

# AI: TensorFlow

---

**CPSC 501: Advanced Programming Techniques  
Fall 2020**

Jonathan Hudson, Ph.D  
Instructor  
Department of Computer Science  
University of Calgary

Wednesday, August 12, 2020



# Machine learning

---

- Mathematical concepts behind deep learning have been around for decades
- Programming libraries dedicated to creating and training these deep models have only been available in recent years
- Often a trade-off
  - Flexible for research
  - Fast and efficient for production
- TensorFlow attempts to solve this problem

# Context (circa 2015)

---

- Deep learning already claiming big successes
- Number of developers/researchers exploding
- A “zoo” of tools and libraries, some of questionable quality...

# TensorFlow

---

# What is TensorFlow?

---

- Created by Google as an internal machine learning tool
- Open sourced under the Apache 2.0 License in November 2015
- An open-source software library for Machine Intelligence
  - numerical computation using data flow graphs
- Google maintains its own internal version.
- One of more than a dozen of machine intelligence libraries developed by big companies
- Other libraries
  - [https://en.wikipedia.org/wiki/Comparison\\_of\\_deep\\_learning\\_software](https://en.wikipedia.org/wiki/Comparison_of_deep_learning_software)

# What is TensorFlow?

- Open source library for numerical computation using **data flow graphs**
- Developed by Google Brain Team to conduct machine learning research
  - Based on DisBelief used internally at Google since 2011
- “TensorFlow is an interface for expressing machine learning algorithms, and an implementation for executing such algorithms”

tensorflow / tensorflow

Watch 7,777 Star 96,717 Fork 61,507

Code Issues 1,313 Pull requests 196 Projects 0 Insights

Computation using data flow graphs for scalable machine learning <https://tensorflow.org>

tensorflow machine-learning python deep-learning deep-neural-networks neural-network ml distributed

31,895 commits 31 branches 54 releases 1,435 contributors Apache-2.0

# Why TensorFlow?

---

## 1. Popularity

- Top are Theano, Torch, and TensorFlow.
- Torch framework is written in Lua, which is a wonderful language but one most popular in game development and AI community.
- Theano has an additional “graph compilation” step that took significant amounts of time and frustration
- TensorFlow has a much cleaner interface as compared to Theano.
- TensorFlow was built with production use in mind

# Why TensorFlow?

---

- The ability to run in mobile environments
  - Easily build models that span multiple GPUs on a single machine
  - Train large-scale networks in a distributed fashion
1. Python API
  2. Portability: CPU/GPU
  3. Flexibility: OS environments
  4. Visualization (TensorBoard is very useful)
  5. Checkpoints (for managing experiments)
  6. Large community.



# Gradient computation: Backpropagation

---

```
train_step = tf.train.GradientDescentOptimizer(0.5).minimize(cross_entropy)
```

`tf.train.GradientDescentOptimizer` is an **Optimizer** object

```
tf.train.GradientDescentOptimizer(lr).minimize(cross_entropy)
```

adds optimization **operation** to computation graph

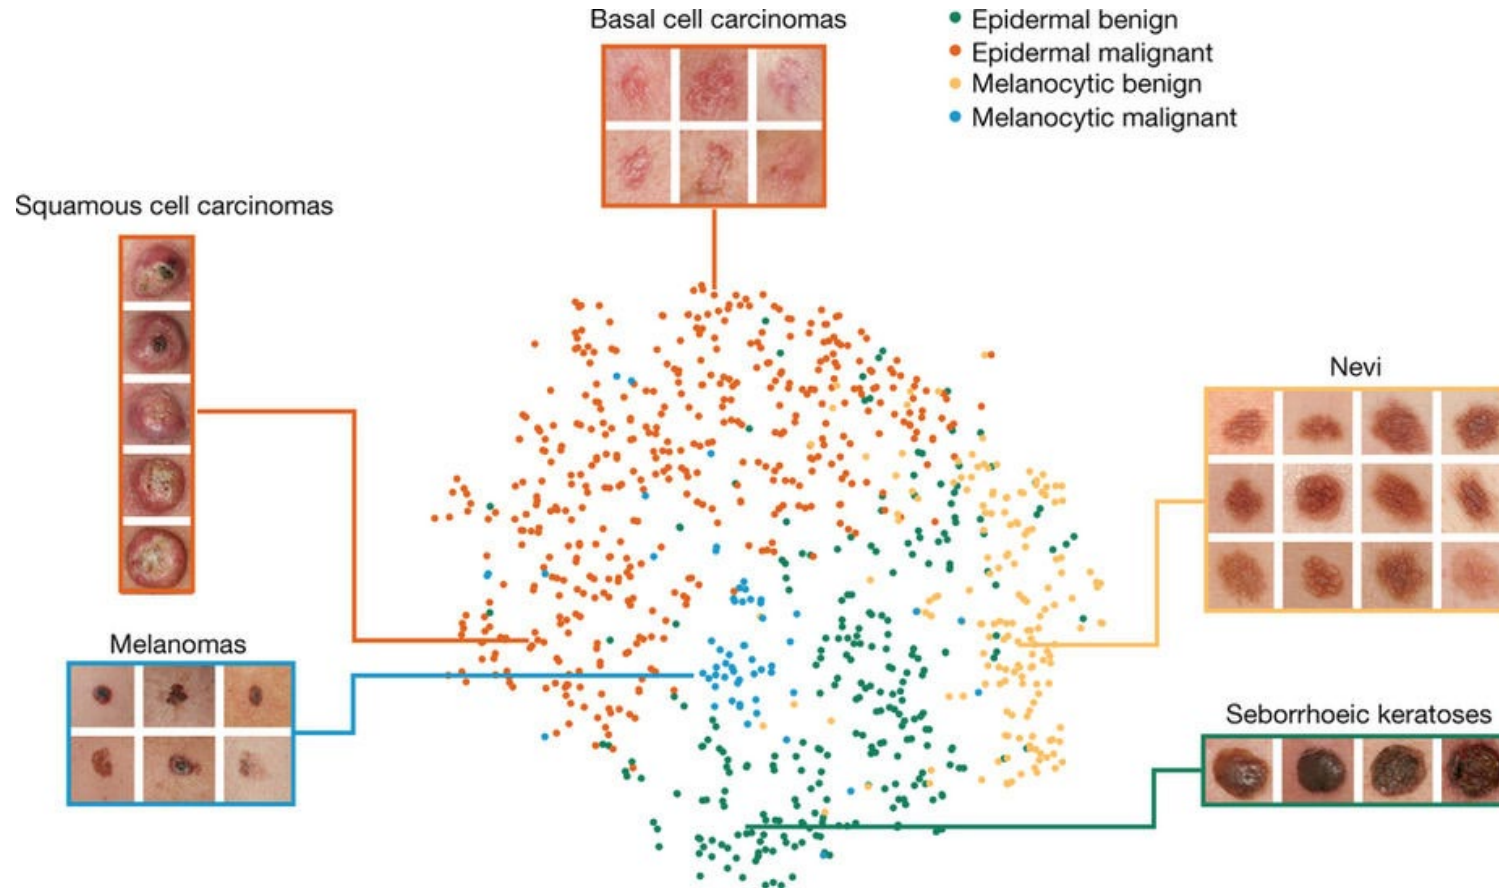
TensorFlow graph **nodes** have **attached gradient operations**

Gradient with respect to **parameters** computed with **backpropagation ... automatically**

# Examples

---

# Examples: Classify skin cancer



Dermatologist-level classification of skin cancer with deep neural networks (Esteva et al., Nature 2017)

# Examples: Neural Style Translation

---



Image Style Transfer Using Convolutional Neural Networks (Gatys et al., 2016) Tensorflow adaptation by Cameron Smith (cysmith@github)

# Goals

---

- Understand TF's approach to moving data through graphs
- Explore what is built into the language as tools
- Learn to build simple models for machine learning

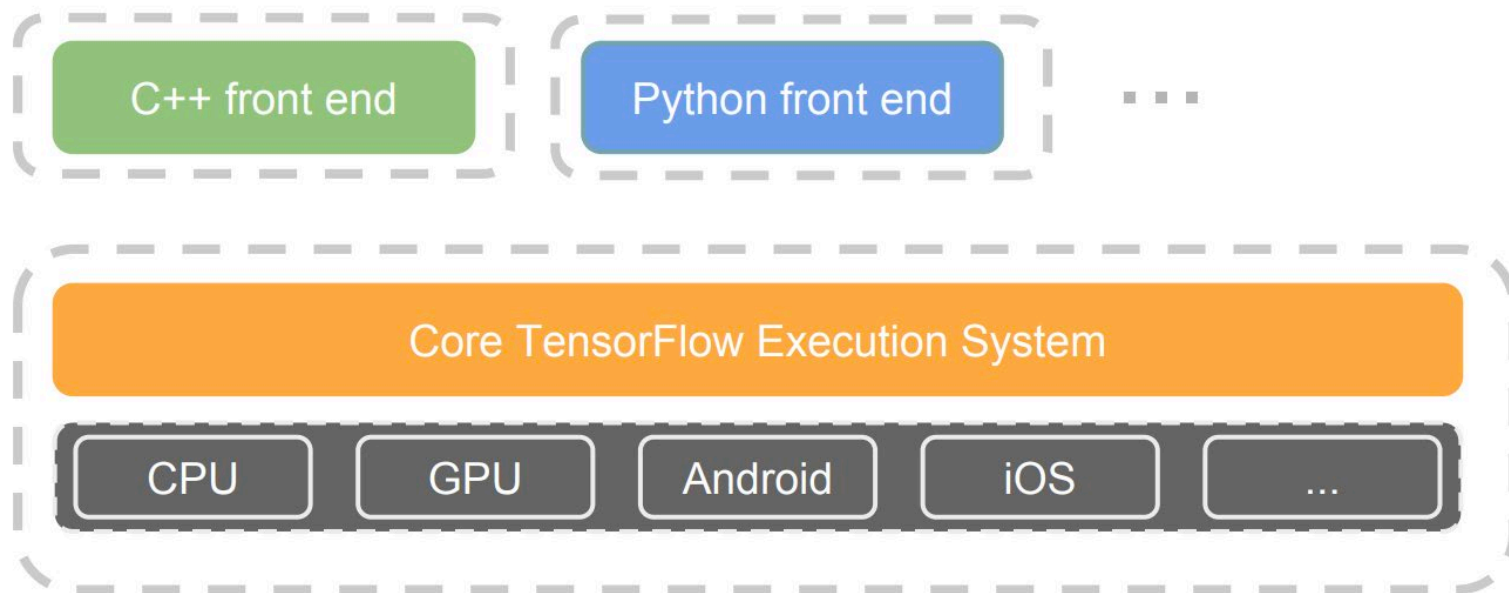
# Architecture

---

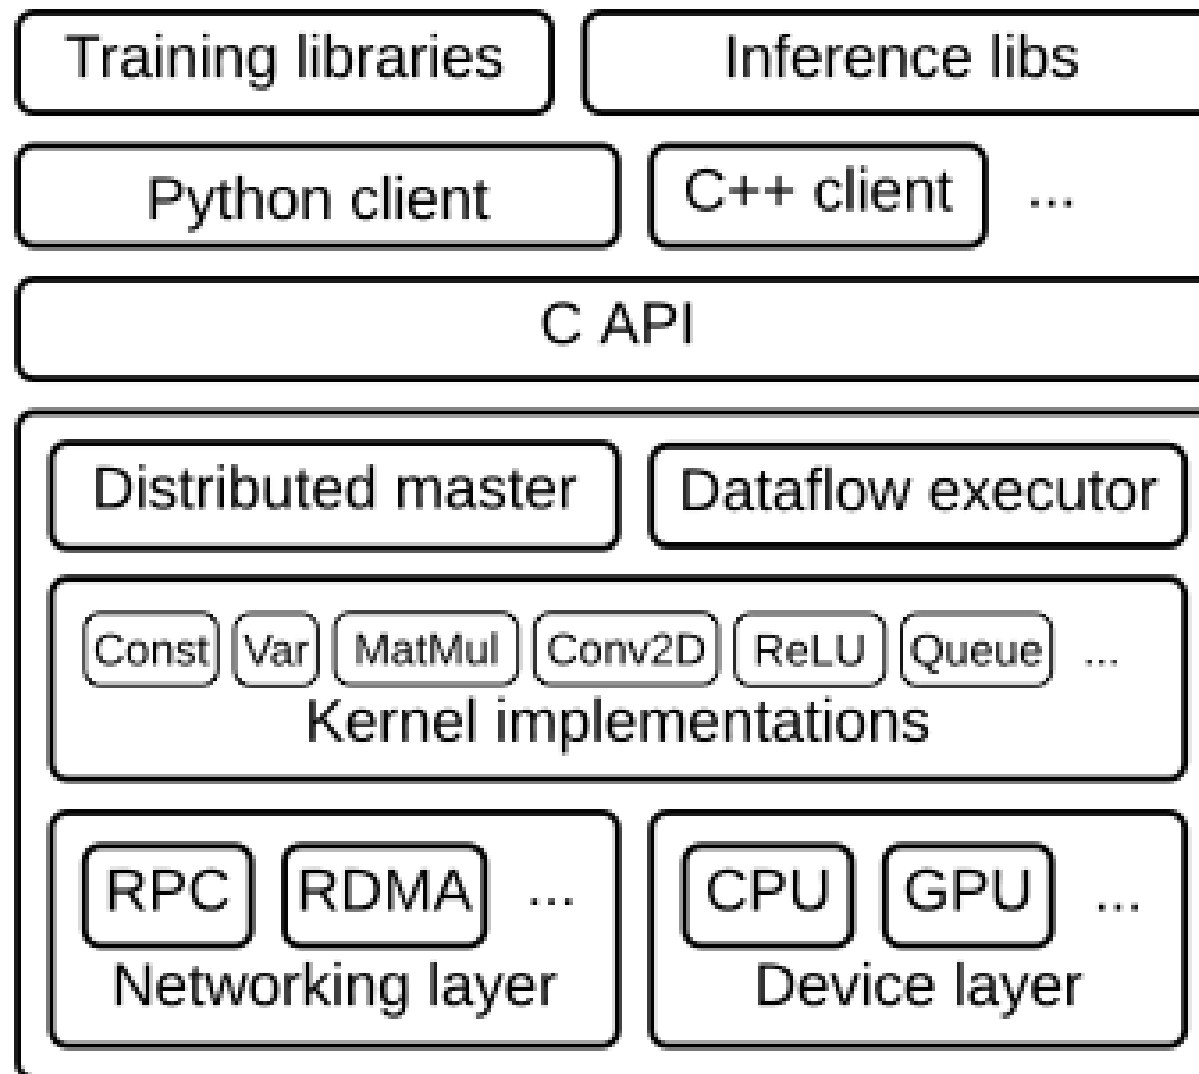
# TensorFlow architecture

---

- Core in C++
  - Very low overhead
- Different front ends for specifying/driving the computation
  - Started with Python and C++



# Detailed architecture






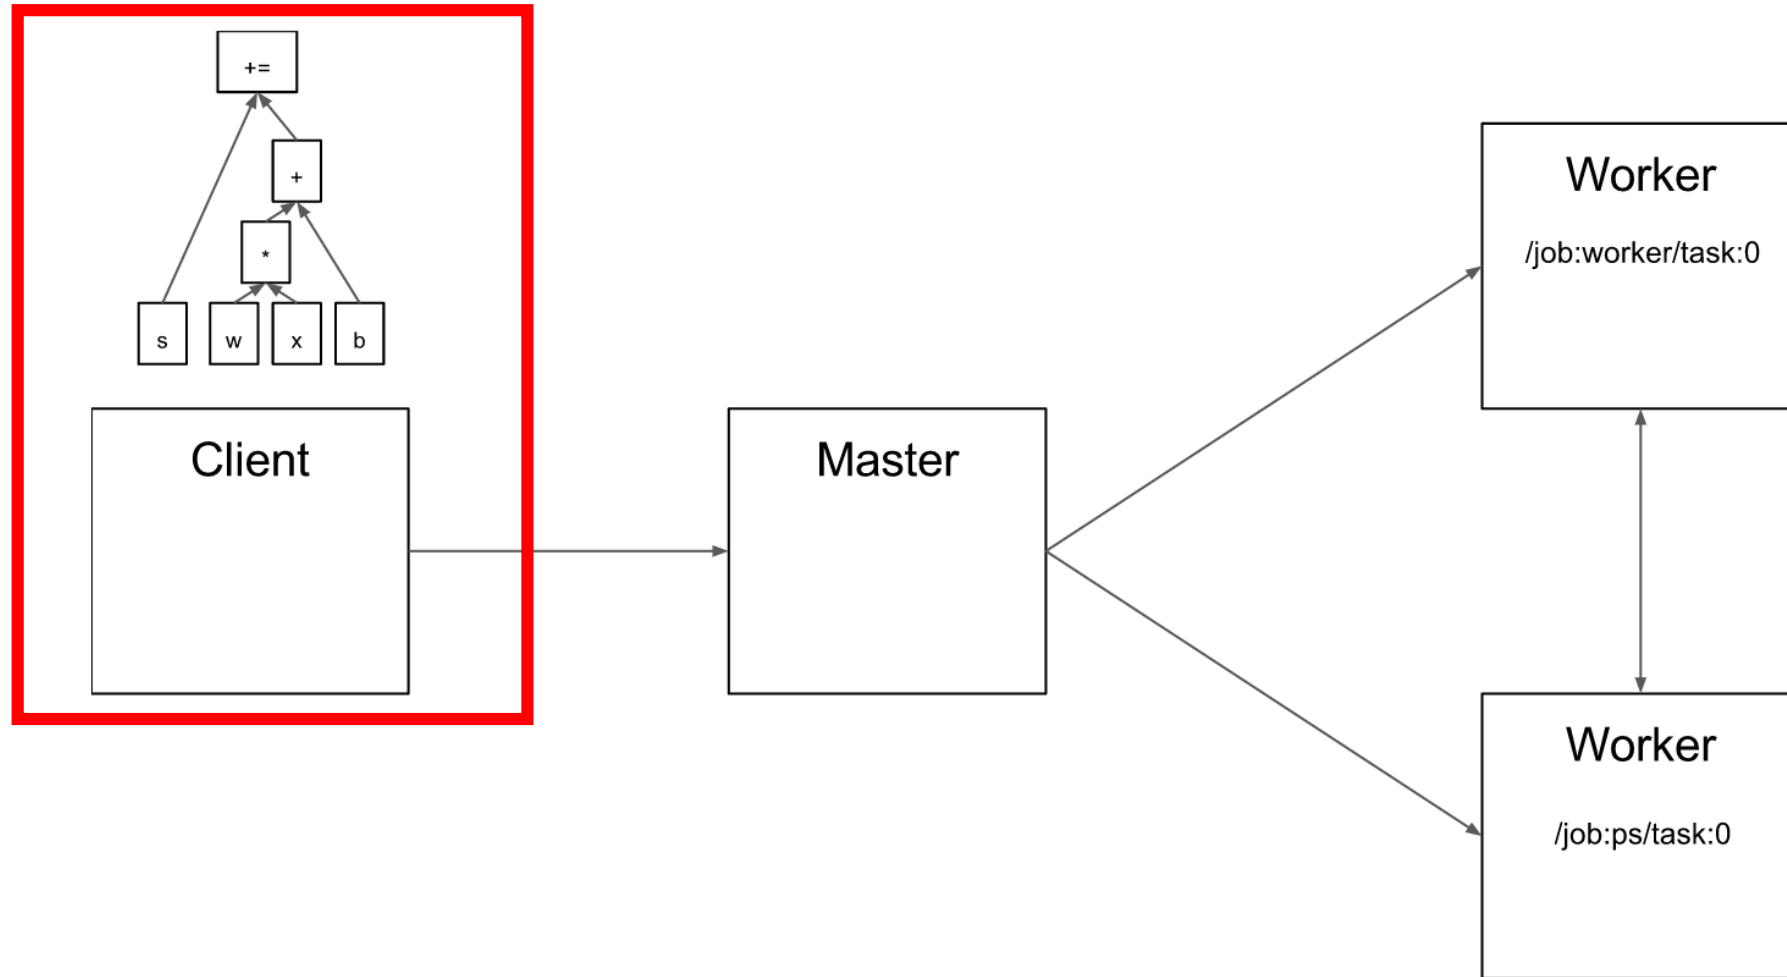
# Key components

---


- Similar to MapReduce, Apache Hadoop, Apache Spark, ...



# Client



# Master




# Computation Graphs

---


# Computation graph partition

PS

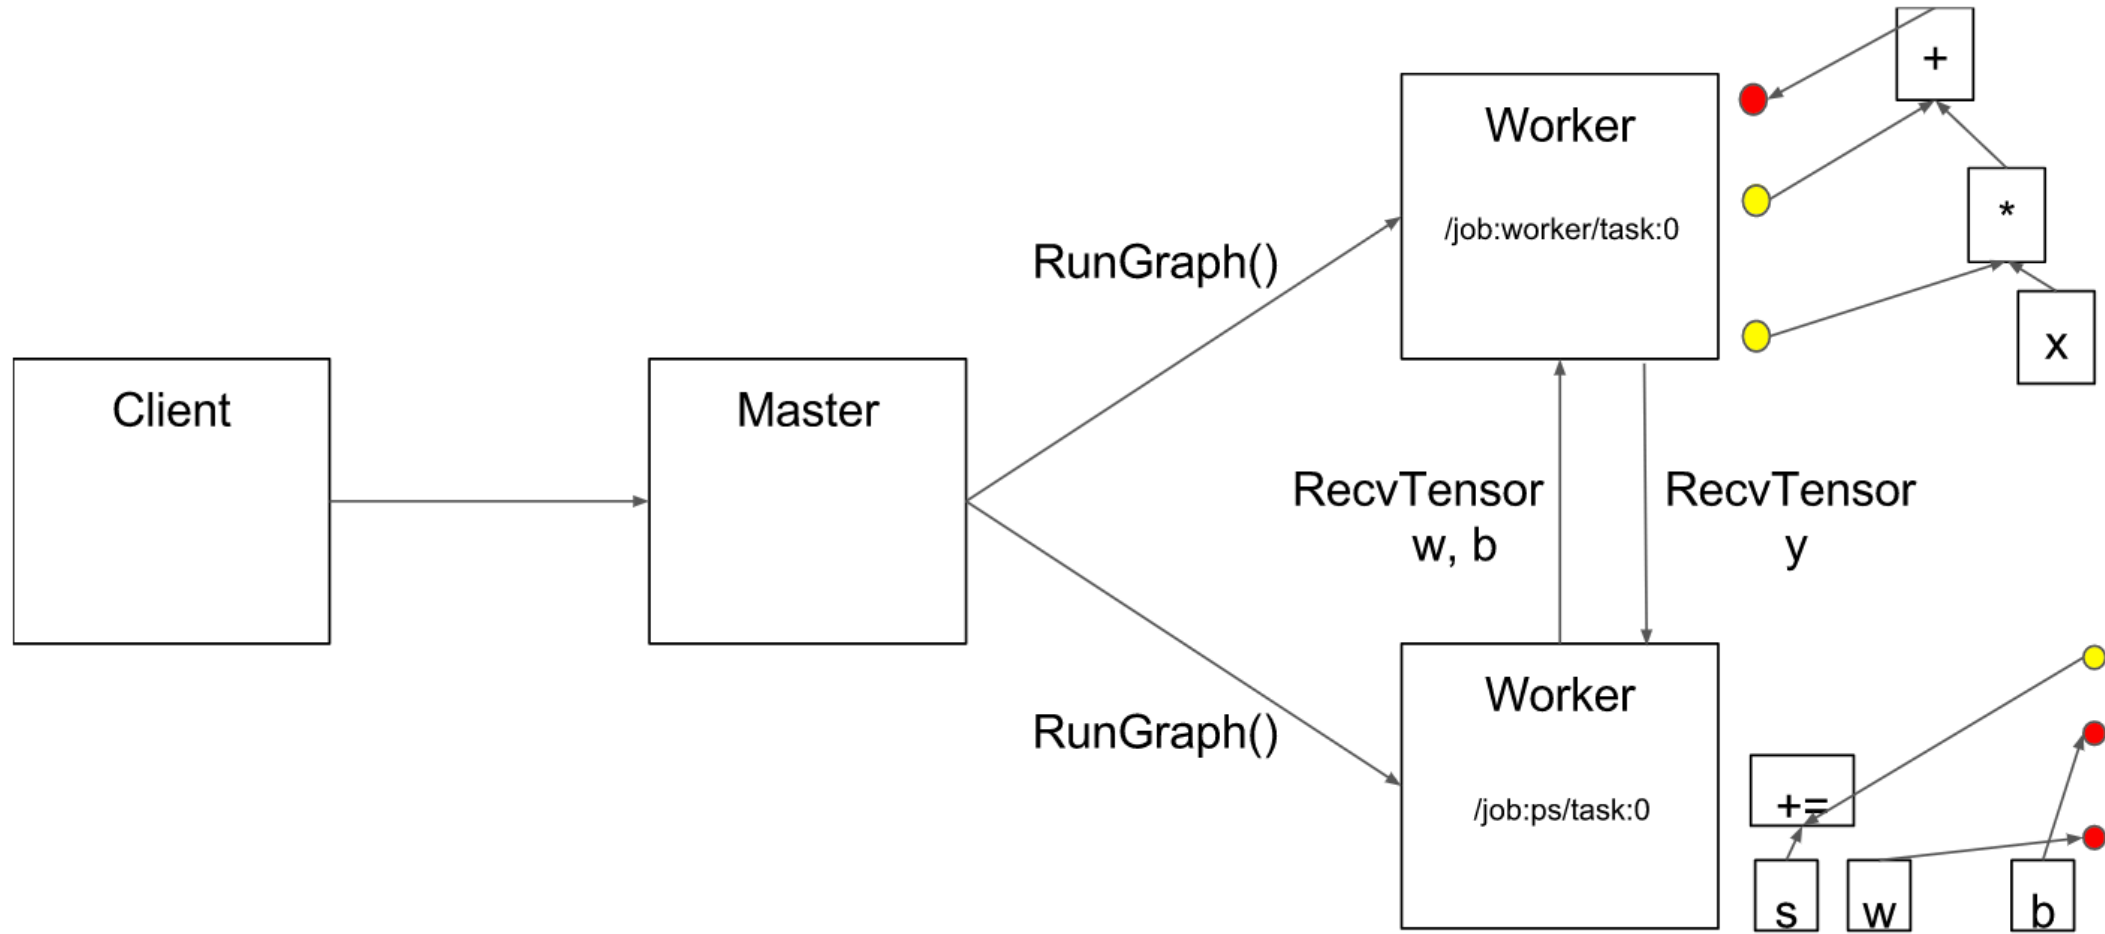
Worker



# Computation graph partition



# Execution



# Onward to ... TensorFlow Applied.

---

Jonathan Hudson  
[jwhudson@ucalgary.ca](mailto:jwhudson@ucalgary.ca)  
<https://pages.cpsc.ucalgary.ca/~hudsonj/>



UNIVERSITY OF  
CALGARY