Joker Conference, October 19-20, 2018 - St. Petersburg, Russia

TensorFlow for Java Developers

Derek Ferguson
Head of Engineering, Chase CB
derek.ferguson@mail.ru
derek.ferguson@chase.com

The Request

- The Business has been reading technology sites
- They want a machine learning solution
- They're looking for you to find a suitable problem… then solve it

Your Starting Point

- You have no prior background in machine learning
- You have no prior background in advanced math / data

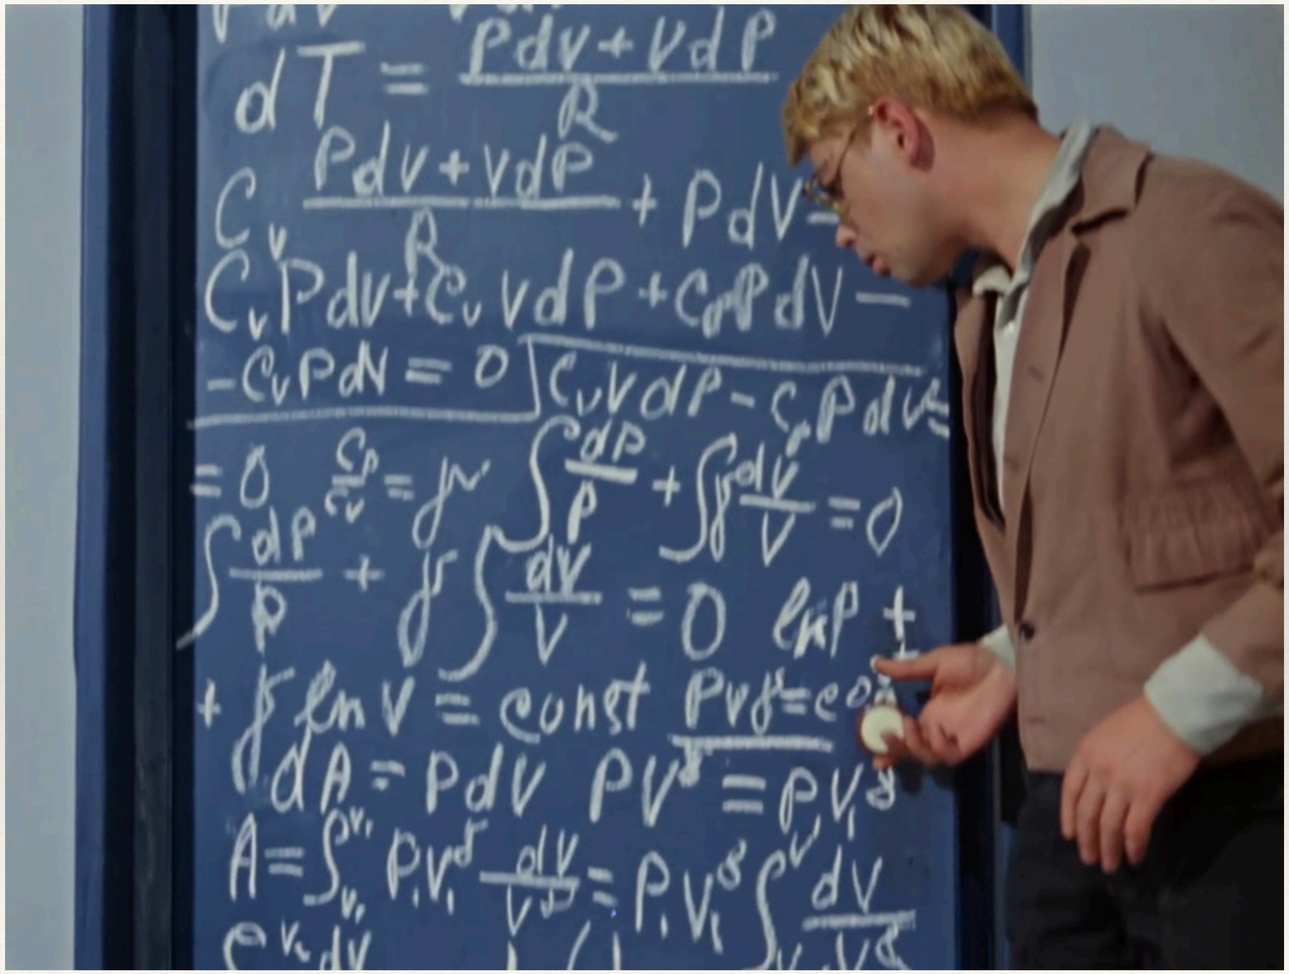

The Solution (aka “the Plan”)

- A very pragmatic, quick intro to machine learning
- An equally pragmatic, quick intro to TensorFlow
- 4 Java coding demonstrations

TensorFlow

What is Machine Learning?

- Teach a computer to solve 2 + 2 - 1 = __ <— arithmetic
- Teach a computer to solve 2 + __ - 1 = 4 <— algebra
- Teach a computer to solve 2 __ 3 __ 1 = 4 <— ML

Supervised Learning

- Start with “tagged” training data
- Use 2/3 of your data to evolve (“train”) an algorithm called a Model
- Use remaining 1/3 of your data to test your model with some predictions

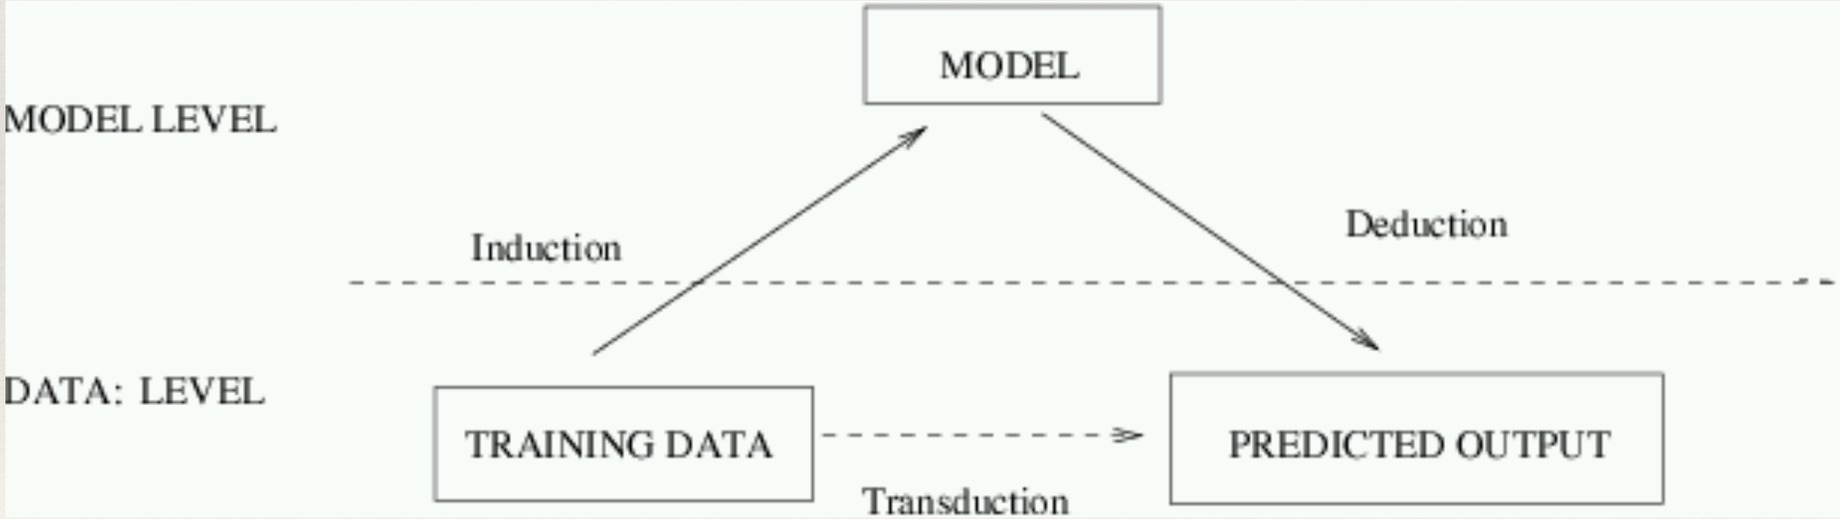

Induction

- $0 \_\_ 0 = 0$
- $1.5 \_\_ 3 = 4.5$
- $6 \_\_ 7 = 42$
- $1 \_\_ 2 \_\_ 1 \_\_ 3 \_\_ 2 \_\_ 1 \_\_ 7 \_\_ 4 \_\_ 3 \_\_ 7 \_\_ 8 \_\_ 4 \_\_ 4 = -2$

No Such Thing as Non-Numeric Data

- Do you smoke? (0 = No, 1 = Yes)
- Images and time data are just number matrices

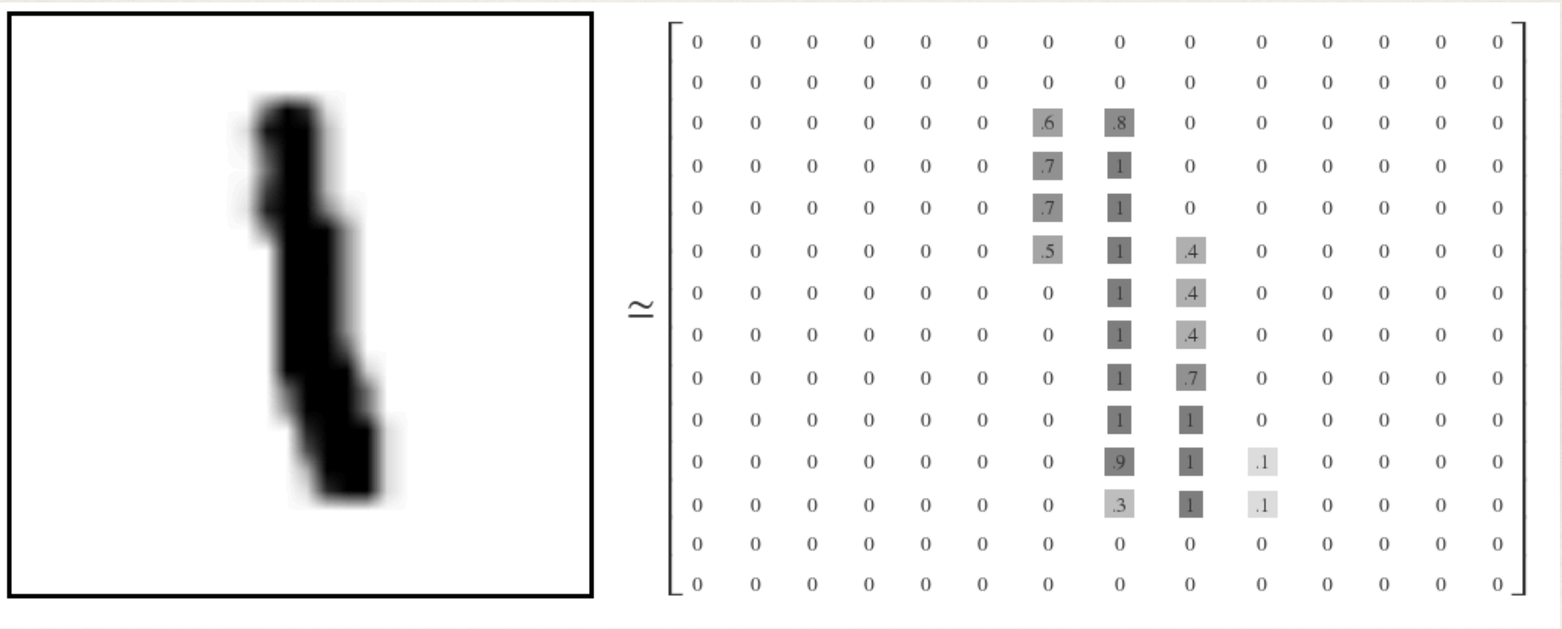

TensorFlow in 1 slide

- TensorFlow is an Open Source, C++ library for ML
- Its primary (think 90%+) wrapper is Python
- Java integration requires… craftiness!

- It provides containers called tensors for constants & variables & ops
- Sequenced collections of tensors are called graphs
- It provides execution environments for graphs to “flow” (train and predict)

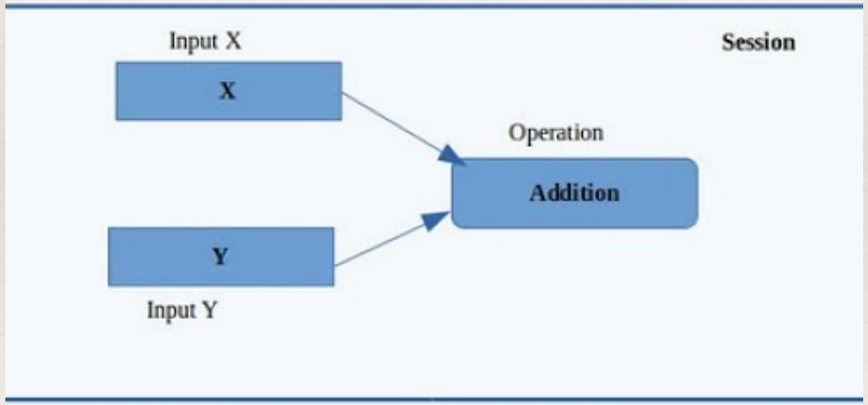

Tensors

The TensorFlow Java SDK

- Available as a maven dependency (and other formats)
- Wraps maybe 20% of the TensorFlow classes

What's in the 20%?

- Basic model creation
- Model training
- Inline predictions
- Predictions across the network

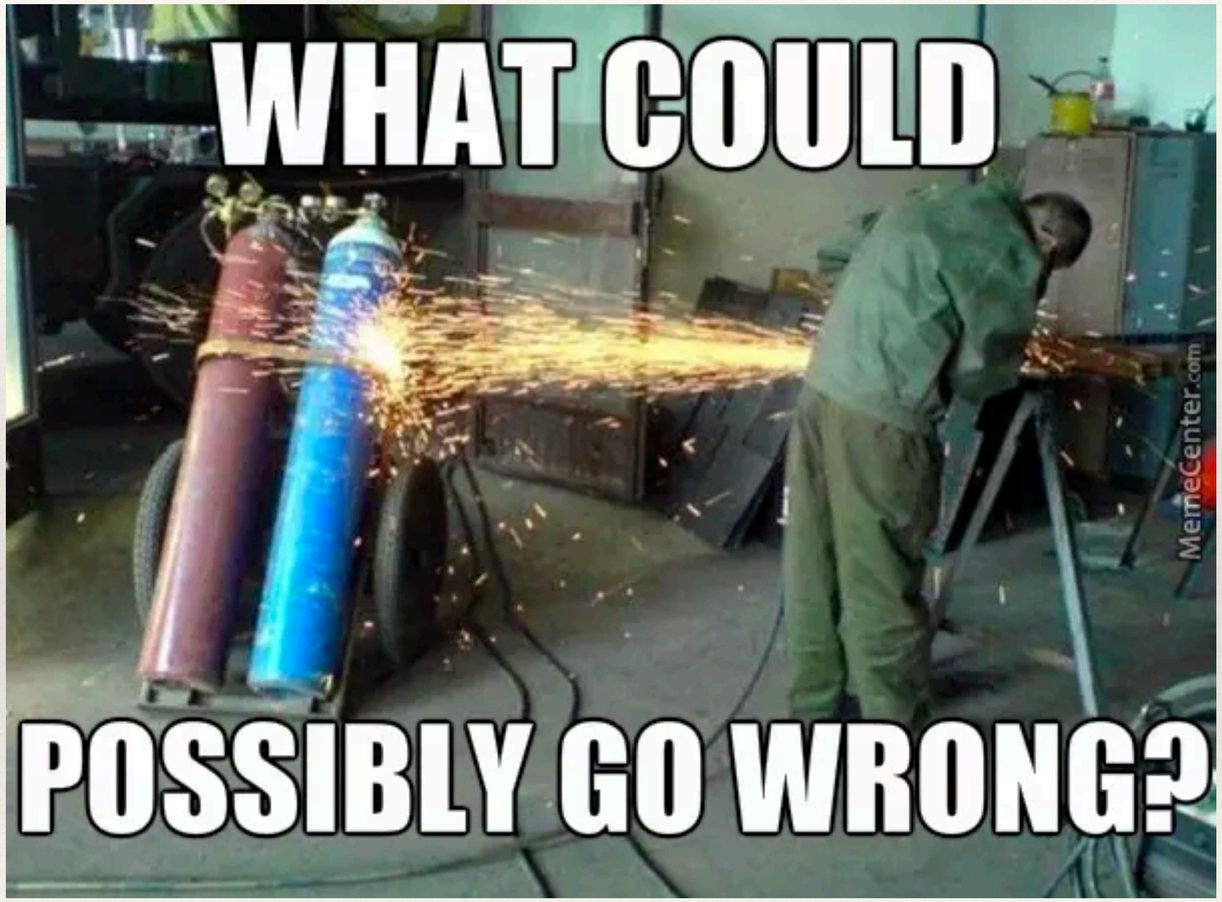

Essential Classes

- `org.tensorflow.Tensor<T>`
- `org.tensorflow.Output`
- `org.tensorflow.Graph`
- `org.tensorflow.Session`

Hands On Lab

Demo 1

Adding 2 numbers

Start with Raw Data

- Let's go back to a prior example: smoking
- We can represent this as cigarettes-per-capita — X
- We can also represent the number of deaths — Y

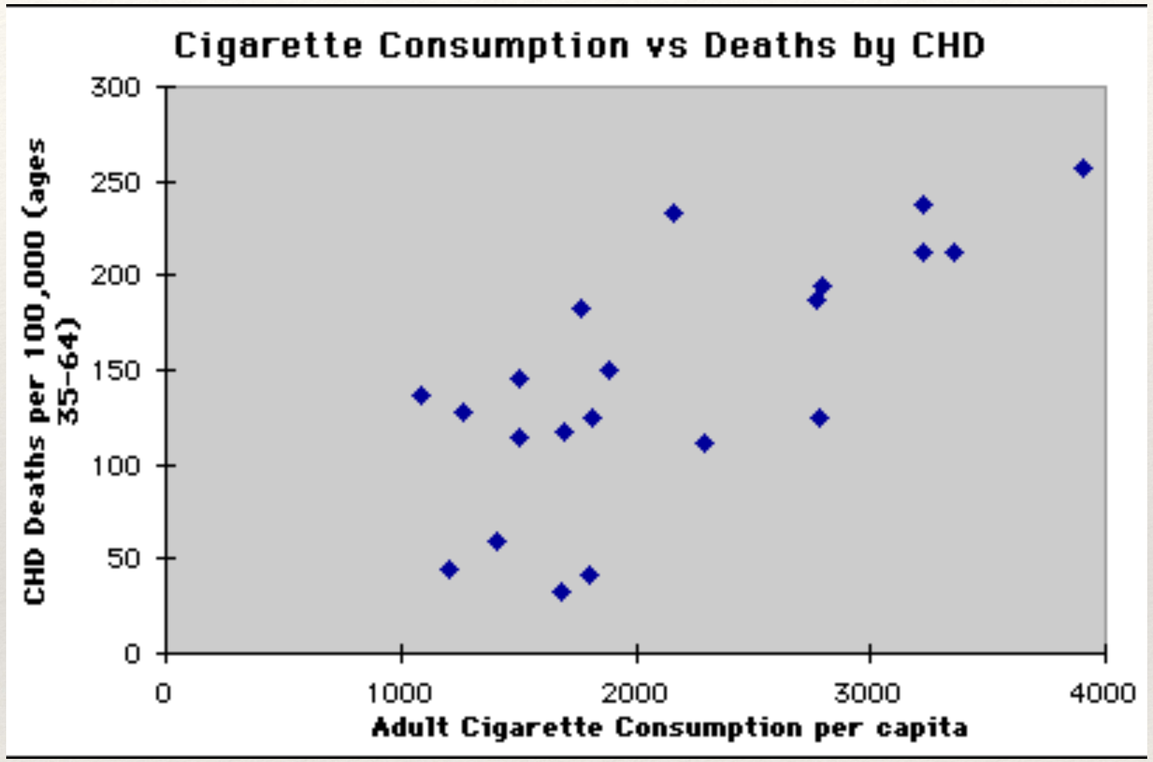

Raw Data Points

Predict via Linear Regression

- y = mX + b (worst math we'll touch, I promise!)

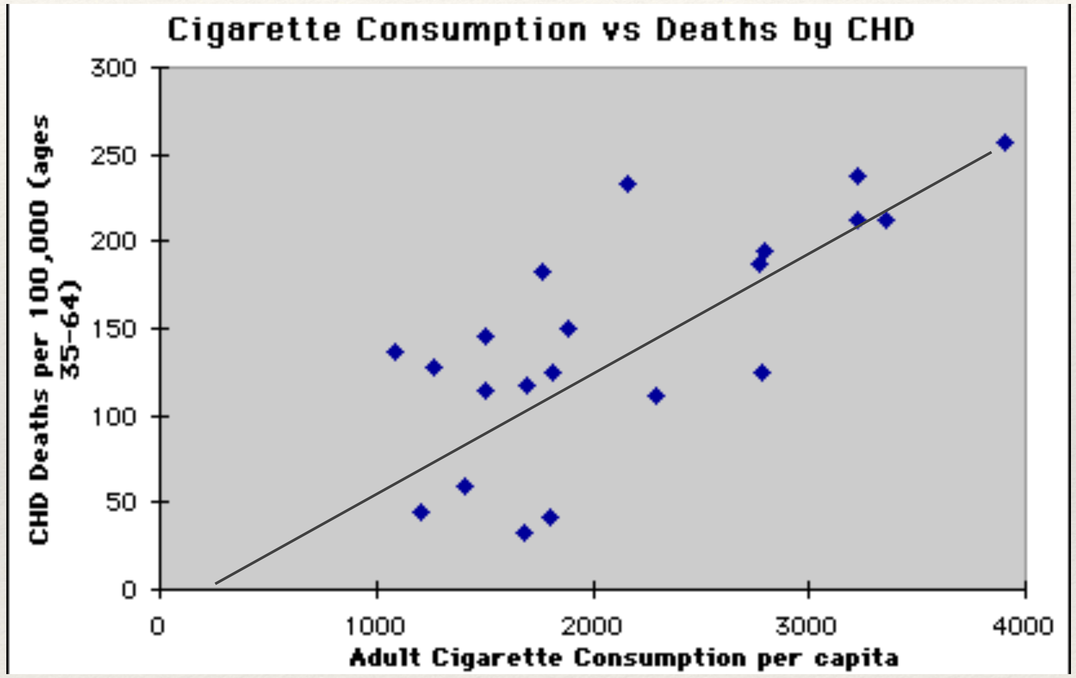

Fit Line

Evolving a model

- We want a class that adjusts the model to fit the data
- TF provides the GradientDescentOptimizer
- Of course - this isn't in the 20% :-(

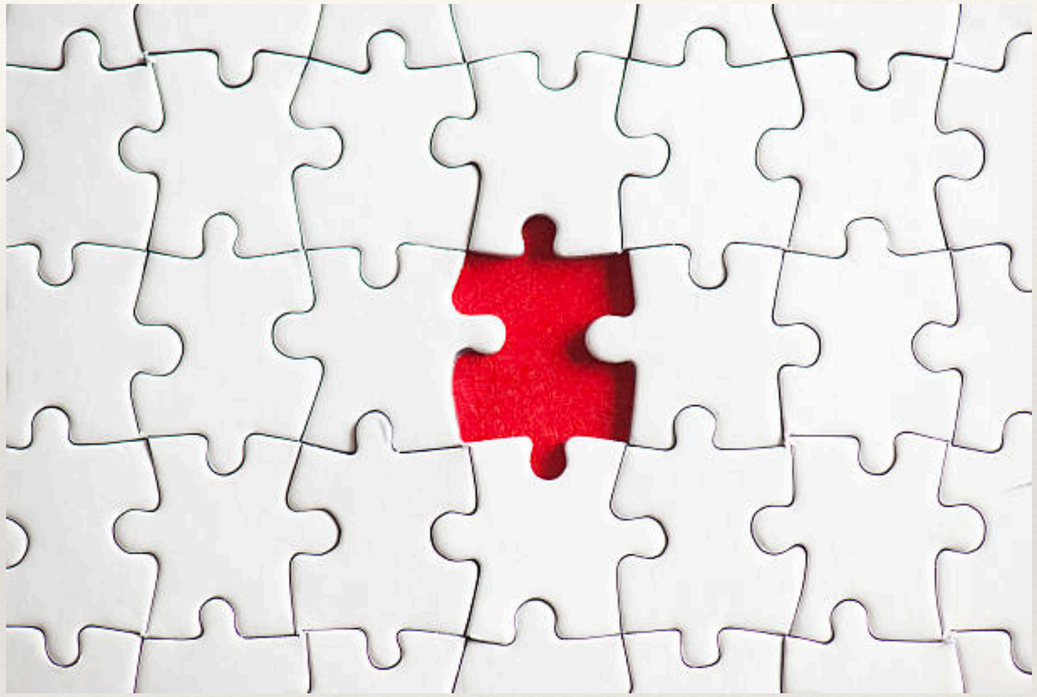

Getting the GradientDescentOptimizer

- You can download a binary of this model online

http://derekferguson-6c9ab059-eval-test.apigee.net/hello?w=3.0&b=2.0&lr=0.1

  - w — Provide the initial “slope”
  - b — Provide the initial “y intercept”
  - lr — Provide the “learning rate”

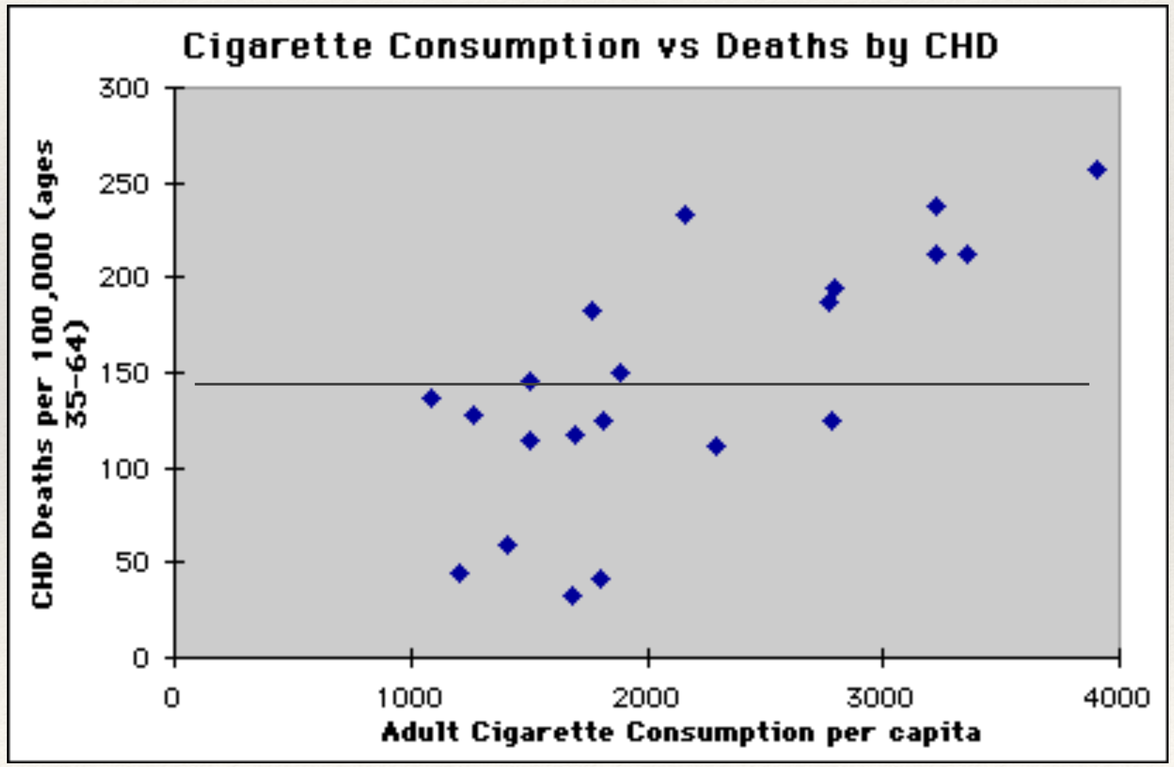

Initial Line

New TensorFlow Concepts

- Graph importation
- Place holders
- Feeds

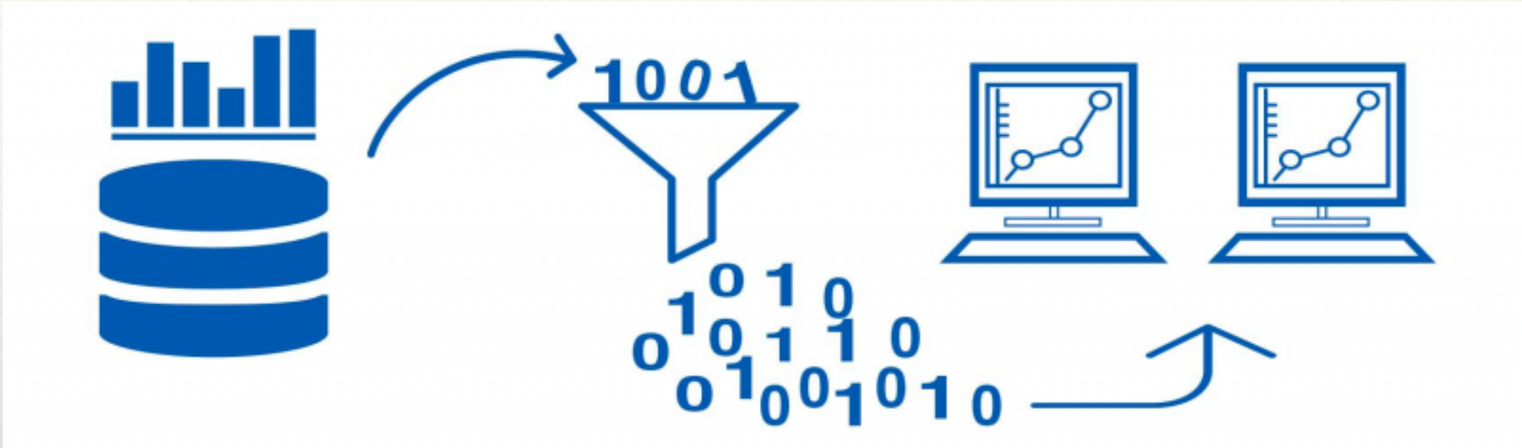

Hands On Lab

Demo 2

Training a model

Is everything a line? (Clue: No)

- At some point, smoking deaths level off
- Participation in the work force peaks by age
- Recognizing images of cars and bicycles - definitely not a line

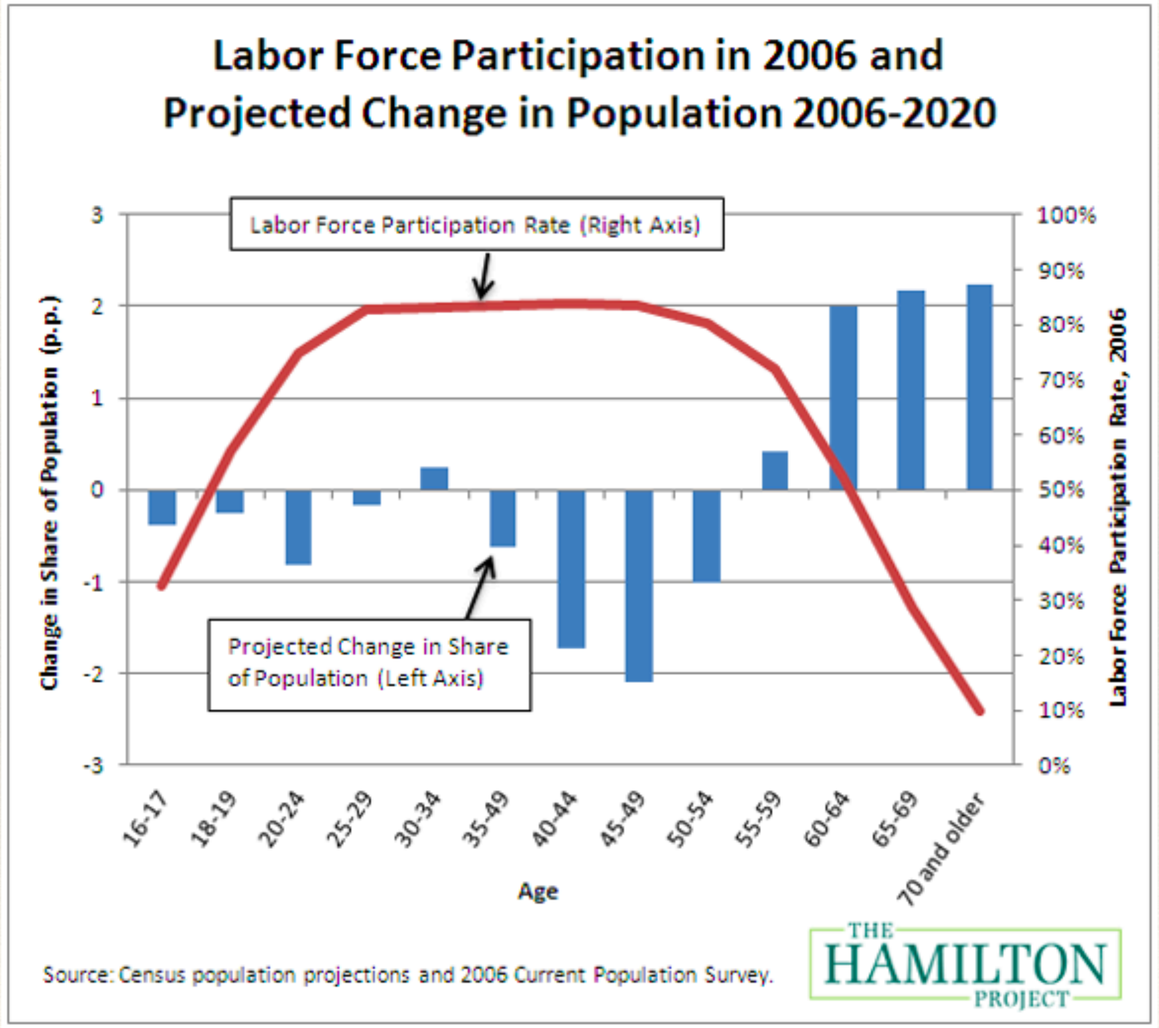

Neural Networks in Nature

- Human nervous system has neurons that specialize
- Specialized neurons work together on larger problems
- Connections between neurons become complex

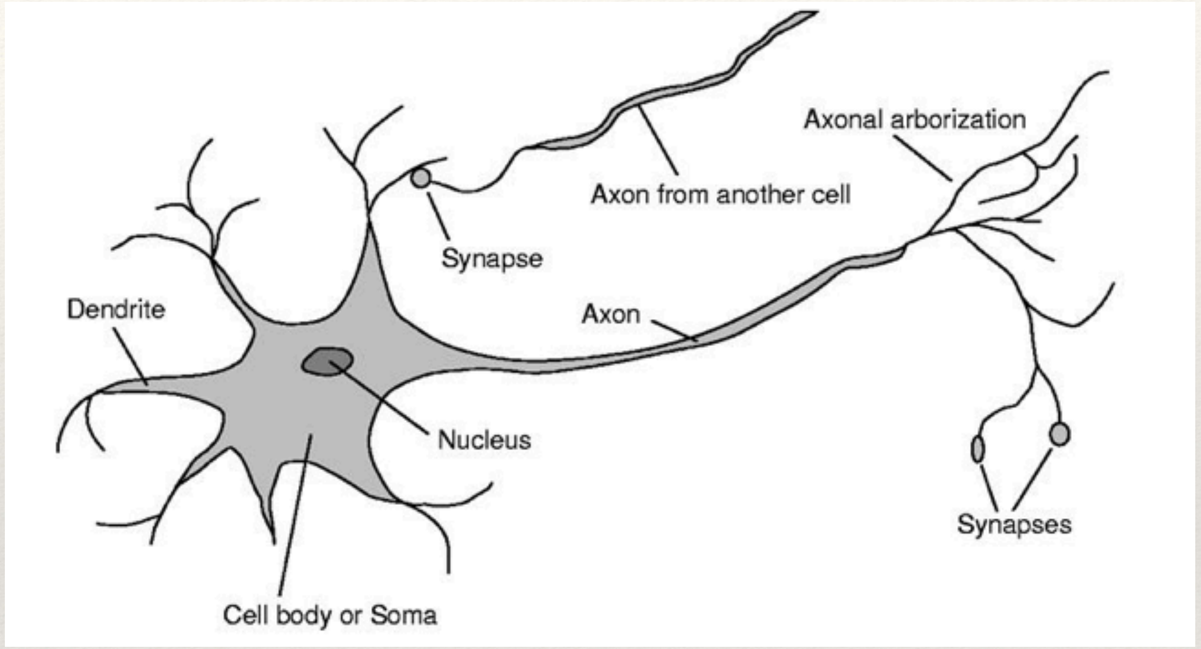

Neural Networks in TensorFlow - pt 1

- In TensorFlow, this is modeled / trained as…
  - Weights - sensitivity to each input (feature or neuron)
  - Biases - constant added to output

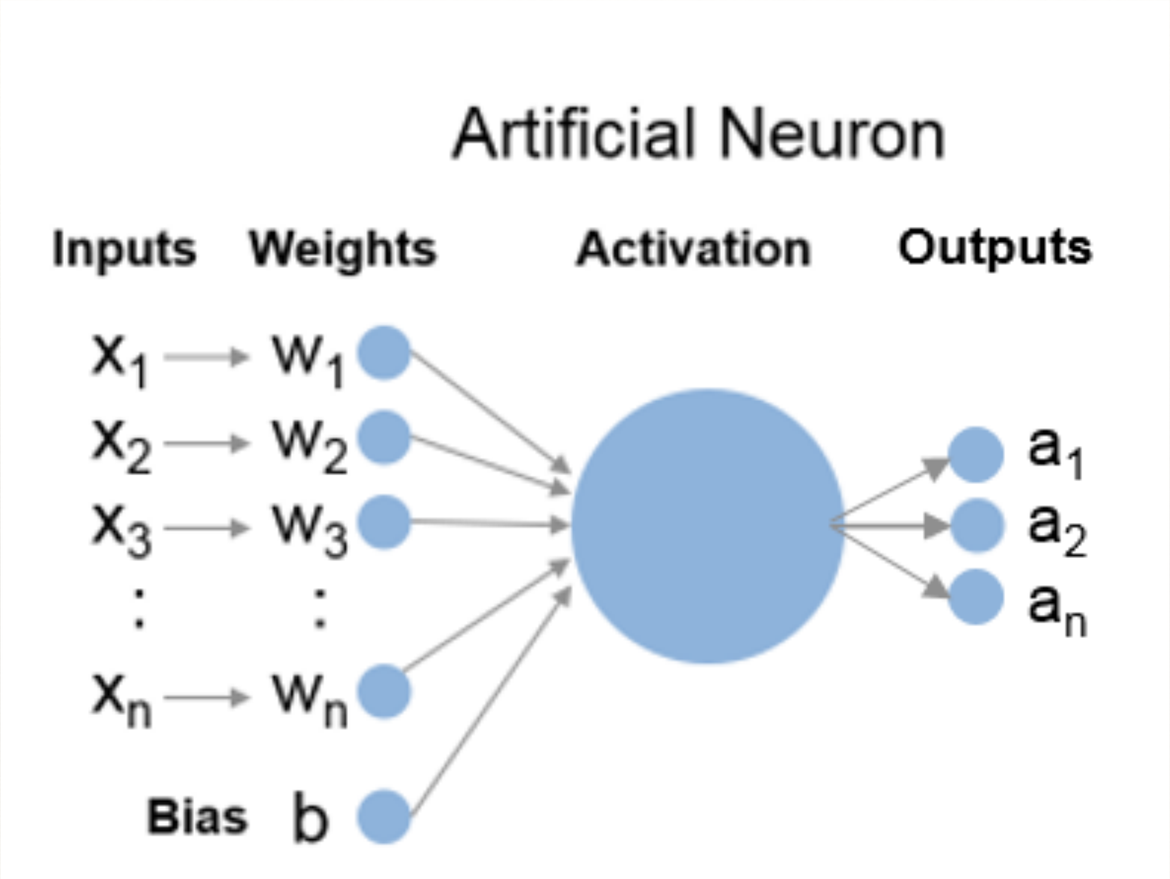

Neural Networks in TensorFlow - pt 2

- TensorFlow provides a DNNClassifier for neural networks
- Of course, this isn't in the 20%, either…

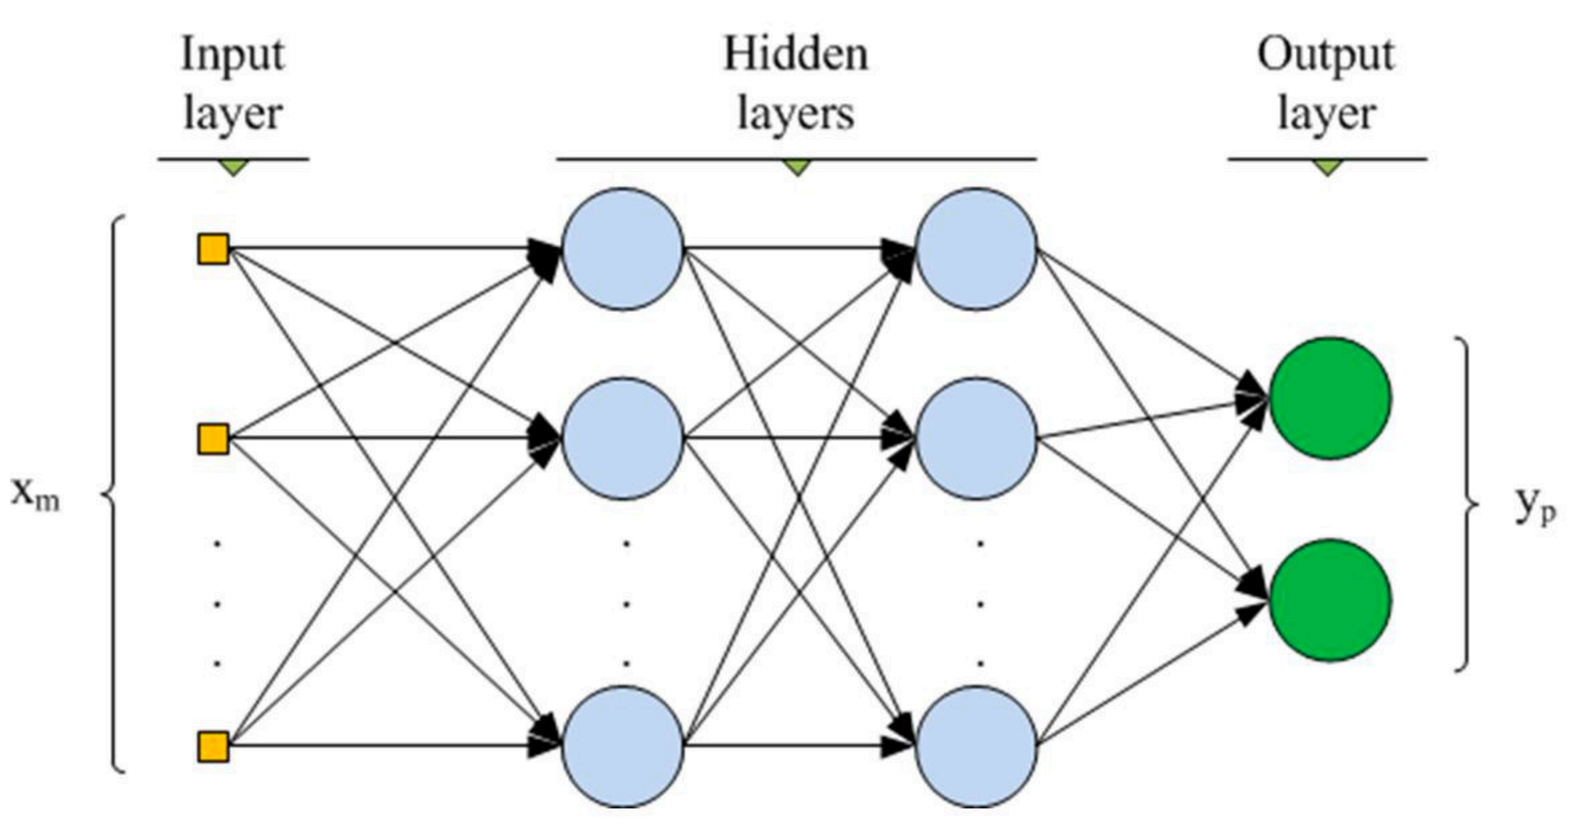

Hands On Lab

Demo 3

Invoking a model inline

Deploying Trained Models

- Before TF 1.8, Python users used Flask
- Since TF 1.8, TensorFlow Serving is available

```
FROM tensorflow/serving

COPY exports/1530401406 /models/model

EXPOSE 8500

EXPOSE 8501
```

Executing your Docker predictor

```
docker build -t predictor .


docker run
  -p 8500:8500 -p 8501:8501
  predictor
```

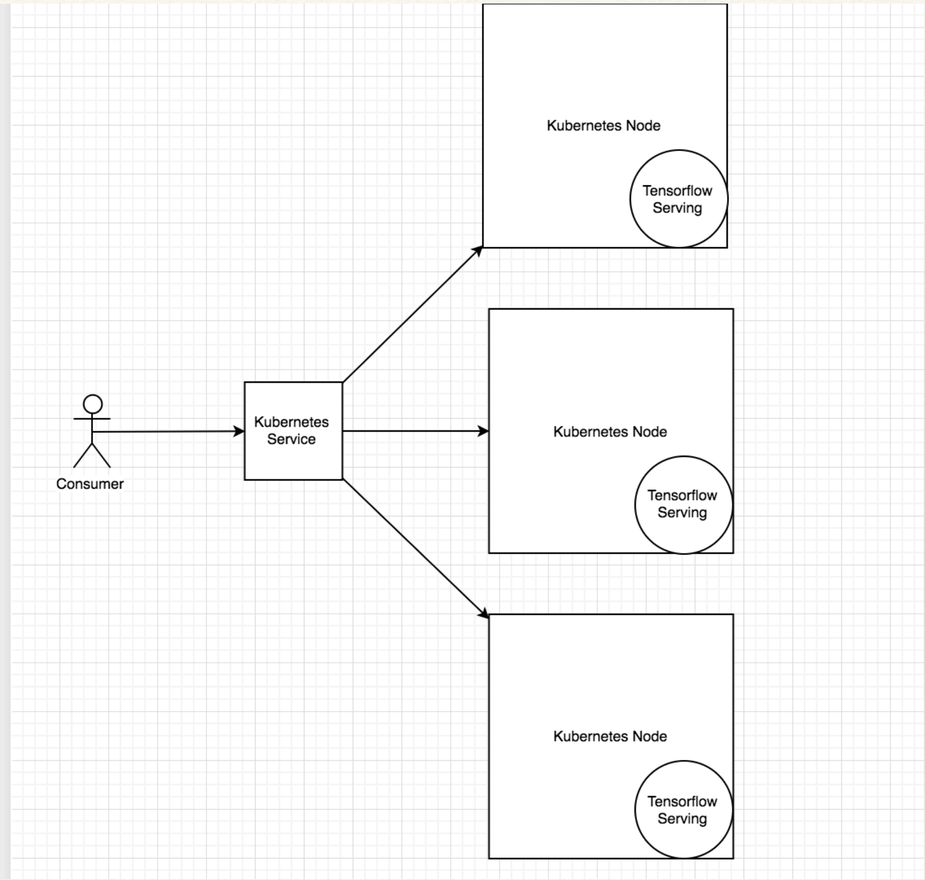

Kubernetes with TF Serving

Protocol Buffers

- Developed by Google
- Intended to be smaller and faster than XML
- Define message formats in .proto files and generate transmission code

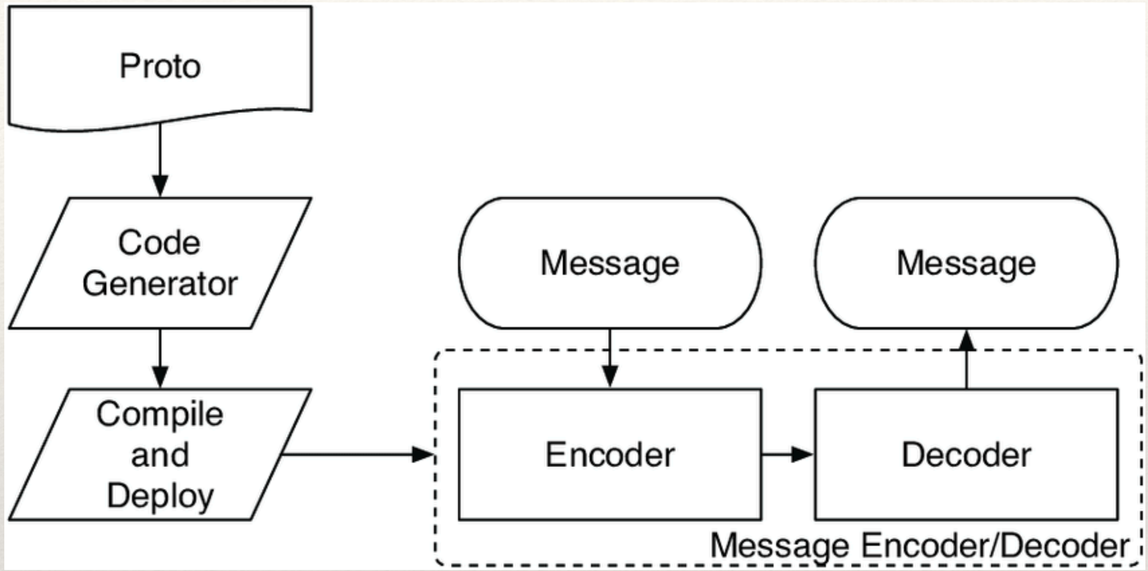

Protocol Buffers and TensorFlow

- ManagedChannelBuilder — communications engine
- Tensor(Shape)Proto — protobuf(s) for Tensors
- ModelSpec — protobuf for an entire model
- PredictResponse — value returned from execution

Better Yet - Let's Just Use REST

```
saved_model_cli.py show –dir modeler –all
```

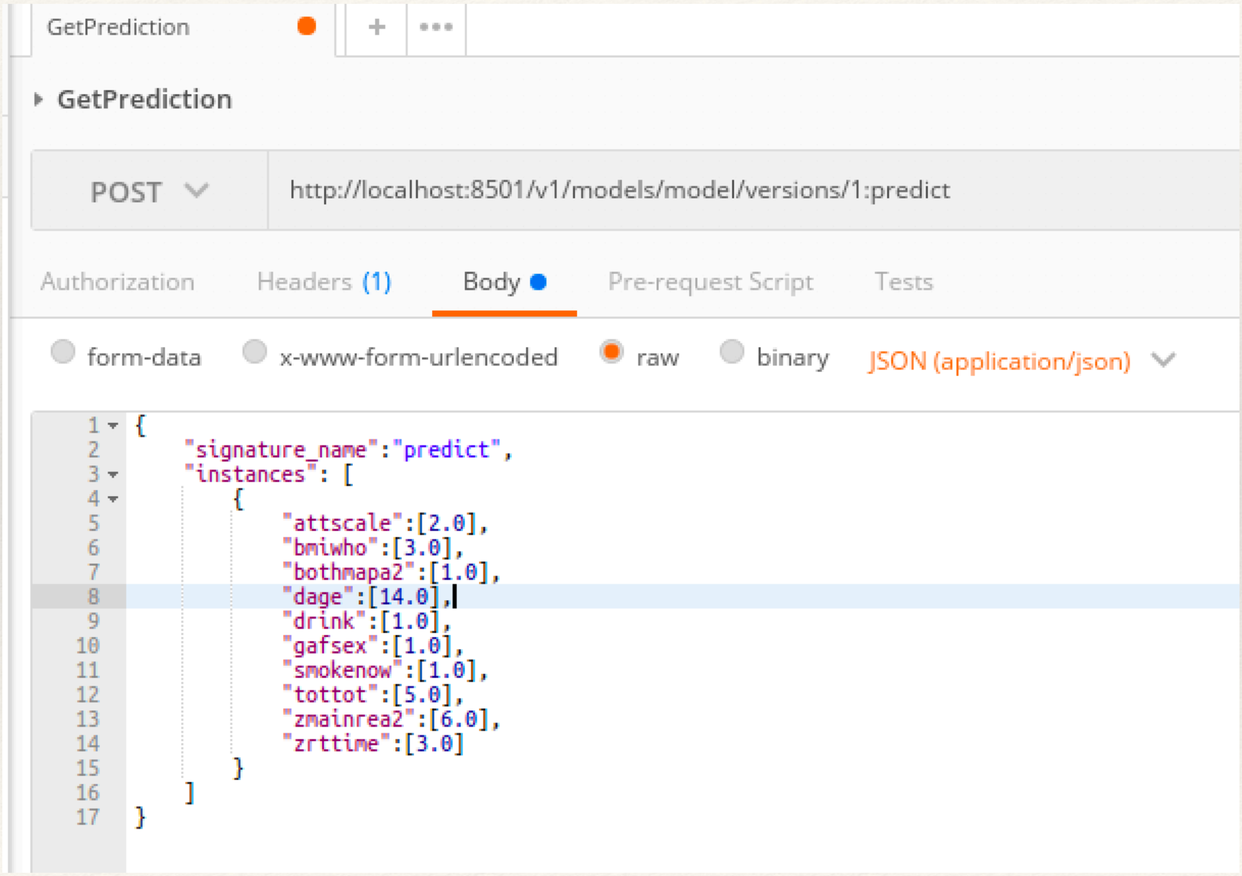

10 Input Features (X)

Hands On Lab

Demo 4

TensorFlow Serving via REST

What we learned…

- How machine learning works
- What TensorFlow is
- Where to get the required libraries
- How to obtain, train and invoke TF models

Where to from here?

- More of the missing 80%
- Advanced model creation
- Unsupervised learning
- Remember: the whole thing is Open Source!

Q&A

Contact me! English или по-русски!

–Derek Ferguson — derek.ferguson@mail.ru