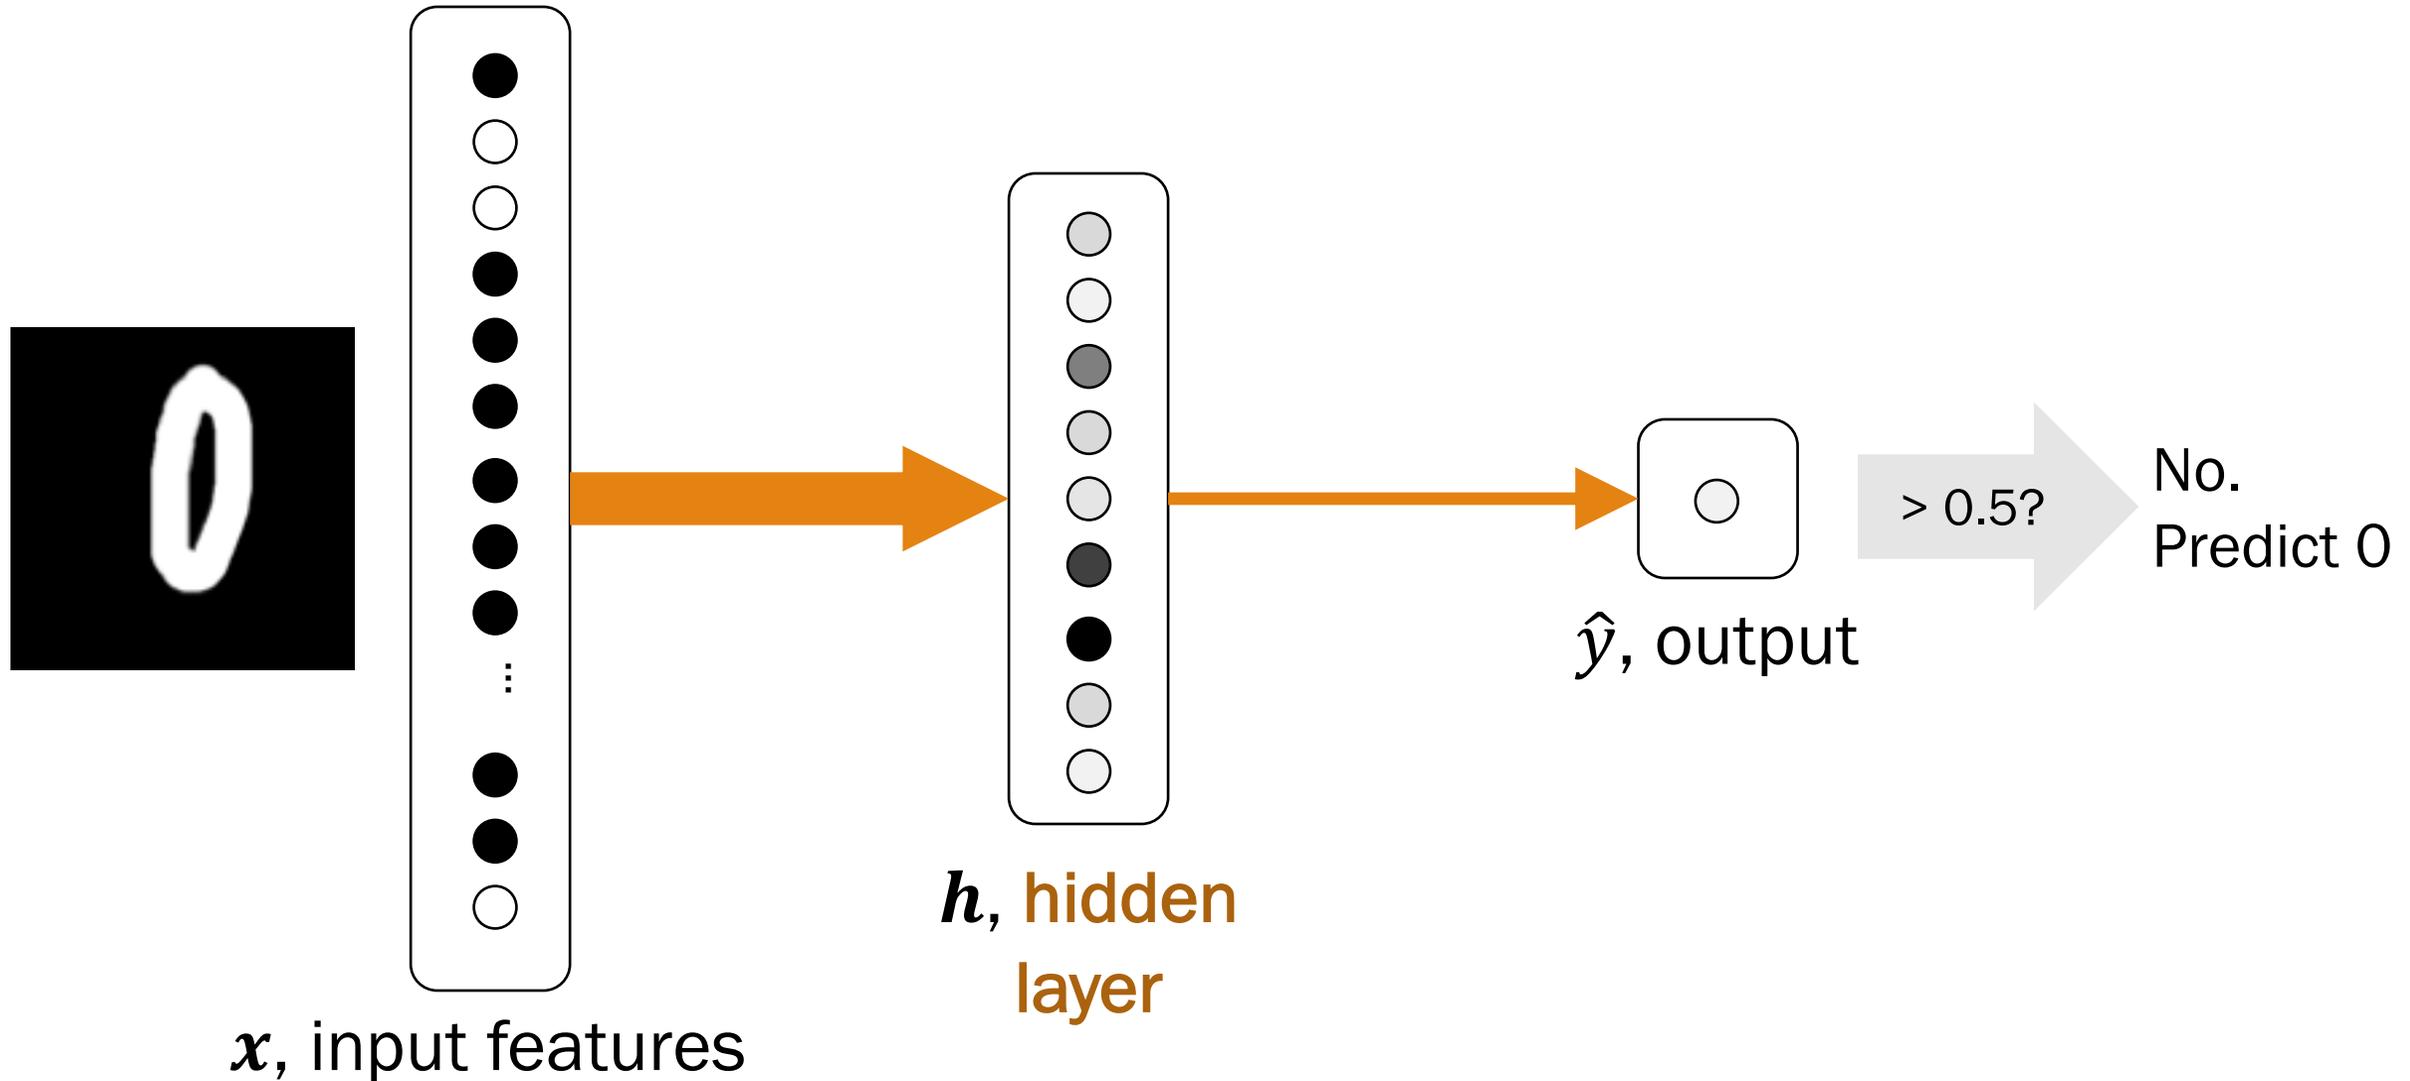
Logistic Regression



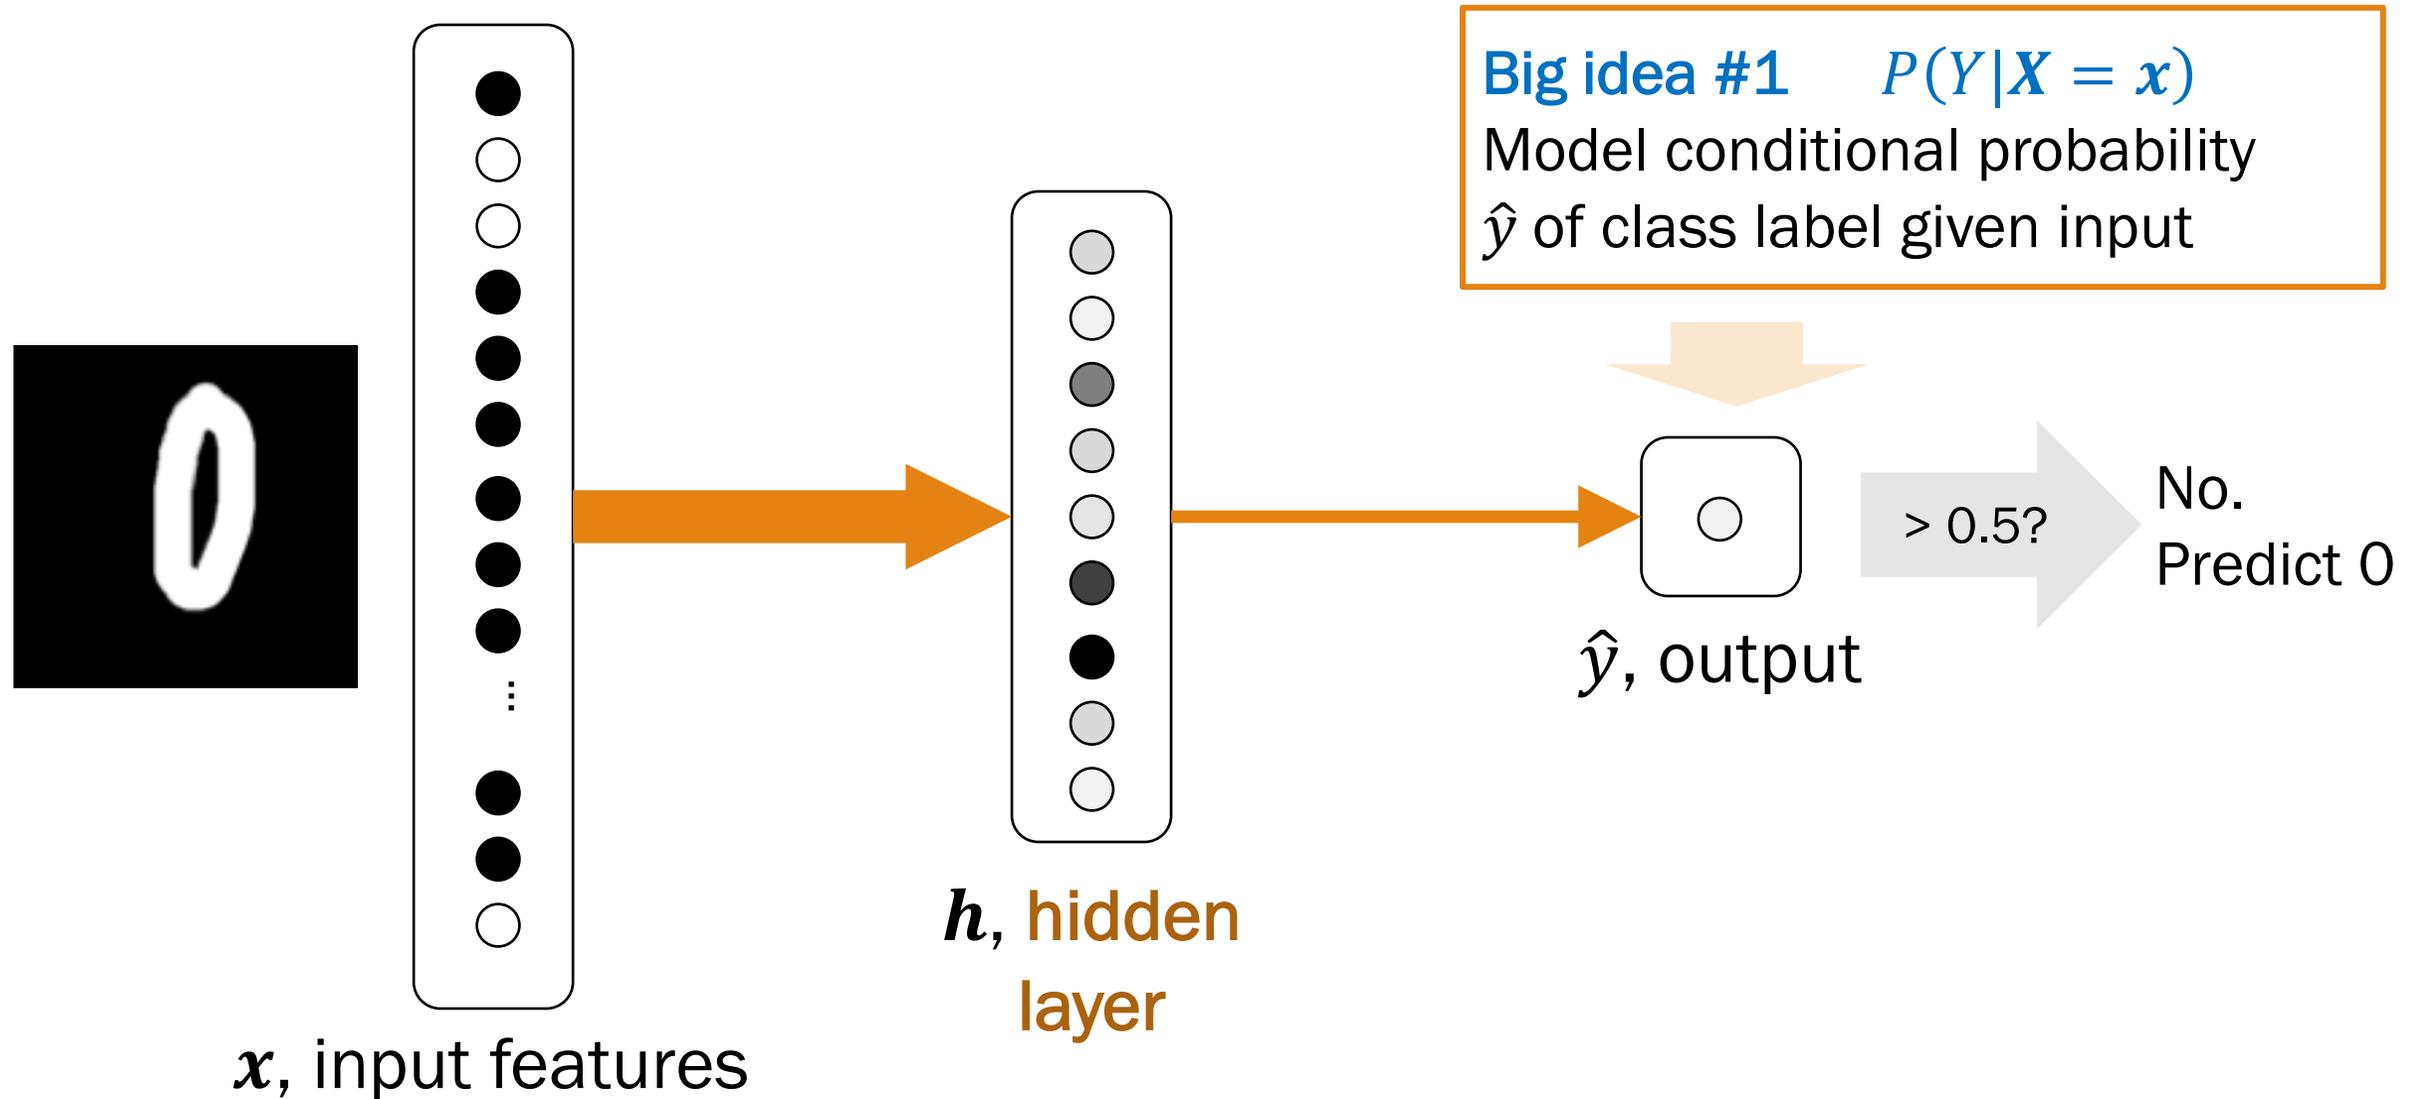
Big ideas from logistic regression



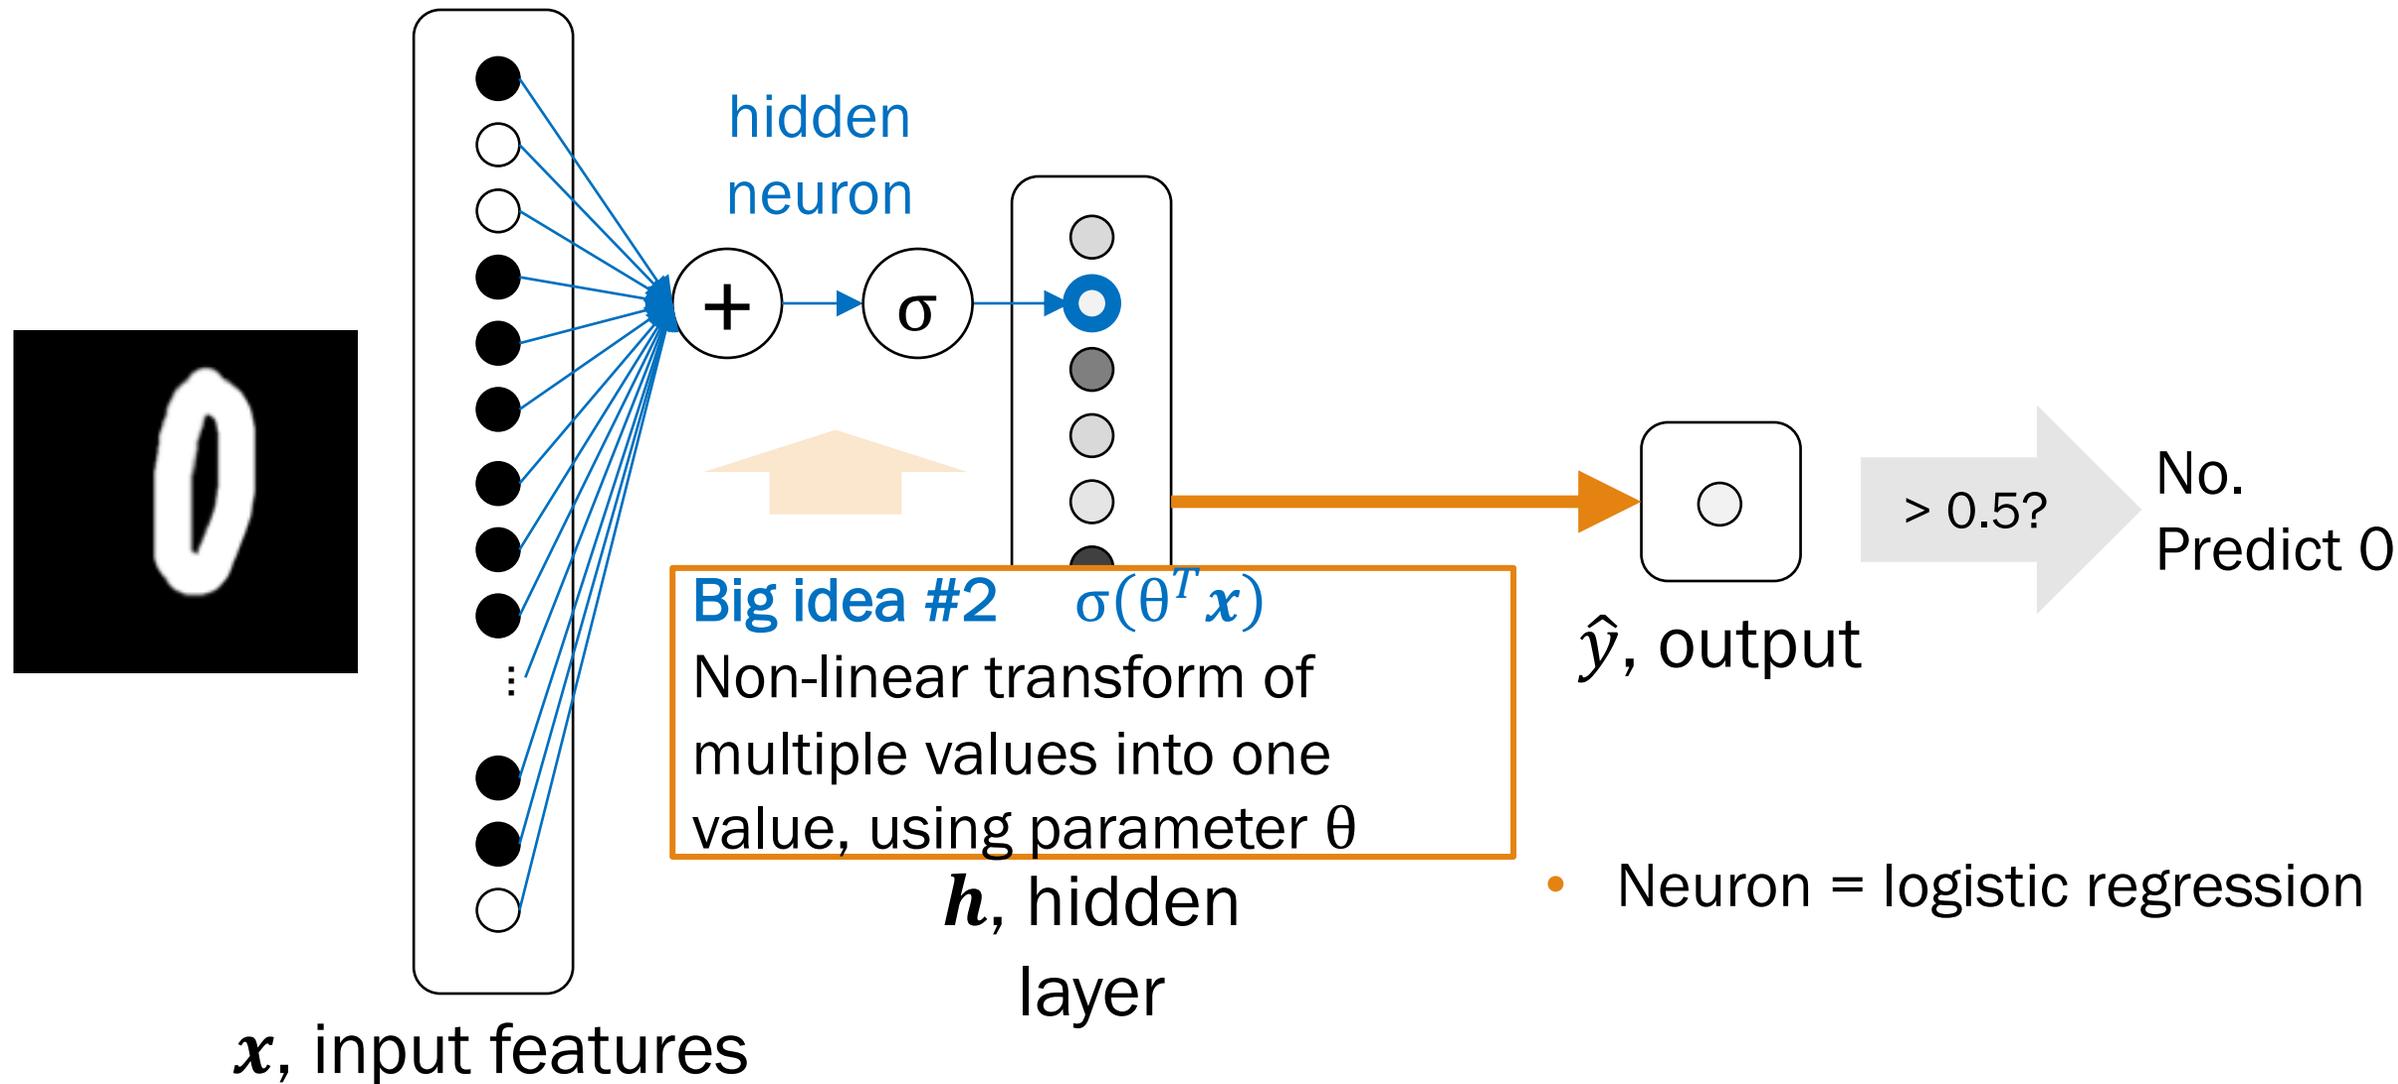
Neural network



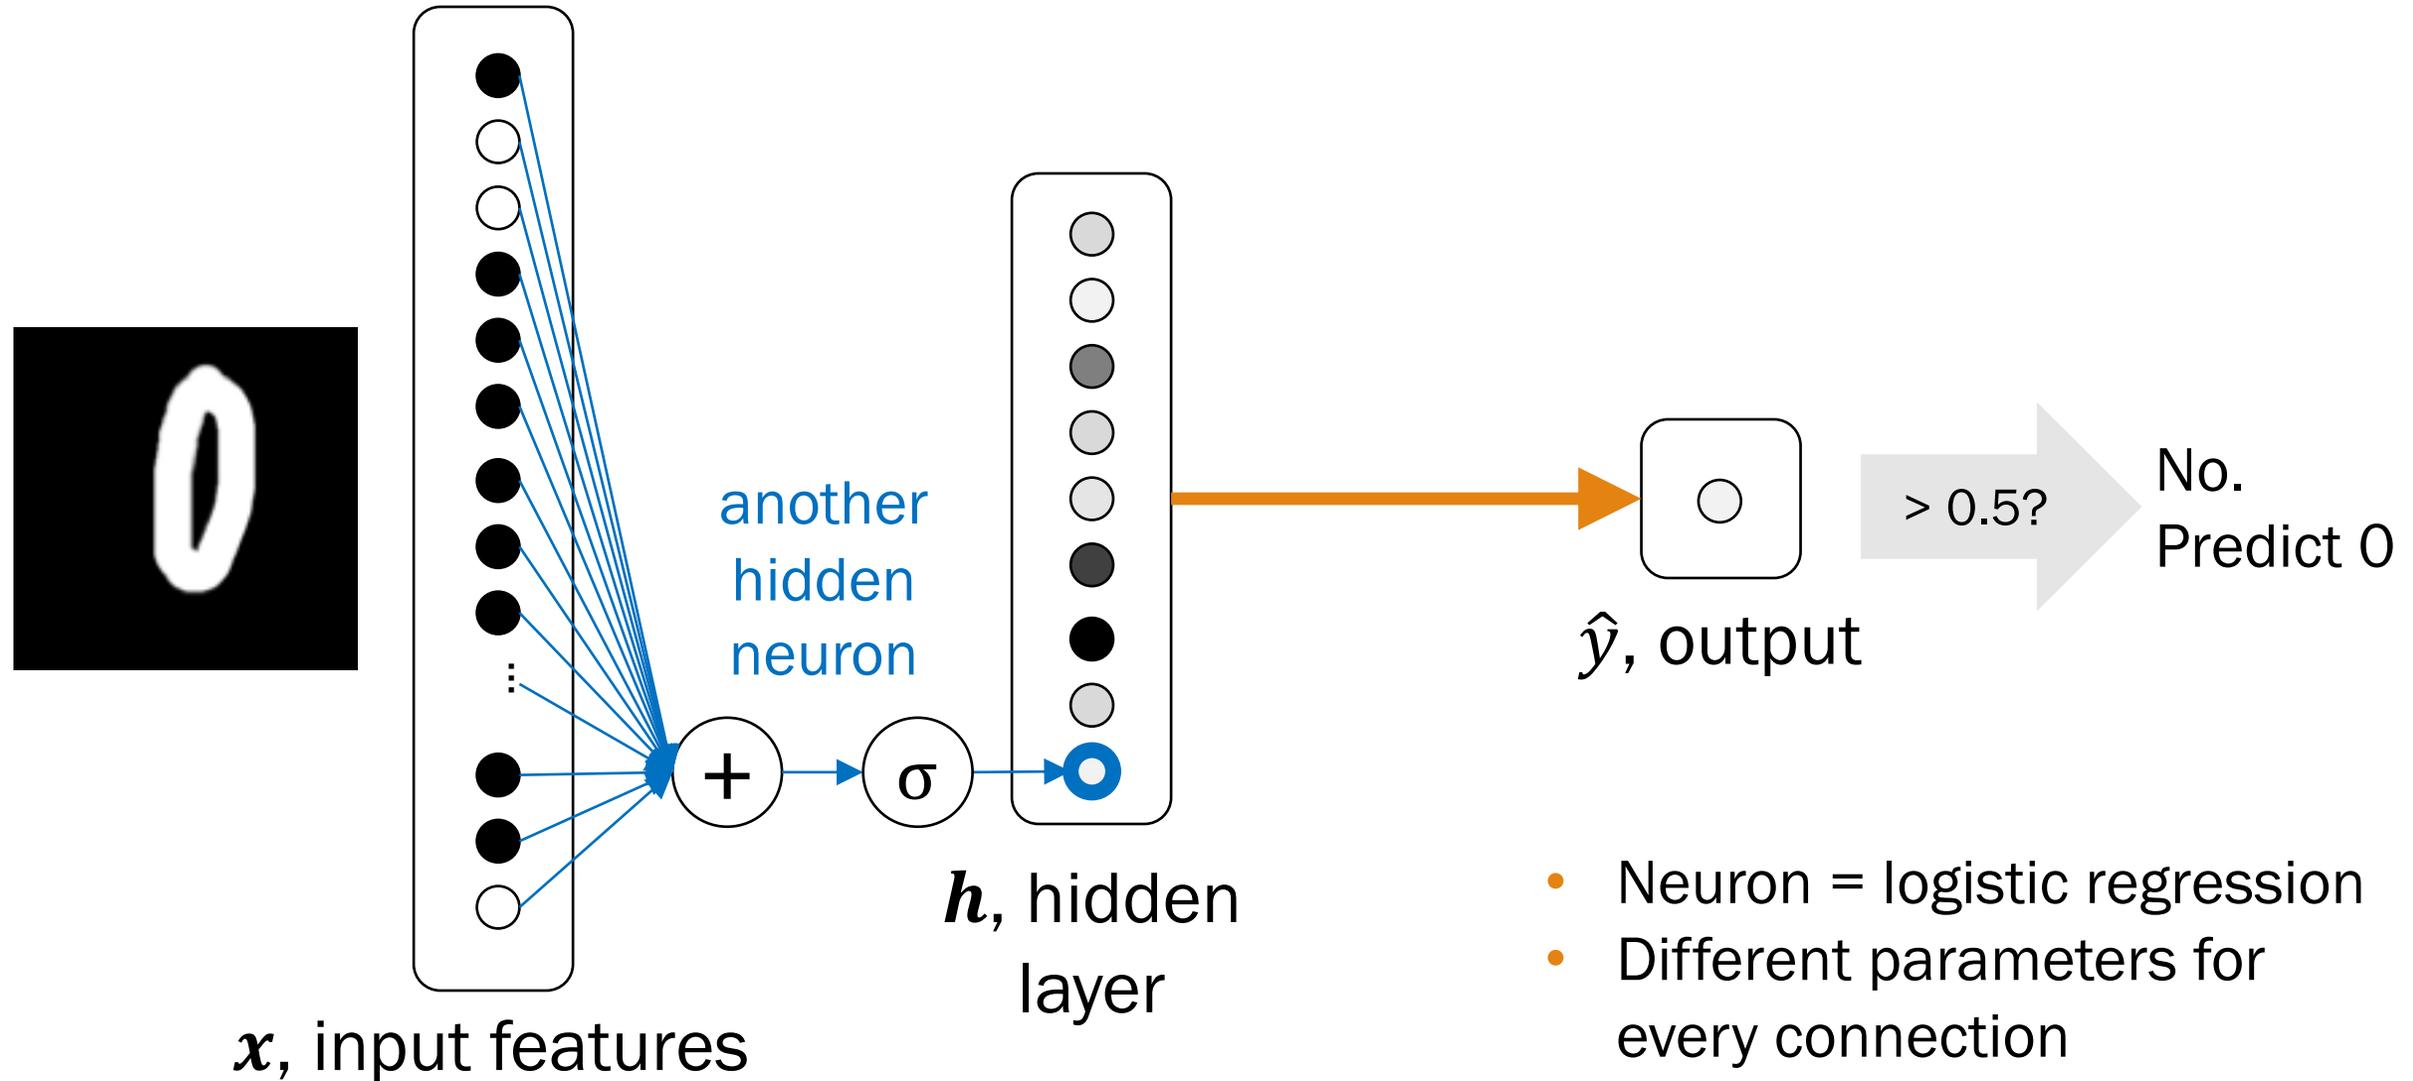
Neural network



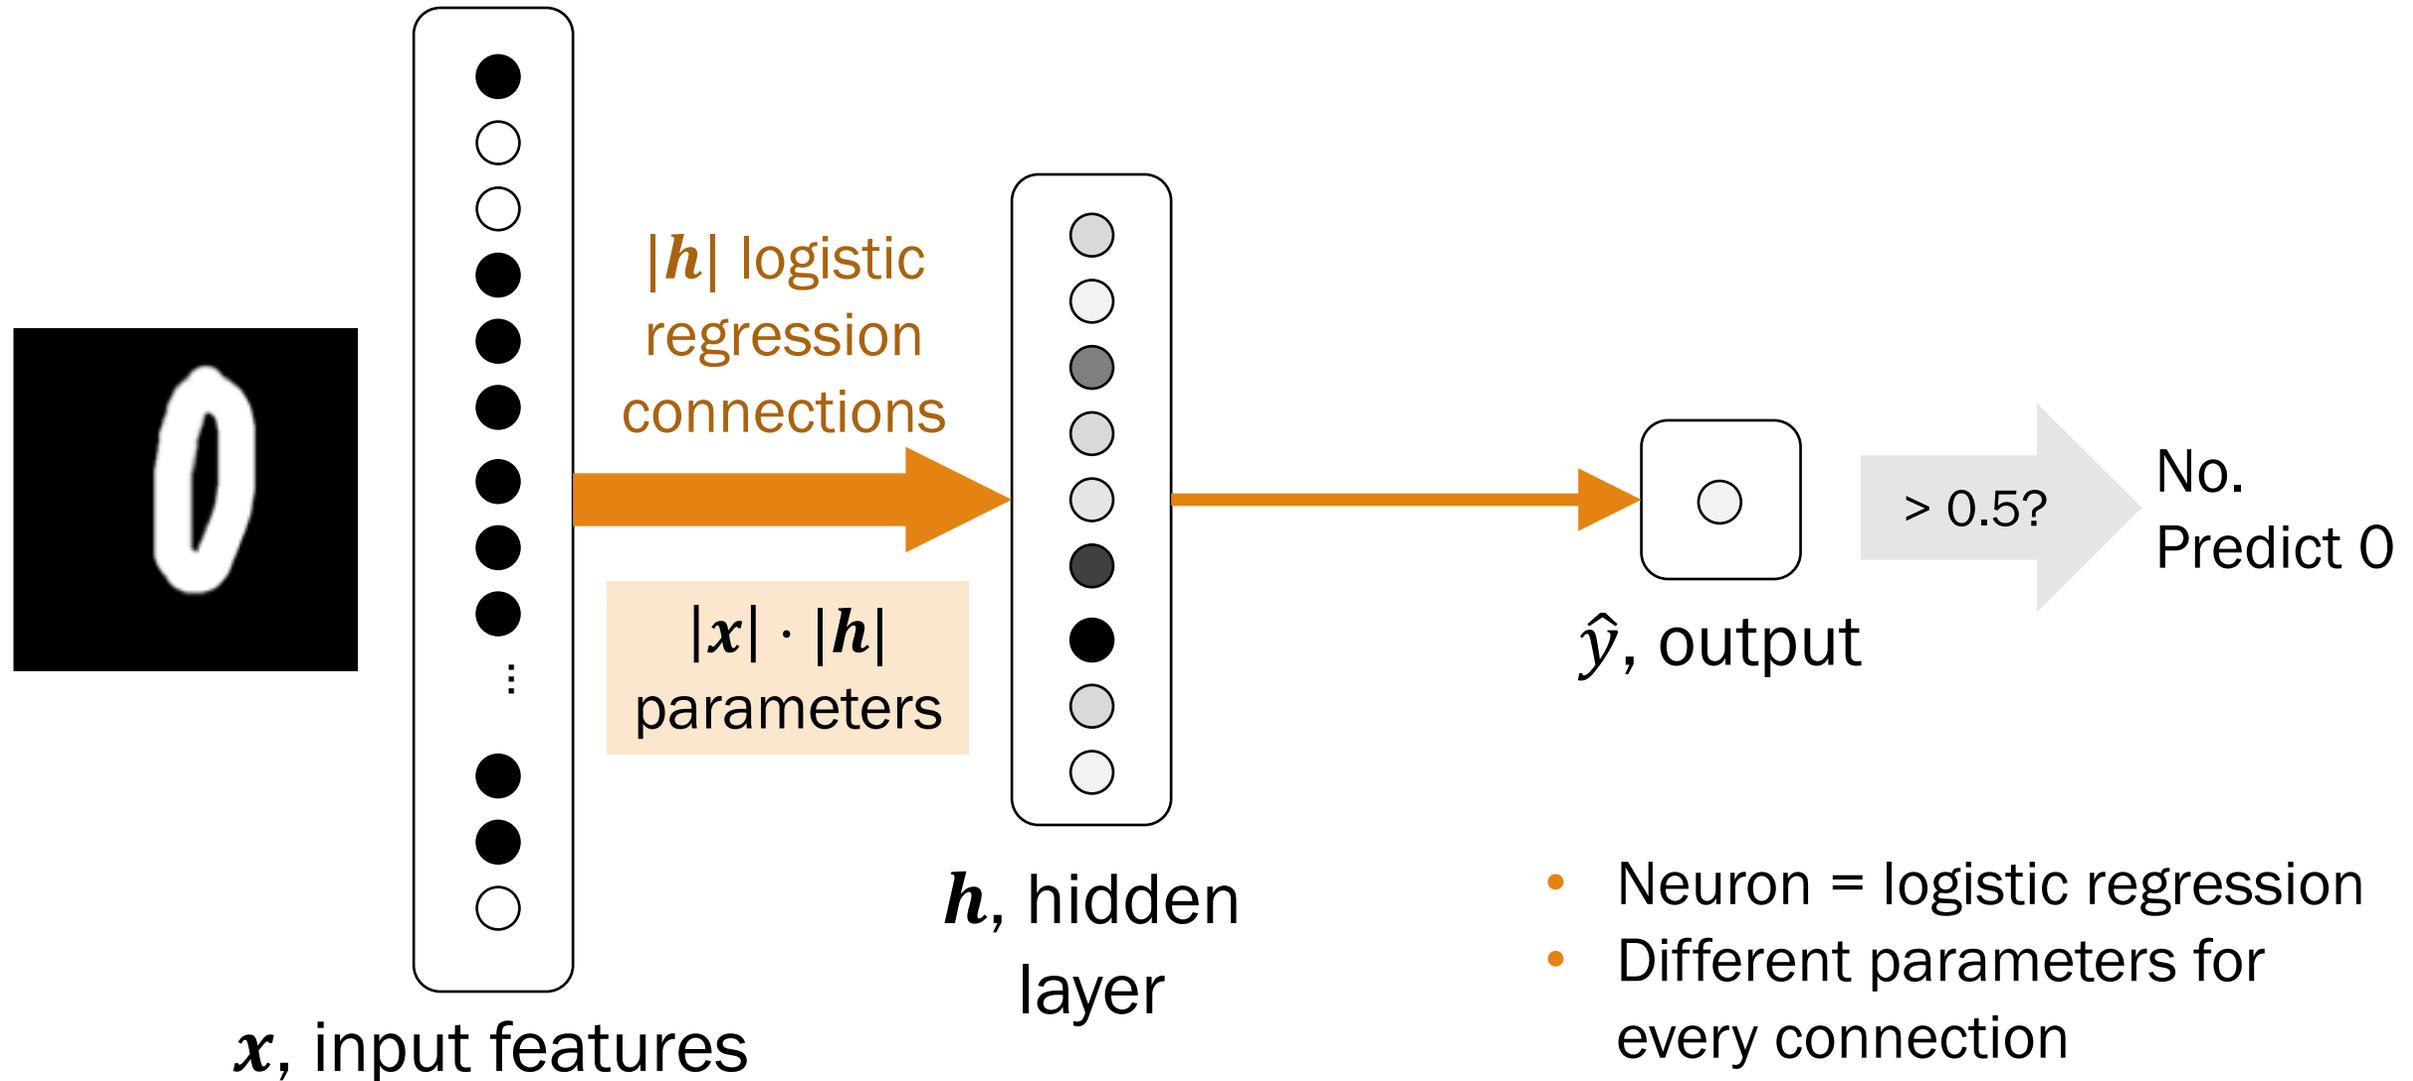
Feed neurons into other neurons



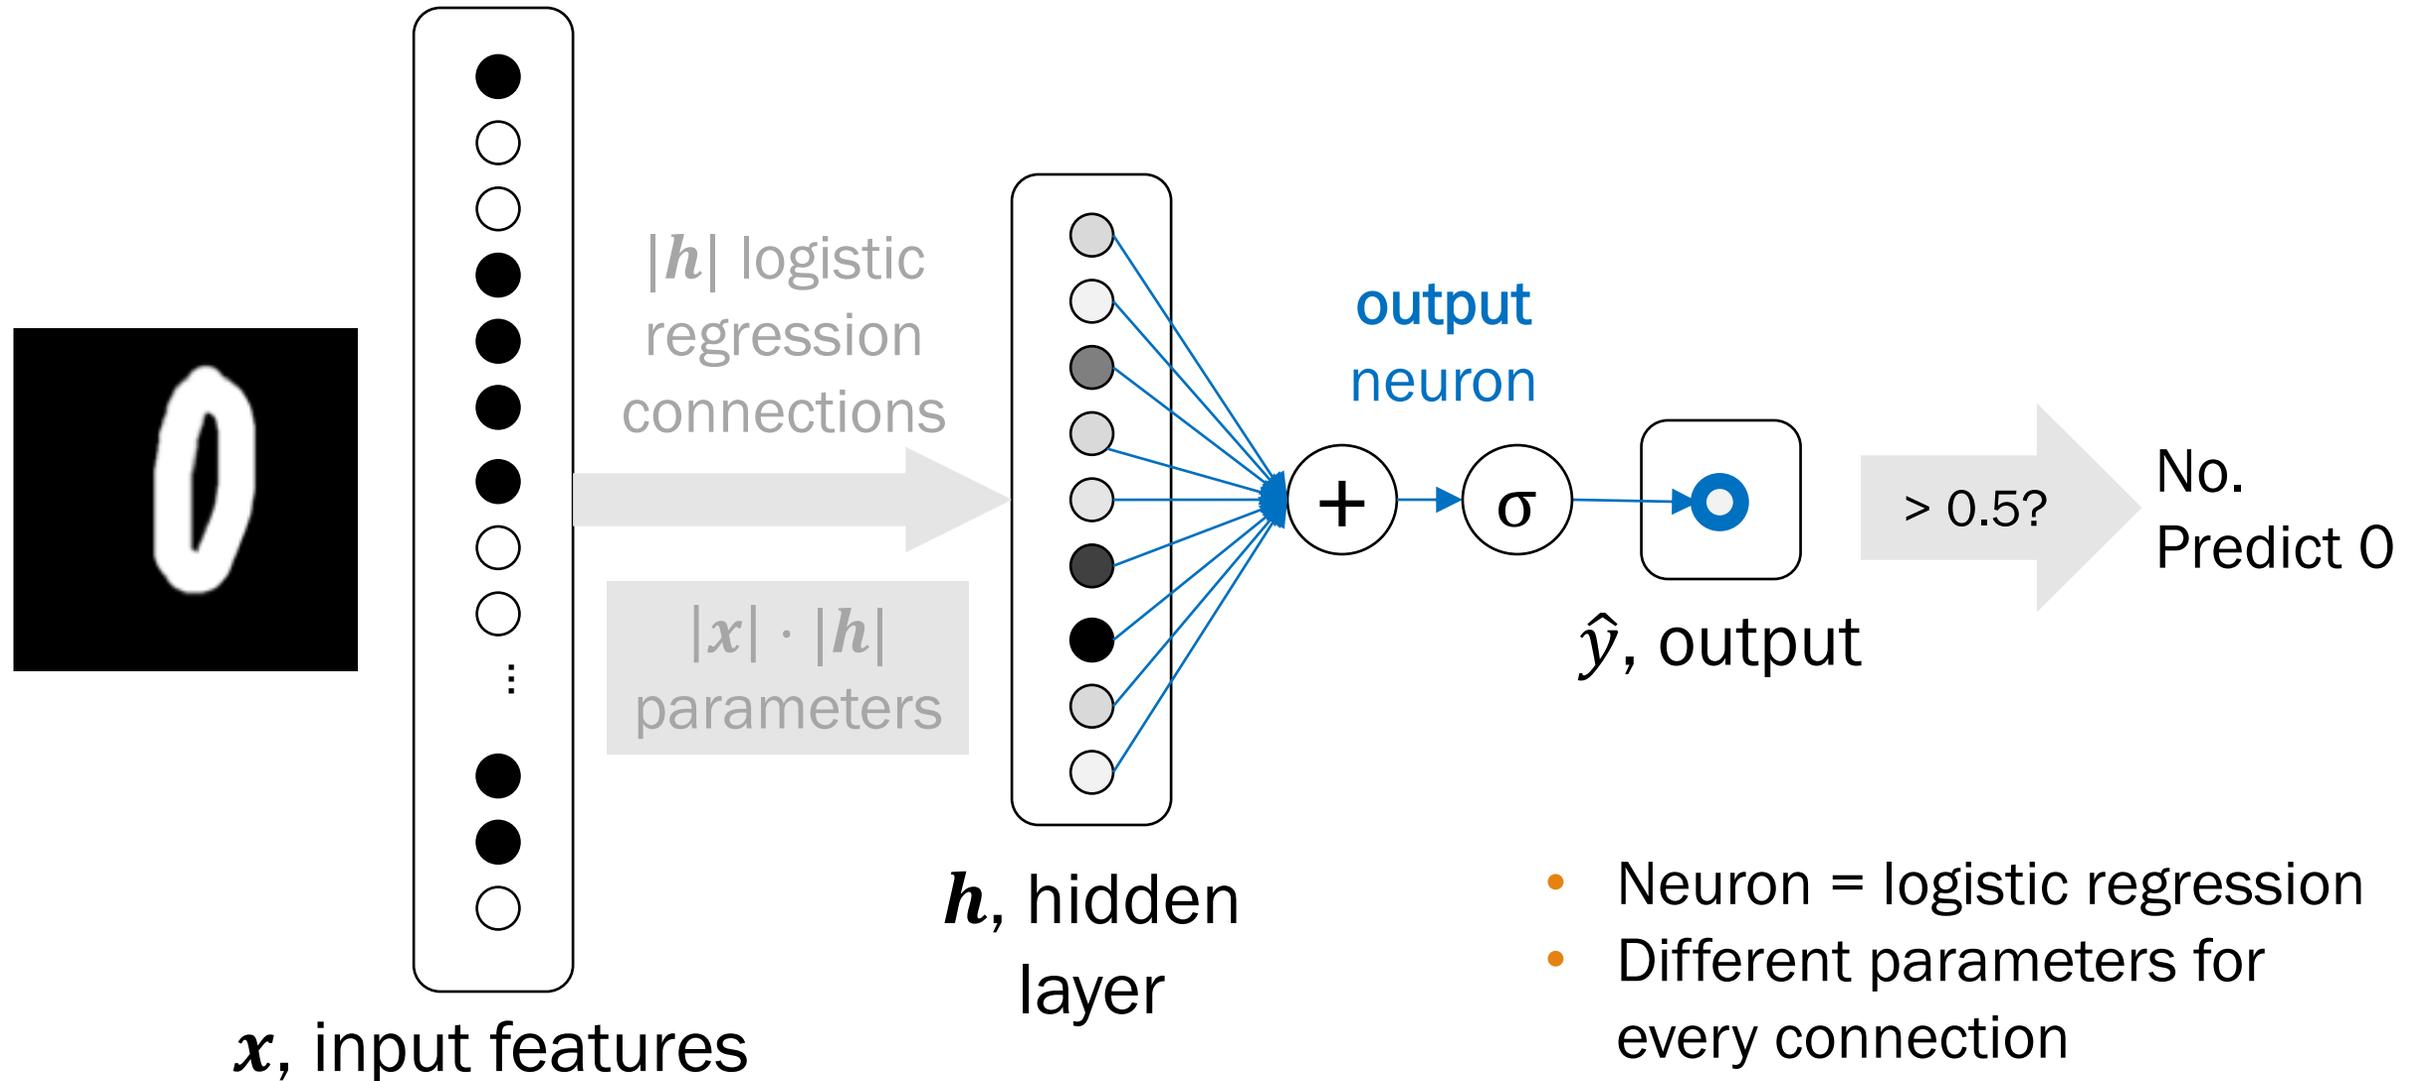
Feed neurons into other neurons



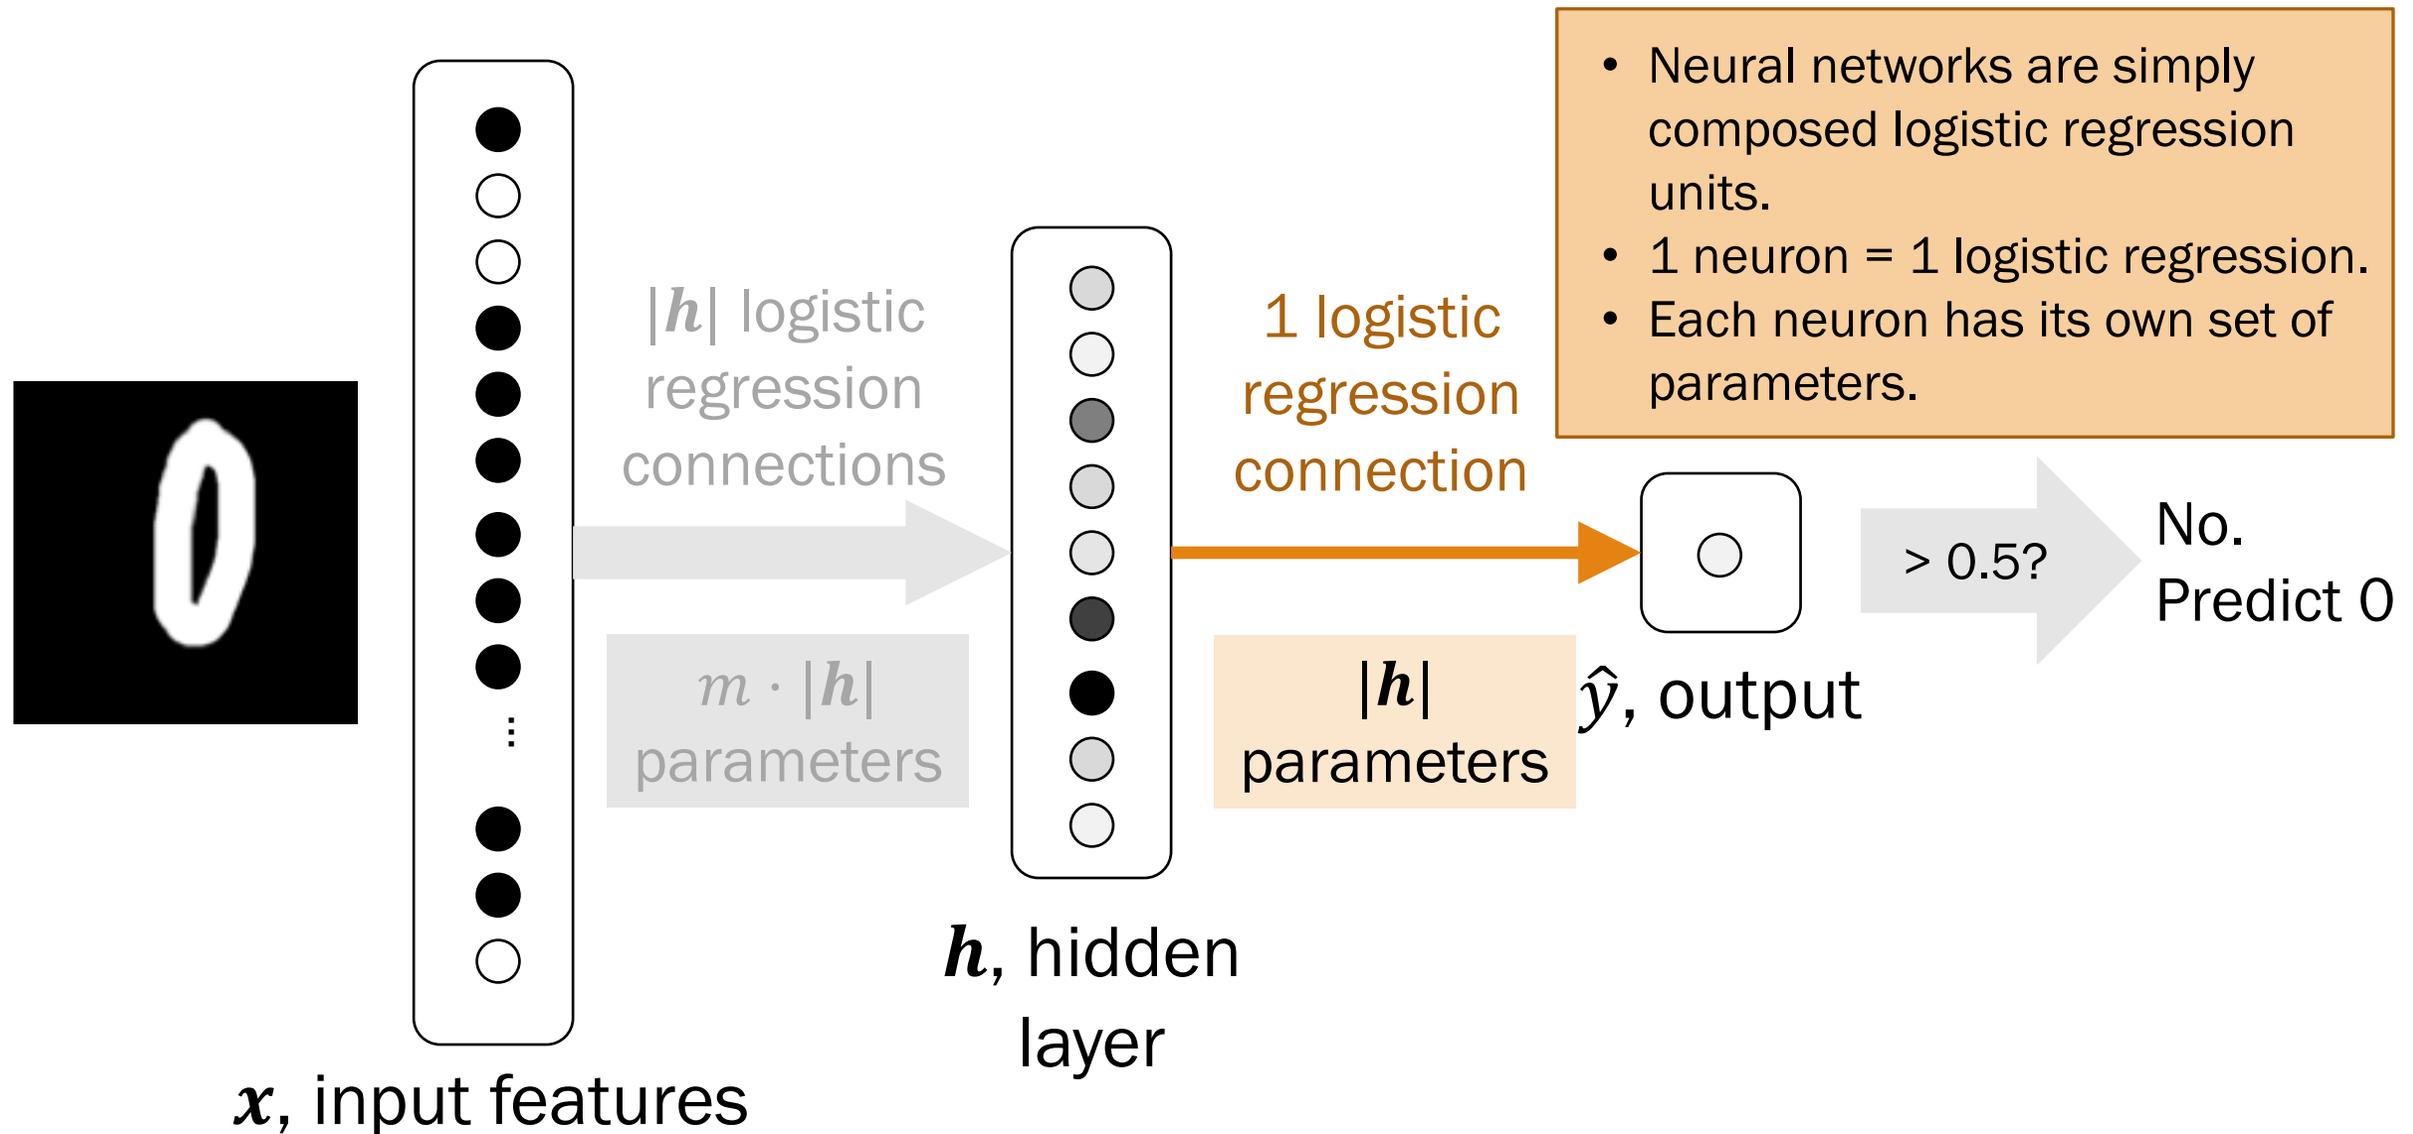
Feed neurons into other neurons



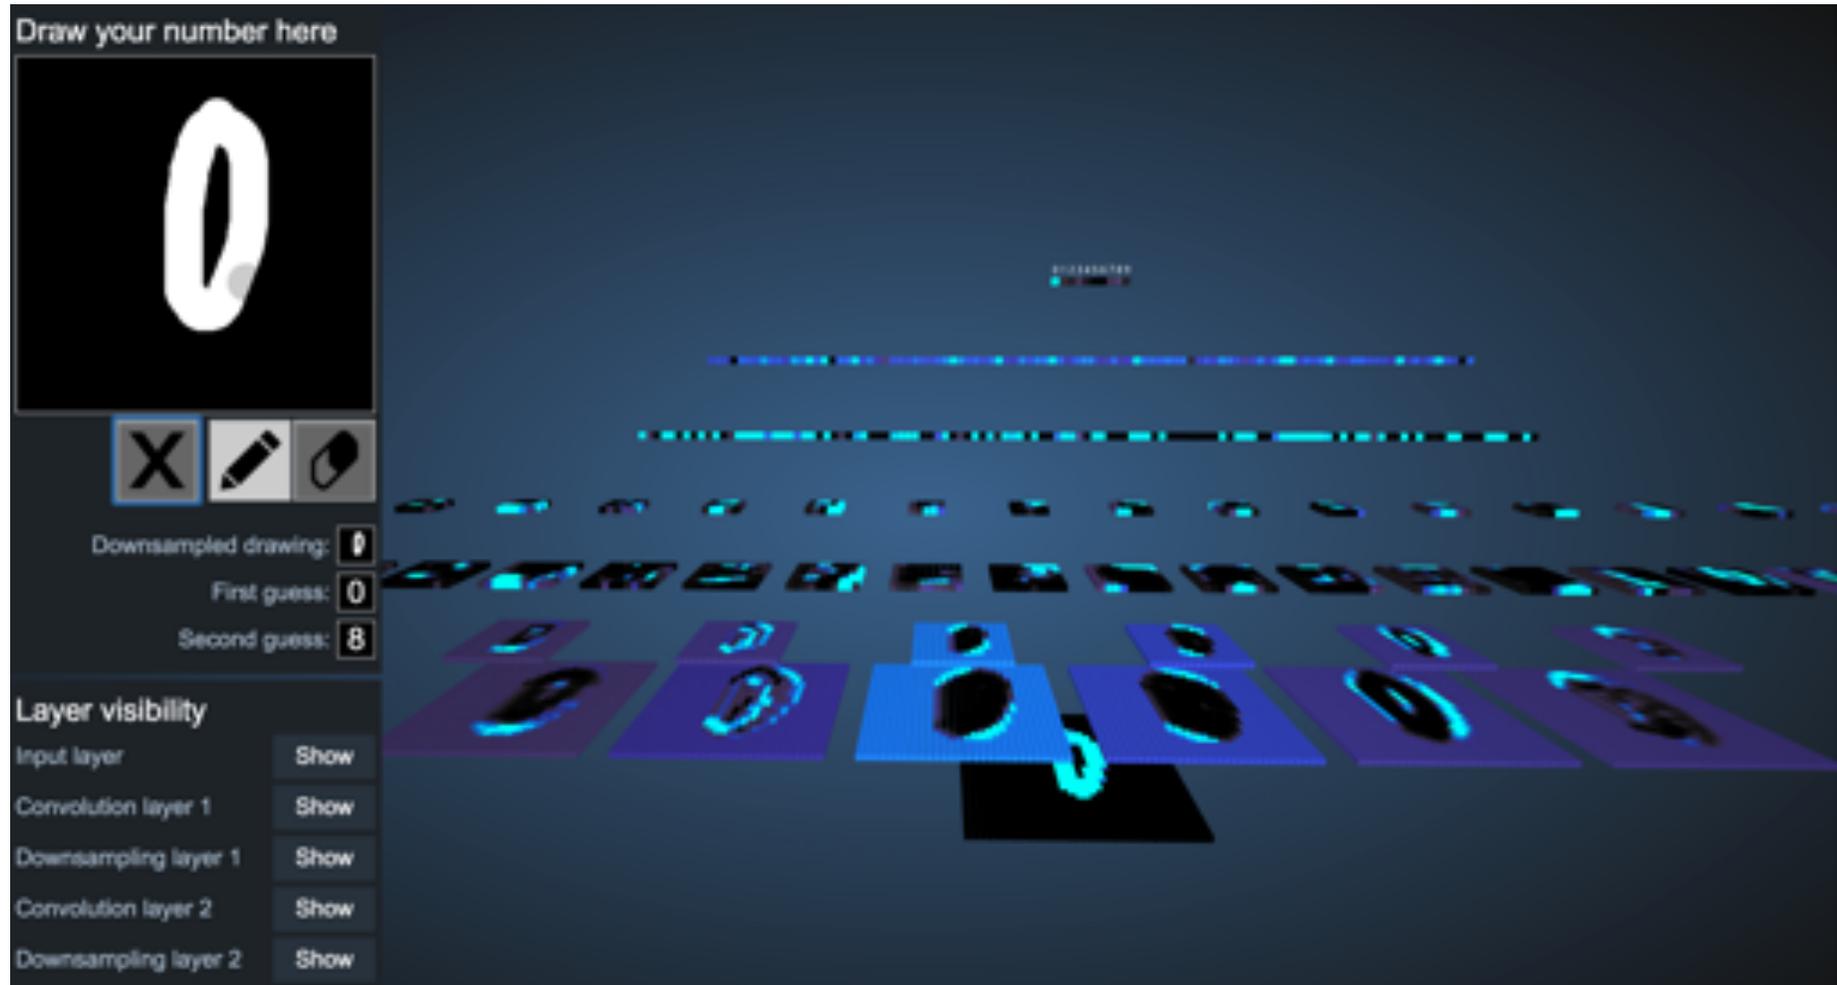
Feed neurons into other neurons



Feed neurons into other neurons



Demonstration



<http://scs.ryerson.ca/~aharley/vis/conv/>

Neural networks

A neural network (like logistic regression) gets intelligence from its parameters θ .

Training

- Learn parameters θ
- Find θ_{MLE} that maximizes likelihood of training data (MLE)

Testing/ Prediction

For input feature vector $\mathbf{X} = \mathbf{x}$:

- Use parameters to compute $\hat{y} = P(Y = 1 | \mathbf{X} = \mathbf{x})$
- Classify instance as:
$$\begin{cases} 1 & \hat{y} > 0.5 \\ 0 & \text{otherwise} \end{cases}$$

Neural networks

A neural network (like logistic regression) gets intelligence from its parameters θ .

Training

- Learn parameters θ
- Find θ_{MLE} that maximizes likelihood of training data (MLE)

How do we learn the $|\mathbf{x}| \cdot |\mathbf{h}| + |\mathbf{h}|$ parameters?

Gradient ascent + chain rule!

1. Optimization problem:

$$\theta_{MLE} = \arg \max_{\theta} \prod_{i=1}^n f(y^{(i)} | \mathbf{x}^{(i)}, \theta) = \arg \max_{\theta} LL(\theta)$$

$$LL(\theta) = \sum_{i=1}^n y^{(i)} \log \hat{y}^{(i)} + (1 - y^{(i)}) \log(1 - \hat{y}^{(i)})$$

$$\hat{y} = \sigma(\theta^T \mathbf{x}^{(i)})$$

2. Compute gradient

$$\frac{\partial LL(\theta)}{\partial \theta_j} = \sum_{i=1}^n [y^{(i)} - \hat{y}] x_j^{(i)}$$

3. Optimize

```
initialize params
repeat many times:
  compute gradient
  params += η * gradient
```

Training a neural net

1. Optimization problem:

$$\theta_{MLE} = \arg \max_{\theta} \prod_{i=1}^n f(y^{(i)} | \mathbf{x}^{(i)}, \theta) = \arg \max_{\theta} LL(\theta)$$

2. Compute gradient

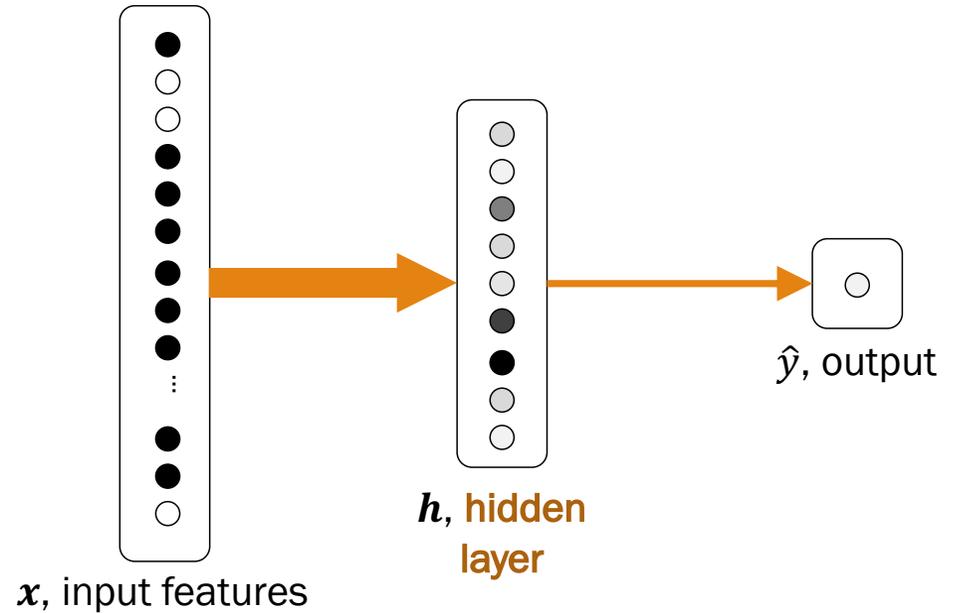
3. Optimize

1. Same output \hat{y} , same conditional log-likelihood

$$P(Y = y | \mathbf{X} = \mathbf{x}) = \begin{cases} \hat{y} & \text{if } y = 1 \\ 1 - \hat{y} & \text{if } y = 0 \end{cases}$$

$$\begin{aligned} L(\theta) &= \prod_{i=1}^n P(Y = y^{(i)} | \mathbf{X} = \mathbf{x}^{(i)}, \theta) \\ &= \prod_{i=1}^n (\hat{y}^{(i)})^{y^{(i)}} (1 - \hat{y}^{(i)})^{1-y^{(i)}} \end{aligned}$$

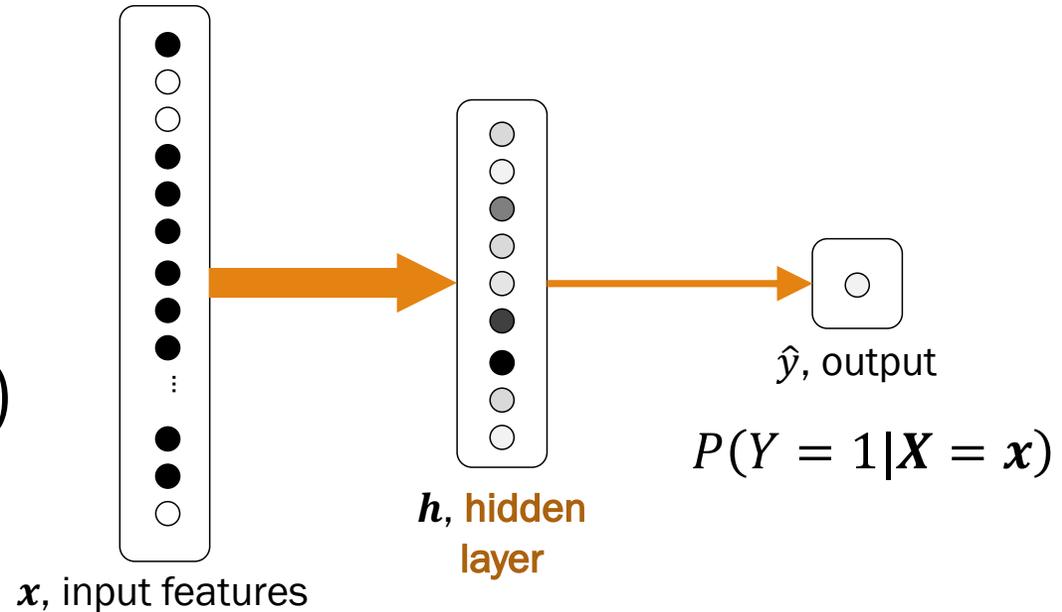
$$LL(\theta) = \sum_{i=1}^n y^{(i)} \log \hat{y}^{(i)} + (1 - y^{(i)}) \log(1 - \hat{y}^{(i)})$$



1. (model is a little more complicated, though)

$$P(Y = y | \mathbf{X} = \mathbf{x}) = \begin{cases} \hat{y} & \text{if } y = 1 \\ 1 - \hat{y} & \text{if } y = 0 \end{cases}$$

$$LL(\theta) = \sum_{i=1}^n y^{(i)} \log \hat{y}^{(i)} + (1 - y^{(i)}) \log(1 - \hat{y}^{(i)})$$



$$h_j = \sigma \left(\theta_j^{(h)T} \mathbf{x} \right) \text{ for } j = 1, \dots, |\mathbf{h}|$$

$$\hat{y} = \sigma \left(\theta^{(\hat{y})T} \mathbf{h} \right)$$

$|\mathbf{h}|$ param vectors, each with dimension $|\mathbf{x}|$

$|\mathbf{h}|$ -dimensional param vector

2. Compute gradient

1. Optimization problem:

$$\theta_{MLE} = \arg \max_{\theta} \prod_{i=1}^n f(y^{(i)} | \mathbf{x}^{(i)}, \theta) = \arg \max_{\theta} LL(\theta)$$

$$LL(\theta) = \sum_{i=1}^n y^{(i)} \log \hat{y}^{(i)} + (1 - y^{(i)}) \log(1 - \hat{y}^{(i)})$$

$$h_j = \sigma \left(\theta_j^{(h)T} \mathbf{x} \right) \quad \text{for } j = 1, \dots, |\mathbf{h}| \quad \hat{y} = \sigma \left(\theta^{(\hat{y})T} \mathbf{h} \right)$$

2. Compute gradient

Take gradient with respect to all θ parameters

3. Optimize

Calculus refresher #1:

Derivative(sum) =
sum(derivative)

Calculus refresher #2:

Chain rule 

3. Optimize

1. Optimization problem:

$$\theta_{MLE} = \arg \max_{\theta} \prod_{i=1}^n f(y^{(i)} | \mathbf{x}^{(i)}, \theta) = \arg \max_{\theta} LL(\theta)$$

$$LL(\theta) = \sum_{i=1}^n y^{(i)} \log \hat{y}^{(i)} + (1 - y^{(i)}) \log(1 - \hat{y}^{(i)})$$

$$h_j = \sigma \left(\theta_j^{(h)T} \mathbf{x} \right) \quad \text{for } j = 1, \dots, |\mathbf{h}| \quad \hat{y} = \sigma \left(\theta^{(\hat{y})T} \mathbf{h} \right)$$

2. Compute gradient

Take gradient with respect to all θ parameters

3. Optimize

```
initialize params
repeat many times:
  compute gradient
  params +=  $\eta$  * gradient
```

Training a neural net

1. Optimization problem:

$$\theta_{MLE} = \arg \max_{\theta} \prod_{i=1}^n f(y^{(i)} | \mathbf{x}^{(i)}, \theta) = \arg \max_{\theta} LL(\theta)$$

$$LL(\theta) = \sum_{i=1}^n y^{(i)} \log \hat{y}^{(i)} + (1 - y^{(i)}) \log(1 - \hat{y}^{(i)})$$

Wait, did we just skip something difficult?

2. Compute

3. Optimize

```
initialize params
repeat many times:
  compute gradient
  params += η * gradient
```

2. Compute gradient via backpropagation

1. Optimization problem:

$$\theta_{MLE} = \arg \max_{\theta} \prod_{i=1}^n f(y^{(i)} | \mathbf{x}^{(i)}, \theta) = \arg \max_{\theta} LL(\theta)$$

$$LL(\theta) = \sum_{i=1}^n y^{(i)} \log \hat{y}^{(i)} + (1 - y^{(i)}) \log(1 - \hat{y}^{(i)})$$

$$h_j = \sigma \left(\theta_j^{(h)T} \mathbf{x} \right) \quad \text{for } j = 1, \dots, |\mathbf{h}| \quad \hat{y} = \sigma \left(\theta^{(\hat{y})T} \mathbf{h} \right)$$

2. Compute gradient

Take gradient with respect to all θ parameters

3. Optimize

initial
repeat
compute
parameters

Learn the tricks behind **backpropagation** in CS229, CS231N, CS224N, etc.

Interlude for jokes

Probability as college students

The Six Probability Distributions You'll Meet in Your Sorority

The One Who Does It All

You see her everywhere. Physics, math, computer science. How is she in all your classes? And she does amazing in all of them, keeping well ahead of the curve. You'd like to be friends, but despite her popularity, she seems to have been progressing towards mean spirited behavior. At least she seems normal.

Gaussian

$$f(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2}$$

The Confidant

You think she's related with The One Who Does It All. But how did she turn out like that? You can tell her anything, and you know she's discreet enough to keep it under wraps. But take care of her. This one will bet all she's got on a handful of coin tosses.

Binomial

$$p(k) = \binom{n}{k} p^k q^{n-k}$$

The Scatterbrain

This girl cannot remember anything. She needs to ask your math every time you meet her. You're pretty sure you were friends during math, but things have dropped off quickly since then.

Exponential

$$f(x) = \lambda e^{-\lambda x}$$

The Background Boyfriend

He started dating The Confidant last semester, but you can't see what they have in common. He's not obnoxious, but he's not particularly charming either. No matter what she says, he always gives the same response: "sure, k." His flat personality might be mistaken for a chill, laid-back attitude.

Uniform

$$f(k) = \frac{1}{b-a}$$

The Ghost

This sister always seems kind of distracted and never shows up to anything. In fact, the last time you saw her was two months ago - standing outside the library building.

Poisson

$$f(k) = e^{-\lambda} \frac{\lambda^k}{k!}$$

The Ride or Die

You always know where she's going to be. Your relationship can get complicated, but she's always got your back when things need to square it. That rock-solid support, you can count on her to never waver in her ~~opinion~~.

Delta

$$\delta(x) = \begin{cases} +\infty, & x = 0 \\ 0, & x \neq 0 \end{cases}$$

(Note: not a valid PDF, but a useful construct nonetheless)

All remaining jokes

(from mid-quarter feedback)

What happens to a frog's car when it breaks down? It gets toad away.

What is similar about ice cream and bees? If you eat too many of either you'll get a stomachache

Knock knock. Whose there. Dishes. Dishes who? Dishes a great class.

Why is your nose in the middle of your face? Because it's the scenter :)

Why does the Norwegian Navy have bar codes on the side of their ships? So when they come back to port they can...

Tell a joke about pizza! Or wait that might be too cheesy ;)

Why do you avoid the top of the bell curve? Because everyone there is mean.

What did the fish say when he swam into a wall?

"Three logicians walk into a bar. The bartender asks, "Will you all have a drink?" The first logician says "I don't know." The second logician says "I don't know." The third logician says "Yes.""

What do you call it when you've been pulling too many all-nighters studying probability, and you see the ghost of Carl Friedrich Gauss? para-NORMAL activity (I'll take my extra credit now please)

Lisa giving us a house tour.

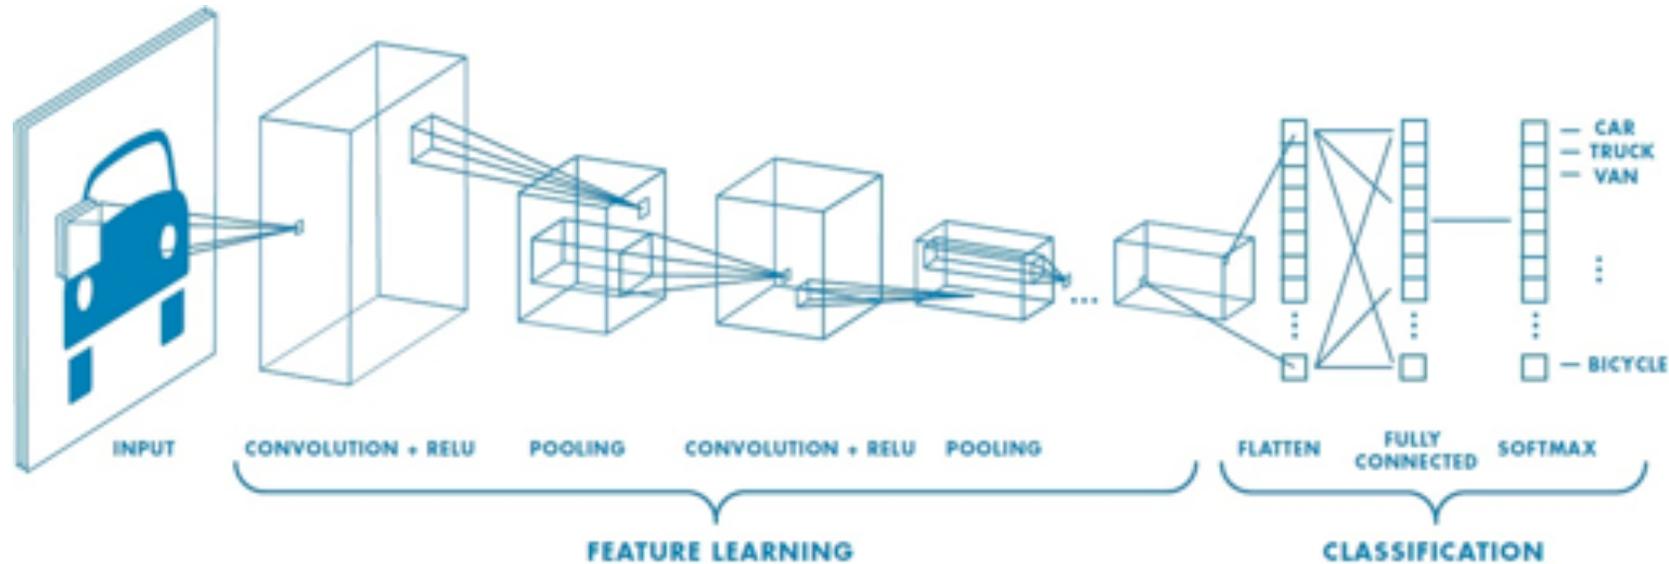
Two random variables were talking in a bar. They thought they were being discreet but I heard their chatter continuously.

what has two butts and kills people?

"A statistics professor plans to travel to a conference by plane. When he passes the security check, they discover a bomb in his carry-on! Of course, he is hauled off immediately for interrogation. "I don't understand!" the

Beyond the basics

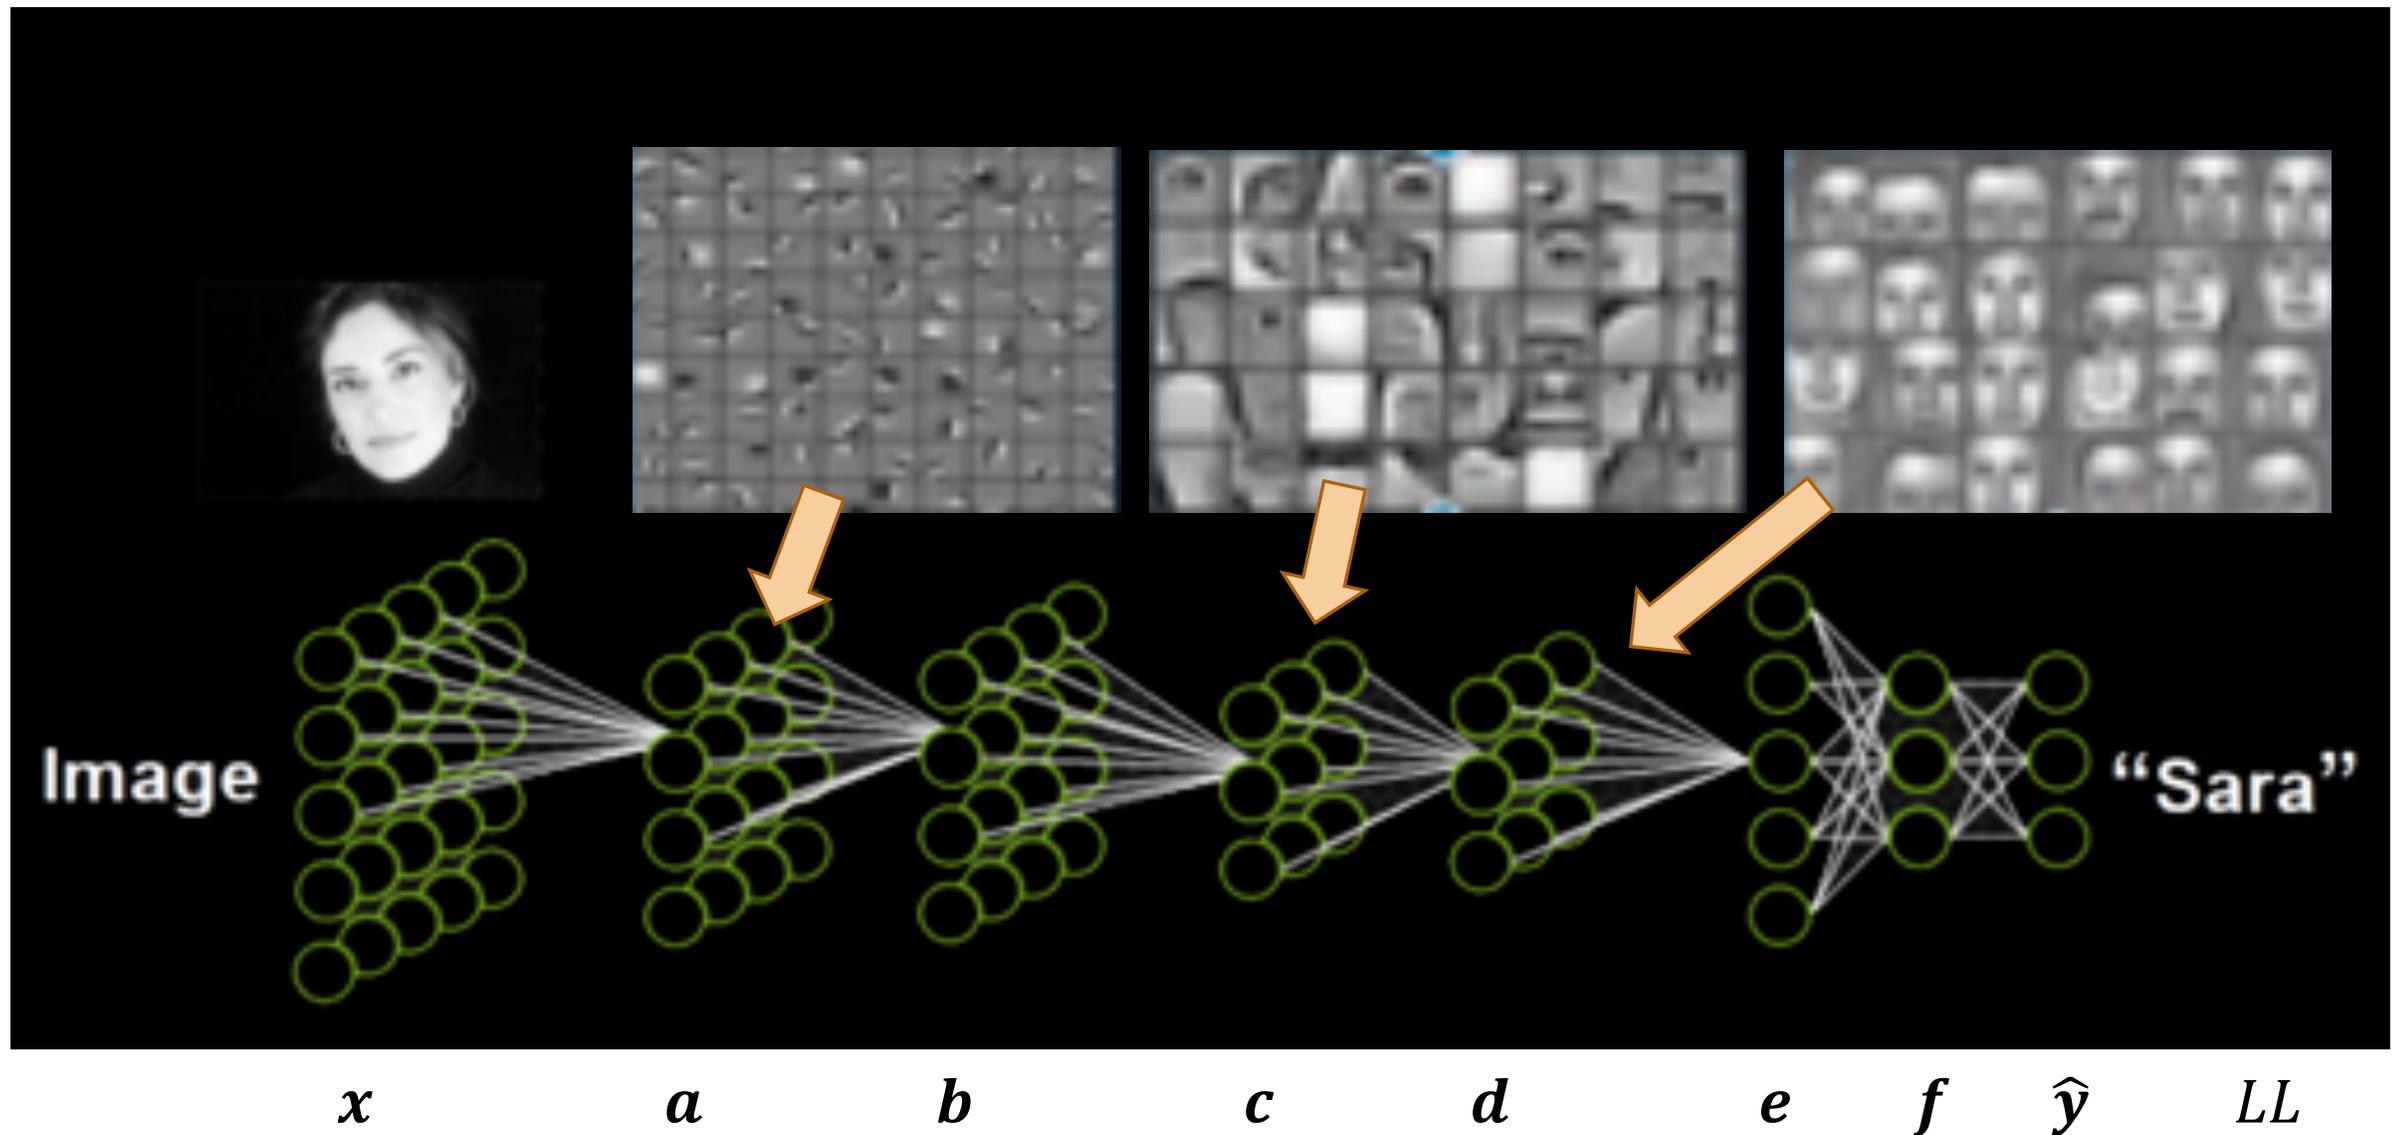
Shared weights?



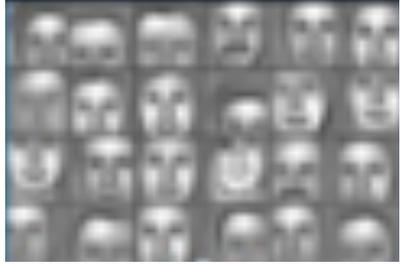
It turns out if you want to force some of your weights to be shared over different neurons, the math isn't much harder.

Convolution is an example of such weight-sharing and is used a lot for vision (Convolutional Neural Networks, CNN).

Neural networks with multiple layers



Neurons learn features of the dataset



Neurons in later layers will respond strongly to high-level features of your **training data**.

If your training data is faces, you will get lots of face neurons.

If your training data is all of YouTube...



...you get a cat neuron.



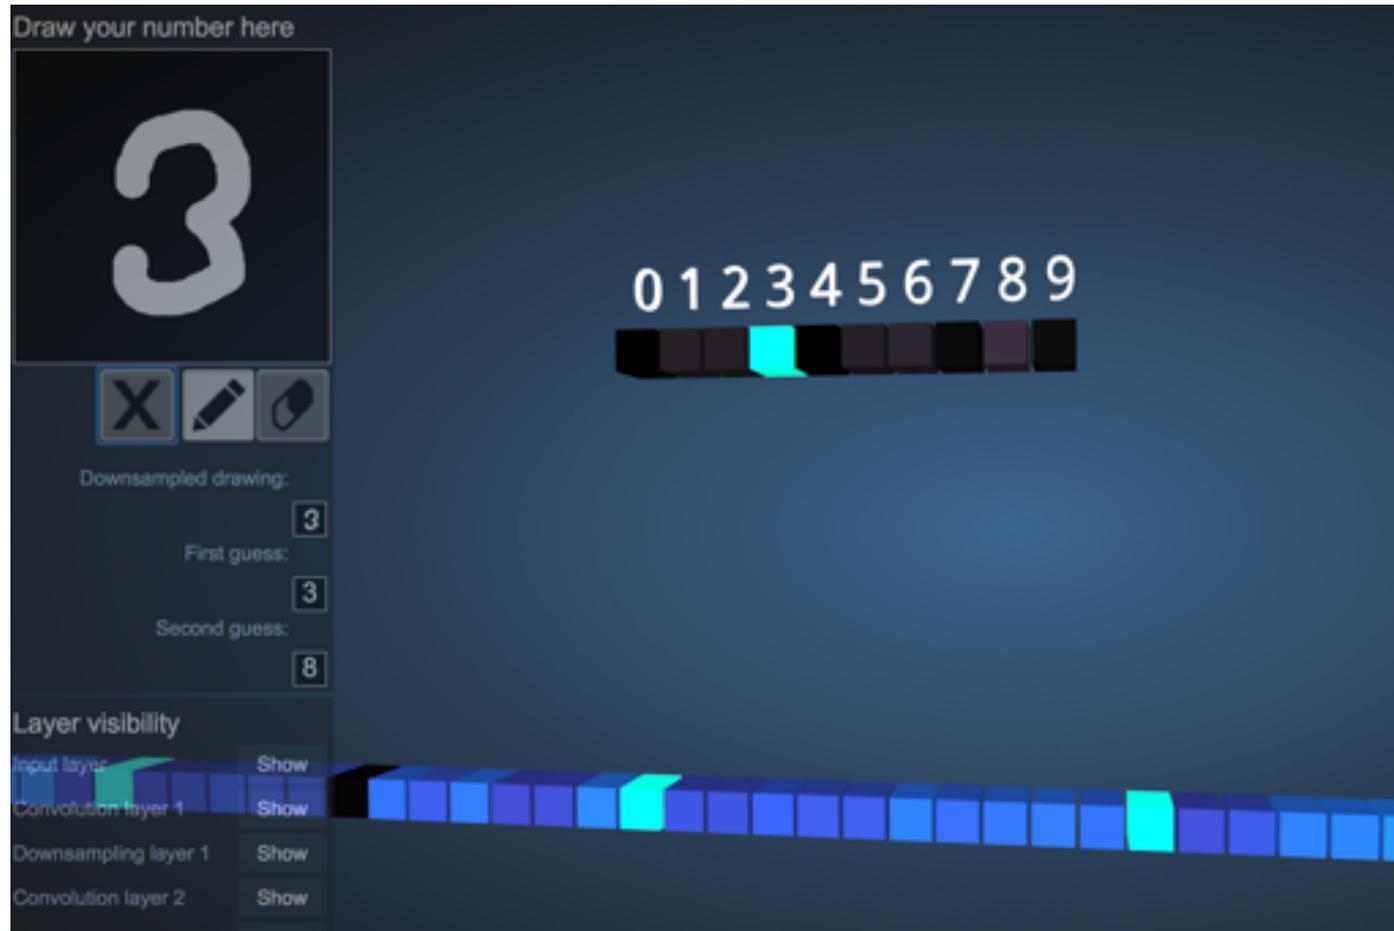
Top stimuli in test set



Optimal stimulus found by numerical optimization



Multiple outputs?



Softmax is a generalization of the sigmoid function.

sigmoid(z): value in range $[0, 1]$

$z \in \mathbb{R}$:

$$P(Y = 1 | \mathbf{X} = \mathbf{x}) = \sigma(z)$$

(equivalent: Bernoulli p)



softmax(z): k -dimensional values in range $[0, 1]$ that add up to 1

$\mathbf{z} \in \mathbb{R}^k$:

$$P(Y = i | \mathbf{X} = \mathbf{x}) = \text{softmax}(\mathbf{z})_i$$

(equivalent: Multinomial p_1, \dots, p_k)

Softmax test metric: Top-5 error

$Y = y$	$P(Y = y X = \mathbf{x})$
5	0.14
8	0.13
7	0.12
2	0.10
9	0.10
4	0.09
1	0.09
0	0.09
6	0.08
3	0.05



Top-5 classification error

What % of datapoints did *not* have the correct class label in the top-5 predictions?



ImageNet classification

22,000 categories

14,000,000 images

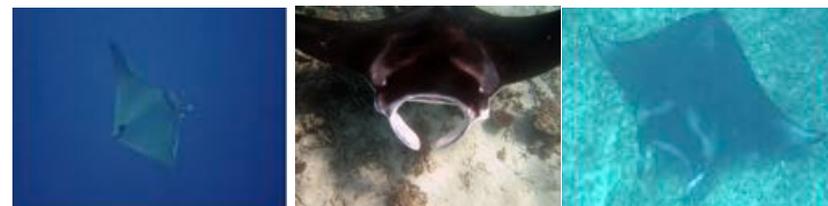
Hand-engineered features
(SIFT, HOG, LBP),
Spatial pyramid,
SparseCoding/Compression

...
smoothhound, smoothhound shark, *Mustelus mustelus*
American smooth dogfish, *Mustelus canis*
Florida smoothhound, *Mustelus norrisi*
whitetip shark, reef whitetip shark, *Triaenodon obseus*
Atlantic spiny dogfish, *Squalus acanthias*
Pacific spiny dogfish, *Squalus suckleyi*
hammerhead, hammerhead shark
smooth hammerhead, *Sphyrna zygaena*
smalleye hammerhead, *Sphyrna tudes*
shovelhead, bonnethead, bonnet shark, *Sphyrna tiburo*
angel shark, angelfish, *Squatina squatina*, monkfish
electric ray, crampfish, numbfish, torpedo
smalltooth sawfish, *Pristis pectinatus*
guitarfish
roughtail stingray, *Dasyatis centroura*
butterfly ray
eagle ray
spotted eagle ray, spotted ray, *Aetobatus narinari*
cownose ray, cow-nosed ray, *Rhinoptera bonasus*
manta, manta ray, devilfish
Atlantic manta, *Manta birostris*
devil ray, *Mobula hypostoma*
grey skate, gray skate, *Raja batis*
little skate, *Raja erinacea*
...

Stingray



Mantaray



ImageNet classification challenge

~~22,000 categories~~

1000 categories

smoothhound shark, *Mustelus mustelus*
dogfish, *Mustelus canis*

Florida smoothhound, *Mustelus norrisi*

14,000,000 images

1,200,000 images in train set

codon obseus

200,000 images in test set

Hand-engineered features
(SIFT, HOG, LBP),
Spatial pyramid,
SparseCoding/Compression

smooth hammerhead, *Sphyrna zygaena*
smalleye hammerhead, *Sphyrna tudes*
shovelhead, bonnethead, bonnet shark, *Sphyrna tiburo*
angel shark, angelfish, *Squatina squatina*, monkfish
electric ray, crampfish, numbfish, torpedo
smalltooth sawfish, *Pristis pectinatus*
guitarfish
rougthead stingray, *Dasyatis centroura*
butterfly ray
eagle ray
spotted eagle ray, spotted ray, *Aetobatus narinari*
cownose ray, cow-nosed ray, *Rhinoptera bonasus*
manta, manta ray, devilfish
Atlantic manta, *Manta birostris*
devil ray, *Mobula hypostoma*
grey skate, gray skate, *Raja batis*
little skate, *Raja erinacea*

...

ImageNet challenge: Top-5 classification error

(lower is better)

99.5%

Random guess

$$P(\text{true class label not in 5 guesses}) = \frac{\binom{999}{5}}{\binom{1000}{5}} = \frac{995}{1000}$$

ImageNet challenge: Top-5 classification error

(lower is better)

99.5%

Random guess

25.8%

Pre-Neural Networks

5.1%

Humans
(2014)

16.4%

GoogLe Net
(2015)

Russakovsky et al., ImageNet Large Scale Visual Recognition Challenge. IJCV 2015

Szegedy et al., Going Deeper With Convolutions. CVPR 2015

Hu et al., Squeeze-and-Excitation Networks. Preprint arXiv 2017

ImageNet challenge: Top-5 classification error

(lower is better)

99.5%

Random guess

25.8%

Pre-Neural Networks

5.1%

Humans
(2014)

16.4%

GoogLe Net
(2015)

2.25%

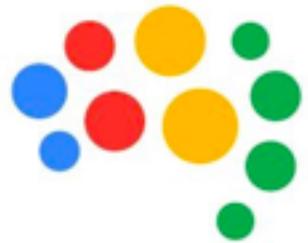
SENet
(2017)

Russakovsky et al., ImageNet Large Scale Visual Recognition Challenge. IJCV 2015

Szegedy et al., Going Deeper With Convolutions. CVPR 2015

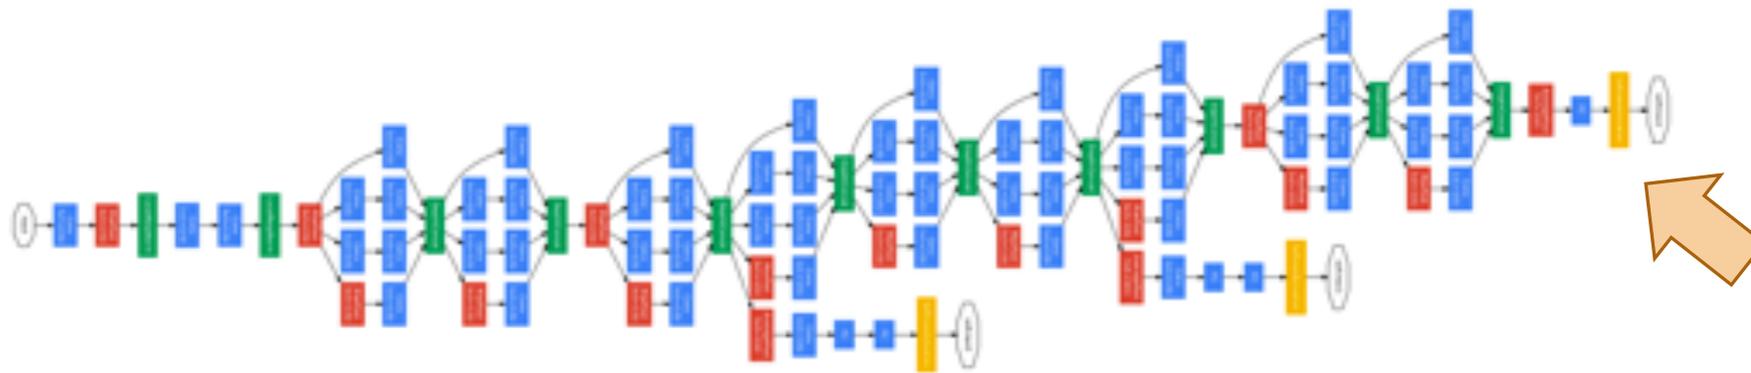
Hu et al., Squeeze-and-Excitation Networks. Preprint arXiv 2017

GoogLeNet (2015)



Google Brain

1 Trillion Artificial Neurons
(btw human brains have 1 billion neurons)



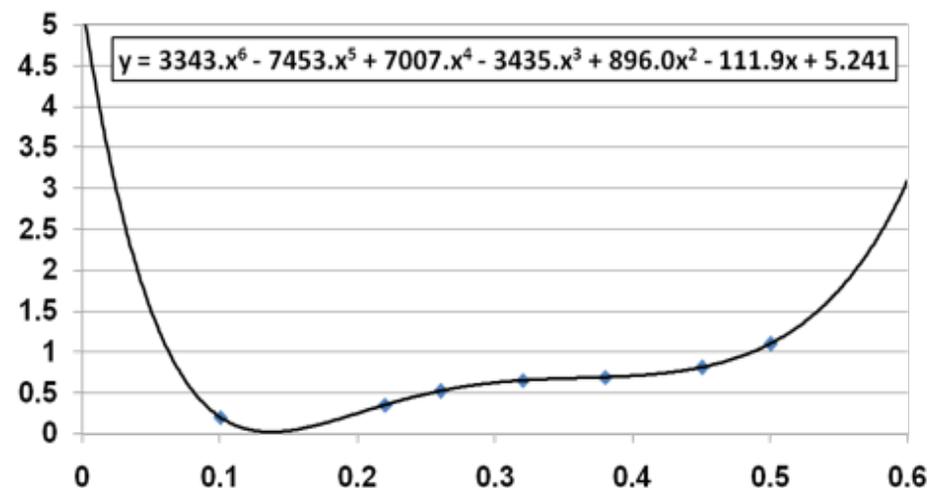
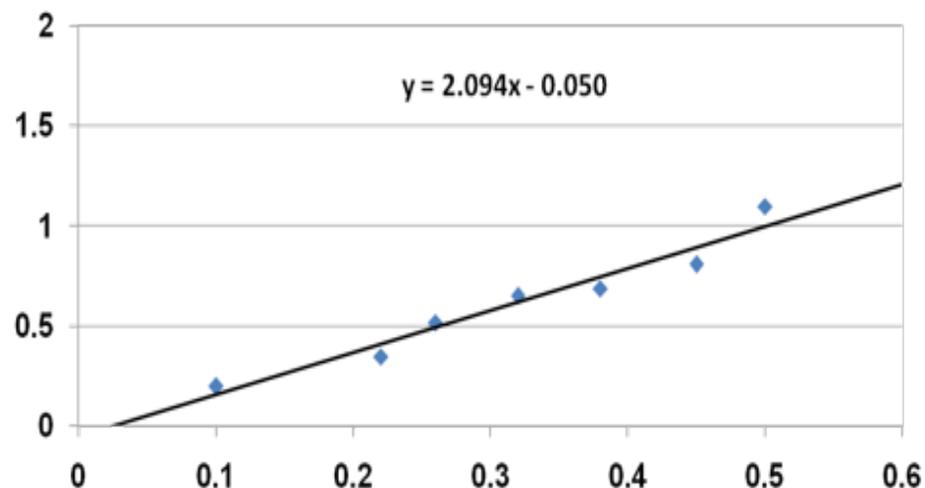
Multiple,
Multi class output

22 layers deep!

Good ML = Generalization

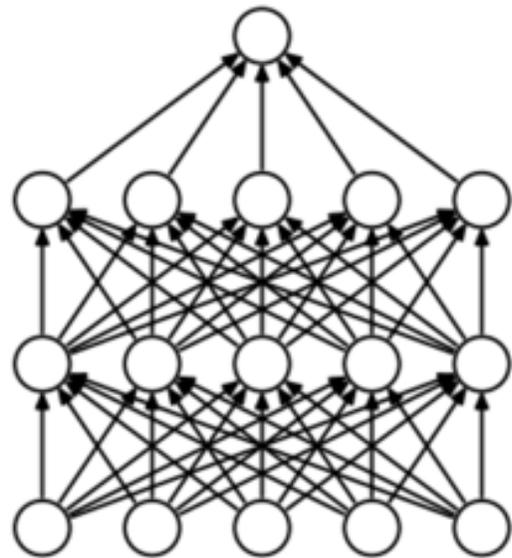
Goal of machine learning: build models that *generalize* well to predicting new data

Overfitting: fitting the training data too well, so we lose generality of model

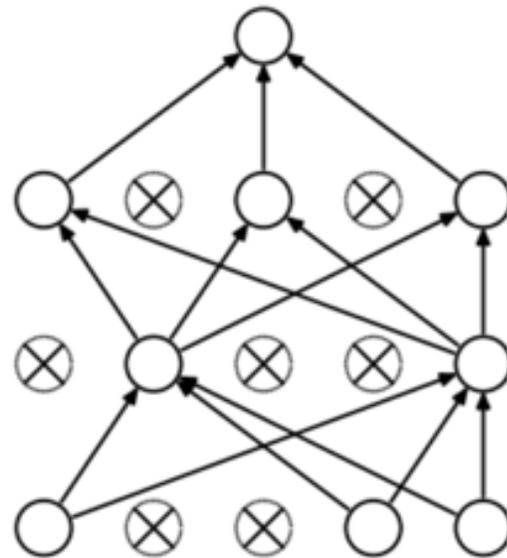


- Polynomial on the right fits training data perfectly!
- Which would you rather use to predict a new data point?

Prevent overfitting?



(a) Standard Neural Net



(b) After applying dropout.

Dropout (training technique)

When your model is training, randomly turn off your neurons with probability 0.5.

It will make your network more robust.

Making decisions?



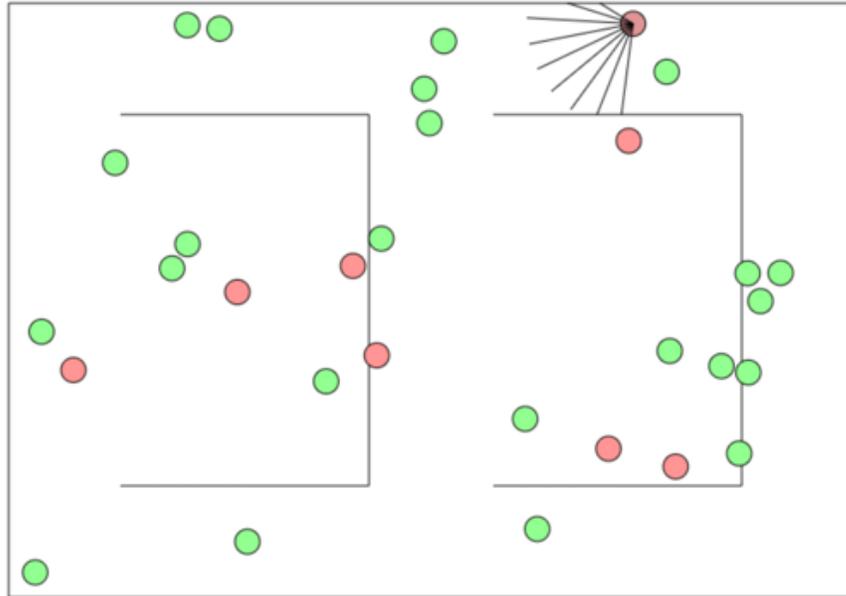
Not everything is classification.

Deep Reinforcement Learning

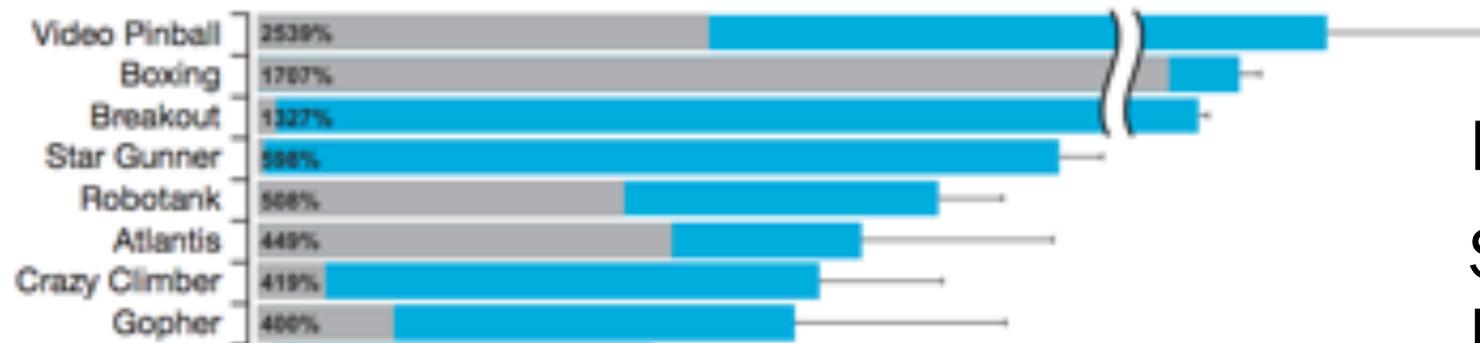
Instead of having the output of a model be a probability, you make output an expectation.

<http://cs.stanford.edu/people/karpathy/convnetjs/demo/rldemo.html>

Deep Reinforcement Learning



<http://cs.stanford.edu/people/karpathy/convnetjs/demo/rldemo.html>

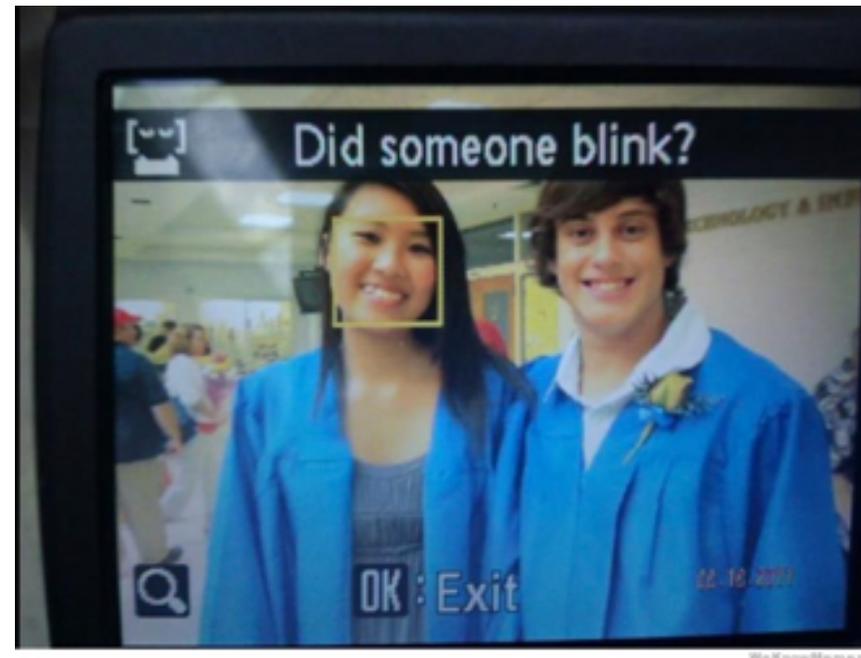


Deep Mind Atari Games
Score compared to best
human

What's missing?

Where is your data
coming from?

Ethics and datasets



Sometimes machine learning feels universally unbiased.

We can even prove our estimators are “unbiased” (mathematically).

Google/Nikon/HP had biased datasets.

Should your data be unbiased?

Dataset: Google News

$$\vec{\text{man}} - \vec{\text{woman}} \approx \vec{\text{king}} - \vec{\text{queen}}$$

$$\vec{\text{man}} - \vec{\text{woman}} \approx \vec{\text{computer programmer}} - \vec{\text{homemaker}}.$$

Should our unbiased data collection reflect society's systemic bias?

How can we explain decisions?



If your task is **image classification**, reasoning about high-level features is relatively easy.

Everything can be visualized.

What if you are trying to classify social outcomes?

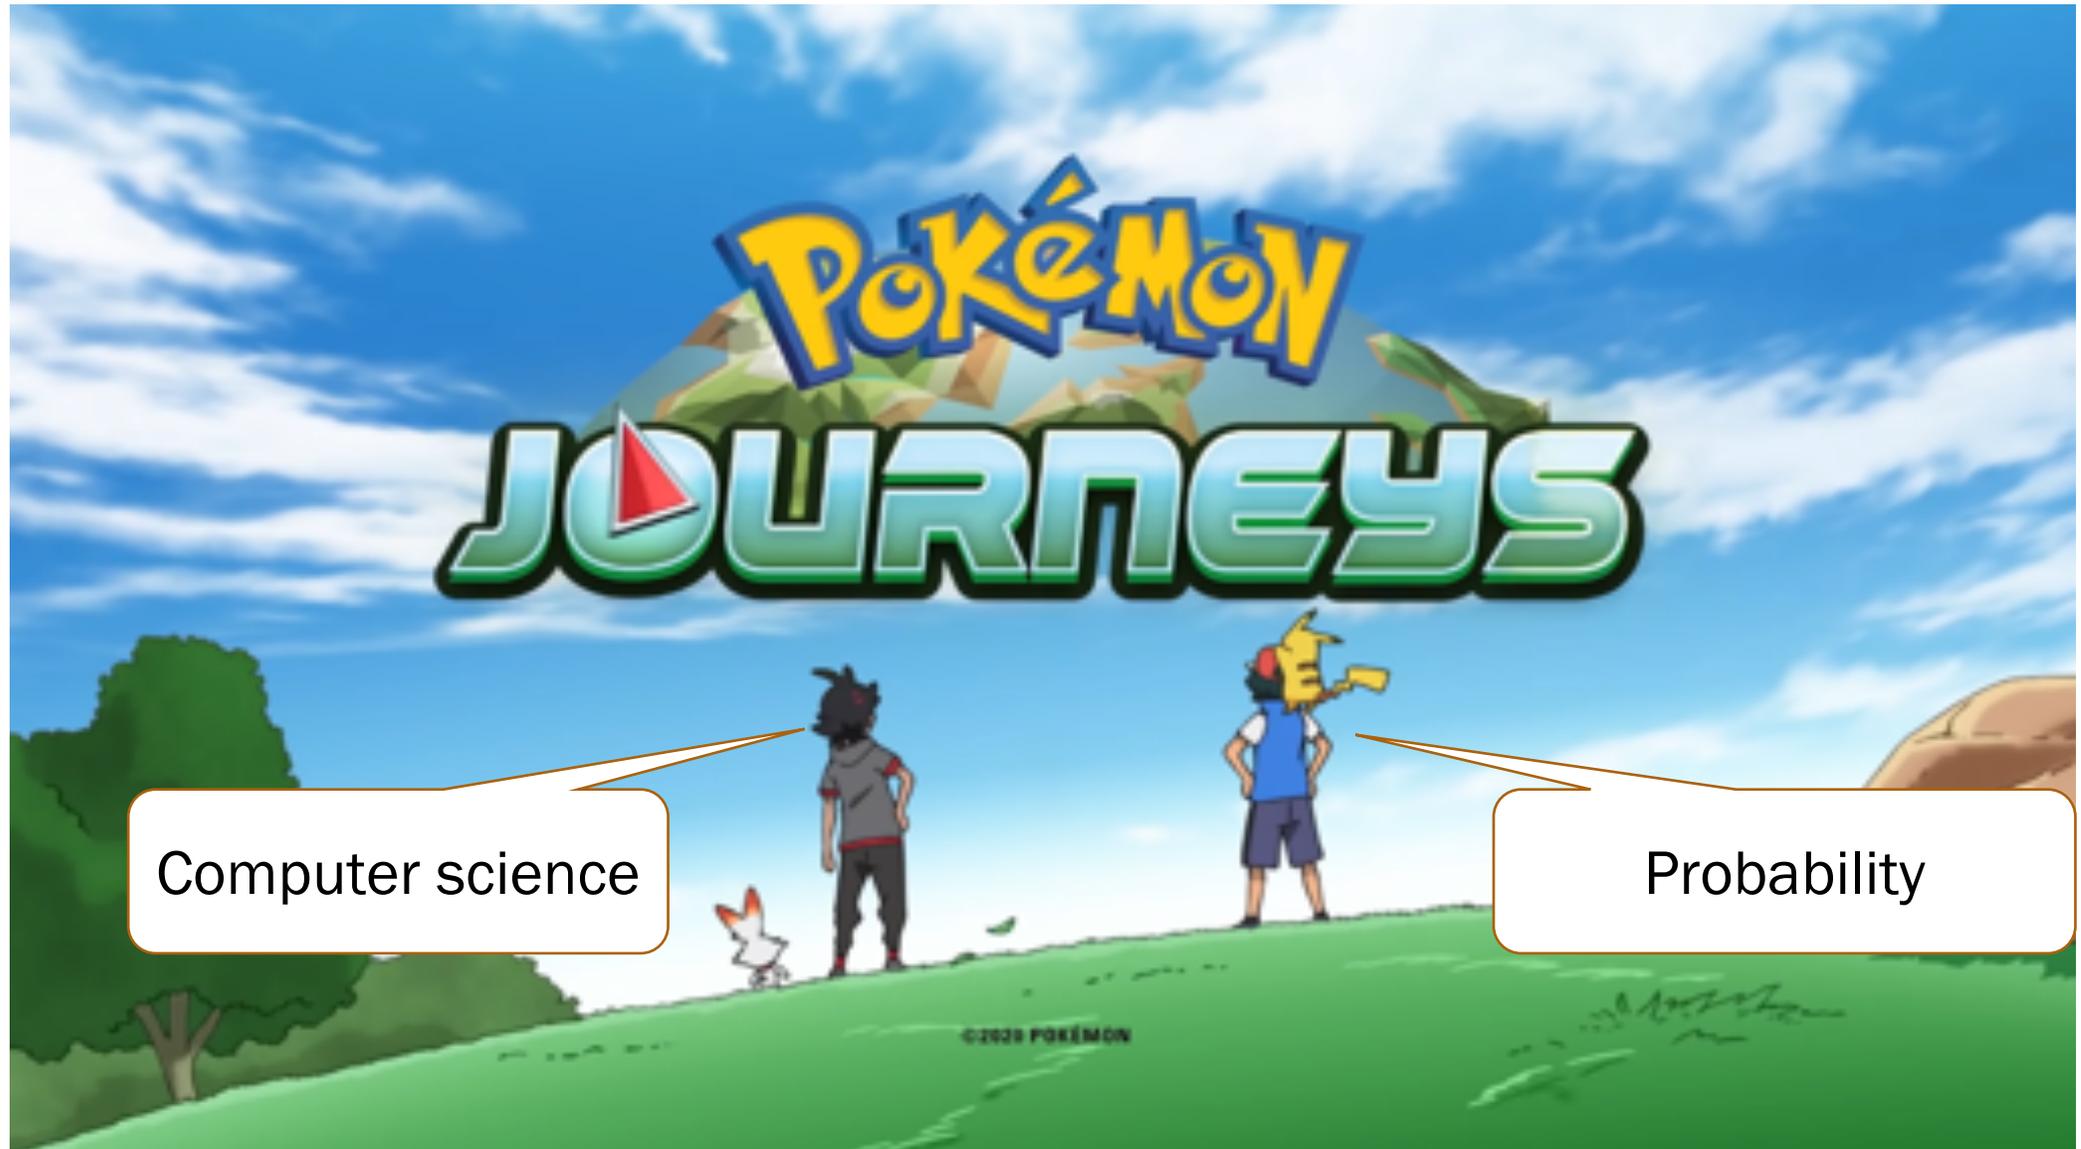
- Criminal recidivism
- Job performance
- Policing
- Terrorist risk
- At-risk kids

Ethics in Machine Learning
is a whole new field. 😊

CS109 Wrap-Up

What have we learned in
CS109?

A wild journey



Computer science

Probability

From combinatorics to probability...



Everything in the world is either



a potato or not a potato.

$$P(E) + P(E^C) = 1$$



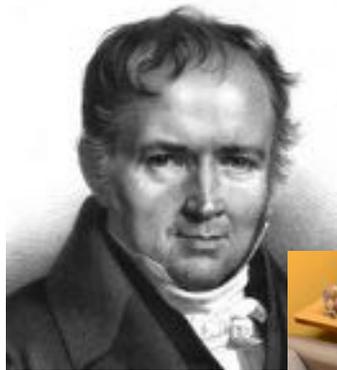
...to random variables and the Central Limit Theorem...



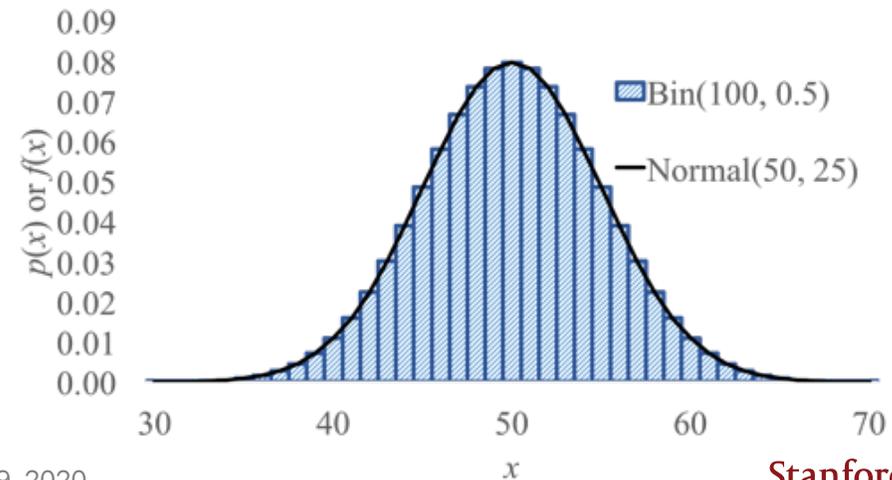
Bernoulli



Gaussian



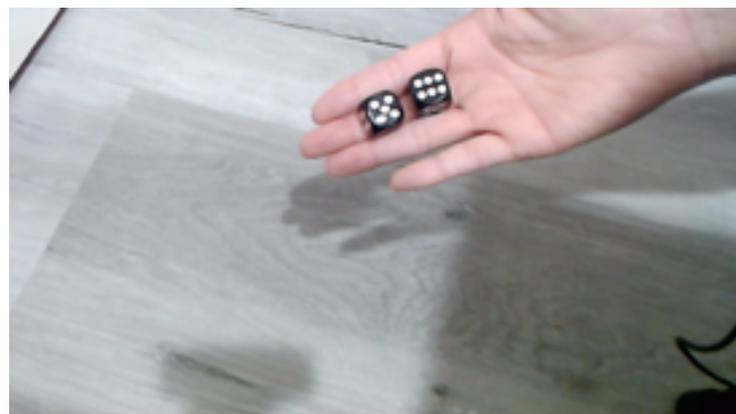
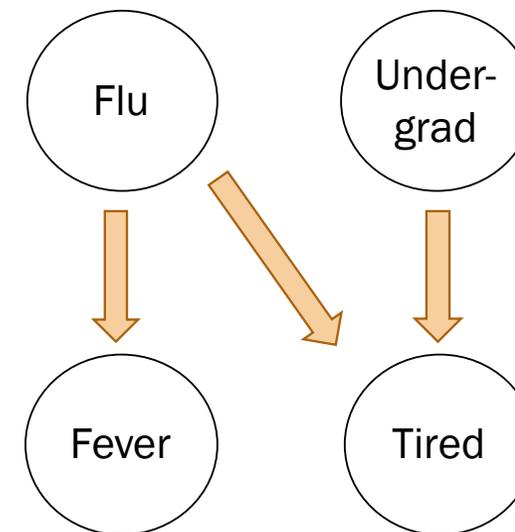
Poisson



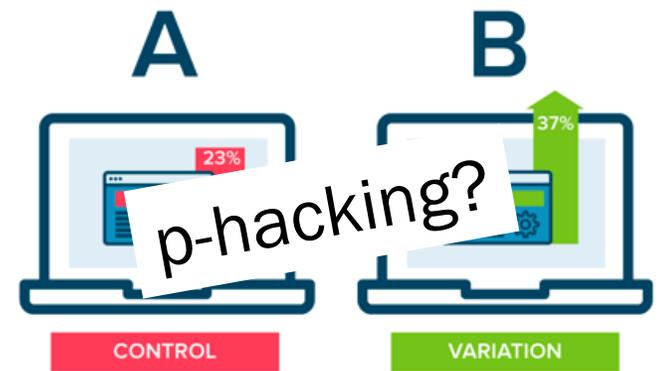
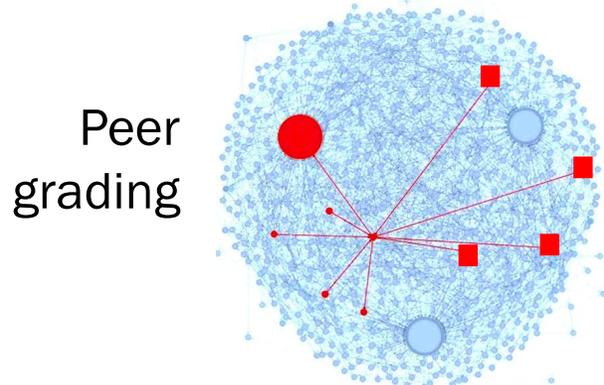
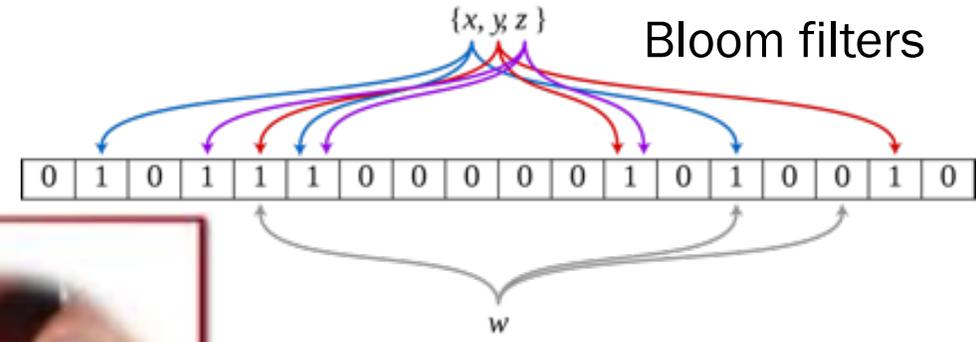
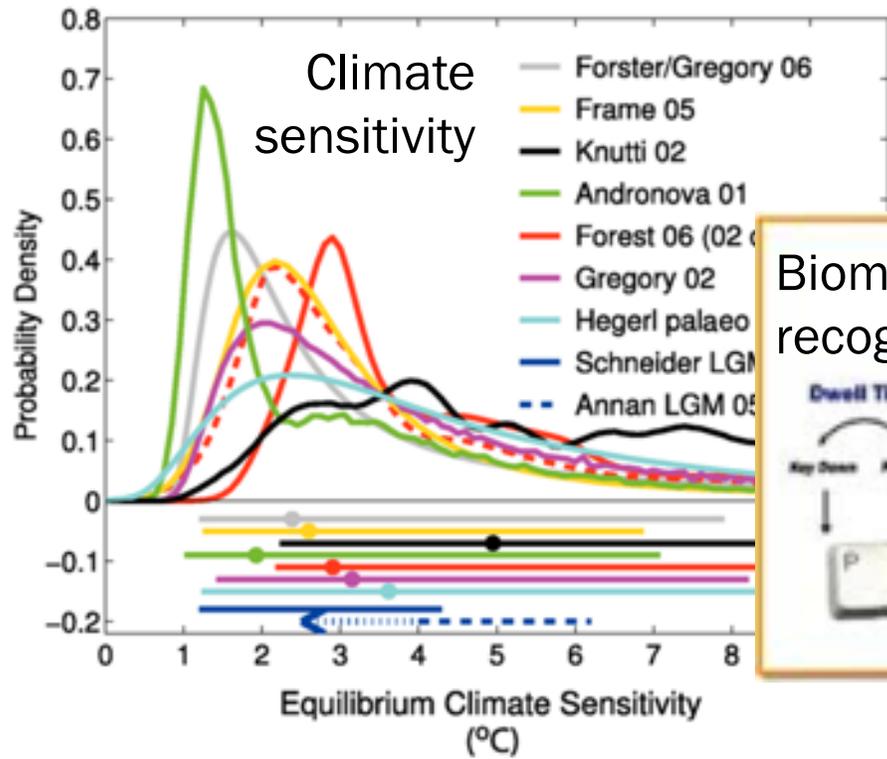
...to statistics, parameter estimation, and machine learning



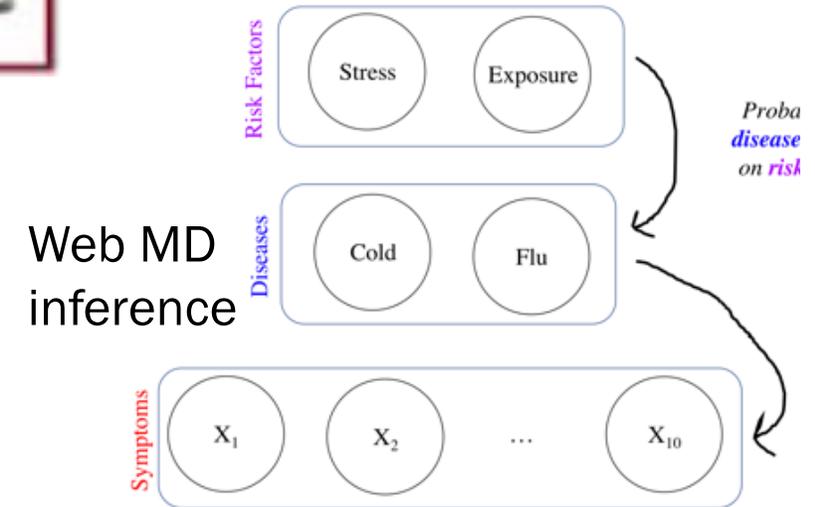
A happy
Bhutanese person



Lots and lots of analysis

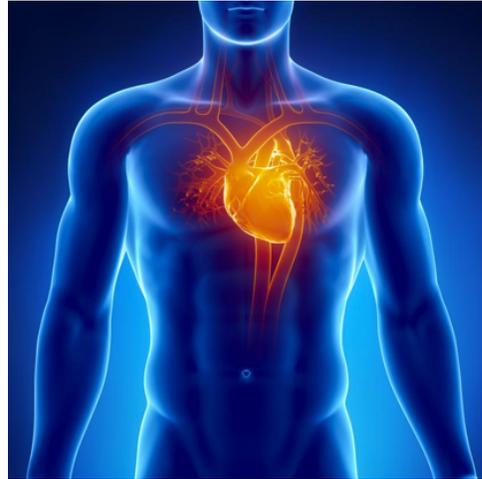


Coursera A/B testing

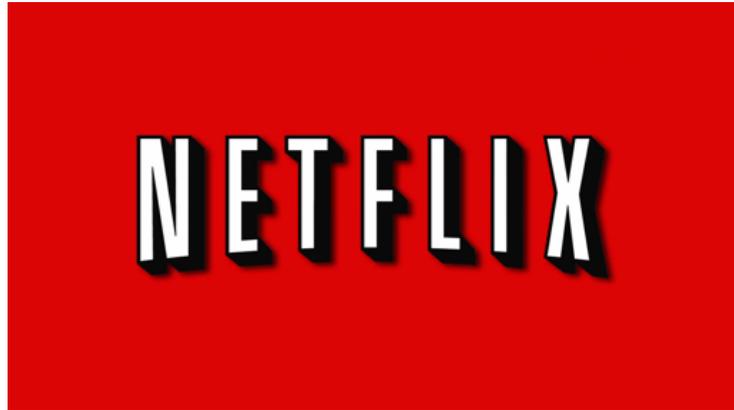


Lots and lots of analysis

Heart



Ancestry



NETFLIX

Netflix

After CS109

Theory

CS161 – Algorithmic analysis

Stats 217 – Stochastic Processes

CS238 – Decision Making Under Uncertainty

CS228 – Probabilistic Graphical Models

Statistics

Stats 200 – Statistical Inference

Stats 208 – Intro to the Bootstrap

Stats 209 – Group Methods/Causal Inference

After CS109

AI

CS 221 – Intro to AI

CS 229 – Machine Learning

CS 230 – Deep Learning

CS 224N – Natural Language Processing

CS 231N – Conv Neural Nets for Visual Recognition

CS 234 – Reinforcement Learning

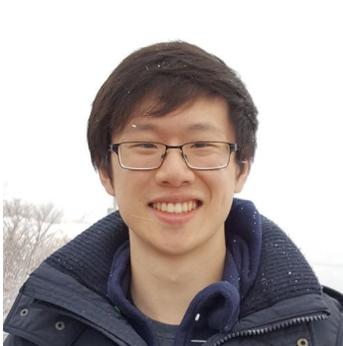
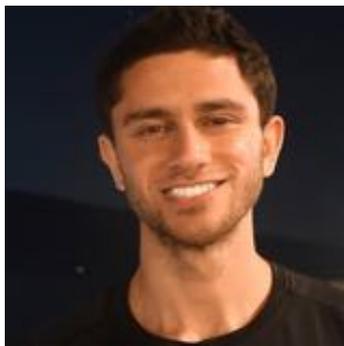
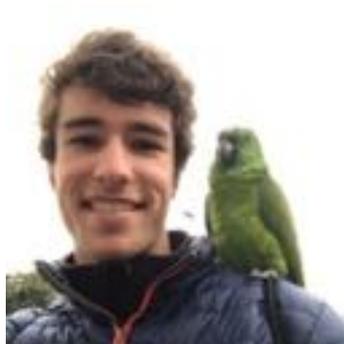
Applications

CS 279 – Bio Computation

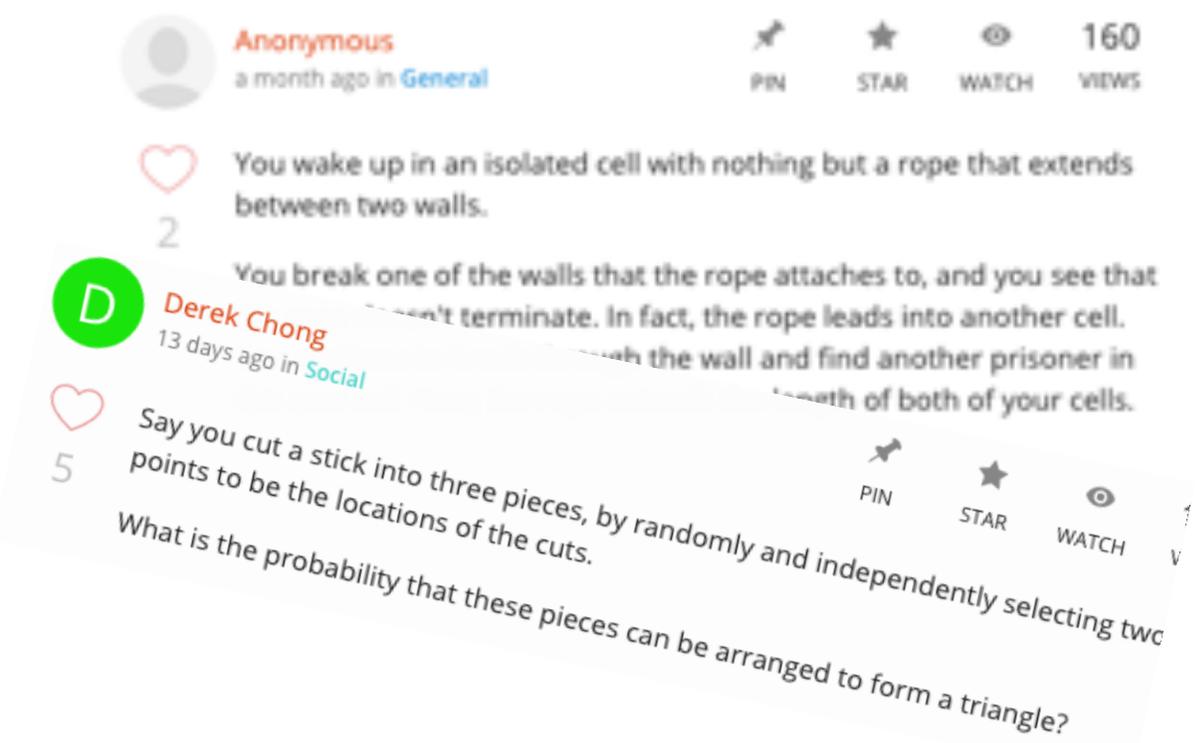
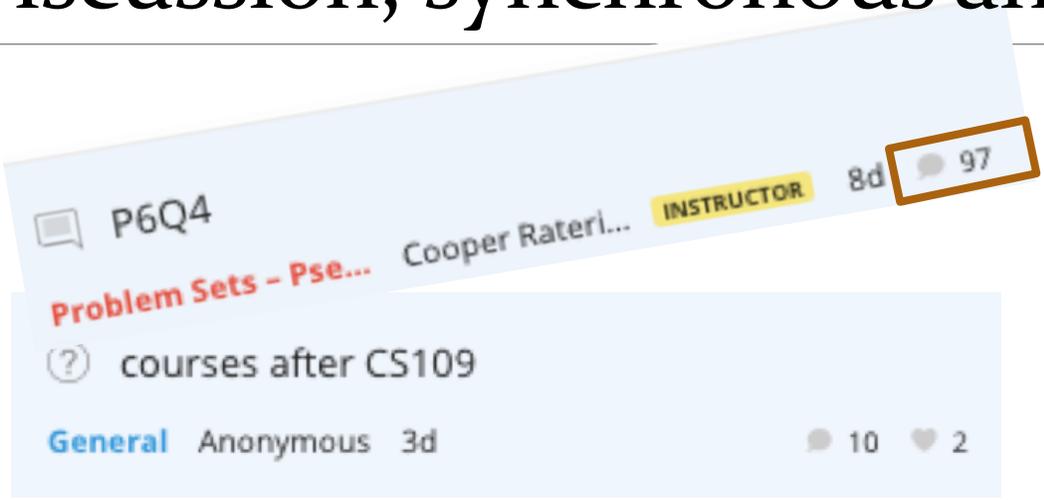
Literally any class with numbers in it

What have we done together
this quarter?

The CS109 teaching team



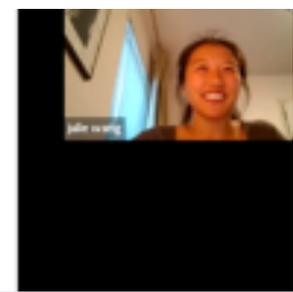
Discussion, synchronous and asynchronous



Negative Binomial $X \sim \text{Neg Bin}(r, p)$
 # of trials until r successes
 PMF: $P(X=k) = \binom{k-1}{r-1} (1-p)^{k-r} p^r$
 k-r failures r successes

Poisson RV discrete / Number of successes over experiment duration
 $X \sim \text{Poi}(\lambda)$ PMF

If $Z \sim \text{Poi}(\lambda=2.3)$
 $P(Z=2.3)$?



Section

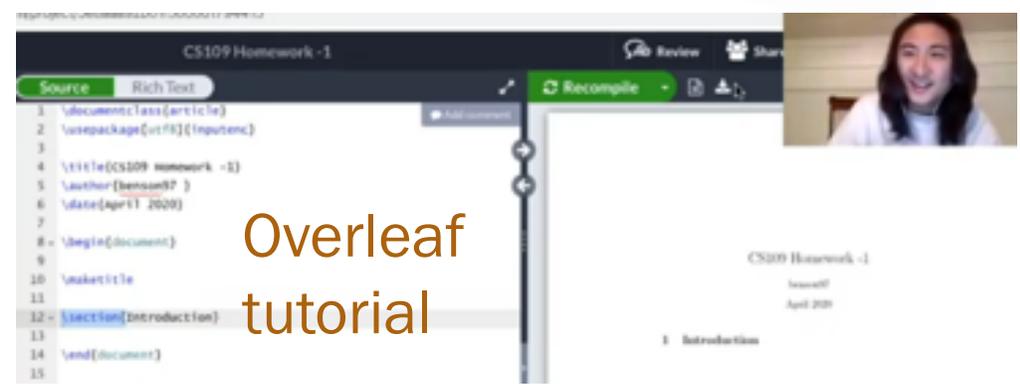
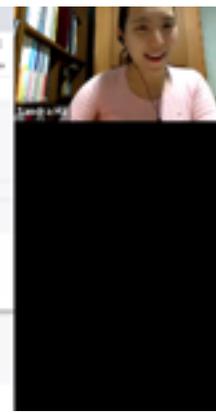
```
#np.random.seed(0) # setting the randomness to be consistent
a = np.random.random()
print(a)
e = np.random.random((2,2)) # Create an array filled with random values
print(e)

0.9428953126593451
[[0.70765589 0.23404159]
 [0.77035459 0.8800494 ]]

f = np.arange(10) # arrange
print(f)

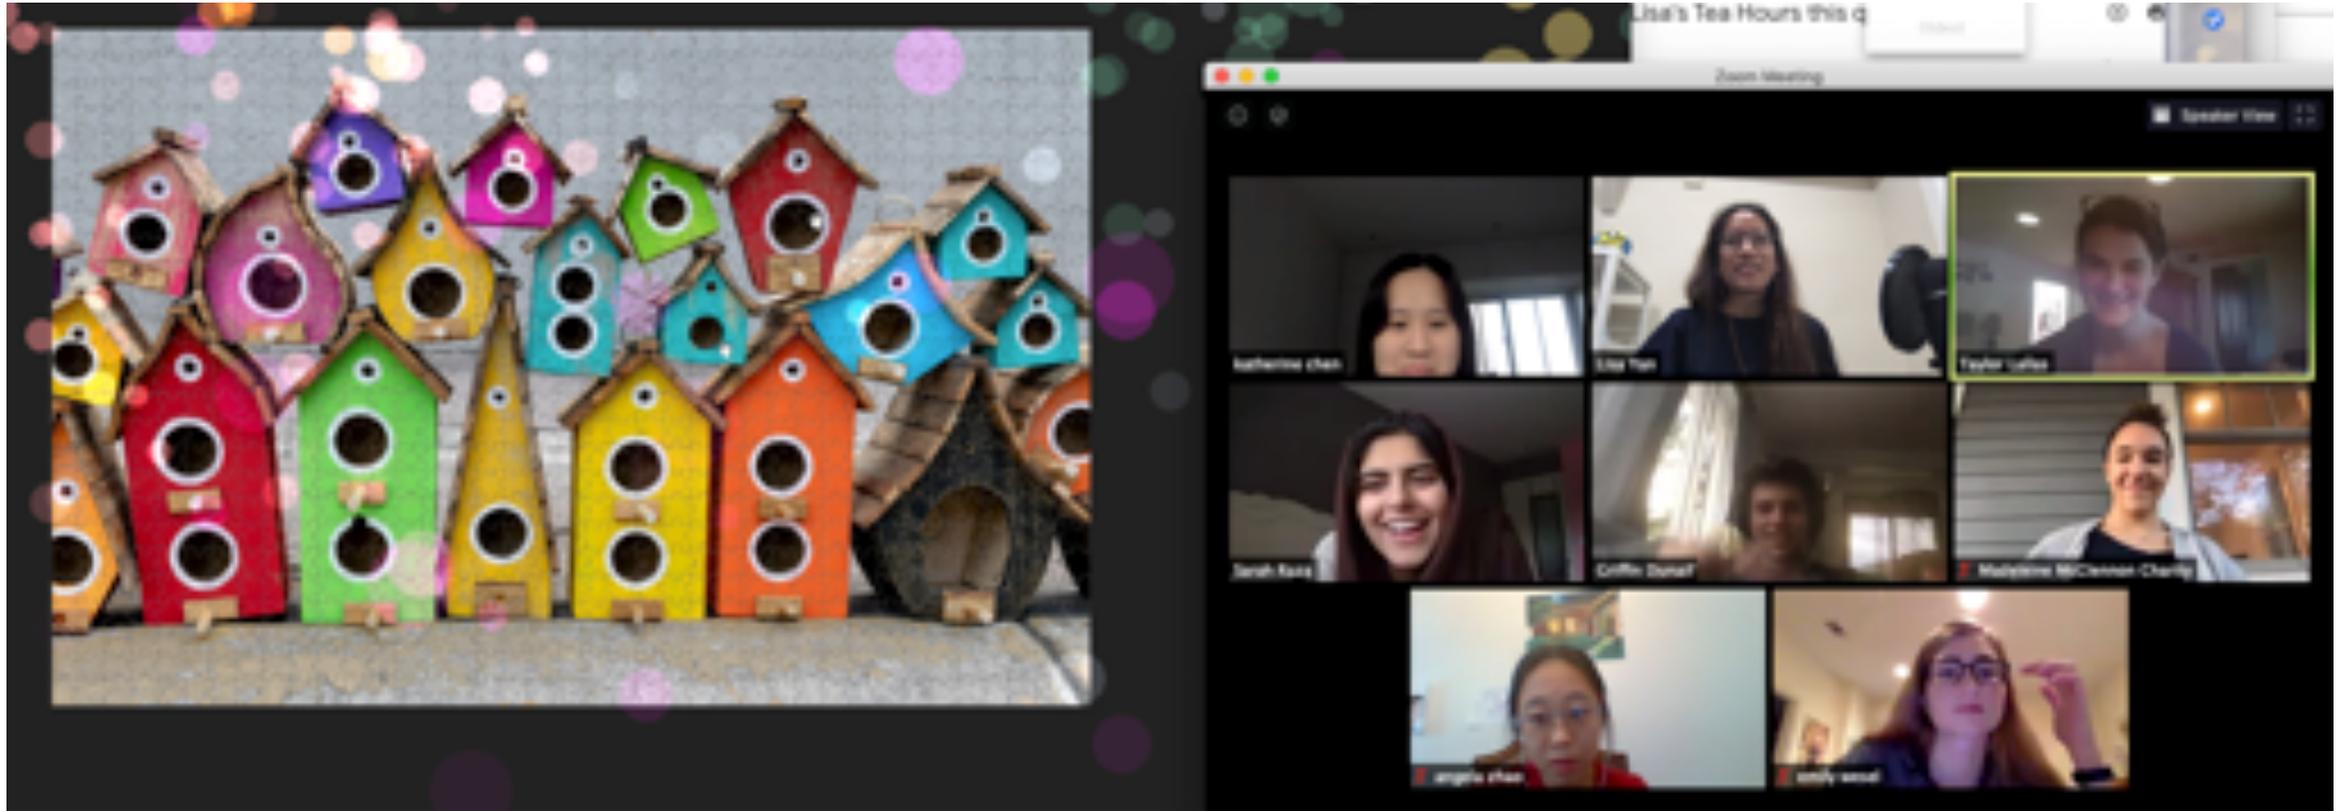
[0 1 2 3 4 5 6 7 8 9]
```

Python session



Overleaf tutorial

Tea hours



Turns out jigsaw puzzles aren't probabilistic!

What have you learned
from CS109?

What do you want to
remember in 5 years?

What do you want to remember in the next 5 years?

Probability has two school of thoughts, frequentist and Bayesian

Bayes Theorem

How to count (lol who knew I'd learn to count during college)

To be honest, counting.

Counting!

How to use random variables, bootstrapping, and random sampling to analyze data.

LOTUS because it helped me realize how I use stats to reason about the world in my everyday life

Always switch the door when playing Monty Hall!

This machine learning stuff! So cool!

Statistics can be p-hacked to misinterpret results

Machine Learning algorithms and Bayesian statistics!
(also "cancellation city")

I want to be able to take data and be able to structure and model it in a way that makes it useful to me and the world!

That I was finally able to code a Naive Bayes Classifier

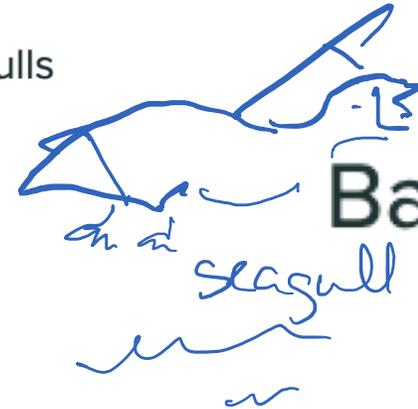
What do you want to remember in the next 5 years?

Poisson!! the shark... but hopefully also the distribution

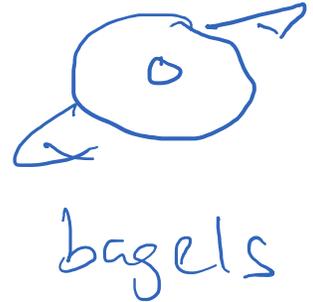
poisson shark



bay-gulls



Bay-gulls



seagulls in the bay area are actually bagels!

Interlude for jokes LOL.

Lisa's jokes

the jokes!

Lisa's jokes, of course! A close second is bootstrapping! I think it is the coolest

the fun times!

What do you want to remember in the next 5 years?

The Probability in the Real World segments

I learned a lot.

How much I struggled.

I want to remember the fundamentals of probability and the commitment of the teaching staff to making this tough quarter as smooth as possible.

I want to remember how probability is so relevant to many different things in the world :

I never knew how much probability was around me!

I'll want to remember how much learning helped me deal with the emotions of the quarter.

I want to remember how I learned about COVID-19 tracking through researching articles for CS109.

I'll remember feeling excited about stats

Thank you so much CS109! You will be missed.

Why study probability + CS?

Why study probability + CS?

Fastest growing occupations: 20 occupations with the highest percent change of employment between 2018-28.

Click on an occupation name to see the full occupational profile.

OCCUPATION	GROWTH RATE, 2018-28	2018 MEDIAN PAY
Physician assistants	31%	\$108,610 per year
Nurse practitioners	28%	\$107,030 per year
Software developers, applications	26%	\$103,620 per year
Mathematicians	26%	\$101,900 per year
Information security analysts	32%	\$98,350 per year
Health specialties teachers, postsecondary	23%	\$97,370 per year
Statisticians	31%	\$87,780 per year
Operations research analysts	26%	\$83,390 per year
Genetic counselors	27%	\$80,370 per year



Source: [US Bureau of Labor Statistics](#)

Why study probability + CS?



Interdisciplinary



Closest thing to magic

Why study probability + CS?



Everyone is welcome!

Technology magnifies.
What do we want
magnified?

You are all one step closer to
improving the world.

(all of you!)

The end



See you soon... 😊