

Introduction to TensorFlow



Alejandro Solano - EuroPython 2017



TensorFlow



TensorFlow



cat



TensorFlow

input



target

??

cat



TensorFlow

input



target

sin
+
log X
exp

cat



TensorFlow

input

target



sin
+
X
log
exp

cat



TensorFlow

Deep Learning



TensorFlow

What is TensorFlow?

- TensorFlow is an open-source library for Deep Learning.
- Developed by the Google Brain team and released in November 2015.
- Version 1.0.0 was launched in February 2017.



Installation



Install TensorFlow (Linux and Mac OS)

- Download Anaconda
- Create an environment with all must-have libraries.

```
$ conda create -n tensorflow python=3.5
```

```
$ source activate tensorflow
```

```
$ conda install pandas matplotlib jupyter notebook scipy scikit
```

```
$ pip install tensorflow
```



TensorFlow

Install TensorFlow (Windows)

- Download Anaconda
- Create an environment with all must-have libraries.

```
$ conda create -n tensorflow python=3.5
```

```
$ activate tensorflow
```

```
$ conda install pandas matplotlib jupyter notebook scipy scikit
```

```
$ pip install tensorflow
```



TensorFlow

Concepts



TensorFlow



cat



TensorFlow



MODEL



cat



TensorFlow



MODEL



non-cat



TensorFlow




MODEL



non-cat

prediction



input

target



cat



TensorFlow



MODEL



non-cat



C
O
S
T



cat



TensorFlow



MODEL



non-cat



C
O
S
T



cat



TensorFlow



MODEL



non-cat



OPTIMIZER



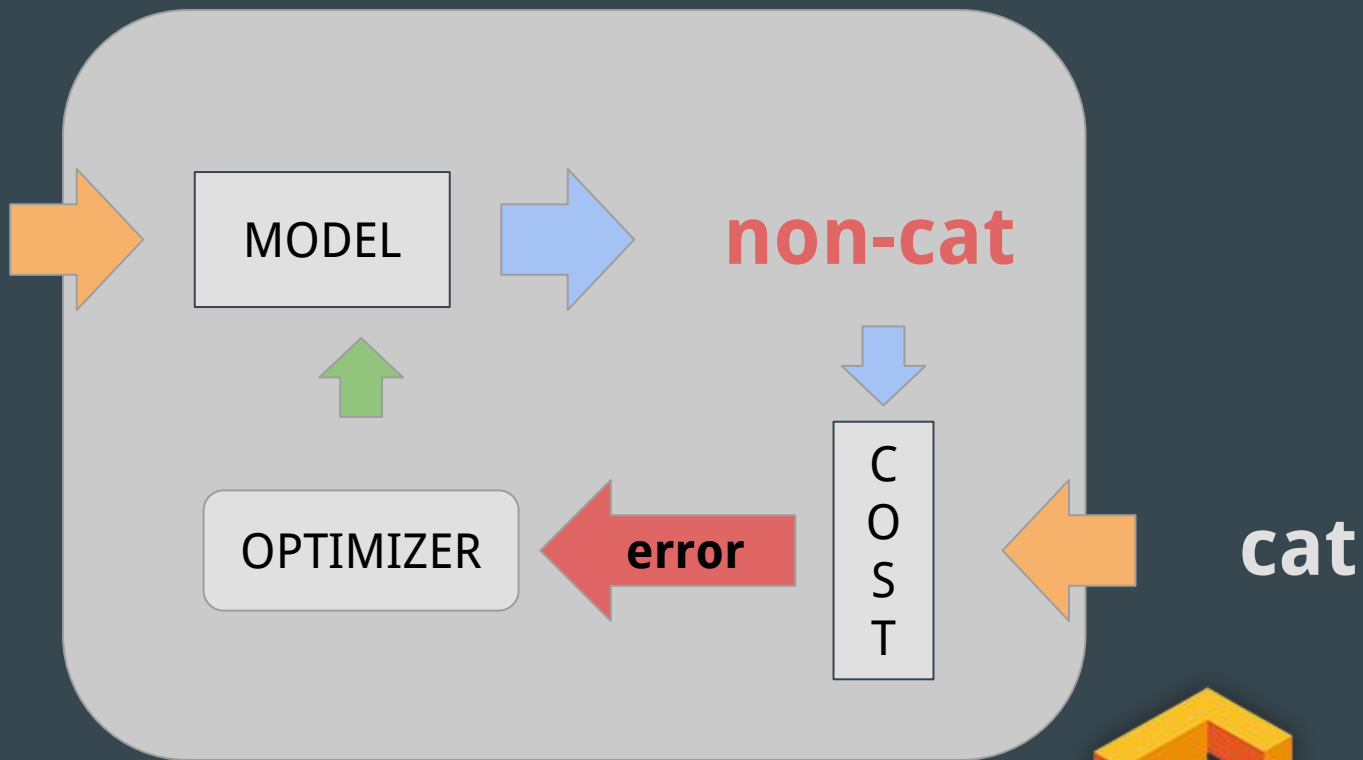
C
O
S
T



cat



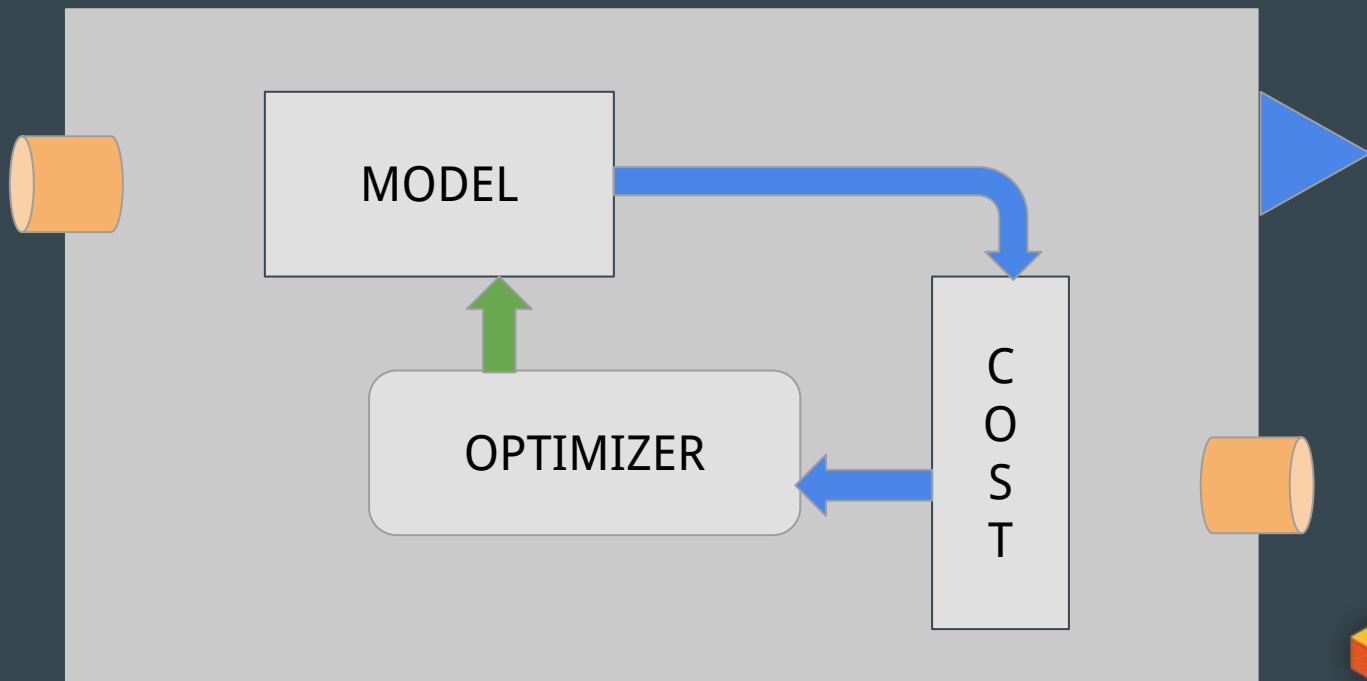
TensorFlow



Graph

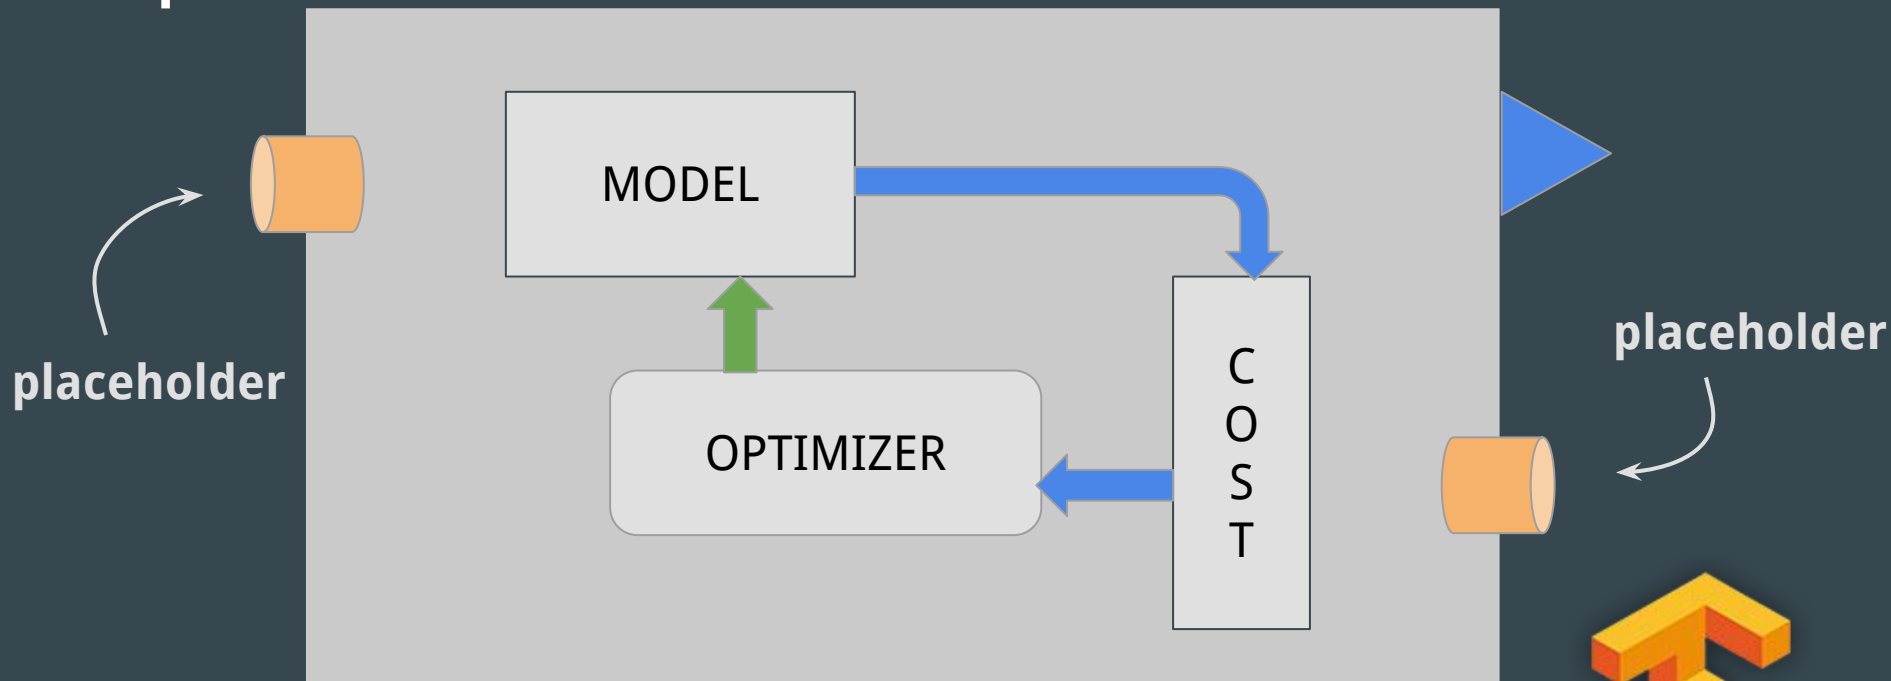


Graph



TensorFlow

Graph



TensorFlow

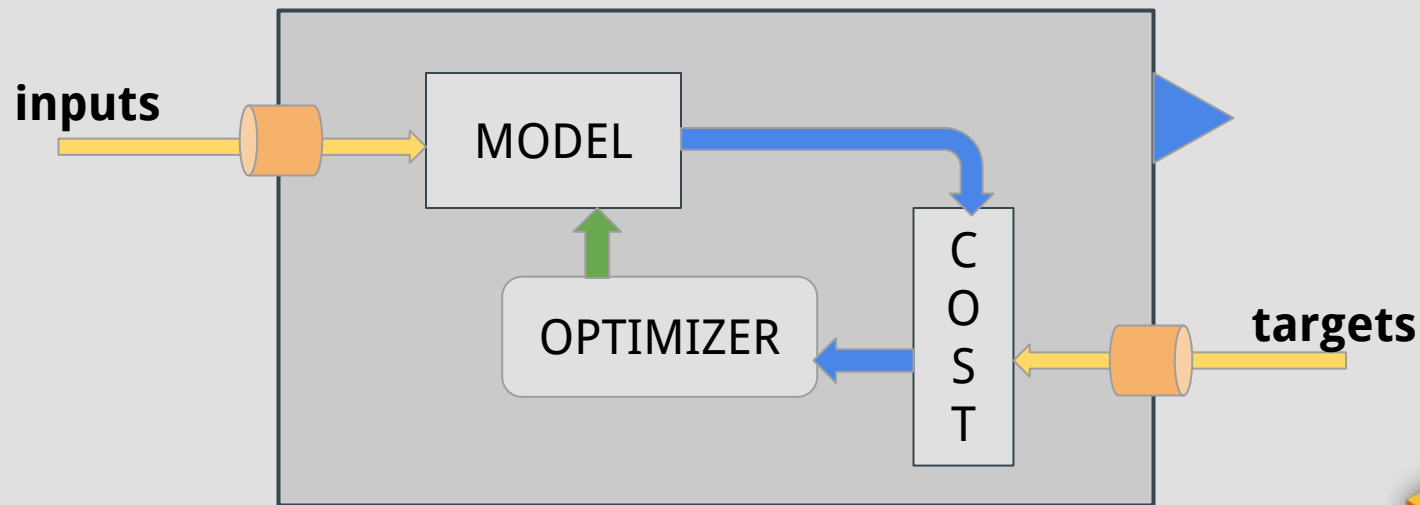
Graph

- **Placeholders:** gates where we introduce example
- **Model:** makes predictions. Set of **variables** and operations
- **Cost function:** function that computes the model error
- **Optimizer:** algorithm that optimizes the variables so the cost would be zero



TensorFlow

Session: Graph + Data



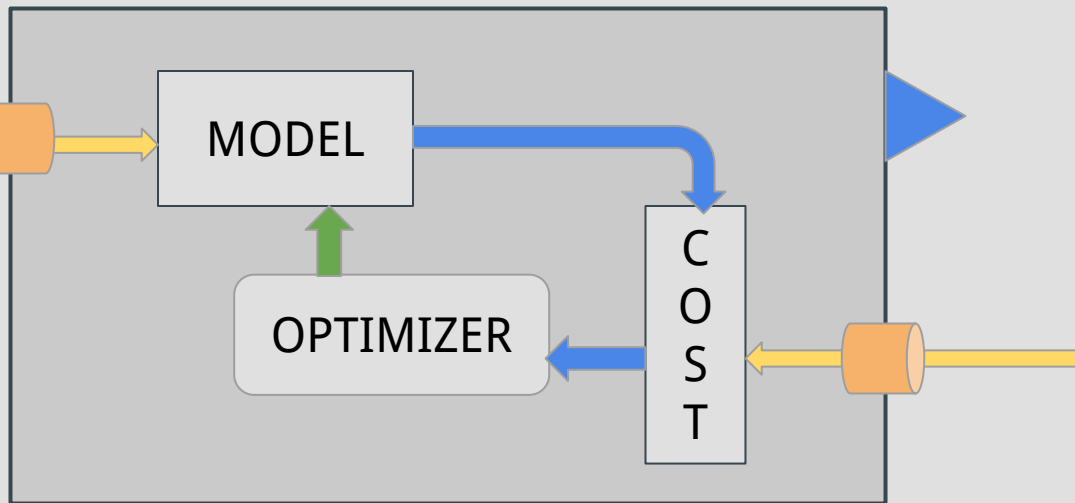
Graph, Data and Session

- **Graph:** Layout of the prediction and learning process. It does not include data.
- **Data:** examples that will train the neural network. It consists on two kinds: inputs and targets.
- **Session:** where everything takes places. Here is where we **feed** the graph with data.



TensorFlow

Session: Graph + Data

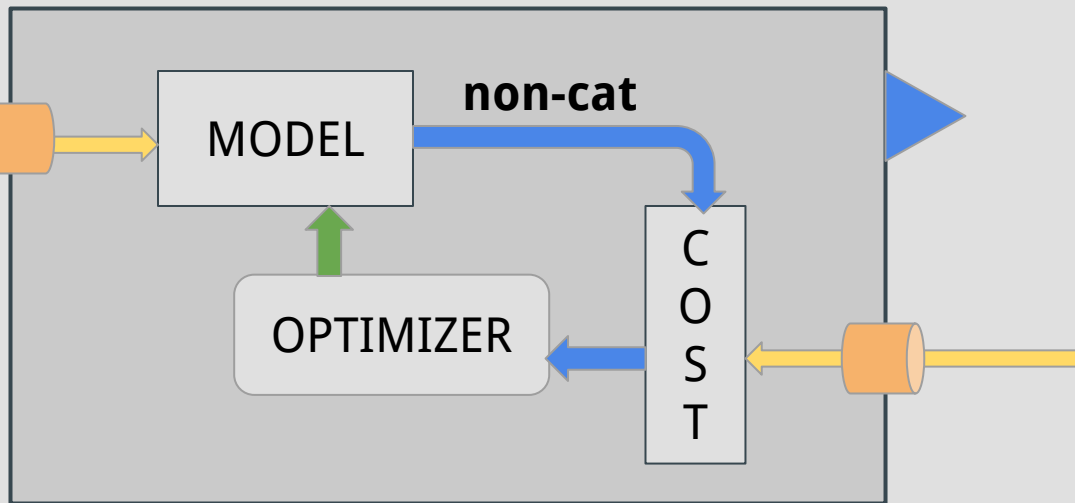



cat



TensorFlow

Session: Graph + Data

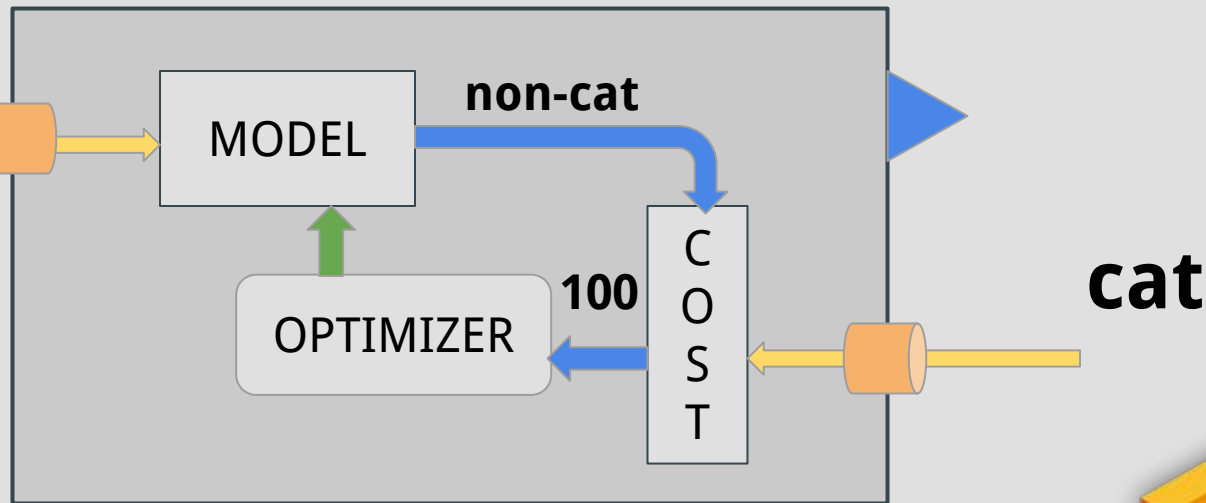


cat

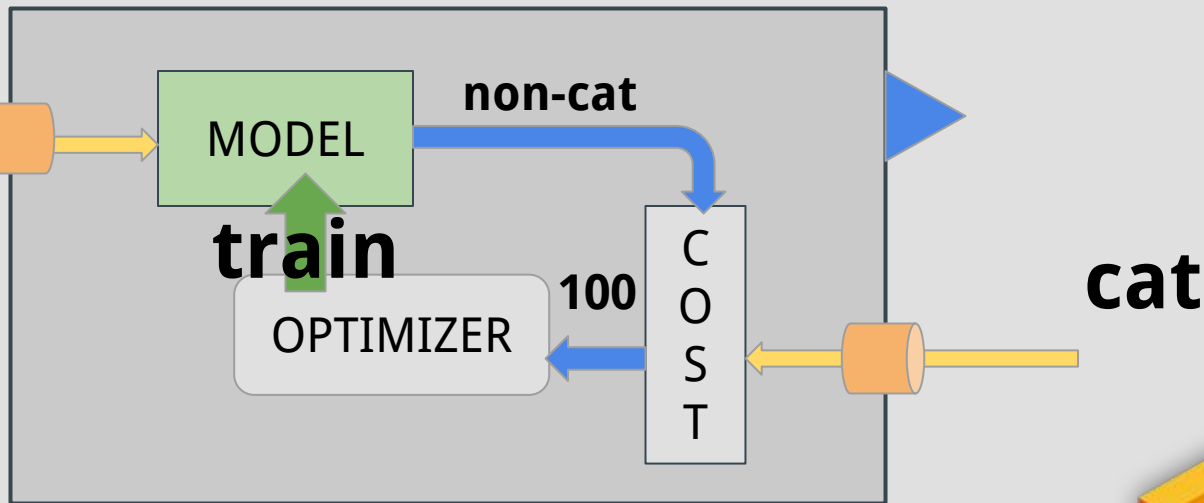


TensorFlow

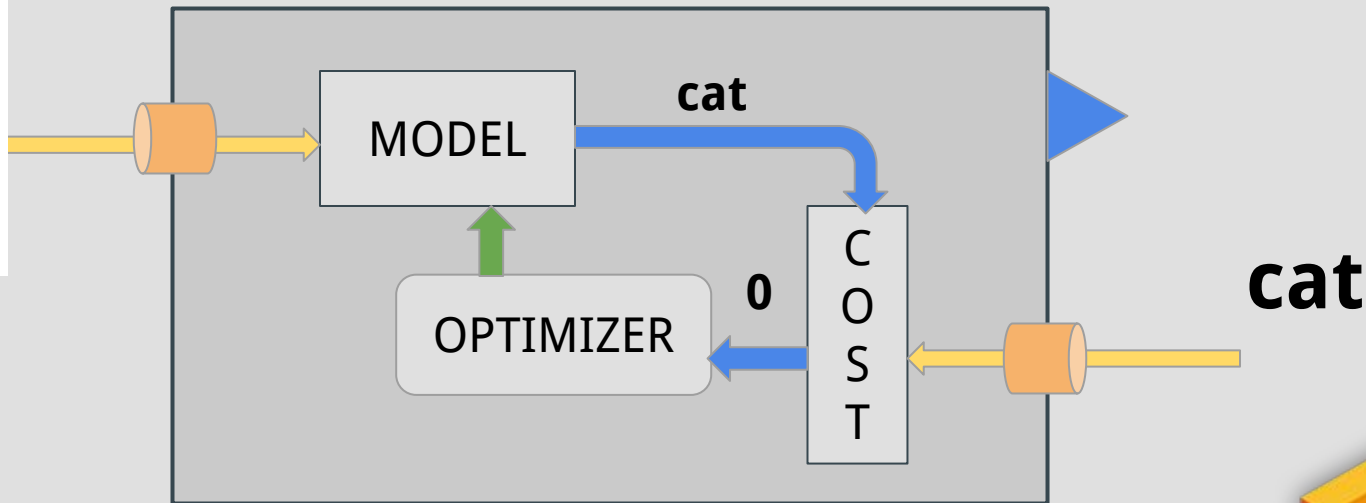
Session: Graph + Data



Session: Graph + Data



Session: Graph + Data



Hello world!




TensorFlow

Hello world!: Sum of two integers

```
import tensorflow as tf
```



Hello world!: Sum of two integers



TensorFlow

Hello world!: Sum of two integers

```
##### GRAPH #####  
a = tf.placeholder(tf.int32)  
b = tf.placeholder(tf.int32)  
sum_graph = tf.add(a, b)
```

```
##### DATA #####  
num1 = 3  
num2 = 8
```



TensorFlow

Hello world!: Sum of two integers

```
##### SESSION #####  
with tf.Session() as sess:  
    sum_outcome = sess.run(sum_graph, feed_dict={  
        a: num1,  
        b: num2  
    })
```



TensorFlow

```
In [5]: with tf.Session() as sess:
        sum_output = sess.run(sum_graph, feed_dict={
            a: num1,
            b: num2
        })

        print("The sum of {} and {} is {}".format(num1, num2, sum_output))
```

The sum of 3 and 8 is 11



TensorFlow

Regression



TensorFlow

TensorFlow for Regression: learning how to sum

- Mission: learn how to sum using 10,000 examples.

$$x_1 + x_2 = y$$



TensorFlow

TensorFlow for Regression: learning how to sum

- Mission: learn how to sum using 10,000 examples.


$$x_1 + x_2 = y$$



TensorFlow

TensorFlow for Regression: learning how to sum

- Mission: learn how to sum using 10,000 examples.



TensorFlow for Regression: learning how to sum

- Mission: learn how to sum using 10,000 examples.

$$x_1 ? x_2 = y$$




250 m



TensorFlow

TensorFlow for Regression: learning how to sum

- Mission: learn how to sum using 10,000 examples.



TensorFlow for Regression: learning how to sum

- Mission: learn how to sum using 10,000 examples.

$$\mathbf{x}_1 \ ? \ \mathbf{x}_2 = \mathbf{y}$$

- We assume the relationship between \mathbf{x} and \mathbf{y} is a linear function.

$$\mathbf{x} \cdot \mathbf{W} + \mathbf{b} = \mathbf{y}$$



TensorFlow

TensorFlow for Regression: learning how to sum


- Mission: learn how to sum using 10,000 examples.

$$\mathbf{x}_1 \ ? \ \mathbf{x}_2 = \mathbf{y}$$

- We assume the relationship between \mathbf{x} and \mathbf{y} is a linear function.

$$\mathbf{x} \cdot \mathbf{W} + \mathbf{b} = \mathbf{y}$$


variables to be learned



TensorFlow

Neural Network


$$x \cdot w_1 + b_1 = y$$



Neural Network

$$x \cdot w_1 + b_1 = y$$

$$(x \cdot w_1 + b_1) \cdot w_2 + b_2 = y$$




Neural Network

$$x \cdot w_1 + b_1 = y$$

$$(x \cdot w_1 + b_1) \cdot w_2 + b_2 = y$$

$$((x \cdot w_1 + b_1) \cdot w_2 + b_2) \cdot w_3 + b_3 = y$$




Neural Network

$$x \cdot W_1 + b_1 = y$$

$$\sigma(x \cdot W_1 + b_1) \cdot W_2 + b_2 = y$$

$$\tanh(\sigma(x \cdot W_1 + b_1) \cdot W_2 + b_2) \cdot W_3 + b_3 = y$$



TensorFlow for Regression: learning how to sum

```
# PLACEHOLDERS  
x = tf.placeholder(tf.float32, [None, 2])  
y = tf.placeholder(tf.float32, [None, 1])
```



TensorFlow

TensorFlow for Regression: learning how to sum

```
# PLACEHOLDERS  
x = tf.placeholder(tf.float32, [None, 2])  
y = tf.placeholder(tf.float32, [None, 1])
```

(we don't know how many examples we'll have, but we do know that each one of them has 2 numbers as input and 1 as target)



TensorFlow for Regression: learning how to sum

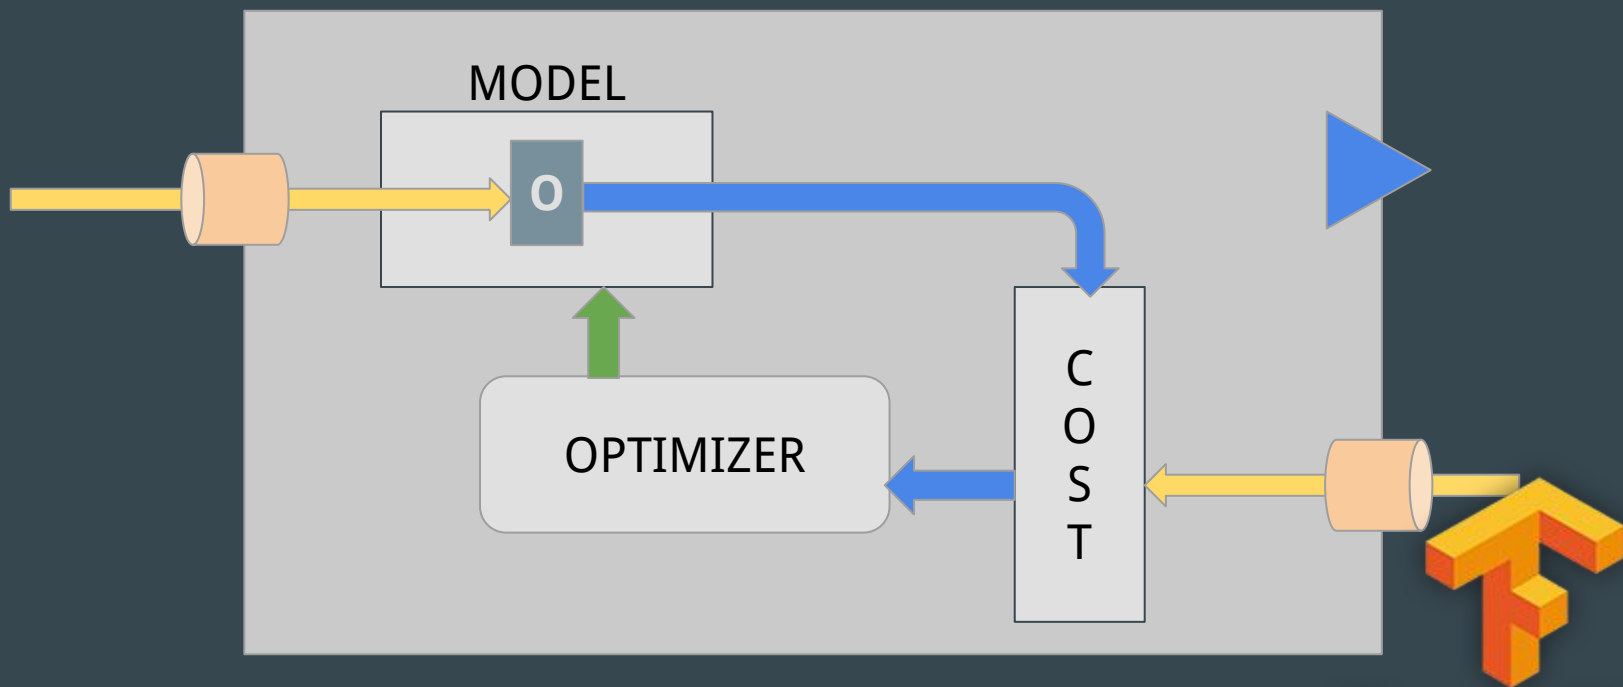
```
# MODEL
W = tf.Variable(tf.truncated_normal([2, 1], stddev=0.05))
b = tf.Variable(tf.random_normal([1]))

output = tf.add(tf.matmul(x, W), b)
```



TensorFlow

TensorFlow for Regression: learning how to sum



Cost (loss) function

$$\underbrace{x \cdot W + b}_{\text{prediction}} = \underset{\text{target}}{y}$$

The diagram illustrates the linear equation $x \cdot W + b = y$. The left side, $x \cdot W + b$, is grouped by a white curly brace and labeled "prediction" with a white arrow pointing to it. The right side, y , is labeled "target" with a white arrow pointing to it.



Cost (loss) function

$$y - (x \cdot W + b)$$



TensorFlow

Cost (loss) function

$$[y - (x \cdot w + b)]^2$$



TensorFlow

Cost (loss) function

$$\Sigma [y_i - (x_i \cdot W + b)]^2$$



TensorFlow

TensorFlow for Regression: learning how to sum

```
cost = tf.reduce_sum(tf.square(output - y))
```




TensorFlow

Gradient Descent

cost function

$$\text{cost} = f(w_1, w_2, b)$$




TensorFlow

Gradient Descent

cost function

$$\text{cost} = f(w_1, w_2, b)$$




TensorFlow

Gradient Descent

cost function

$$\text{cost} = f(w_1, w_2, b)$$




TensorFlow

Gradient Descent

cost function

$$\text{cost} = f(w_1, w_2, b)$$



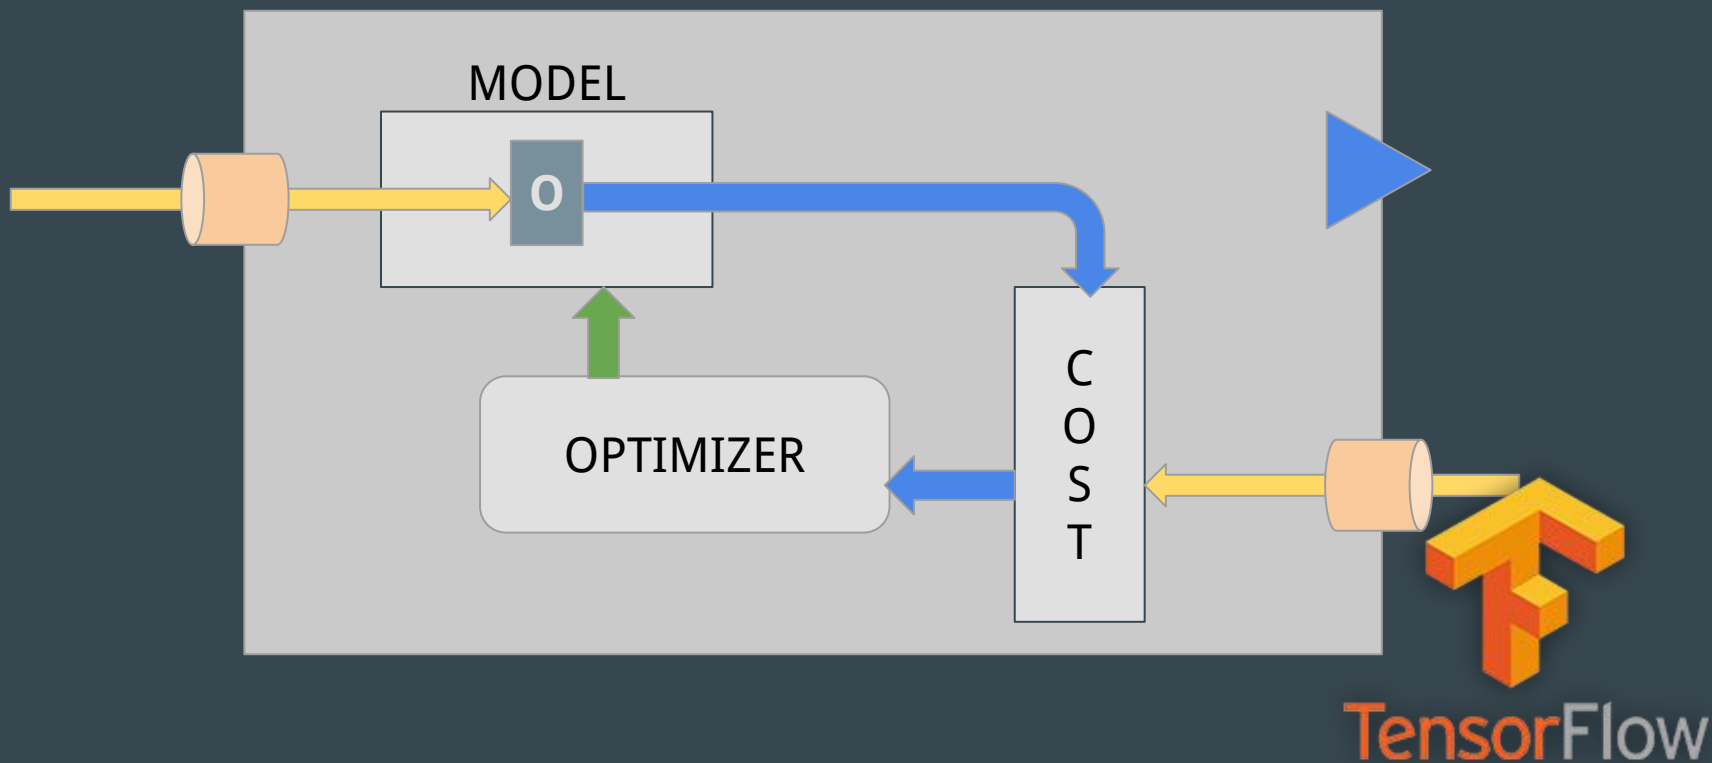
TensorFlow

TensorFlow for Regression: learning how to sum

```
optimizer =  
tf.train.GradientDescentOptimizer(Learning_rate=0.00001)  
  
optimizer = optimizer.minimize(cost)
```



TensorFlow for Regression: learning how to sum



Data split

data



TensorFlow

Data split

data



train data test data



TensorFlow

TensorFlow for Regression: learning how to sum

```
from helper import get_data, split_data

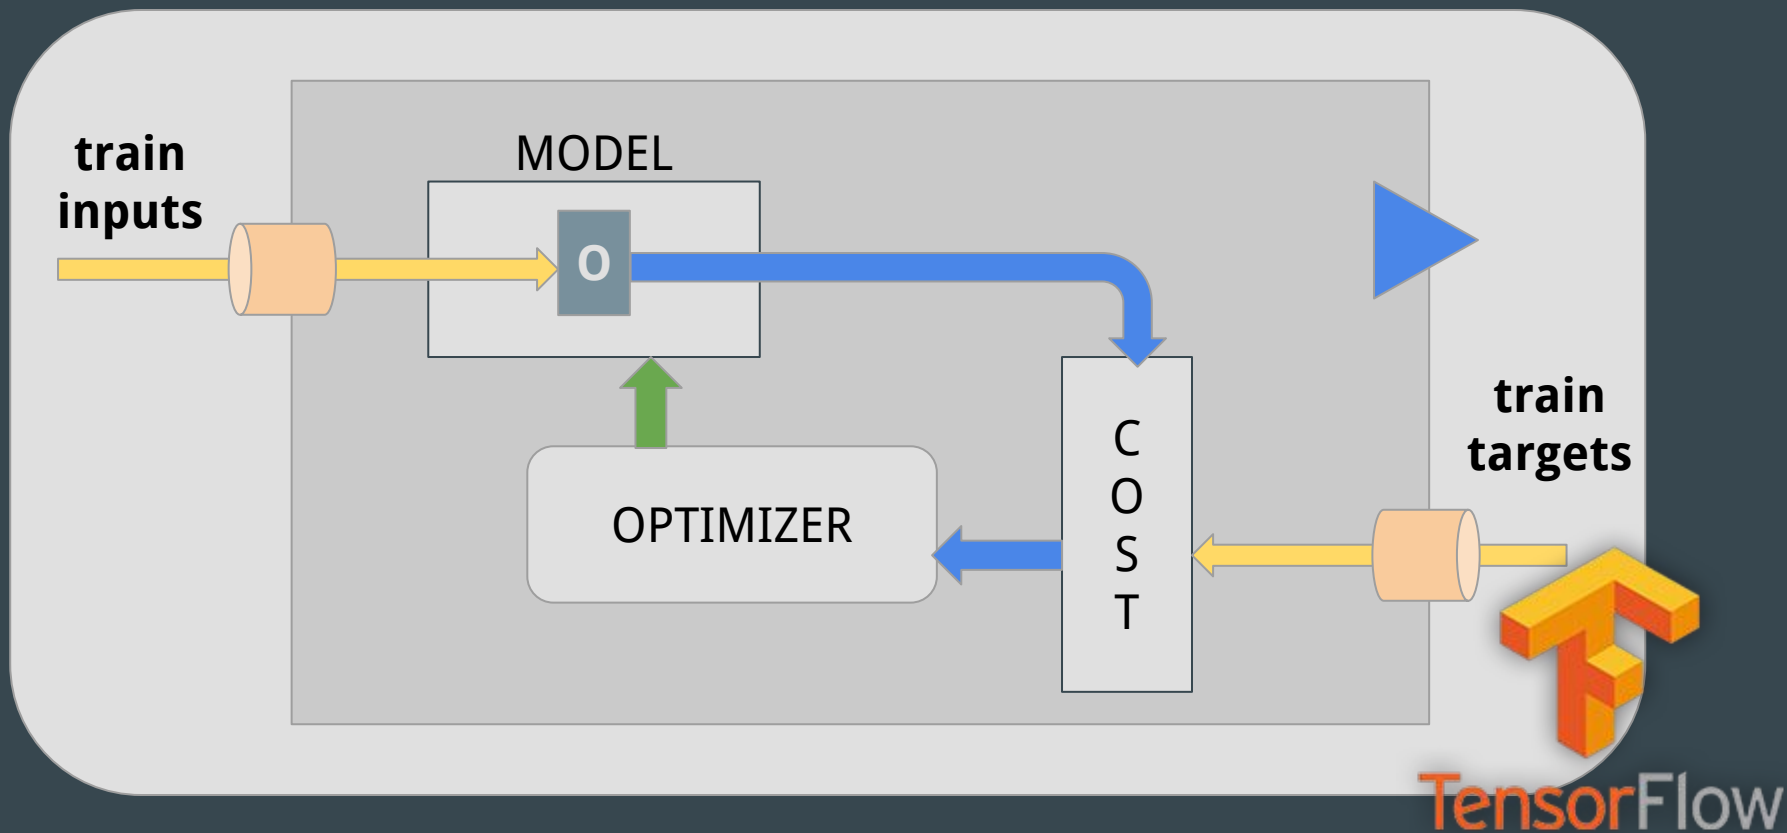
# DATA
inputs, targets = get_data(max_int=10, size=10000)

# split train and test data
train_inputs, test_inputs, train_targets, test_targets =
split_data(inputs, targets)
```



TensorFlow

TensorFlow for Regression: learning how to sum



TensorFlow for Regression: learning how to sum

```
with tf.Session() as sess:  
    sess.run(tf.global_variables_initializer())  
  
    for epoch in range(epochs):  
        sess.run(optimizer, feed_dict={  
            x: train_inputs,  
            y: train_targets  
        })
```



TensorFlow

```
Testing Accuracy: 0.9568796753883362
```

```
The sum of 5 plus 7 is 11.991448402404785
```

```
The weights are: [[ 0.88418758]  
 [ 0.8903569 ]]  
and the bias is: [ 1.33801162]
```



TensorFlow

Classification



TensorFlow



cat



non-cat



TensorFlow



$[0, 1]$



$[1, 0]$



TensorFlow



MODEL



[0.175,
0.825]



TensorFlow



MODEL



[0.457,
0.543]



TensorFlow



MODEL



[0.457,
+
0.543]

1.000



TensorFlow

TensorFlow for Classification

- Mission: learn if the sum of two numbers is higher than 10.

if ($x_1 + x_2 > 10$) then $y = [0; 1]$

else $y = [1; 0]$



TensorFlow

TensorFlow for Classification

- Mission: learn if the sum of two numbers is higher than 10.

$$x_1 \text{ ?? } x_2 = y$$




TensorFlow for Classification

- Mission: learn if the sum of two numbers is higher than 10

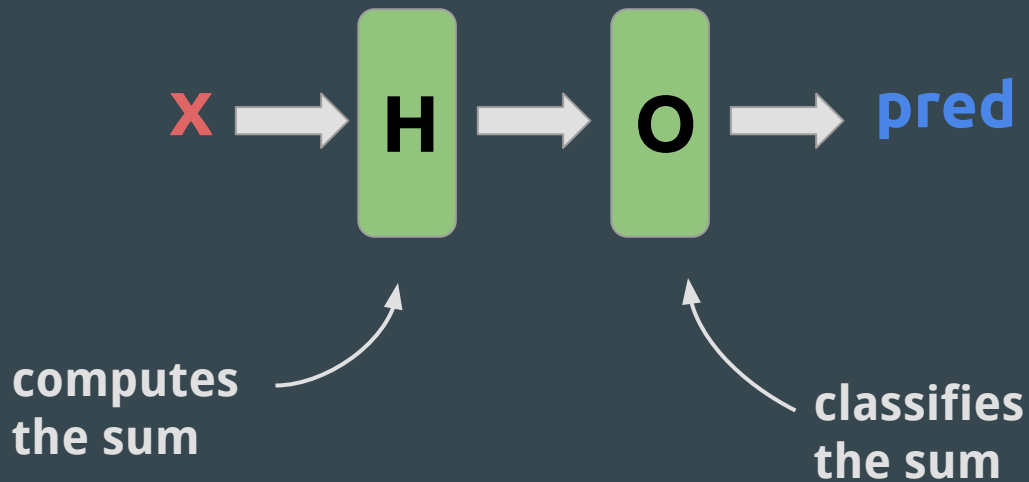
$$x_1 \text{ ?? } x_2 = y$$

- More complexity: we add a new layer



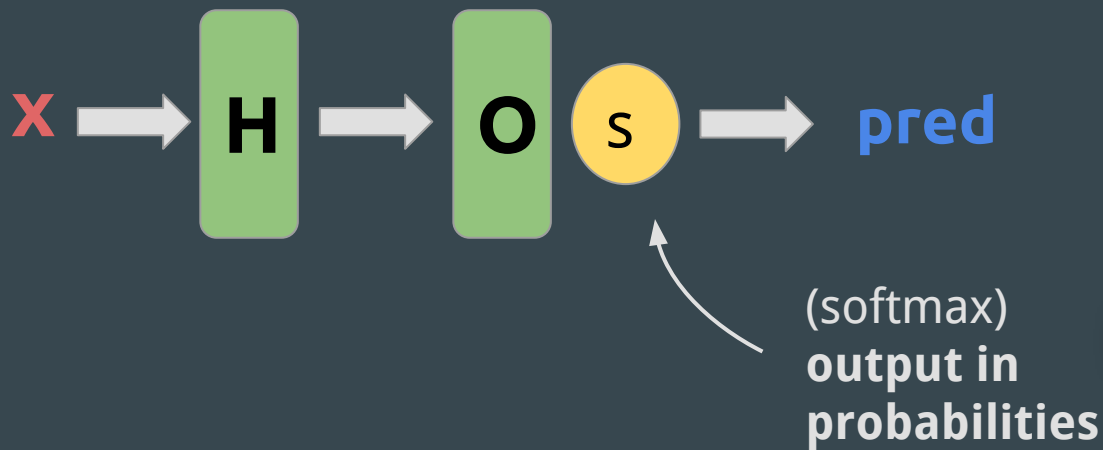
Neural Networks: intuition

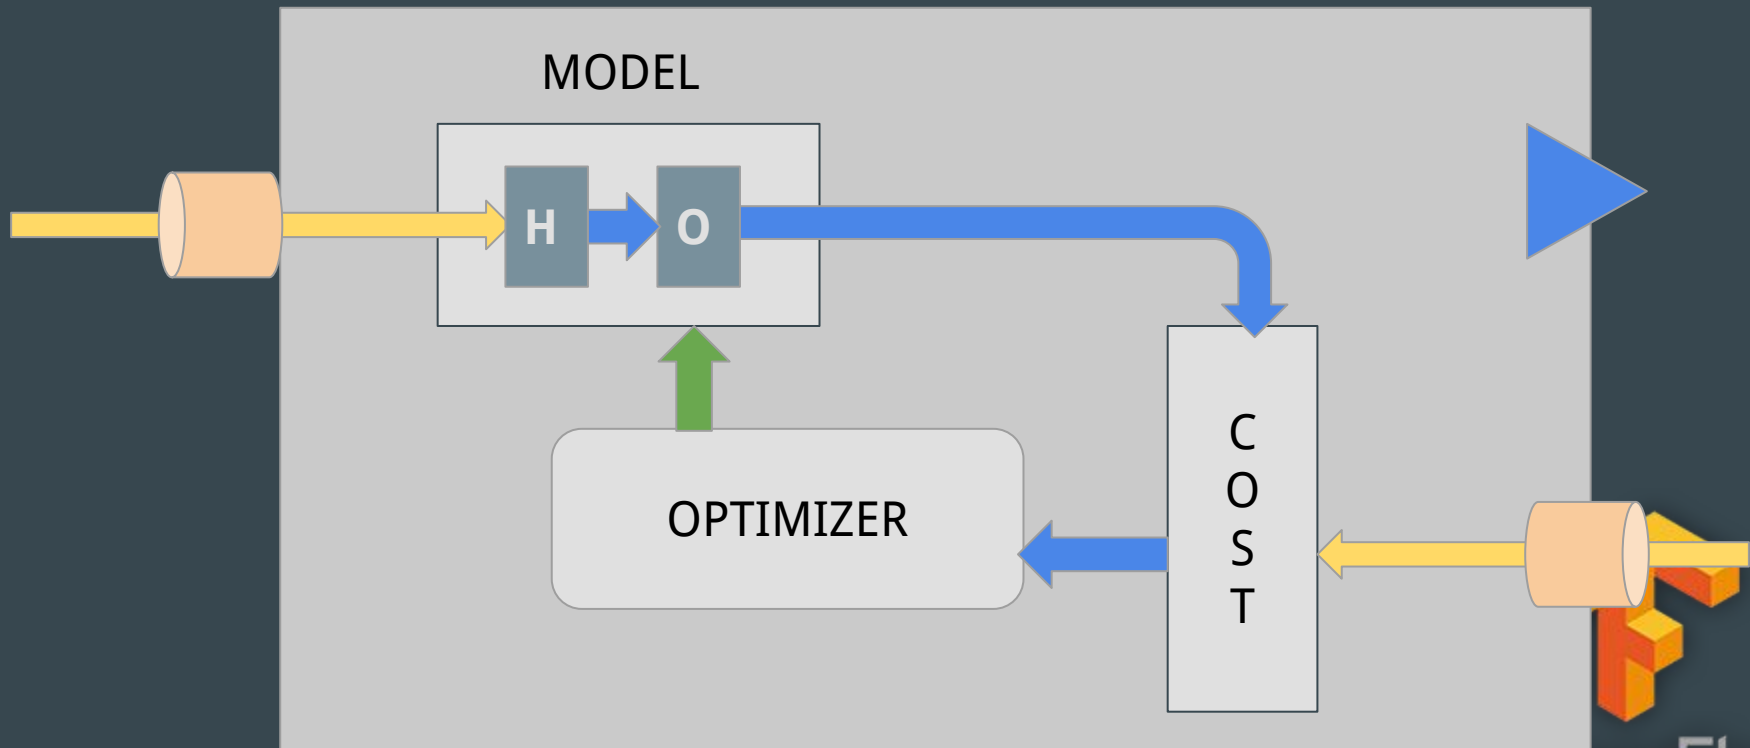
First layers extract the more basic features, while the next ones will work from this information.



Neural Networks: intuition

First layers extract the more basic features, while the next ones will work from this information.





```
Testing Accuracy: 1.0
```

```
Final test (5 + 3), (7 + 6), (10 + 10):
```

```
[[ 8.93008471e-01  1.06991500e-01]
 [ 1.24206737e-01  8.75793278e-01]
 [ 4.71629581e-04  9.99528408e-01]]
```

```
Hidden layer weights and bias:
```

```
[[ -1.64704931]
 [ -1.64078069]]
 [ 0.19933932]
```

```
Output layer weights and bias:
```

```
[[ 1.21974468 -1.26005733]]
 [-0.2669307]
```



TensorFlow

To know more...

Deep learning

- Neural Networks and Deep Learning - **Michael Nielsen**
- Stanford's CS231n - **Andrej Karpathy**

Tensorflow

- Tensorflow Tutorials - **Hvass Laboratories**
- Deep Learning Foundations Nanodegree - **Udacity**



TensorFlow

To start to know more...

Basics

- Intro to Data Science - Udacity
- Intro to Machine Learning - Udacity





`alesolano/mastering_tensorflow`



TensorFlow



+

A.I.



TensorFlow

