

# Distributed TensorFlow



**CSE545 - Spring 2022**  
**Stony Brook University**


H. Andrew Schwartz

# Big Data Analytics, The Class

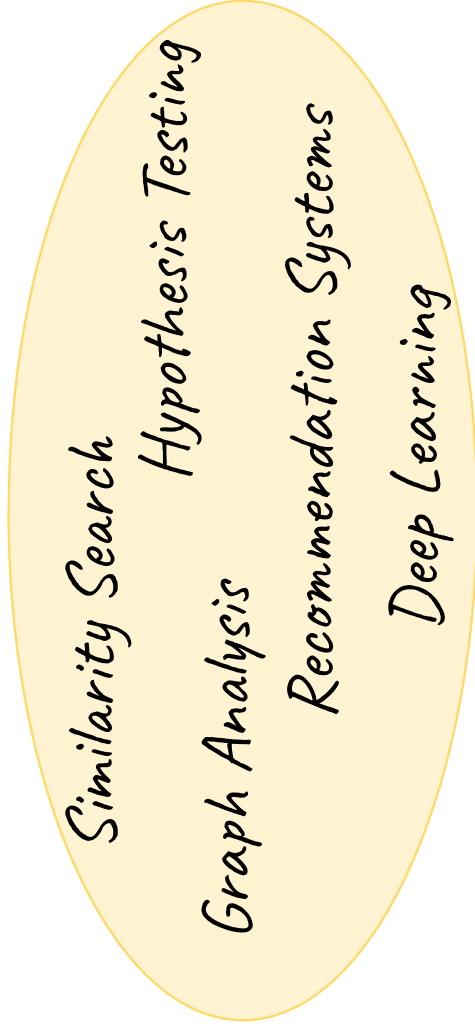
**Goal:** Generalizations  
A model or summarization of the data.



*Data Frameworks*



*Algorithms and Analyses*



# Limitations of Spark

Spark is fast for being so flexible

- Fast: RDDs in memory + Lazy evaluation: optimized chain of operations.
- Flexible: Many transformations -- can contain any custom code.

# Limitations of Spark

Spark is fast for being so flexible

- Fast: RDDs in memory + Lazy evaluation: optimized chain of operations.
- Flexible: Many transformations -- can contain any custom code.

However:

- Hadoop MapReduce can still be better for extreme IO, data that will not fit in memory across cluster.


# Limitations of Spark

Spark is fast for being so flexible

- Fast: RDDs in memory + Lazy evaluation: optimized chain of operations.
- Flexible: Many transformations -- can contain any custom code.

However:

- Hadoop MapReduce can still be better for extreme IO, data that will not fit in memory across cluster.



# Limitations of Spark

Spark is fast for being so flexible

- Fast: RDDs in memory + Lazy evaluation: optimized chain of operations.
- Flexible: Many transformations -- can contain any custom code.

However:

- Hadoop MapReduce can still be better for extreme IO, data that will not fit in memory across cluster.
- Modern machine learning (esp. Deep learning), a common big data task, requires heavy numeric computation.

IO Bound

*MapReduce*

*Spark*

Compute Bound

(large files: TBs or PBs)

(1s of TBs, 100s of GBs) (many numeric computations)

\* this is the subjective approximation of the instructor as of February 2020. A lot of factors at play.

# Limitations of Spark

Spark is fast for being so flexible

- Fast: RDDs in memory + Lazy evaluation: optimized chain of operations.
- Flexible: Many transformations -- can contain any custom code.

However:

- Hadoop MapReduce can still be better for extreme IO, data that will not fit in memory across cluster.
- Modern machine learning (esp. Deep learning), a common big data task, requires heavy numeric computation.

IO Bound

*MapReduce*

*Spark*

*TensorFlow*

Compute Bound

(large files: TBs or PBs)

(1s of TBs, 100s of GBs) (many numeric computations)

\* this is the subjective approximation of the instructor as of February 2020. A lot of factors at play.

# Learning Objectives

- Understand TensorFlow as a data workflow system.
  - Know the key components of TensorFlow.
  - Understand the key concepts of *distributed* TensorFlow.
- Execute basic distributed tensorflow program.
- Establish a foundation to distribute deep learning models:
  - Convolutional Neural Networks
  - Recurrent Neural Network (or LSTM, GRU)



# What is TensorFlow?


A workflow system catered to numerical computation.

One view: Like Spark, but uses *tensors* instead of *RDDs*.

# What is TensorFlow?

A workflow system catered to numerical computation.

One view: Like Spark, but uses *tensors* instead of *RDDs*.



A multi-dimensional matrix

(i.stack.imgur.com)

# What is TensorFlow?


A workflow system catered to numerical computation.

One view: Like Spark, but uses *tensors* instead of *RDDs*.

A 2-d tensor is just a matrix.

1-d: vector

0-d: a constant / scalar



(i.stack.imgur.com)

**Note: Linguistic ambiguity:  
Dimensions of a Tensor  $\neq$   
Dimensions of a Matrix**

# What is TensorFlow?

A workflow system catered to numerical computation.

One view: Like Spark, but uses *tensors* instead of *RDDs*.



Examples > 2-d :

Image definitions in terms of RGB per pixel

Image[row][column][rgb]

Subject, Verb, Object representation of language:

Counts[verb][subject][object]

# What is TensorFlow?

A workflow system catered to numerical computation.

One view: Like Spark, but uses *tensors* instead of *RDDs*.



Technically, less abstract than *RDDs* which could hold tensors as well as many other data structures (dictionaries/HashMaps, Trees, ...etc...).

Then, why TensorFlow?

# What is TensorFlow?

Efficient, high-level built-in **linear algebra** and **machine learning optimization operations** (i.e. transformations).

enables complex models, like deep learning

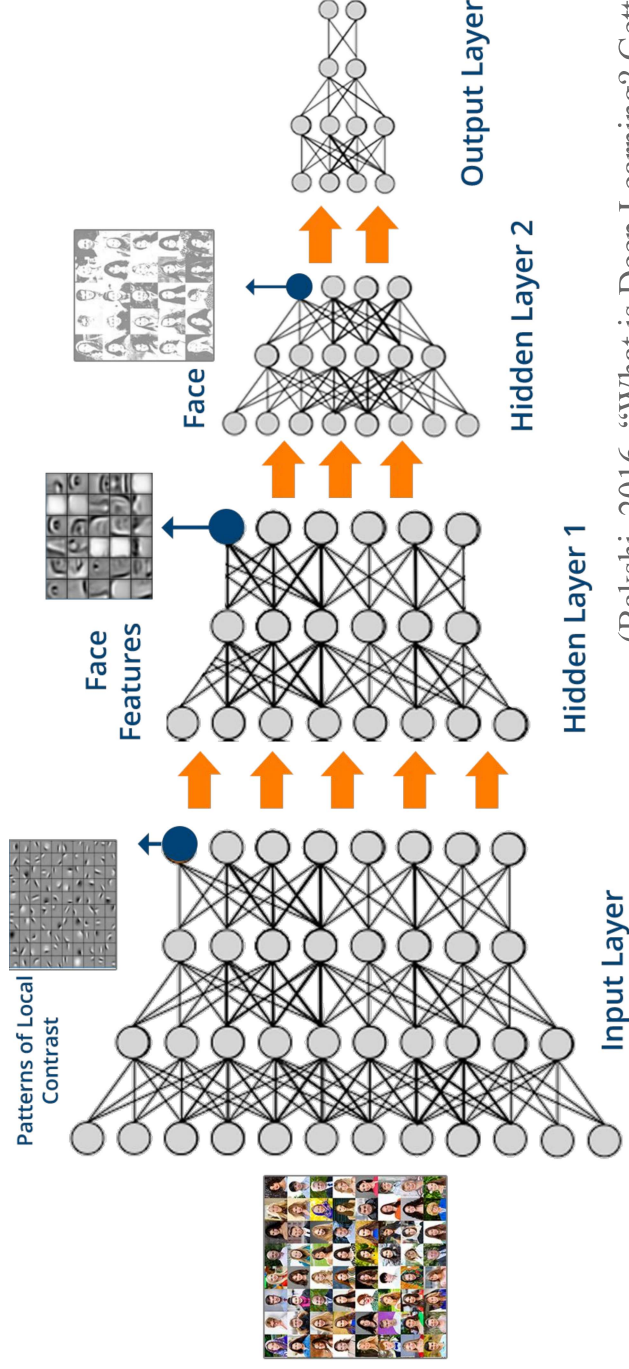


**Then, why TensorFlow?**

# What is TensorFlow?

Efficient, high-level built-in **linear algebra** and **machine learning optimization operations**.

enables complex models, like deep learning



(Bakshi, 2016, "What is Deep Learning? Getting Started With Deep Learning")

# What is TensorFlow?

## Efficient, high-level built-in linear algebra and machine learning operations.

```
import tensorflow as tf

b = tf.Variable(tf.zeros([100]))
W = tf.Variable(tf.random_uniform([784, 100], -1, 1))
x = tf.placeholder(name="x")
relu = tf.nn.relu(tf.matmul(W, x) + b)
C = [...]

s = tf.Session()
for step in xrange(0, 10):
    input = ...construct 100-D input array ...
    result = s.run(C, feed_dict={x: input})
    print step, result

# 100-d vector, init to zeroes
# 784x100 matrix w/rnd vals
# Placeholder for input
# ReLU(Wx+b)
# Cost computed as a function
# of ReLU

# Create 100-d vector for input
# Fetch cost, feeding x=input
```

(Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., ... & Ghemawat, S. (2016). TensorFlow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*.)




# TensorFlow

Operations on tensors are often conceptualized as **graphs**:

```
import tensorflow as tf

b = tf.Variable(tf.zeros([100]))
W = tf.Variable(tf.random_uniform([784, 100], -1, 1))
x = tf.placeholder(name="x")
relu = tf.nn.relu(tf.matmul(W, x) + b)
C = [...]

s = tf.Session()
for step in xrange(0, 10):
    input = ...construct 100-D input array ...
    result = s.run(C, feed_dict={x: input})
    print step, result
```




(Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., ... & Ghemawat, S. (2016). TensorFlow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*.)

# TensorFlow

Operations on tensors are often conceptualized as graphs:

A simpler example:

```
c = tensorflow.matmul(a, b)
```



# TensorFlow


Operations on tensors are often conceptualized as graphs:

example:

$$d = b + c$$

$$e = c + 2$$

$$a = d * e$$



(Adventures in Machine Learning.  
*Python TensorFlow Tutorial*, 2017)

# Ingredients of a TensorFlow

## **tensors\***

*variables* - persistent  
mutable tensors  
*constants* - constant  
*placeholders* - from data

\* technically, still *operations*

**operations**  
an abstract computation  
(e.g. matrix multiply, add)  
executed by device *kernels*

**graph**

## **session**

defines the environment in  
which operations *run*.  
(like a Spark context)

## **devices**

the specific devices (cpus or  
gpus) on which to run the  
session.

# Ingredients of a TensorFlow

## **tensors\***

- variables* - persistent mutable tensors
- constants* - constant
- placeholders* - from data

- `tf.Variable(initial_value, name)`
  - `tf.constant(value, type, name)`
  - `tf.placeholder(type, shape, name)`
- executed by device kernels

\* technically, still *operations*

**graph**

## **session**

defines the environment in which operations *run*.  
(like a Spark context)

## **devices**

the specific devices (cpus or gpus) on which to run the session.

# Ingredients of a TensorFlow

## *tensors\**

*variables* - persistent  
mutable tensors  
*constants* - constant  
*placeholders* - from data

## *operations*

an abstract computation  
(e.g. matrix multiply, add)  
executed by device *kernels*

## Category

Element-wise mathematical operations  
Array operations  
Matrix operations  
Stateful operations  
Neural-net building blocks  
Checkpointing operations  
Queue and synchronization operations  
Control flow operations

## Examples

Add, Sub, Mul, Div, Exp, Log, Greater, Less, Equal, ...  
Concat, Slice, Split, Constant, Rank, Shape, Shuffle, ...  
MatMul, MatrixInverse, MatrixDeterminant, ...  
Variable, Assign, AssignAdd, ...  
SoftMax, Sigmoid, ReLU, Convolution2D, MaxPool, ...  
Save, Restore  
Enqueue, Dequeue, MutexAcquire, MutexRelease, ...  
Merge, Switch, Enter, Leave, NextIteration

# Ingredients of a TensorFlow

- **Places operations on devices**  
*variables - persistent*
- **Stores the values of variables (when not distributed)**  
*constants - constant placeholders from data*
- **Carries out execution: eval() or run()**

## *operations*

an abstract computation  
(e.g. `matrix_multiply`, `add`)  
executed by device kernels

## *graph*

## *session*

defines the environment in  
which operations *run*.  
(like a Spark context)

## *devices*

the specific devices (cpus or  
gpus) on which to run the  
session.



# Ingredients of a TensorFlow

## **tensors\***

*variables* - persistent  
mutable tensors  
*constants* - constant  
*placeholders* - from data

## **operations**

an abstract computation  
(e.g. matrix multiply, add)  
executed by device *kernels*

**graph**

## **session**

defines the environment in  
which operations *run*.  
(like a Spark context)

## **devices**

the specific devices (cpus or  
gpus) on which to run the  
session.



# Distributed TensorFlow


Typical use-case: (Supervised Machine Learning)

Determine weights,  $\mathcal{W}$ , of a function,  $f$ , such that  $\mathcal{E}$  is minimized:  $f(\mathcal{X}|\mathcal{W}) = \mathcal{Y} + \mathcal{E}$

# Distributed TensorFlow

Typical use-case:


Determine weights,  $\mathcal{W}$ , of a function,  $f$ , such that  $\epsilon$  is minimized:  $f(X|\mathcal{W}) = \mathcal{Y} + \epsilon$



# Distributed TensorFlow

Typical use-case:


Determine weights,  $\mathcal{W}$ , of a function,  $f$ , such that  $\epsilon$  is minimized:  $f(X|\mathcal{W}) = \mathcal{Y} + \epsilon$



# Distributed TensorFlow

Typical use-case:

Determine weights,  $\mathcal{W}$ , of a function,  $f$ , such that  $\mathcal{E}$  is minimized:  $f(X|\mathcal{W}) = \mathcal{Y} + \mathcal{E}$




# Distributed TensorFlow

Typical use-case:

Determine weights,  $\mathcal{W}$ , of a function,  $f$ , such that  $|\varepsilon|$  is minimized:

$$\begin{aligned} f(X/W) &= \hat{Y} \\ Y &= (X/W) + \varepsilon \\ Y &= \hat{Y} + \varepsilon \\ \varepsilon &= \hat{Y} - Y \end{aligned}$$



# Distributed TensorFlow

Typical use-case:



Determine weights,  $\mathcal{W}$ , of a function,  $f$ , such that  $\mathcal{E}$  is minimized:

$$f(X/W) = \hat{Y}$$

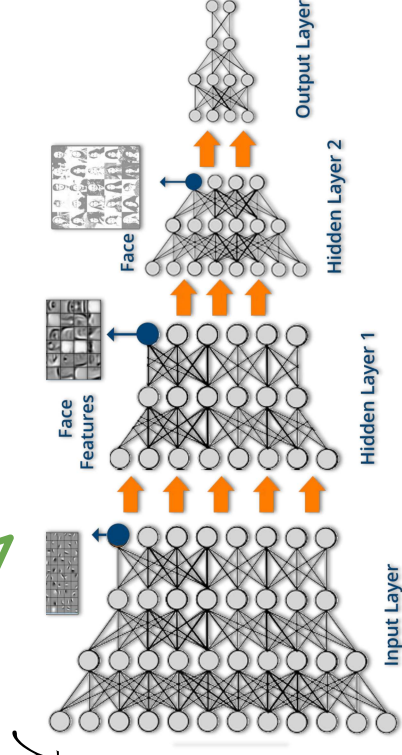
$$Y = (X/W) + \mathcal{E}$$

$$Y = \hat{Y} + \mathcal{E}$$

$$\mathcal{E} = \hat{Y} - Y$$



Typically, very complex!



# Distributed TensorFlow

Typical use-case:

Determine weights,  $\mathcal{W}$ , of a function,  $f$ , such that  $\mathcal{E}$  is minimized:

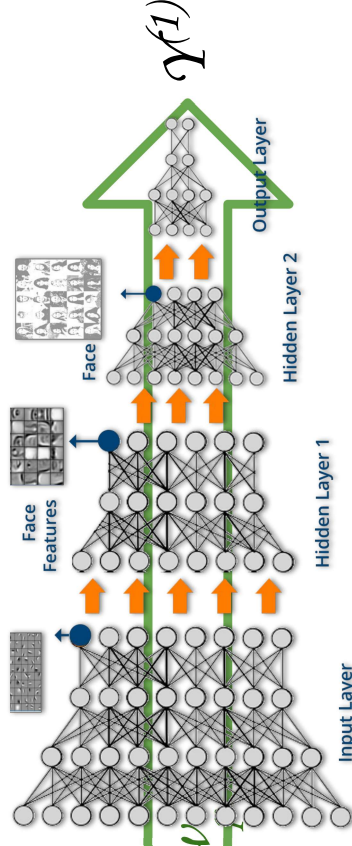
$\mathcal{W}$  determined through *gradient descent*:

*back propagating* error across the network that defines  $f$ .

$$\begin{aligned} f(X/W) &= \hat{Y} \\ Y &= (X/W) + \mathcal{E} \\ Y &= \hat{Y} + \mathcal{E} \\ \mathcal{E} &= \hat{Y} - Y \end{aligned}$$

$$\begin{array}{cccccccc} X_1^{(1)} & X_2 & X_3 & X_4 & X_5 & X_6 & & \\ X_7 & X_8 & X_9 & & & & & \\ X_{13} & X_{14} & X_{15} & & & & & \end{array}$$

$f$  given  $w_1, w_2, \dots, w_p$   
(typically,  $p \geq m$ )



# Distributed TensorFlow


Typical use-case:

Determine weights,  $\mathcal{W}$ , of a function,  $f$ , such that  $\epsilon$  is minimized:

$$\begin{aligned} f(X/W) &= \hat{Y} \\ Y &= (X/W) + \epsilon \\ Y &= \hat{Y} + \epsilon \\ \epsilon &= \hat{Y} - Y \end{aligned}$$

$\mathcal{W}$  determined through **gradient descent**:

back propagating error across the network that defines  $f$ .



minimizes  $\epsilon$  on  $M$  training examples