



PennState
College of the
Liberal Arts



Day 4/5 - Introduction to Neural Nets / Deep Learning for NLP

Advanced Text as Data: Natural Language Processing
Essex Summer School in Social Science Data Analysis

Burt L. Monroe (Instructor) & Sam Bestvater (TA)
Pennsylvania State University


July 29-30, 2021

Why Deep Learning?

Hand engineered features are time consuming, brittle, and not scalable in practice

Can we learn the **underlying features** directly from data?

Low Level Features



Lines & Edges

Mid Level Features



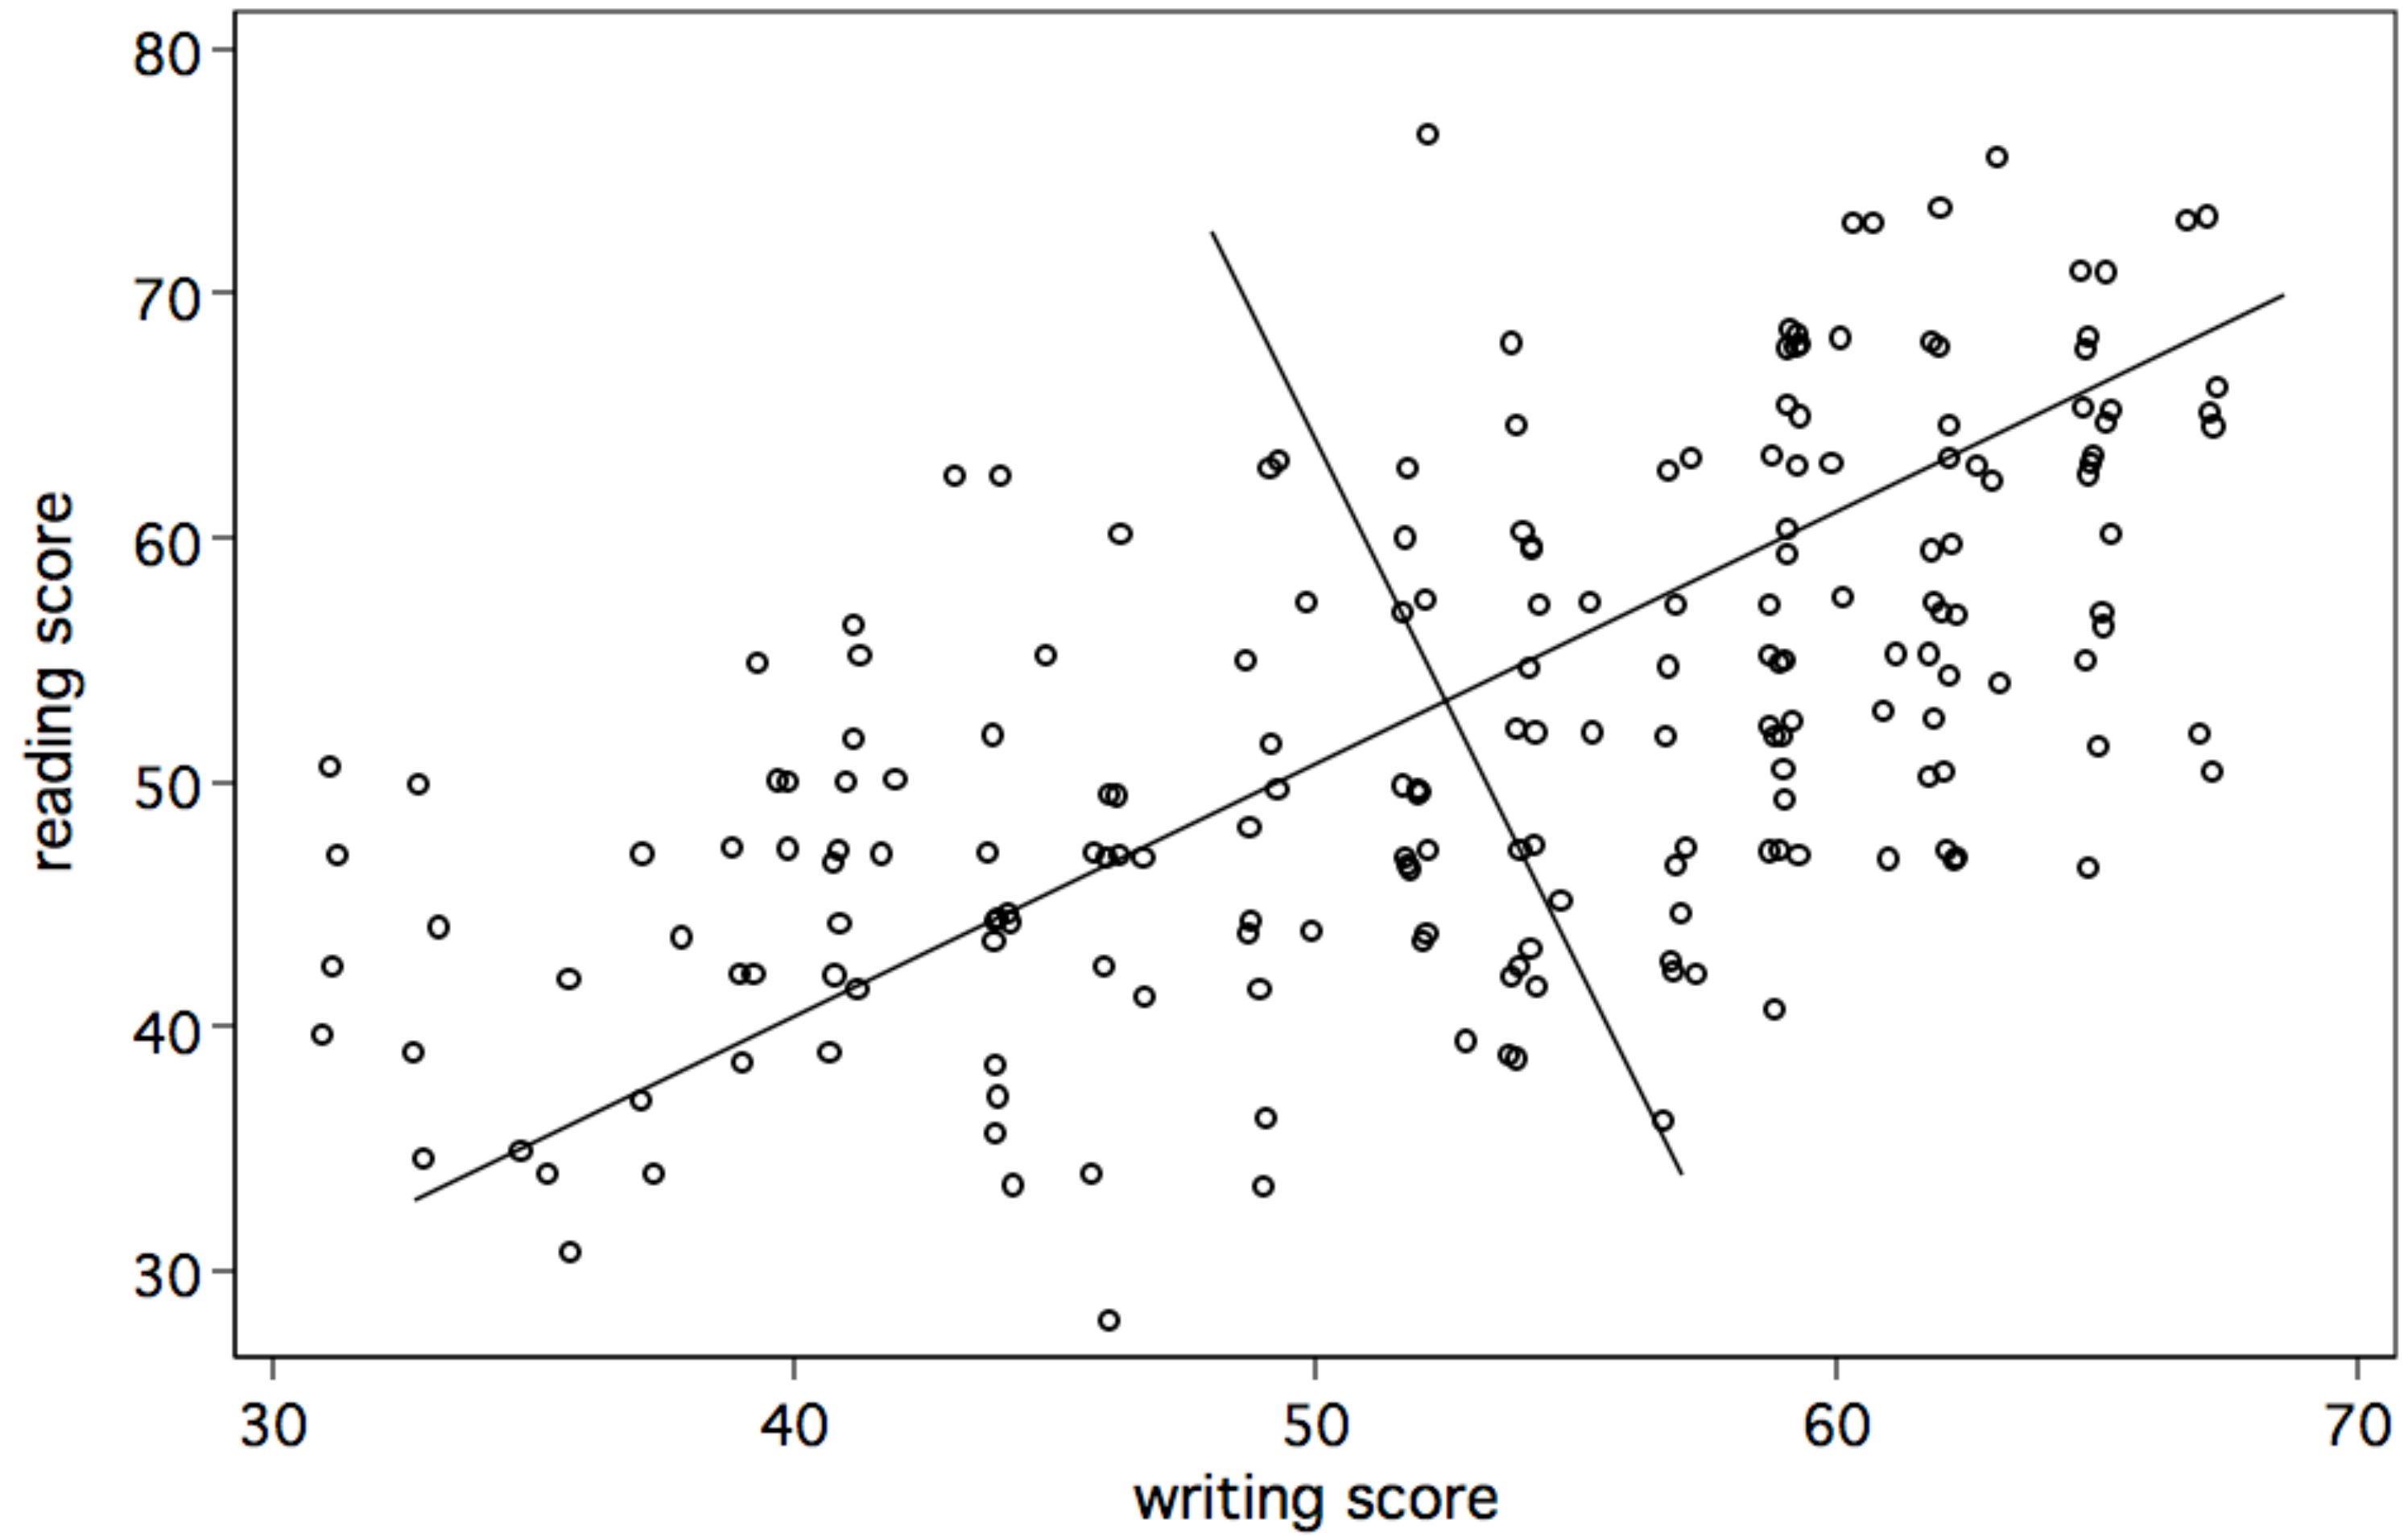
Eyes & Nose & Ears

High Level Features




Facial Structure

You've probably seen a lot of **linear** approaches to finding new **feature representations**:




You've probably seen a lot of **linear** approaches to finding new **feature representations**:

We might do this to find interpretable or intuitive latent concepts.




You've probably seen a lot of **linear** approaches to finding new **feature representations**:

We might do this to make computing more efficient (e.g., orthogonalization).





You've probably seen a lot of **linear** approaches to finding new **feature representations**:

We might do this to reduce dimensionality for generalizability or compression

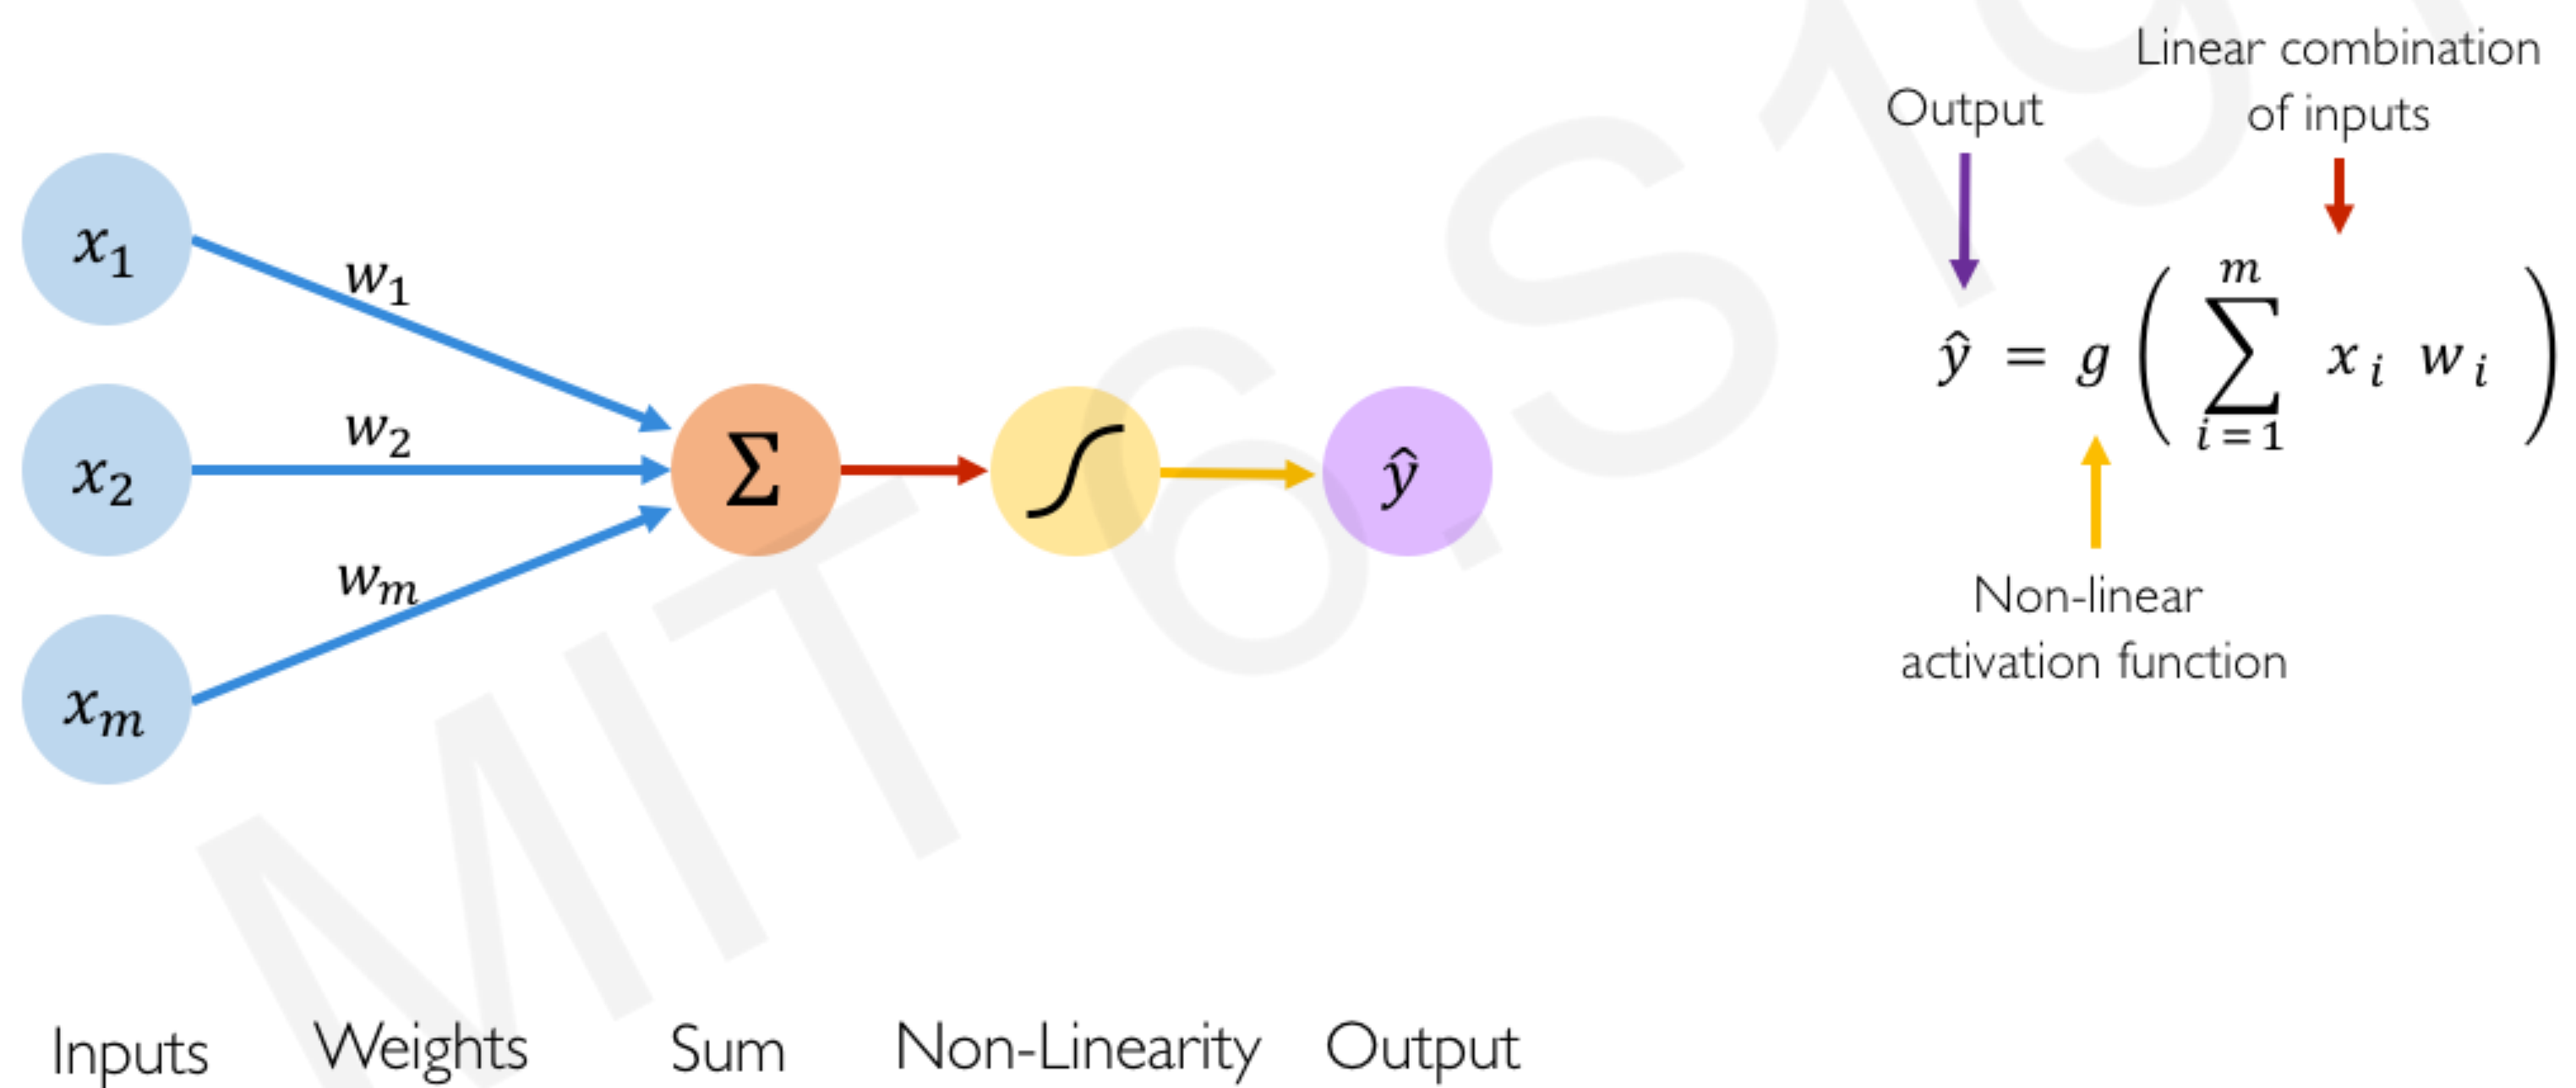


Domain knowledge may allow us to do successful *feature engineering*.


Figure 4.3. Feature engineering for reading the time on a clock

Raw data: pixel grid		
Better features: clock hands' coordinates	{x1: 0.7, y1: 0.7} {x2: 0.5, y2: 0.0}	{x1: 0.0, y2: 1.0} {x2: -0.38, 2: 0.32}
Even better features: angles of clock hands	theta1: 45 theta2: 0	theta1: 90 theta2: 140

The Perceptron: Forward Propagation



The Perceptron: Forward Propagation



Output

$$\hat{y} = g \left(w_0 + \sum_{i=1}^m x_i w_i \right)$$

Linear combination of inputs

Non-linear activation function

Bias


The Perceptron: Forward Propagation

“It was **great.**”

great

banana

worst



Output

$$\hat{y} = g \left(w_0 + \sum_{i=1}^m x_i w_i \right)$$

Linear combination of inputs


Non-linear activation function

Bias

The Perceptron: Forward Propagation

“The **worst.**”

great
banana
worst



Output

Linear combination of inputs

$$\hat{y} = g \left(w_0 + \sum_{i=1}^m x_i w_i \right)$$


Non-linear activation function

Bias

The Perceptron: Forward Propagation

“Acting was **great**.
Worst script ever.”

great
banana
worst



Output

Linear combination of inputs


$$\hat{y} = g \left(w_0 + \sum_{i=1}^m x_i w_i \right)$$

Non-linear activation function

Bias

The Perceptron: Forward Propagation

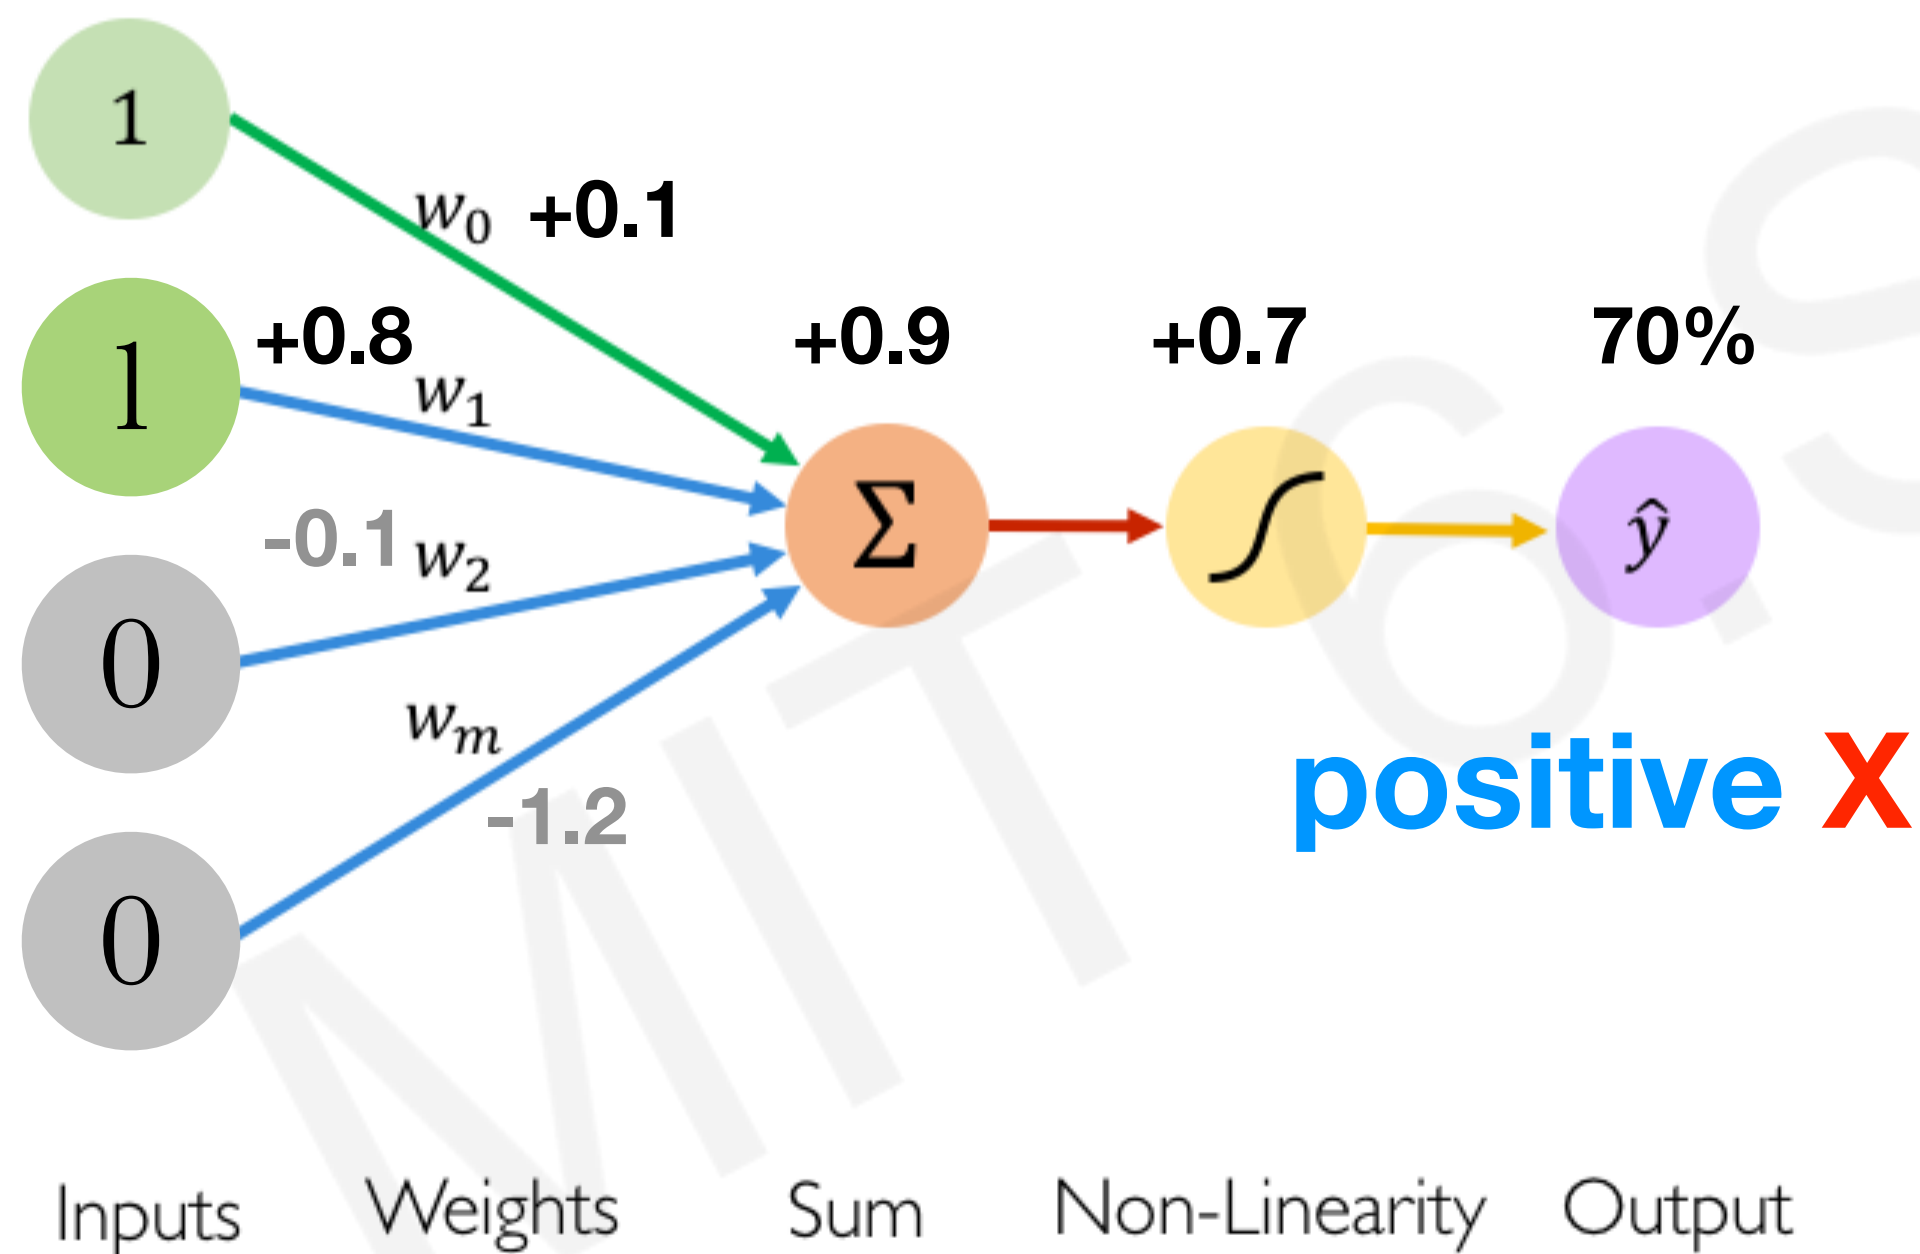
“I want a **banana**.”



The Perceptron: Forward Propagation

“Not great.”

great
banana
worst



Output


Linear combination of inputs

$$\hat{y} = g \left(w_0 + \sum_{i=1}^m x_i w_i \right)$$

Non-linear activation function

Bias

The Perceptron: Forward Propagation




$$\hat{y} = g \left(w_0 + \sum_{i=1}^m x_i w_i \right)$$

$$\hat{y} = g (w_0 + \mathbf{X}^T \mathbf{W})$$

$$\text{where: } \mathbf{X} = \begin{bmatrix} x_1 \\ \vdots \\ x_m \end{bmatrix} \text{ and } \mathbf{W} = \begin{bmatrix} w_1 \\ \vdots \\ w_m \end{bmatrix}$$

The Perceptron: Forward Propagation




Activation Functions

$$\hat{y} = g(w_0 + \mathbf{X}^T \mathbf{W})$$


- Example: sigmoid function

$$g(z) = \sigma(z) = \frac{1}{1 + e^{-z}}$$



This is *exactly the same* as the logit / logistic regression “inverse link function.”

The Perceptron: Forward Propagation




Activation Functions

$$\hat{y} = g(w_0 + \mathbf{X}^T \mathbf{W})$$

- Example: sigmoid function

$$g(z) = \sigma(z) = \frac{1}{1 + e^{-z}}$$




Machine learning in one slide

Social science (inference)	Machine learning (prediction)
GLM inverse link function	Activation function
$\mathbb{E}(y) = f(\mathbf{x}'\boldsymbol{\beta})$	$\mathbb{E}(y) = f(\mathbf{x}'\boldsymbol{\beta})$
Preferred objective function	
Log-likelihood	Cross-entropy
$\log \mathcal{L} = \sum_{i=1}^n \log P(y_i \mathbf{x}_i, \boldsymbol{\beta})$	$-\log \mathcal{L} = -\sum_{i=1}^n \log P(y_i \mathbf{x}_i, \boldsymbol{\beta})$
Solving algorithm	
Newton-Raphson	Gradient descent
$\boldsymbol{\beta}_t := \boldsymbol{\beta}_{t-1} - [\mathbf{H} \log \mathcal{L}]^{-1} \nabla \log \mathcal{L}$	$\boldsymbol{\beta}_t := \boldsymbol{\beta}_{t-1} - \eta \nabla (-\log \mathcal{L})$
Quantities of interest	
$\hat{\boldsymbol{\beta}}; \text{Var}(\hat{\boldsymbol{\beta}})$	$\hat{y}; \sum \mathbf{1}(\hat{y} = y)/n$


Common Activation Functions

Sigmoid Function




$$g(z) = \frac{1}{1 + e^{-z}}$$

$$g'(z) = g(z)(1 - g(z))$$


 `tf.math.sigmoid(z)`

Hyperbolic Tangent




$$g(z) = \frac{e^z - e^{-z}}{e^z + e^{-z}}$$

$$g'(z) = 1 - g(z)^2$$


 `tf.math.tanh(z)`

Rectified Linear Unit (ReLU)



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

$$g'(z) = \begin{cases} 1, & z > 0 \\ 0, & \text{otherwise} \end{cases}$$

 `tf.nn.relu(z)`

 TensorFlow code blocks

NOTE: All activation functions are non-linear

Common Activation Functions

```
model = models.Sequential()
model.add(layers.Dense(16, activation = 'relu', input_shape=(5000,)))
model.add(layers.Dense(16, activation = 'relu'))
model.add(layers.Dense(1, activation= 'sigmoid'))

model.compile(optimizer='adam',
              loss='binary_crossentropy',
              metrics=[ 'accuracy' ])

history = model.fit(partial_x_train,
                    partial_y_train,
                    epochs=4,
                    batch_size=512,
                    validation_data=(x_val,y_val))
```


```

{r}
model <- keras_model_sequential() %>%
  layer_dense(units = 16, activation = "relu", input_shape = c(5000)) %>%
  layer_dense(units = 16, activation = "relu") %>%
  layer_dense(units = 1, activation = "sigmoid")

model %>% compile(
  optimizer = "adam",
  loss = "binary_crossentropy",
  metrics = c("accuracy")
)

model %>% fit(x_train, y_train, epochs = 4, batch_size = 512)
results <- model %>% evaluate(x_test, y_test)

```



$$g'(z) = g(z)(1 - g(z))$$


```
tf.math.sigmoid(z)
```

$$g'(z) = 1 - g(z)^2$$

```
tf.math.tanh(z)
```

$$g'(z) = \begin{cases} 1, & z > 0 \\ 0, & \text{otherwise} \end{cases}$$


```
tf.nn.relu(z)
```




Building Neural Networks with Perceptrons

The Perceptron: Simplified

$$\hat{y} = g(w_0 + \mathbf{X}^T \mathbf{W})$$




The Perceptron: Simplified



$$z = w_0 + \sum_{j=1}^m x_j w_j$$

Multi Output Perceptron


Because all inputs are densely connected to all outputs, these layers are called **Dense** layers



$$z_i = w_{0,i} + \sum_{j=1}^m x_j w_{j,i}$$

Multi Output Perceptron


Because all inputs are densely connected to all outputs, these layers are called **Dense** layers



```
import tensorflow as tf
layer = tf.keras.layers.Dense(
    units=2)
```

$$z_i = w_{0,i} + \sum_{j=1}^m x_j w_{j,i}$$

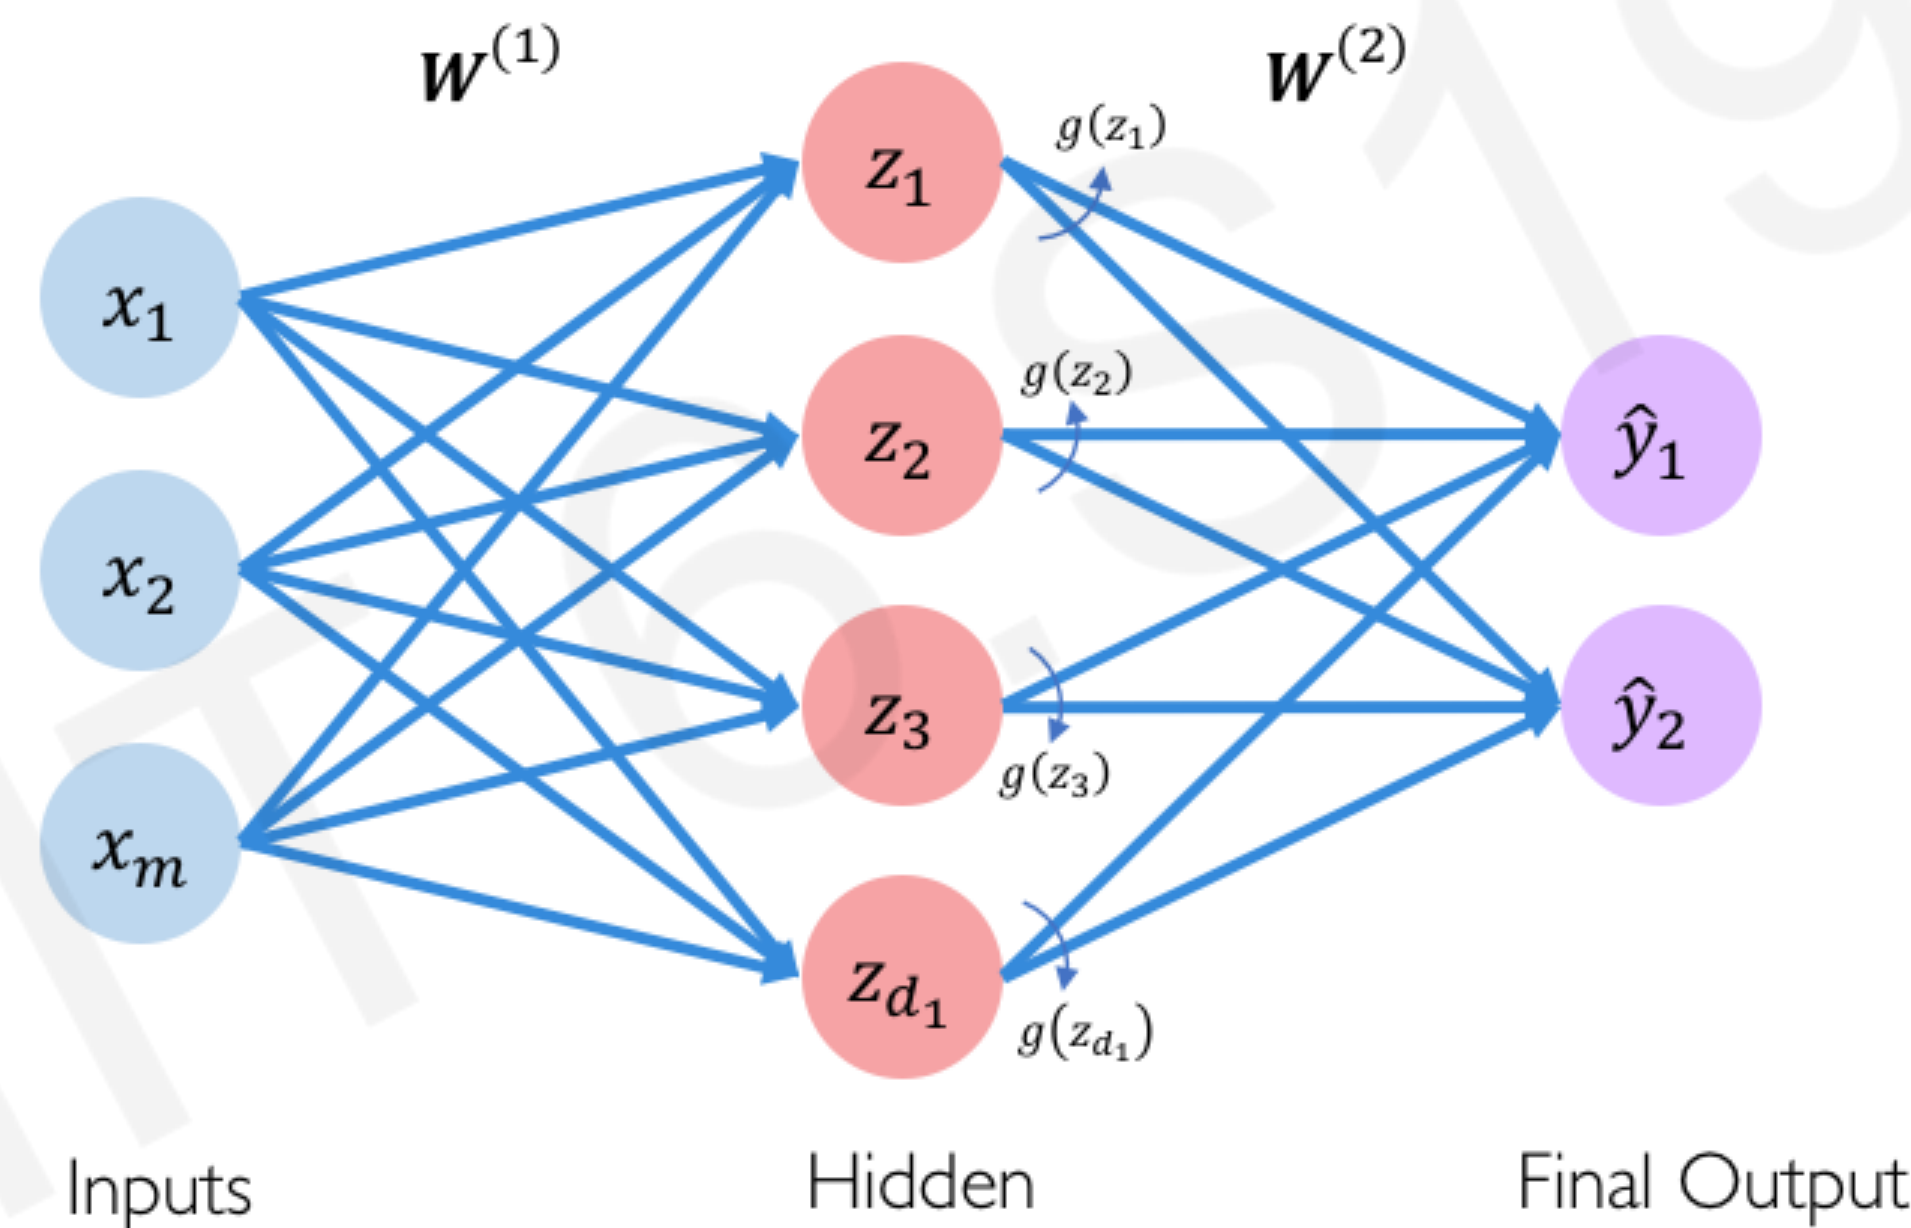
Single Layer Neural Network



$$z_i = w_{0,i}^{(1)} + \sum_{j=1}^m x_j w_{j,i}^{(1)} \quad \hat{y}_i = g \left(w_{0,i}^{(2)} + \sum_{j=1}^{d_1} g(z_j) w_{j,i}^{(2)} \right)$$

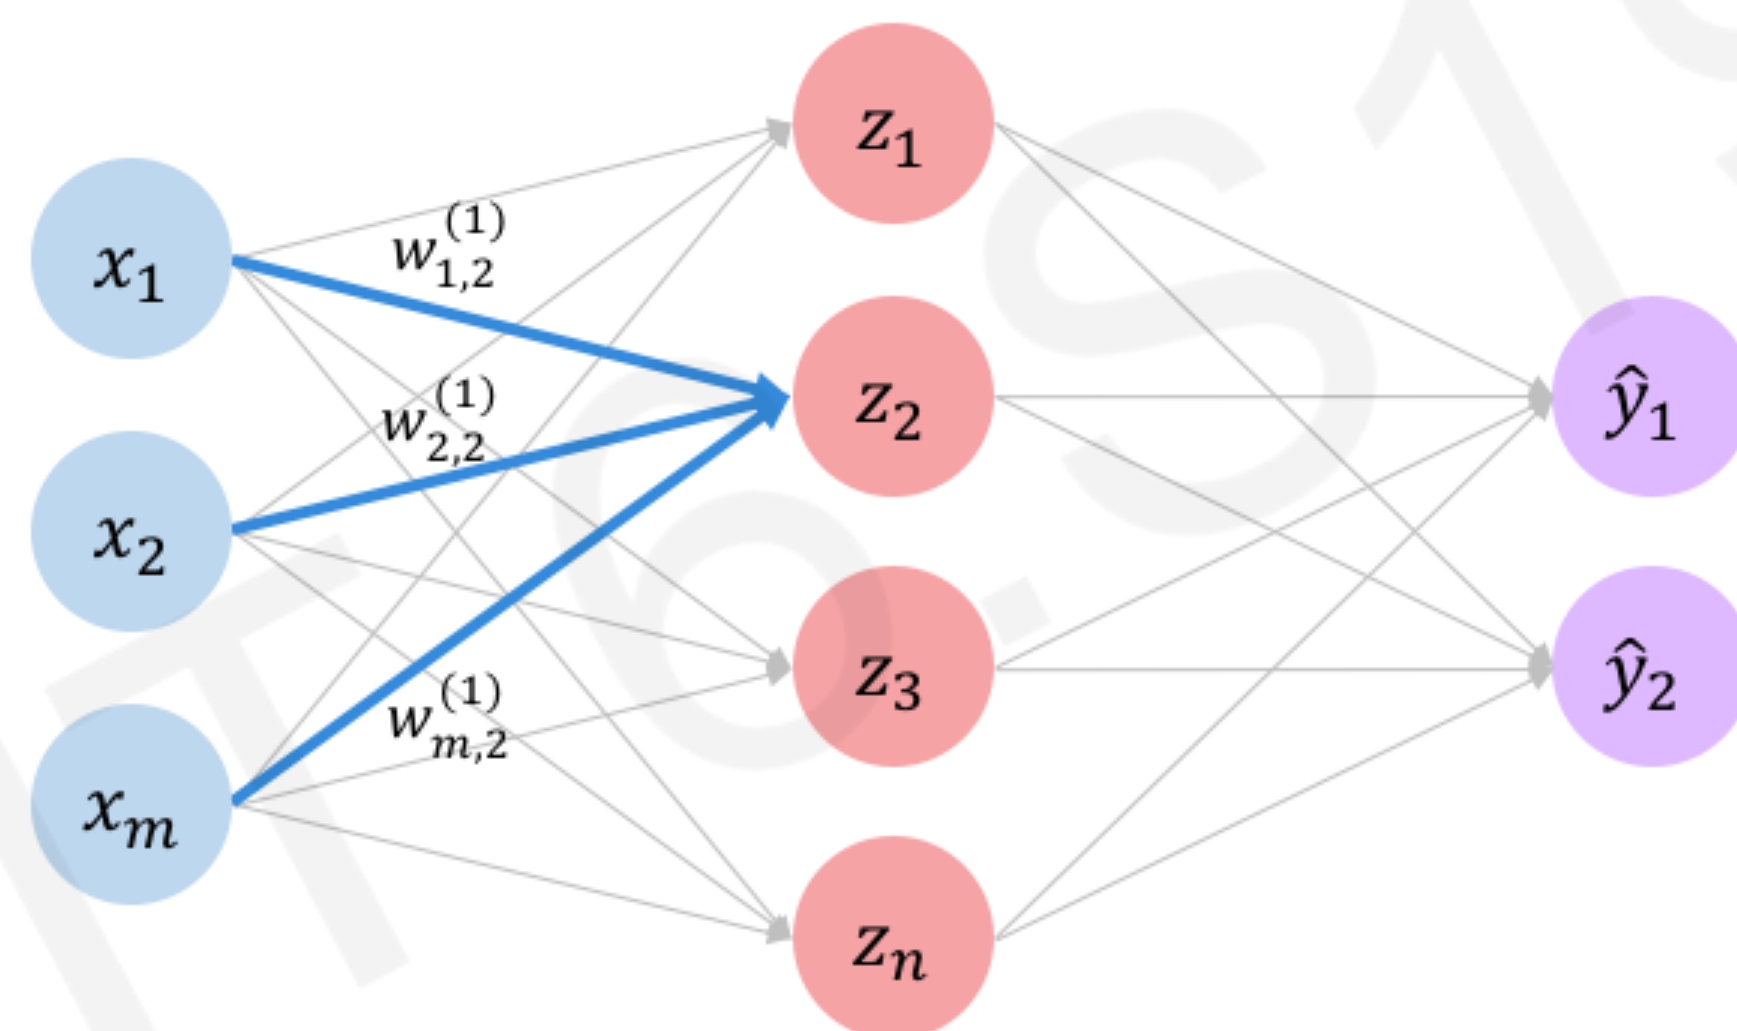
There are important contexts — like modeling with Keras — where we would refer to this as a “two-layer” neural network

Single Layer Neural Network




$$z_i = w_{0,i}^{(1)} + \sum_{j=1}^m x_j w_{j,i}^{(1)} \quad \hat{y}_i = g \left(w_{0,i}^{(2)} + \sum_{j=1}^{d_1} g(z_j) w_{j,i}^{(2)} \right)$$

Single Layer Neural Network



$$\begin{aligned} z_2 &= w_{0,2}^{(1)} + \sum_{j=1}^m x_j w_{j,2}^{(1)} \\ &= w_{0,2}^{(1)} + x_1 w_{1,2}^{(1)} + x_2 w_{2,2}^{(1)} + x_m w_{m,2}^{(1)} \end{aligned}$$

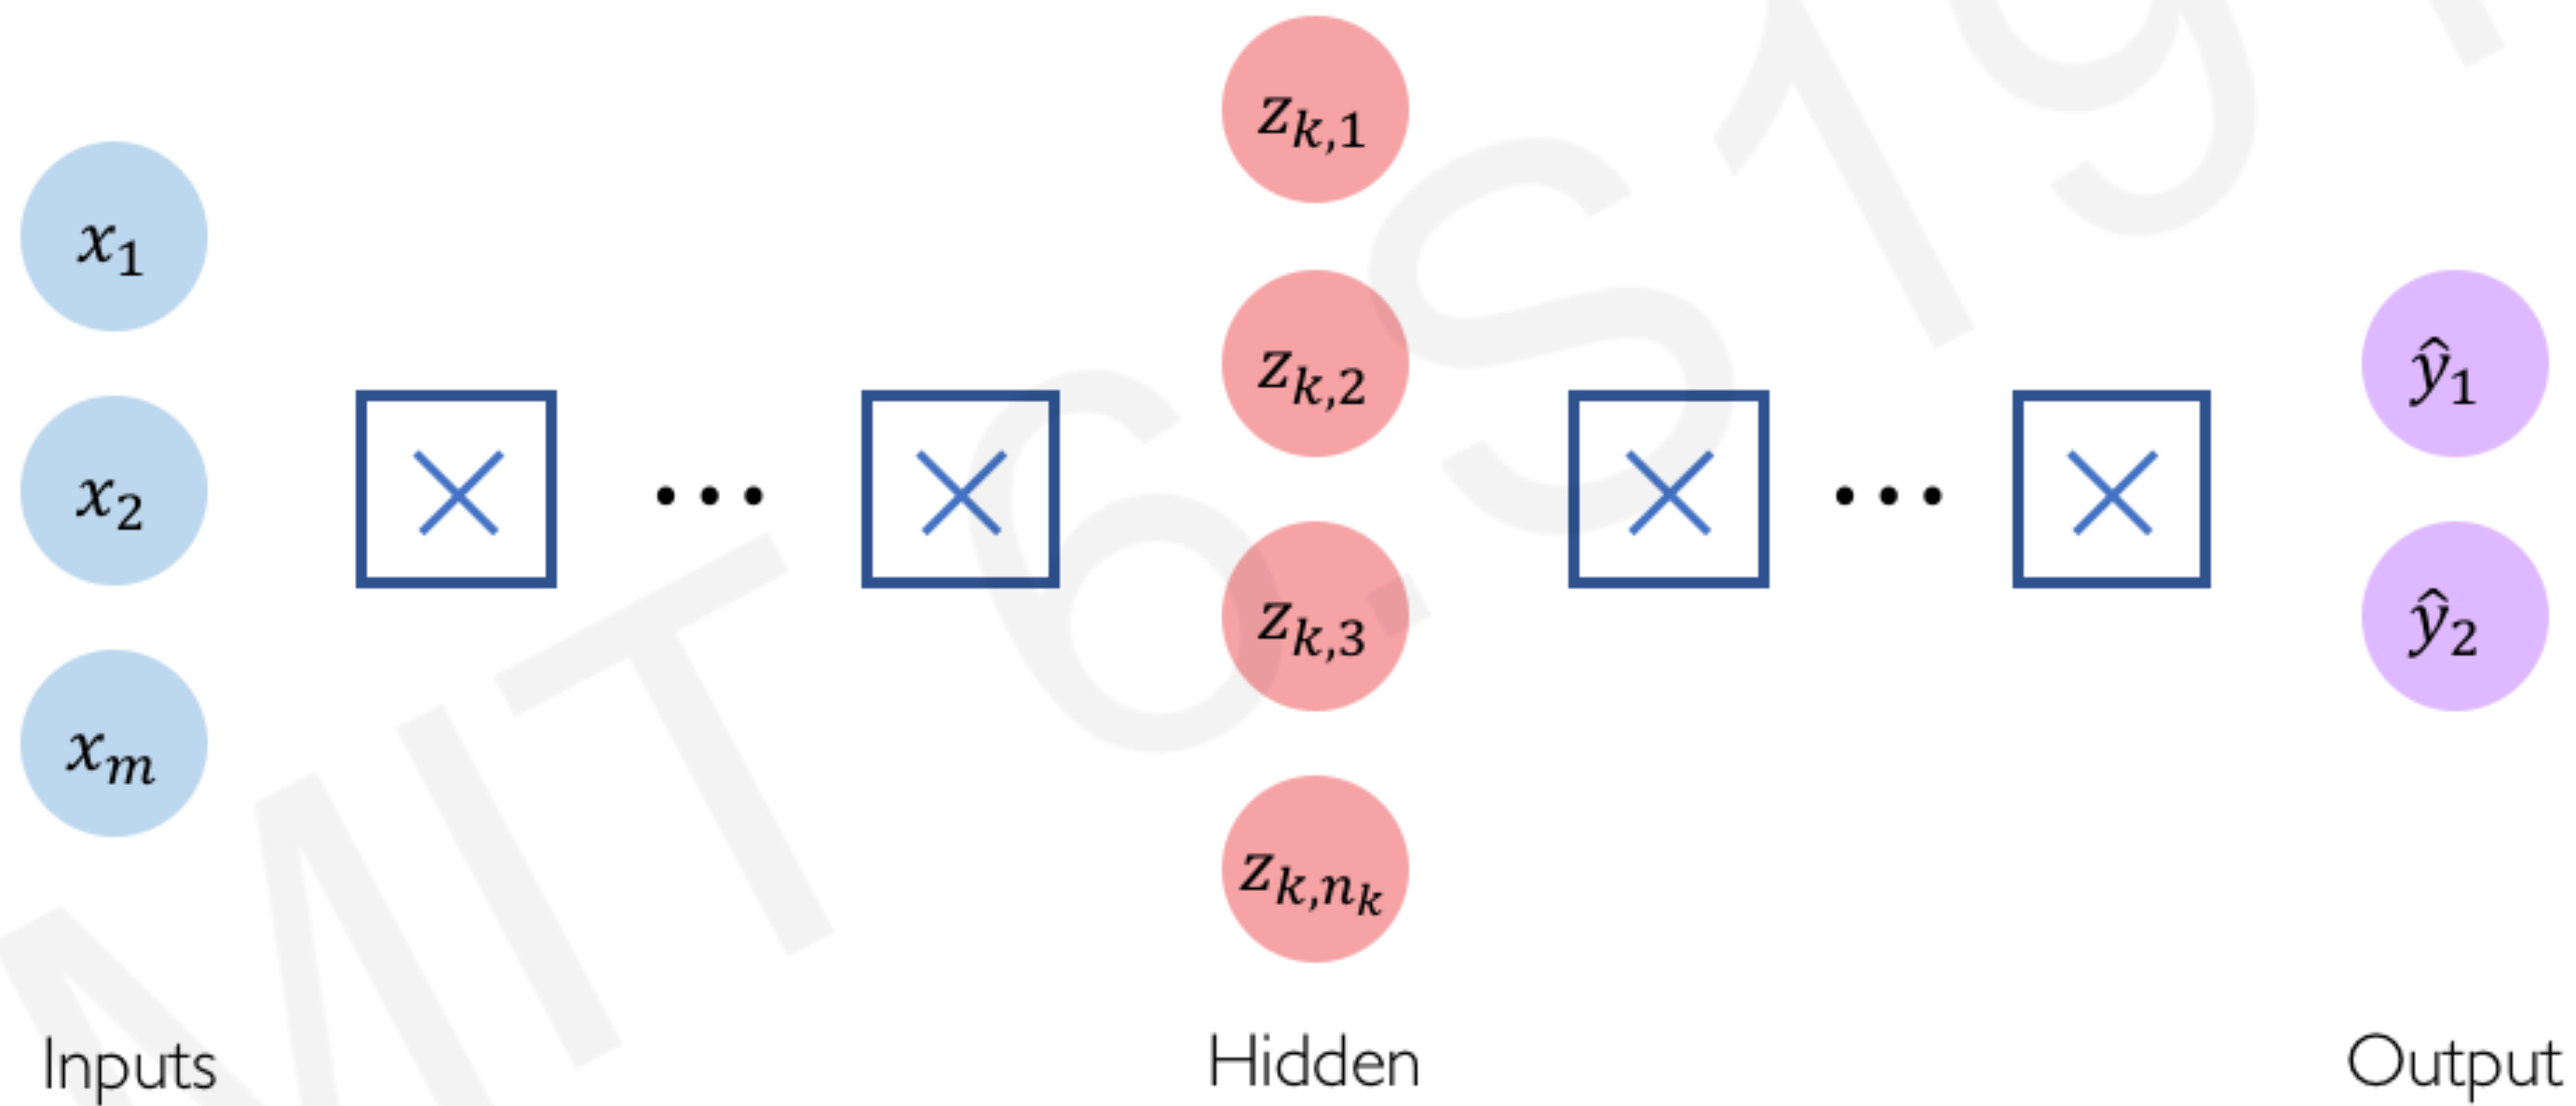
Multi Output Perceptron



```
import tensorflow as tf

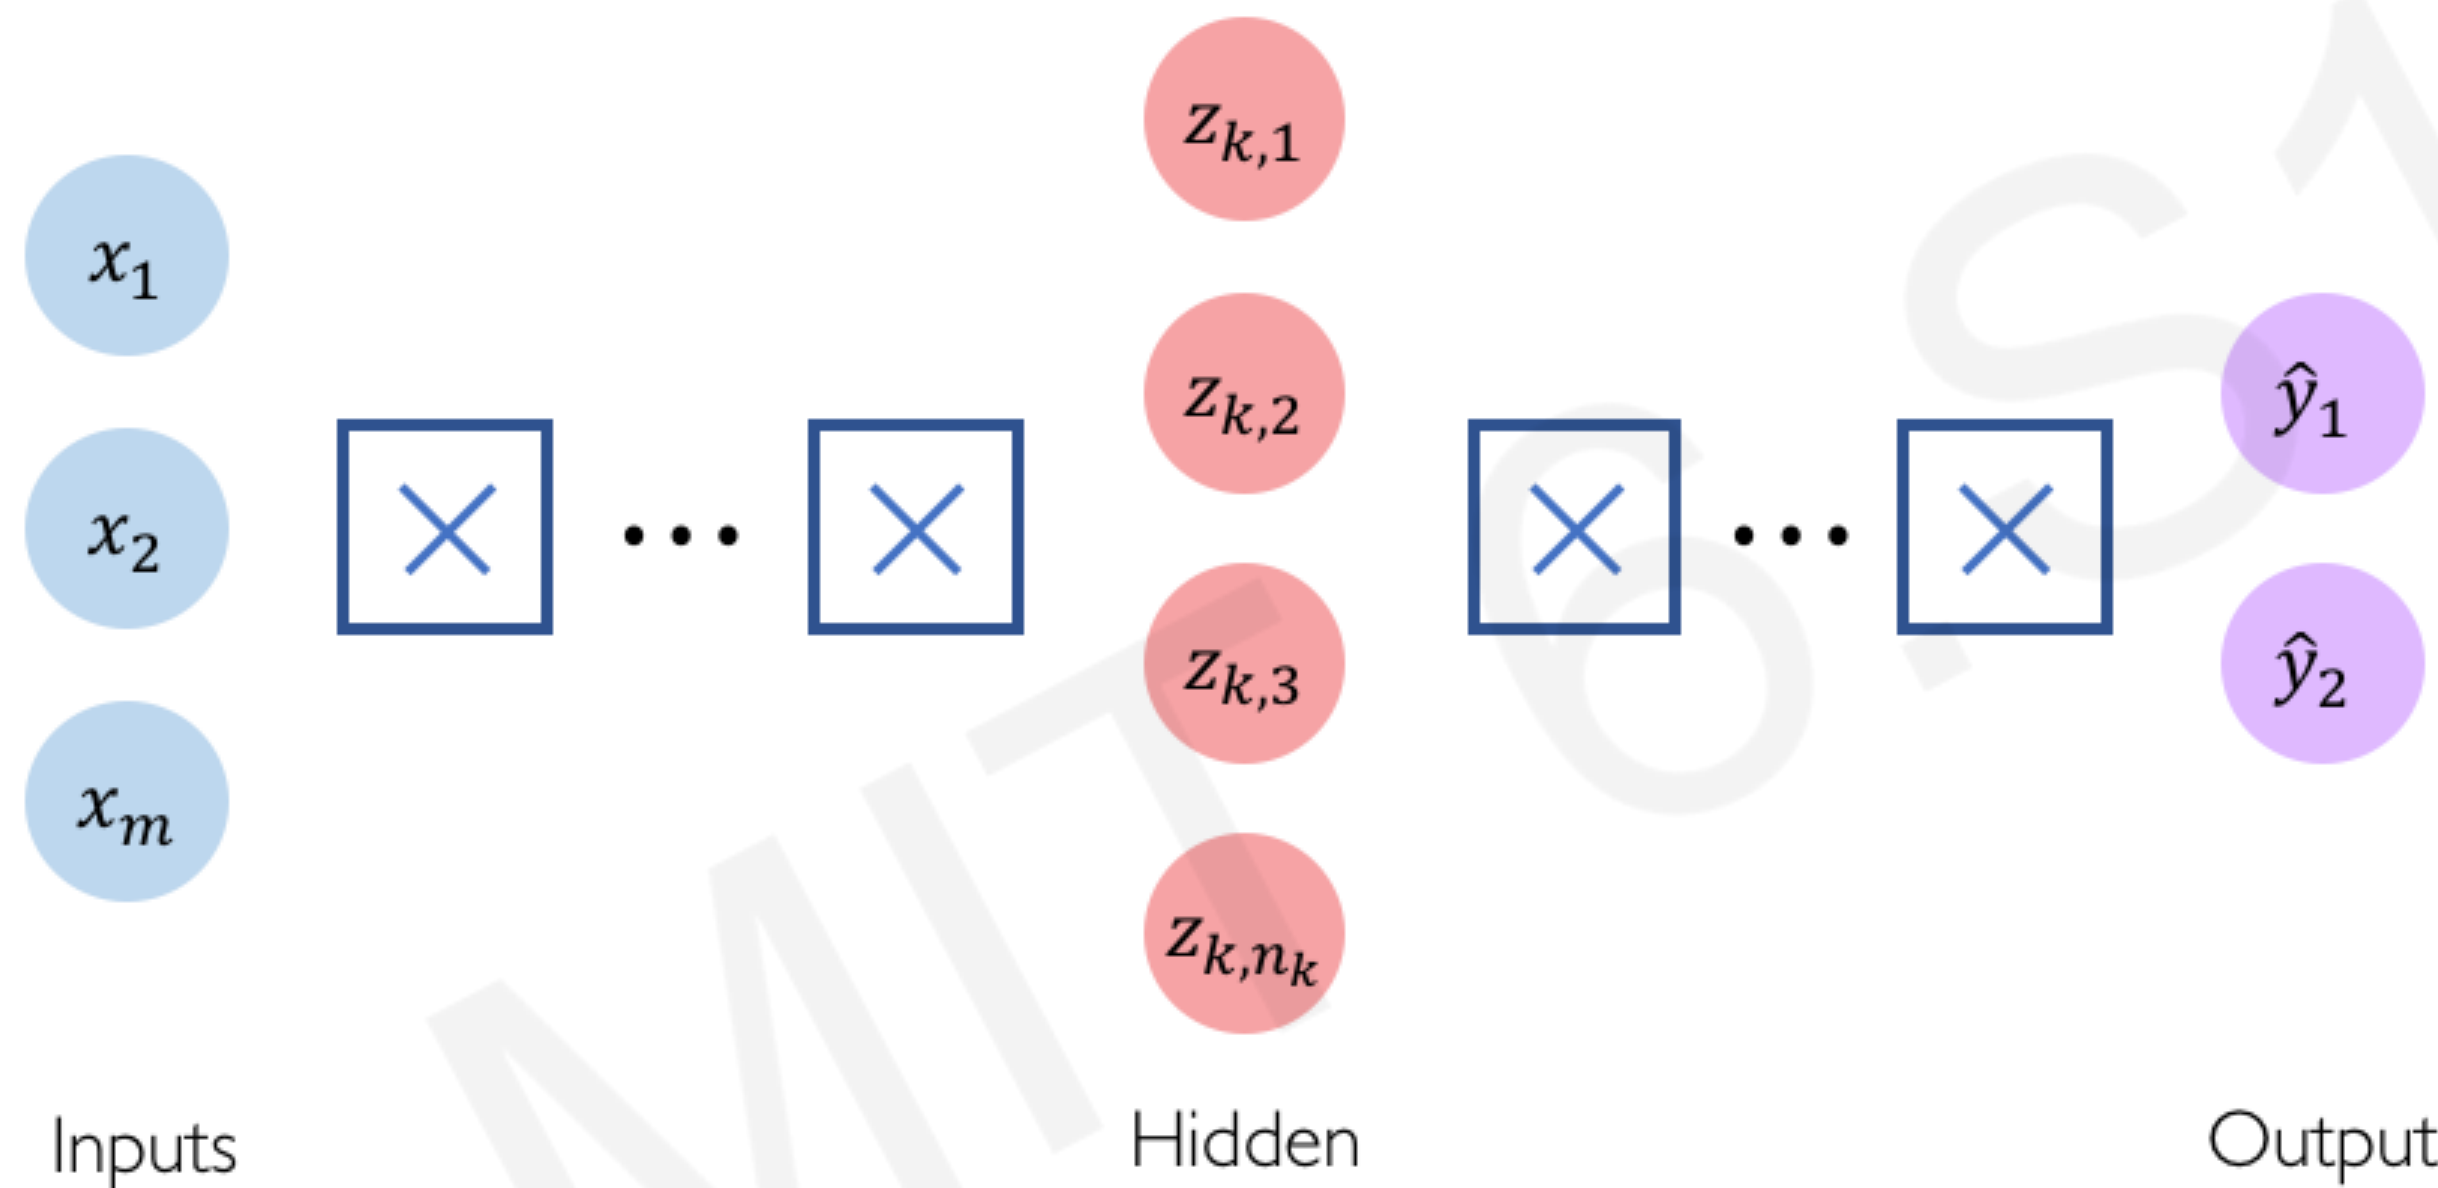
model = tf.keras.Sequential([
    tf.keras.layers.Dense(n),
    tf.keras.layers.Dense(2)
])
```

Deep Neural Network



$$z_{k,i} = w_{0,i}^{(k)} + \sum_{j=1}^{n_{k-1}} g(z_{k-1,j}) w_{j,i}^{(k)}$$

Deep Neural Network



```
import tensorflow as tf

model = tf.keras.Sequential([
    tf.keras.layers.Dense(n1),
    tf.keras.layers.Dense(n2),
    ...
    tf.keras.layers.Dense(2)
])
```

$$z_{k,i} = w_{0,i}^{(k)} + \sum_{j=1}^{n_{k-1}} g(z_{k-1,j}) w_{j,i}^{(k)}$$

Deep Neural Network

```
model = models.Sequential()  
model.add(layers.Dense(16, activation = 'relu', input_shape=(5000,)))  
model.add(layers.Dense(16, activation = 'relu'))  
model.add(layers.Dense(1, activation= 'sigmoid'))
```

x_1

```
model.compile(optimizer='adam',  
              loss='binary_crossentropy',  
              metrics=['accuracy'])
```

x_2

x_m

```
history = model.fit(partial_x_train,  
                   partial_y_train,  
                   epochs=4,  
                   batch_size=512,  
                   validation_data=(x_val,y_val))
```

Input

$$z_{k,i} = w_{0,i}^{(k)} + \sum_{j=1}^{n_{k-1}} g(z_{k-1,j}) w_{j,i}^{(k)}$$

Deep Neural Network

```
```\r}\nmodel <- keras_model_sequential() %>%\n  layer_dense(units = 16, activation = "relu", input_shape = c(5000)) %>%\n  layer_dense(units = 16, activation = "relu") %>%\n  layer_dense(units = 1, activation = "sigmoid")\n\nmodel %>% compile(\n  optimizer = "adam",\n  loss = "binary_crossentropy",\n  metrics = c("accuracy")\n)\n\nmodel %>% fit(x_train, y_train, epochs = 4, batch_size = 512)\nresults <- model %>% evaluate(x_test, y_test)\n```\nit
```



```
import tensorflow as tf\n\nmodel = tf.keras.Sequential([\n tf.keras.layers.Dense(n1),\n tf.keras.layers.Dense(n2),\n :\n tf.keras.layers.Dense(2)\n])
```

$$z_{k,i} = w_{0,i}^{(k)} + \sum_{j=1}^{n_{k-1}} g(z_{k-1,j}) w_{j,i}^{(k)}$$

# Applying Neural Networks

# Example Problem


Will I pass this class?

Let's start with a simple two feature model


$x_1$  = Number of lectures you attend

$x_2$  = Hours spent on the final project


# Example Problem: Will I pass this class?




# Example Problem: Will I pass this class?



# Example Problem: Will I pass this class?




# Example Problem: Will I pass this class?





# Quantifying Loss


The **loss** of our network measures the cost incurred from incorrect predictions



$$\mathcal{L}(\underbrace{f(x^{(i)}; \mathbf{W})}_{\text{Predicted}}, \underbrace{y^{(i)}}_{\text{Actual}})$$

# Empirical Loss

The *empirical loss* measures the total loss over our entire dataset



Also known as:

- Objective function
- Cost function
- Empirical Risk

$$J(\mathbf{W}) = \frac{1}{n} \sum_{i=1}^n \mathcal{L}(\underbrace{f(x^{(i)}; \mathbf{W})}_{\text{Predicted}}, \underbrace{y^{(i)}}_{\text{Actual}})$$

# Binary Cross Entropy Loss

```
model = models.Sequential()
model.add(layers.Dense(16, activation = 'relu', input_shape=(5000,)))
model.add(layers.Dense(16, activation = 'relu'))
model.add(layers.Dense(1, activation= 'sigmoid'))

model.compile(optimizer='adam',
 loss='binary_crossentropy',
 metrics=['accuracy'])

history = model.fit(partial_x_train,
 partial_y_train,
 epochs=4,
 batch_size=512,
 validation_data=(x_val,y_val))
```


Actual Predicted Actual Predicted



```
loss = tf.reduce_mean(tf.nn.softmax_cross_entropy_with_logits(y, predicted))
```

# Binary Cross Entropy Loss

Cross entropy loss can be used with models that output a probability between 0 and 1



$$J(W) = \frac{1}{n} \sum_{i=1}^n \underbrace{y^{(i)}}_{\text{Actual}} \log \left( \underbrace{f(x^{(i)}; W)}_{\text{Predicted}} \right) + (1 - \underbrace{y^{(i)}}_{\text{Actual}}) \log \left( 1 - \underbrace{f(x^{(i)}; W)}_{\text{Predicted}} \right)$$



```
loss = tf.reduce_mean(tf.nn.softmax_cross_entropy_with_logits(y, predicted))
```

# Binary Cross Entropy Loss

```
```\r}
model <- keras_model_sequential() %>%
  layer_dense(units = 16, activation = "relu", input_shape = c(5000)) %>%
  layer_dense(units = 16, activation = "relu") %>%
  layer_dense(units = 1, activation = "sigmoid")

model %>% compile(
  optimizer = "adam",
  loss = "binary_crossentropy",
  metrics = c("accuracy")
)

model %>% fit(x_train, y_train, epochs = 4, batch_size = 512)
results <- model %>% evaluate(x_test, y_test)
```\r
```

Actual Predicted Actual Predicted




```
loss = tf.reduce_mean(tf.nn.softmax_cross_entropy_with_logits(y, predicted))
```

This is exactly equivalent to the negative log-likelihood.  
 So this is so far identical to logit/logistic regression.

## Binary Cross Entropy Loss

*Cross entropy loss can be used with models that output a probability between 0 and 1*



$$J(W) = \frac{1}{n} \sum_{i=1}^n \underbrace{y^{(i)}}_{\text{Actual}} \log \left( \underbrace{f(x^{(i)}; W)}_{\text{Predicted}} \right) + (1 - \underbrace{y^{(i)}}_{\text{Actual}}) \log \left( 1 - \underbrace{f(x^{(i)}; W)}_{\text{Predicted}} \right)$$



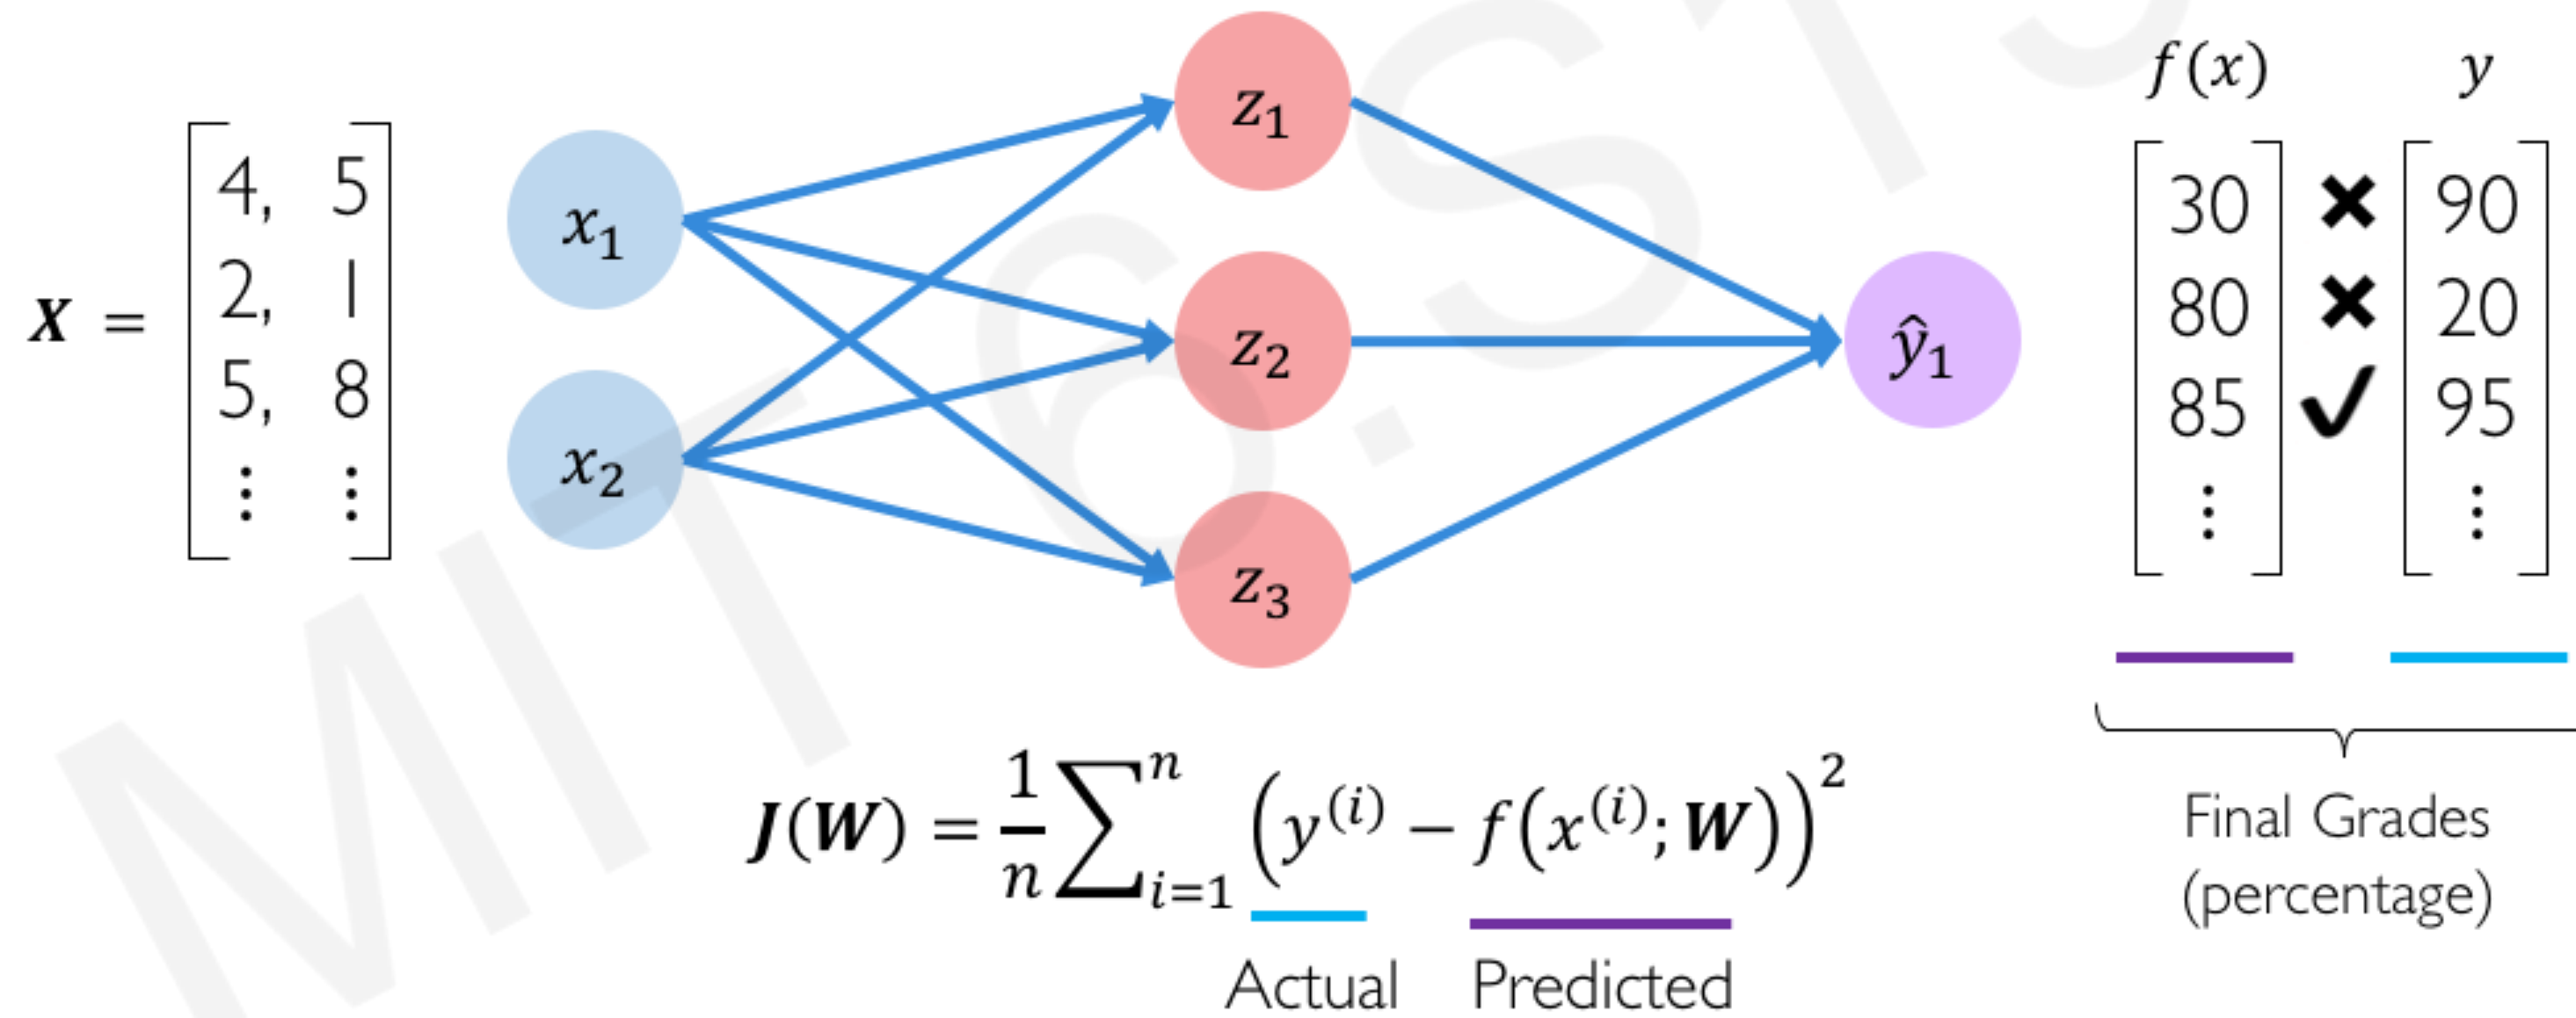
```
loss = tf.reduce_mean(tf.nn.softmax_cross_entropy_with_logits(y, predicted))
```

## Machine learning in one slide

Social science (inference)	Machine learning (prediction)
GLM inverse link function	Activation function
$\mathbb{E}(y) = f(\mathbf{x}'\boldsymbol{\beta})$	$\mathbb{E}(y) = f(\mathbf{x}'\boldsymbol{\beta})$
Preferred objective function	
Log-likelihood	Cross-entropy
$\log \mathcal{L} = \sum_{i=1}^n \log P(y_i \mathbf{x}_i, \boldsymbol{\beta})$	$-\log \mathcal{L} = -\sum_{i=1}^n \log P(y_i \mathbf{x}_i, \boldsymbol{\beta})$
Solving algorithm	
Newton-Raphson	Gradient descent
$\boldsymbol{\beta}_t := \boldsymbol{\beta}_{t-1} - [\mathbf{H} \log \mathcal{L}]^{-1} \nabla \log \mathcal{L}$	$\boldsymbol{\beta}_t := \boldsymbol{\beta}_{t-1} - \eta \nabla (-\log \mathcal{L})$
Quantities of interest	
$\hat{\boldsymbol{\beta}}; \text{Var}(\hat{\boldsymbol{\beta}})$	$\hat{y}; \sum \mathbf{1}(\hat{y} = y)/n$

# Mean Squared Error Loss

Mean squared error loss can be used with regression models that output continuous real numbers



```
loss = tf.reduce_mean(tf.square(tf.subtract(y, predicted)))
```



# Training Neural Networks

# Loss Optimization

We want to find the network weights that **achieve the lowest loss**

$$\mathbf{W}^* = \operatorname{argmin}_{\mathbf{W}} \frac{1}{n} \sum_{i=1}^n \mathcal{L}(f(x^{(i)}; \mathbf{W}), y^{(i)})$$

$$\mathbf{W}^* = \operatorname{argmin}_{\mathbf{W}} J(\mathbf{W})$$

# Loss Optimization

We want to find the network weights that **achieve the lowest loss**

$$\mathbf{W}^* = \operatorname{argmin}_{\mathbf{W}} \frac{1}{n} \sum_{i=1}^n \mathcal{L}(f(x^{(i)}; \mathbf{W}), y^{(i)})$$

$$\mathbf{W}^* = \operatorname{argmin}_{\mathbf{W}} J(\mathbf{W})$$


Remember:

$$\mathbf{W} = \{\mathbf{W}^{(0)}, \mathbf{W}^{(1)}, \dots\}$$

# Loss Optimization


$$W^* = \operatorname{argmin}_W J(W)$$

Remember:  
*Our loss is a function of  
the network weights!*




# Loss Optimization

Randomly pick an initial  $(w_0, w_1)$




# Loss Optimization

Compute gradient,  $\frac{\partial J(\mathbf{w})}{\partial \mathbf{w}}$




# Loss Optimization

Take small step in opposite direction of gradient



# Gradient Descent

Repeat until convergence






Contrast with Newton-Raphson, which also uses *second derivative* (Hessian)

# Gradient Descent

Repeat until convergence



## Machine learning in one slide

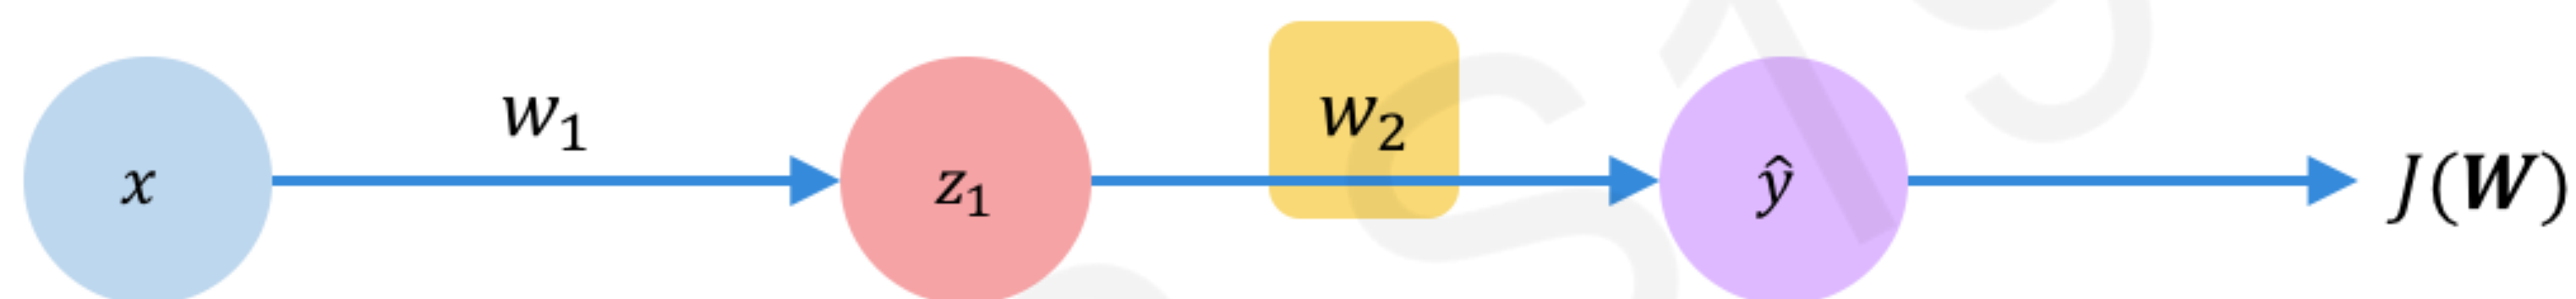
Social science (inference)	Machine learning (prediction)
GLM inverse link function	Activation function
$\mathbb{E}(y) = f(\mathbf{x}'\boldsymbol{\beta})$	$\mathbb{E}(y) = f(\mathbf{x}'\boldsymbol{\beta})$
Preferred objective function	
Log-likelihood	Cross-entropy
$\log \mathcal{L} = \sum_{i=1}^n \log P(y_i \mathbf{x}_i, \boldsymbol{\beta})$	$-\log \mathcal{L} = -\sum_{i=1}^n \log P(y_i \mathbf{x}_i, \boldsymbol{\beta})$
Solving algorithm	
Newton-Raphson	Gradient descent
$\boldsymbol{\beta}_t := \boldsymbol{\beta}_{t-1} - [\mathbf{H} \log \mathcal{L}]^{-1} \nabla \log \mathcal{L}$	$\boldsymbol{\beta}_t := \boldsymbol{\beta}_{t-1} - \eta \nabla (-\log \mathcal{L})$
Quantities of interest	
$\hat{\boldsymbol{\beta}}; \text{Var}(\hat{\boldsymbol{\beta}})$	$\hat{y}; \sum \mathbf{1}(\hat{y} = y)/n$

# Gradient Descent

## Algorithm

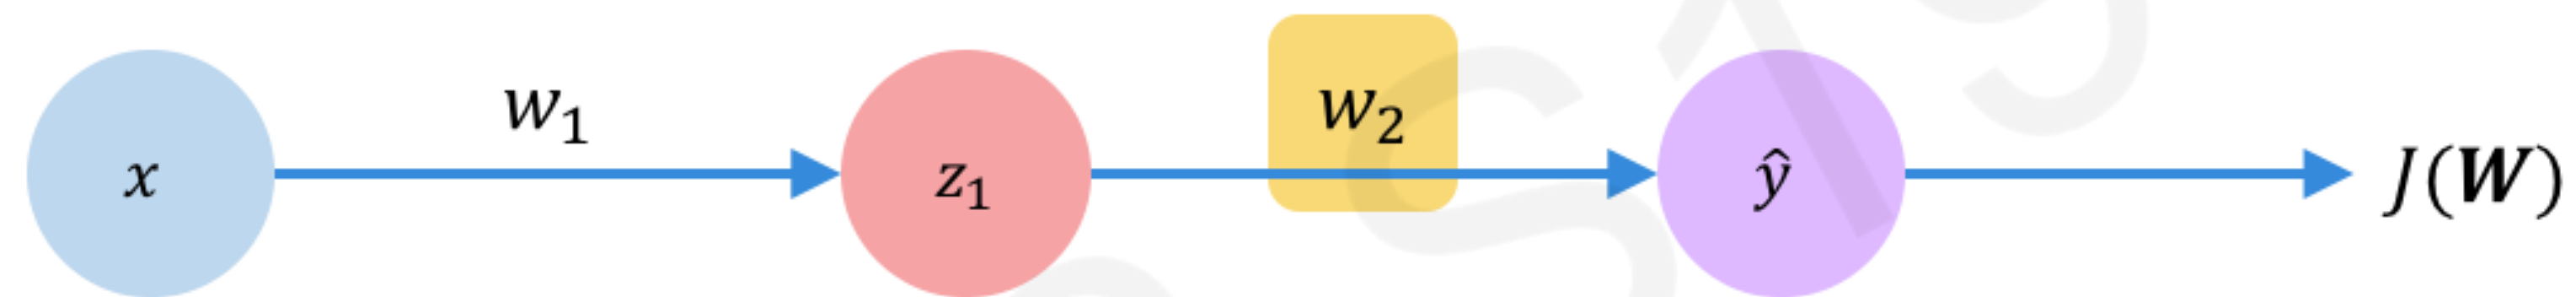
1. Initialize weights randomly  $\sim \mathcal{N}(0, \sigma^2)$
2. Loop until convergence:
3. Compute gradient,  $\frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
4. Update weights,  $\mathbf{W} \leftarrow \mathbf{W} - \eta \frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
5. Return weights

# Computing Gradients: Backpropagation



*How does a small change in one weight (ex.  $w_2$ ) affect the final loss  $J(\mathbf{W})$ ?*


# Computing Gradients: Backpropagation



$$\frac{\partial J(W)}{\partial w_2} =$$

Let's use the chain rule!


# Computing Gradients: Backpropagation



$$\frac{\partial J(W)}{\partial w_2} = \frac{\partial J(W)}{\partial \hat{y}} * \frac{\partial \hat{y}}{\partial w_2}$$

The equation shows the chain rule for backpropagation. The term  $\frac{\partial J(W)}{\partial \hat{y}}$  is highlighted with a purple bar underneath, and the term  $\frac{\partial \hat{y}}{\partial w_2}$  is highlighted with a red bar underneath.


# Computing Gradients: Backpropagation



$$\frac{\partial J(W)}{\partial w_1} = \frac{\partial J(W)}{\partial \hat{y}} * \frac{\partial \hat{y}}{\partial w_1}$$

Apply chain rule!      Apply chain rule!

# Computing Gradients: Backpropagation




$$\frac{\partial J(W)}{\partial w_1} = \frac{\partial J(W)}{\partial \hat{y}} * \frac{\partial \hat{y}}{\partial z_1} * \frac{\partial z_1}{\partial w_1}$$

█   █   █



# Computing Gradients: Backpropagation




$$\frac{\partial J(W)}{\partial w_1} = \underbrace{\frac{\partial J(W)}{\partial \hat{y}}}_{\text{purple}} * \underbrace{\frac{\partial \hat{y}}{\partial z_1}}_{\text{red}} * \underbrace{\frac{\partial z_1}{\partial w_1}}_{\text{blue}}$$

Repeat this for **every weight in the network** using gradients from later layers

# Neural Networks in Practice: Optimization

# Training Neural Networks is Difficult



*"Visualizing the loss landscape of neural nets". Dec 2017.*

# Loss Functions Can Be Difficult to Optimize

## Remember:

Optimization through gradient descent

$$\mathbf{W} \leftarrow \mathbf{W} - \eta \frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$$

# Loss Functions Can Be Difficult to Optimize

## Remember:


Optimization through gradient descent

$$\mathbf{W} \leftarrow \mathbf{W} - \eta \frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$$

How can we set the learning rate?


# Setting the Learning Rate

*Small learning rate converges slowly and gets stuck in false local minima*




# Setting the Learning Rate

*Large learning rates overshoot, become unstable and diverge*



# Setting the Learning Rate

*Stable learning rates converge smoothly and avoid local minima*





# How to deal with this?

## Idea 1:

Try lots of different learning rates and see what works “just right”

# How to deal with this?

## Idea 1:

Try lots of different learning rates and see what works “just right”

## Idea 2:






Do something smarter!

Design an adaptive learning rate that “adapts” to the landscape

# Adaptive Learning Rates

- Learning rates are no longer fixed
- Can be made larger or smaller depending on:
  - how large gradient is
  - how fast learning is happening
  - size of particular weights
  - etc...

# Gradient Descent Algorithms

Algorithm	TF Implementation	Reference
• SGD	 <code>tf.keras.optimizers.SGD</code>	Kiefer & Wolfowitz. "Stochastic Estimation of the Maximum of a Regression Function." 1952.
• Adam	 <code>tf.keras.optimizers.Adam</code>	Kingma et al. "Adam: A Method for Stochastic Optimization." 2014.
• Adadelta	 <code>tf.keras.optimizers.Adadelta</code>	Zeiler et al. "ADADELTA: An Adaptive Learning Rate Method." 2012.
• Adagrad	 <code>tf.keras.optimizers.Adagrad</code>	Duchi et al. "Adaptive Subgradient Methods for Online Learning and Stochastic Optimization." 2011.
• RMSProp	 <code>tf.keras.optimizers.RMSProp</code>	

Additional details: <http://ruder.io/optimizing-gradient-descent/>

# Gradient Descent Algorithms

```
model = models.Sequential()
model.add(layers.Dense(16, activation = 'relu', input_shape=(5000,)))
model.add(layers.Dense(16, activation = 'relu'))
model.add(layers.Dense(1, activation= 'sigmoid'))

model.compile(optimizer='adam',
 loss='binary_crossentropy',
 metrics=['accuracy'])

history = model.fit(partial_x_train,
 partial_y_train,
 epochs=4,
 batch_size=512,
 validation_data=(x_val,y_val))
```

Rate

Online

Additional details: <http://ruder.io/optimizing-gradient-descent/>

# Gradient Descent Algorithms

Alg

- S
- A
- A
- A
- R

```
```\r}
model <- keras_model_sequential() %>%
  layer_dense(units = 16, activation = "relu", input_shape = c(5000)) %>%
  layer_dense(units = 16, activation = "relu") %>%
  layer_dense(units = 1, activation = "sigmoid")
model %>% compile(
  optimizer = "adam",
  loss = "binary_crossentropy",
  metrics = c("accuracy")
)
model %>% fit(x_train, y_train, epochs = 4, batch_size = 512)
results <- model %>% evaluate(x_test, y_test)
```\r}
```

n of the  
52.  
nastic  
earning Rate  
ods for Online  
.011.


Additional details: <http://ruder.io/optimizing-gradient-descent/>

# Neural Networks in Practice: Mini-batches

# Gradient Descent

## Algorithm

1. Initialize weights randomly  $\sim \mathcal{N}(0, \sigma^2)$
2. Loop until convergence:
3.     Compute gradient,  $\frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
4.     Update weights,  $\mathbf{W} \leftarrow \mathbf{W} - \eta \frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
5. Return weights






# Gradient Descent

## Algorithm

1. Initialize weights randomly  $\sim \mathcal{N}(0, \sigma^2)$
2. Loop until convergence:
3. Compute gradient,  $\frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
4. Update weights,  $\mathbf{W} \leftarrow \mathbf{W} - \eta \frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
5. Return weights




Can be very  
computationally  
intensive to compute!

# Stochastic Gradient Descent

## Algorithm


1. Initialize weights randomly  $\sim \mathcal{N}(0, \sigma^2)$
2. Loop until convergence:
3.     Pick single data point  $i$
4.     Compute gradient,  $\frac{\partial J_i(\mathbf{W})}{\partial \mathbf{W}}$
5.     Update weights,  $\mathbf{W} \leftarrow \mathbf{W} - \eta \frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
6. Return weights



# Stochastic Gradient Descent

## Algorithm

1. Initialize weights randomly  $\sim \mathcal{N}(0, \sigma^2)$
2. Loop until convergence:
3. Pick single data point  $i$
4. Compute gradient,  $\frac{\partial J_i(W)}{\partial W}$
5. Update weights,  $\mathbf{W} \leftarrow \mathbf{W} - \eta \frac{\partial J(W)}{\partial W}$
6. Return weights




Easy to compute but  
**very noisy** (stochastic)!

# Stochastic Gradient Descent

## Algorithm


1. Initialize weights randomly  $\sim \mathcal{N}(0, \sigma^2)$
2. Loop until convergence:
3. Pick batch of  $B$  data points
4. Compute gradient,  $\frac{\partial J(\mathbf{W})}{\partial \mathbf{W}} = \frac{1}{B} \sum_{k=1}^B \frac{\partial J_k(\mathbf{W})}{\partial \mathbf{W}}$
5. Update weights,  $\mathbf{W} \leftarrow \mathbf{W} - \eta \frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
6. Return weights



# Stochastic Gradient Descent

## Algorithm

1. Initialize weights randomly  $\sim \mathcal{N}(0, \sigma^2)$
2. Loop until convergence:
3. Pick batch of  $B$  data points
4. Compute gradient,  $\frac{\partial J(\mathbf{W})}{\partial \mathbf{W}} = \frac{1}{B} \sum_{k=1}^B \frac{\partial J_k(\mathbf{W})}{\partial \mathbf{W}}$
5. Update weights,  $\mathbf{W} \leftarrow \mathbf{W} - \eta \frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
6. Return weights



Fast to compute and a much better estimate of the true gradient!

# Stochastic Gradient Descent

```
model = models.Sequential()
model.add(layers.Dense(16, activation = 'relu', input_shape=(5000,)))
model.add(layers.Dense(16, activation = 'relu'))
model.add(layers.Dense(1, activation= 'sigmoid'))

model.compile(optimizer='adam',
 loss='binary_crossentropy',
 metrics=['accuracy'])

history = model.fit(partial_x_train,
 partial_y_train,
 epochs=4,
 batch_size=512,
 validation_data=(x_val,y_val))
```

Fast to compute and a much better  
estimate of the true gradient!

# Stochastic Gradient Descent

Algo

1. lr

2. L

3.

4.


5.

6. R

```
```\r\
model <- keras_model_sequential() %>%
  layer_dense(units = 16, activation = "relu", input_shape = c(5000)) %>%
  layer_dense(units = 16, activation = "relu") %>%
  layer_dense(units = 1, activation = "sigmoid")

model %>% compile(
  optimizer = "adam",
  loss = "binary_crossentropy",
  metrics = c("accuracy")
)

model %>% fit(x_train, y_train, epochs = 4, batch_size = 512)
results <- model %>% evaluate(x_test, y_test)
```\r
```



Fast to compute and a much better  
estimate of the true gradient!

# Mini-batches while training

## More accurate estimation of gradient

- Smother convergence
- Allows for larger learning rates



# Mini-batches while training

## More accurate estimation of gradient

Smoother convergence  
Allows for larger learning rates

## Mini-batches lead to fast training!

Can parallelize computation + achieve significant speed increases on GPU's

So, SGD is different from Newton-Raphson in derivative information used,  
*and* in its optimization over small subsets of the data at a time and in parallel.

## Mini-batches while training

More accurate estimation of gradient

Smoother convergence

Allows for larger learning rates

**Mini-batches lead to fast training!**


Can parallelize computation + achieve significant speed increases on GPU's

<https://playground.tensorflow.org>

# Core Foundation Review


## The Perceptron

- Structural building blocks
- Nonlinear activation functions



## Neural Networks

- Stacking Perceptrons to form neural networks
- Optimization through backpropagation



## Training in Practice

- Adaptive learning
- Batching
- Regularization

