CLOUD NATIVE COMPUTING FOUNDATION

Machine learning and AI in the datacenter and how it will affect you

Nick Chase
Head of Content, Mirantis
Editor in Chief, Open Cloud Digest
Author, *Machine Learning for Mere Mortals*

Who am I?

Nick Chase

- Long-time programmer
- Head of Technical and Marketing Content for Mirantis
- Editor in Chief of *Open Cloud Digest*
- Author of *Machine Learning for Mere Mortals*
- Absolutely **NOT** a math major
- nchase at mirantis dot com

Shhhhhh.....

Turns out, this is not as tough as you think.

Special Bonus

- Free early access to *Machine Learning for Mere Mortals* for live attendees

http://bit.ly/machine_learning_course

Agenda

- What is machine learning?
- Types of machine learning
- Problems facing the datacenter and how ML can solve them
- You, yes, YOU, can do this (and I'll prove it)
- Intelligent Delivery datacenter and how to start
- Q&A

Introduction

Just what is machine learning, anyway?

What is machine learning?

Machine learning is the study of enabling a computer to perform in situations it may not have seen before.

Where can we see ML today?

- Analytics
- Advertising
- Bots
- Pattern recognition
- NLP
- Search

Where can we see ML today?

- Analytics
- Advertising
- Bots
- Pattern recognition
- NLP
- Search
- Image recognition
- Video analysis
- Threat detection
- Medicine
- Self driving cars

Lots of hype

- Existing companies
  http://indeedhi.re/2nFRKY0

- Amazon.com (1252)
- JPMorgan Chase (733)
- Goldman Sachs (674)
- Microsoft (580)
- Apple (322)
- Google (217)
- Facebook (217)
- KPMG (148)
- Jobspring Partners (124)
- Capital One (118)
- NVIDIA (116)
- Strategic IT Staffing (111)
- Oracle (110)
- Lockheed Martin (106)
- IBM (101)

Lots of hype

- New companies (http://bit.ly/2E2qFoU)
- Newly accessible
  - Tensorflow
  - SciKit-Learn
  - etc.

Machine learning vs artificial intelligence

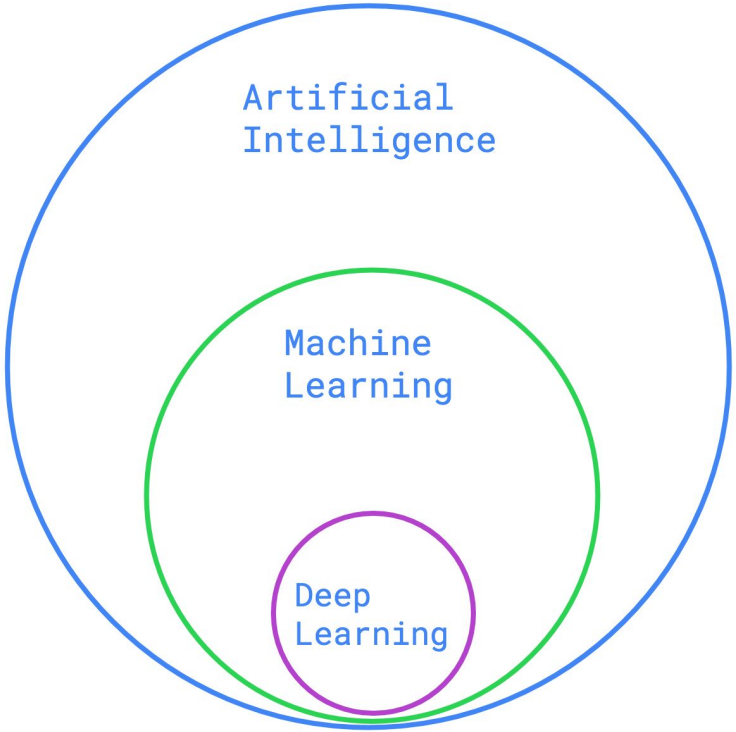

What does ML have to do with "cloud native"?

- Dynamic environments

What does ML have to do with "cloud native"?

- Dynamic environments
- High-velocity management

What does ML have to do with "cloud native"?

- Dynamic environments
- High-velocity management
- Continuous automation

What does ML have to do with "cloud native"?

- Dynamic environments
- High-velocity management
- Continuous automation
- Observability

What does ML have to do with "cloud native"?

- Dynamic environments
- High-velocity management
- Continuous automation
- Observability
- Efficiency

What can ML do for the datacenter?

Continuous Delivery infrastructure

What can ML do for the datacenter?

Continuous Delivery infrastructure

Intelligent Delivery infrastructure

Intelligent Delivery infrastructure

- Defined architecture

Intelligent Delivery infrastructure

- Defined architecture
- Flexible but controllable infrastructure

Intelligent Delivery infrastructure

- Defined architecture
- Flexible but controllable infrastructure
- Intelligent oversight

Intelligent Delivery infrastructure

- Defined architecture
- Flexible but controllable infrastructure
- Intelligent oversight
- Secure footing

Types of Machine Learning

Supervised learning

Training a computer to understand
data it hasn't seen by teaching it about data
that has already been labeled

Supervised learning: Examples

- Image recognition

Supervised learning: Examples

- Image recognition
- Speech recognition

Supervised learning: Examples

- Image recognition
- Speech recognition
- Medical outcome prediction

Supervised learning: Examples

- Image recognition
- Speech recognition
- Medical outcome prediction
- Spam filtering

Supervised learning: Examples

- Image recognition
- Speech recognition
- Medical outcome prediction
- Spam filtering
- Stock market/crypto pricing

Unsupervised learning

Letting the computer discover patterns and relationships in unlabeled data.

Unsupervised learning: example

- Clustering

Unsupervised learning: example

- Clustering
  - Customers into categories

Unsupervised learning: example

- Clustering
  - Customers into categories
  - Chemical compounds

Unsupervised learning: example

- Clustering
  - Customers into categories
  - Chemical compounds
- Anomaly detection

Unsupervised learning: example

- Clustering
  - Customers into categories
  - Chemical compounds
- Anomaly detection
- Signal separation

Reinforcement learning

A combination of supervised and unsupervised learning, where the computer is given an objective and attempts to reach that objective, discovering tactics as it goes along.

Reinforcement learning: example

- Just about anything humans do

Reinforcement learning: example

- Just about anything humans do
- Self-driving cars

Reinforcement learning: example

- Just about anything humans do
- Self-driving cars
- Strategy development

Reinforcement learning: example

- Just about anything humans do
- Self-driving cars
- Strategy development
  - Games

Reinforcement learning: example

- Just about anything humans do
- Self-driving cars
- Strategy development
  - Games
  - Scheduling

Problems Facing the Datacenter

Configuration: potential issues

- Self-service

Configuration: potential issues

- Self-service
- Microservices architectures

Configuration: potential issues

- Self-service
- Microservices architectures
- Dynamic environments and drift issues

Configuration: how ML/AI can help

- App configuration based on past history

Configuration: how ML/AI can help

- App configuration based on past history
- Reinforcement learning to retroactively improve configuration

Performance optimization: Potential issues

- Lots of factors in play

Performance optimization: Potential issues

- Lots of factors
- Dynamic environments create a changing landscape

Performance optimization: how ML/AI can help

- Analyze multiple factors

Performance optimization: how ML/AI can help

- Analyze multiple factors
- Predict future load and move or scale workloads based on patterns

Performance optimization: how ML/AI can help

- Analyze multiple factors
- Predict future load and move or scale workloads based on patterns
- Move workloads to closer geographical resources based on patterns

Cost optimization: potential issues

- Complex environments

Cost optimization: potential issues

- Complex environments
- Changing costs

Cost optimization: potential issues

- Complex environments
- Changing costs
- Soaring storage requirements

Cost optimization: how AI/ML can help

- Keep track of multiple environments and costs

Cost optimization: how AI/ML can help

- Keep track of multiple environments and costs
- Adapt to changing conditions

Cost optimization: how AI/ML can help

- Keep track of multiple environments and costs
- Adapt to changing conditions
- Optimize what data is kept

Fault detection: potential issues

- Hardware failure

Fault detection: potential issues

- Hardware failure
- Software configuration drift

Fault detection: potential issues

- Hardware failure
- Software configuration drift
- Data corruption

Fault detection: potential issues

- Hardware failure
- Software configuration drift
- Data corruption

No, I am **not** making that up: http://bit.ly/cosmic-corruption

Fault detection: potential issues

- Hardware failure
- Software configuration drift
- Data corruption
- Resource overrun

Fault detection: how AI/ML can help

- Autoscaling

Fault detection: how AI/ML can help

- Autoscaling
- Pattern recognition

Fault detection: how AI/ML can help

- Autoscaling
- Pattern recognition
- Anomaly detection

Security: potential issues

- DDOS

Security: potential issues

- Distributed Denial Of Service attacks
- Advanced Persistent Threats

Security: potential issues

- DDOS
- APT
- Inside jobs

Security: how ML/AI can help

- Anomaly detection to find APTs and inside jobs

Security: how ML/AI can help

- Anomaly detection to find APTs and inside jobs
- Pattern recognition to predict DDOS attacks before they're in full swing

Try It Yourself

Once upon a time...

Once upon a time...

- Self-configuring

Once upon a time...

- Self-configuring
- Self-optimizing

Once upon a time...

- Self-configuring
- Self-optimizing
- Self-healing

Once upon a time...

- Self-configuring
- Self-optimizing
- Self-healing
- Self-protecting

Once upon a time...

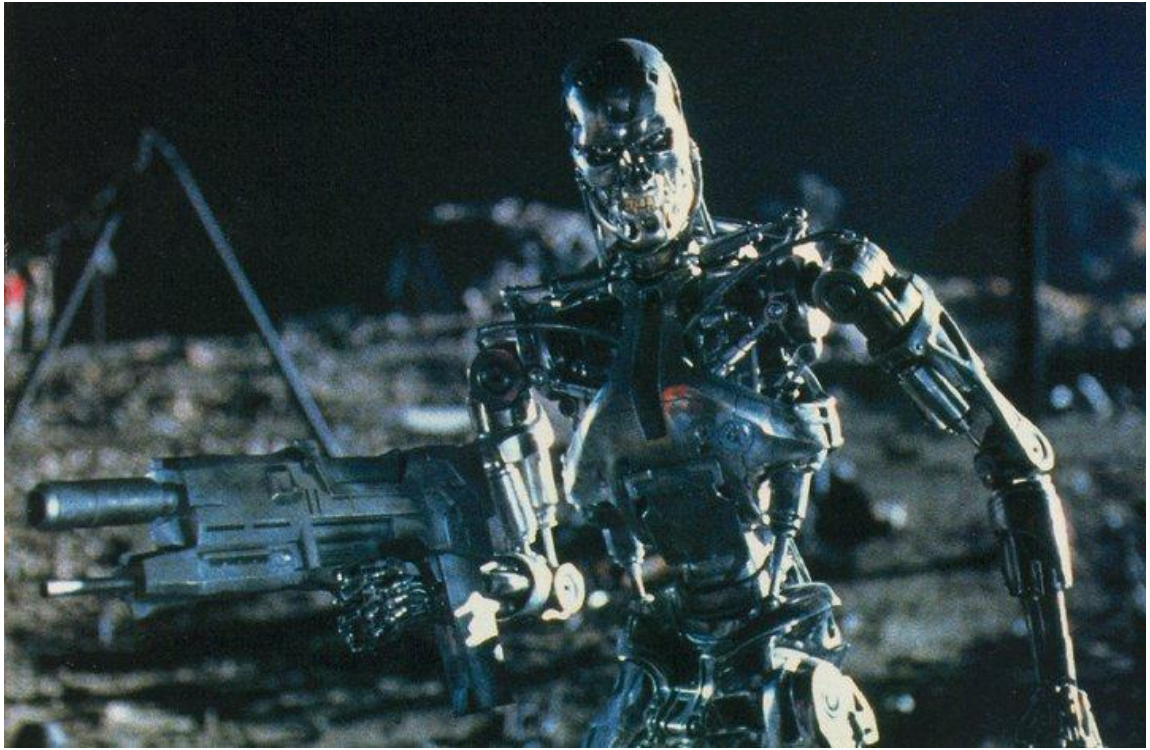

What are we trying to do?

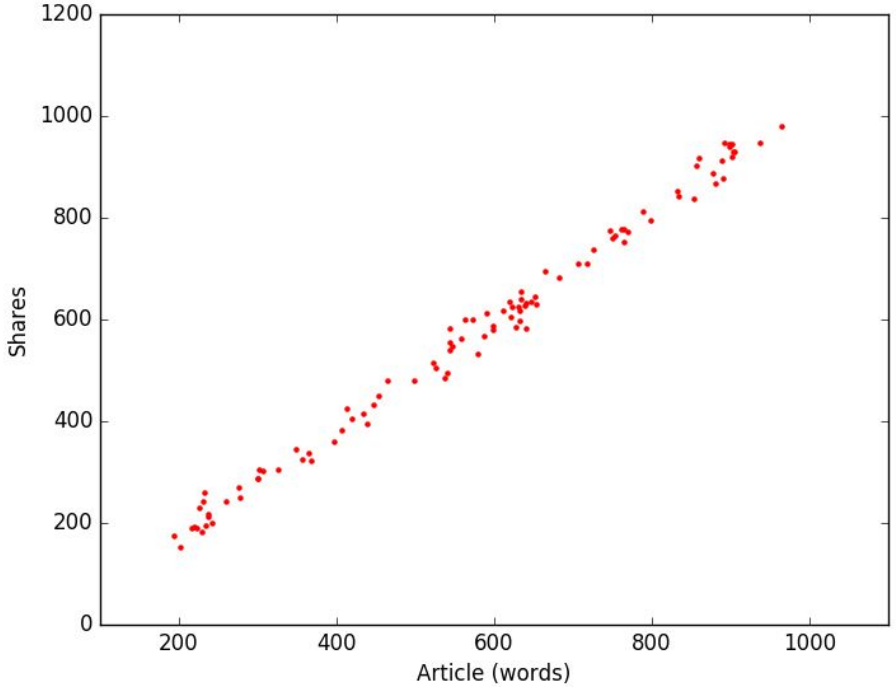

What are we trying to do?

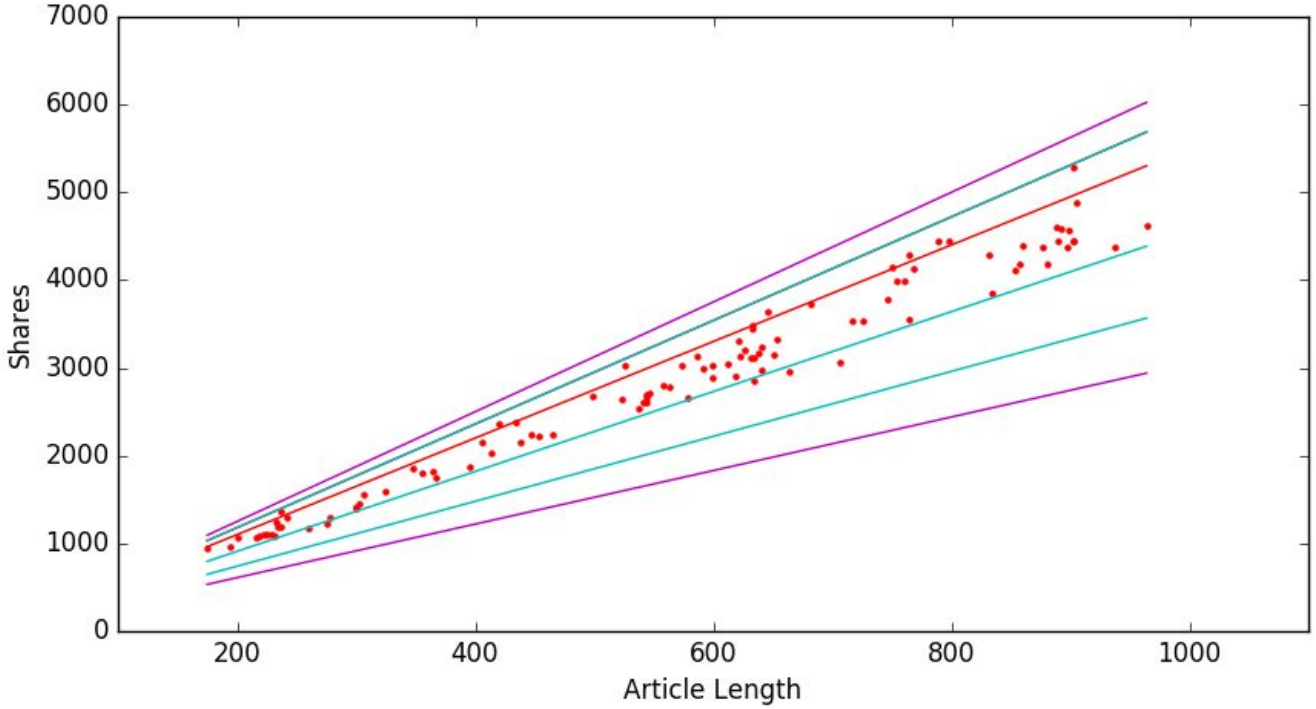

Defining the error

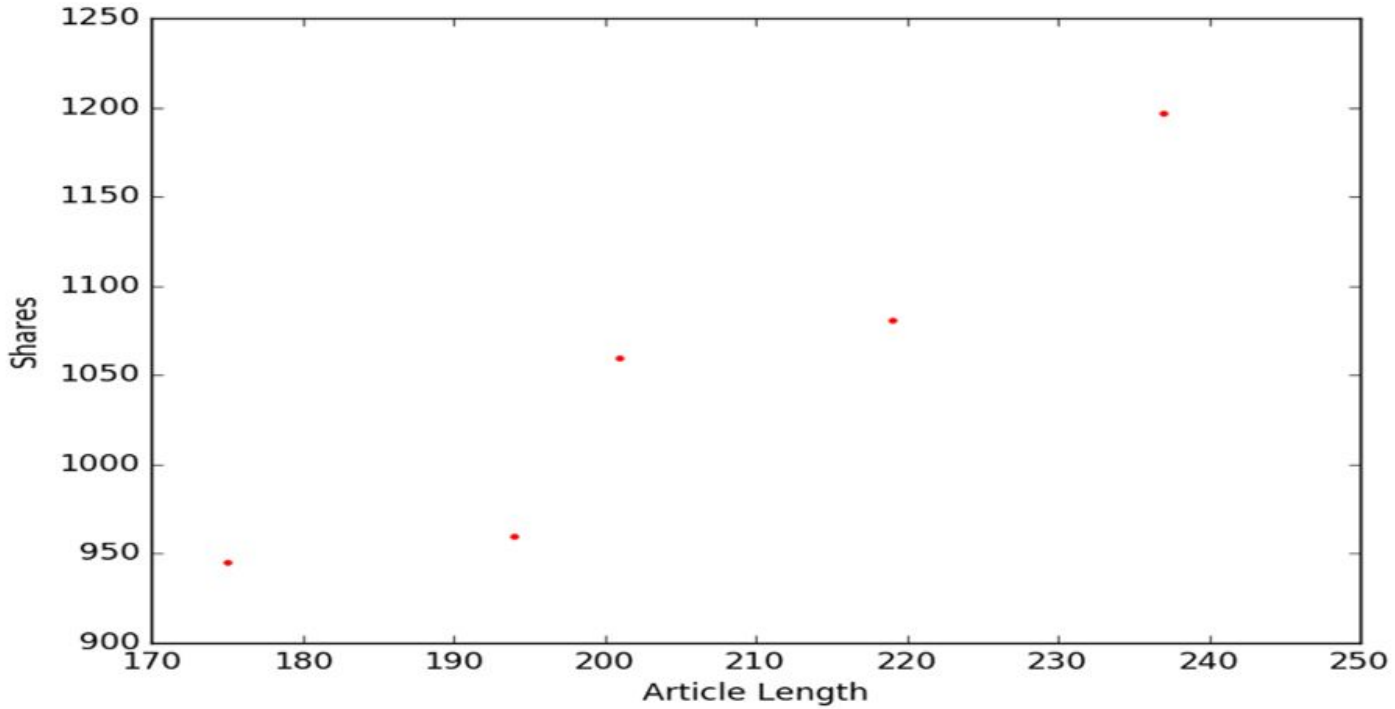

Defining the error

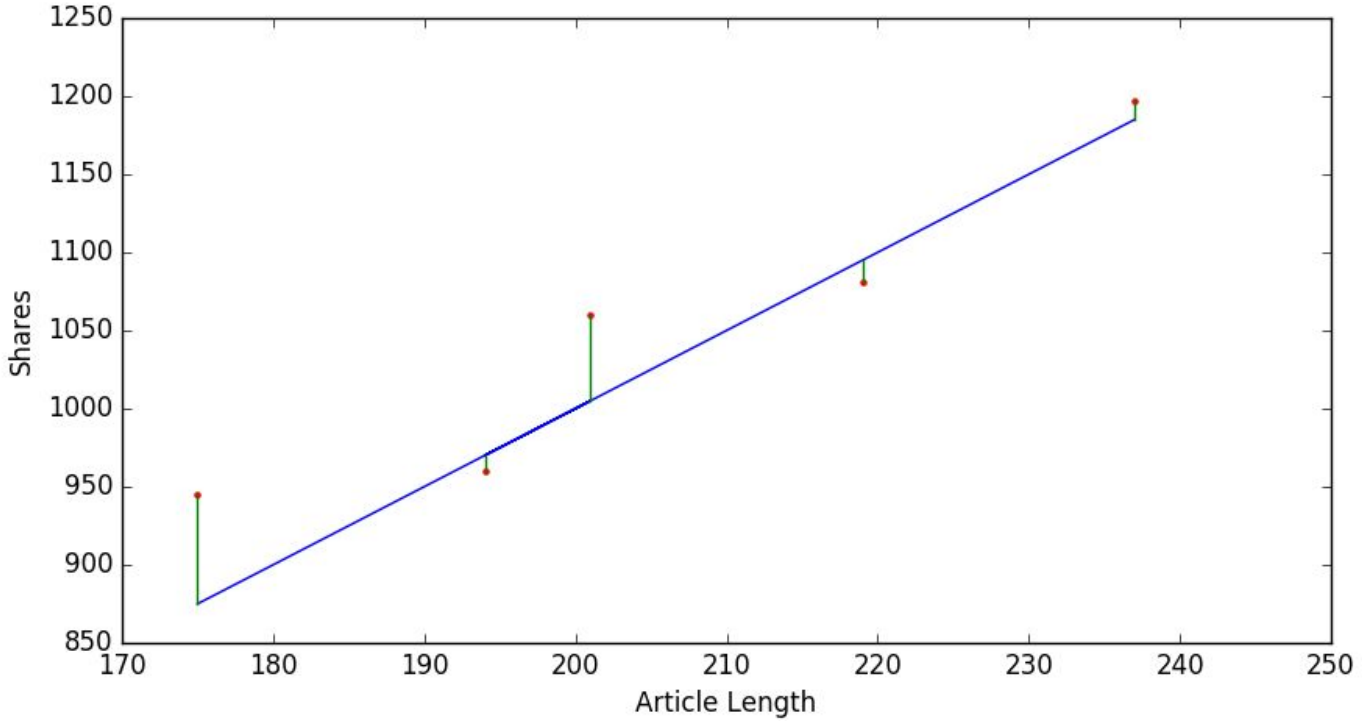

Defining the error

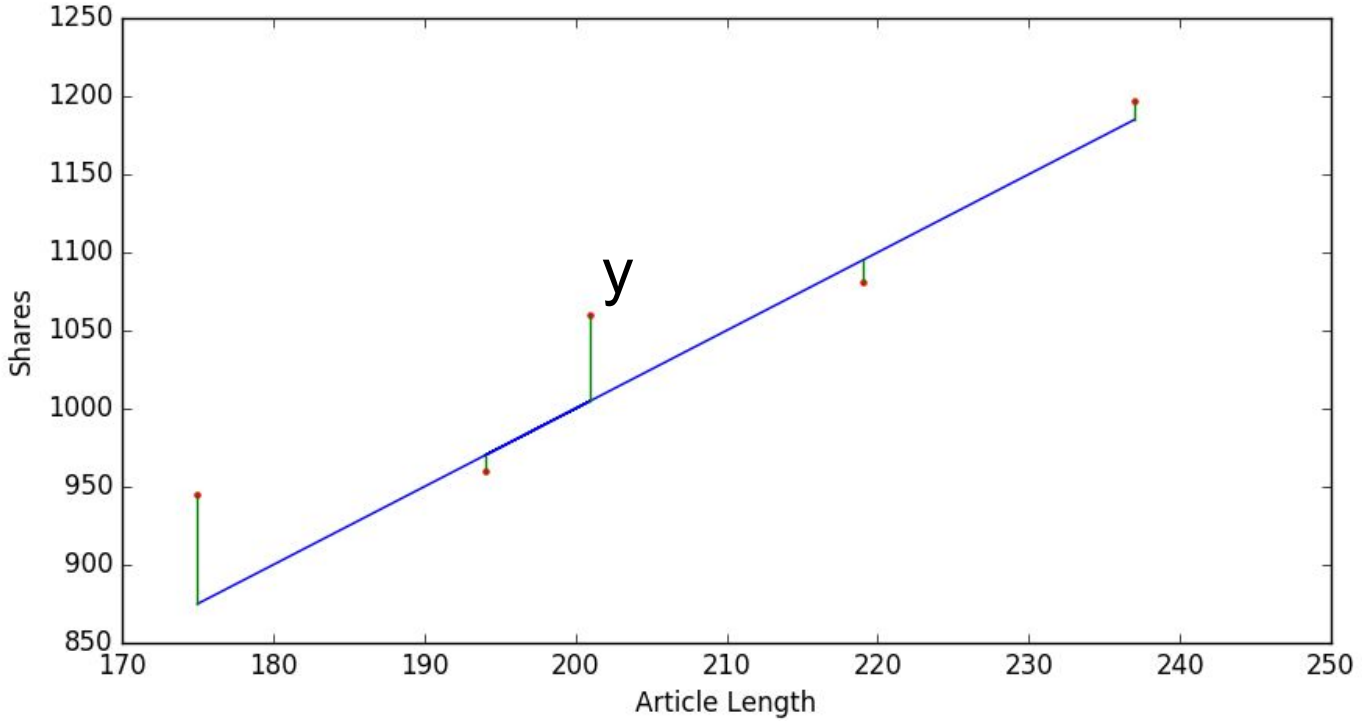

Defining the error

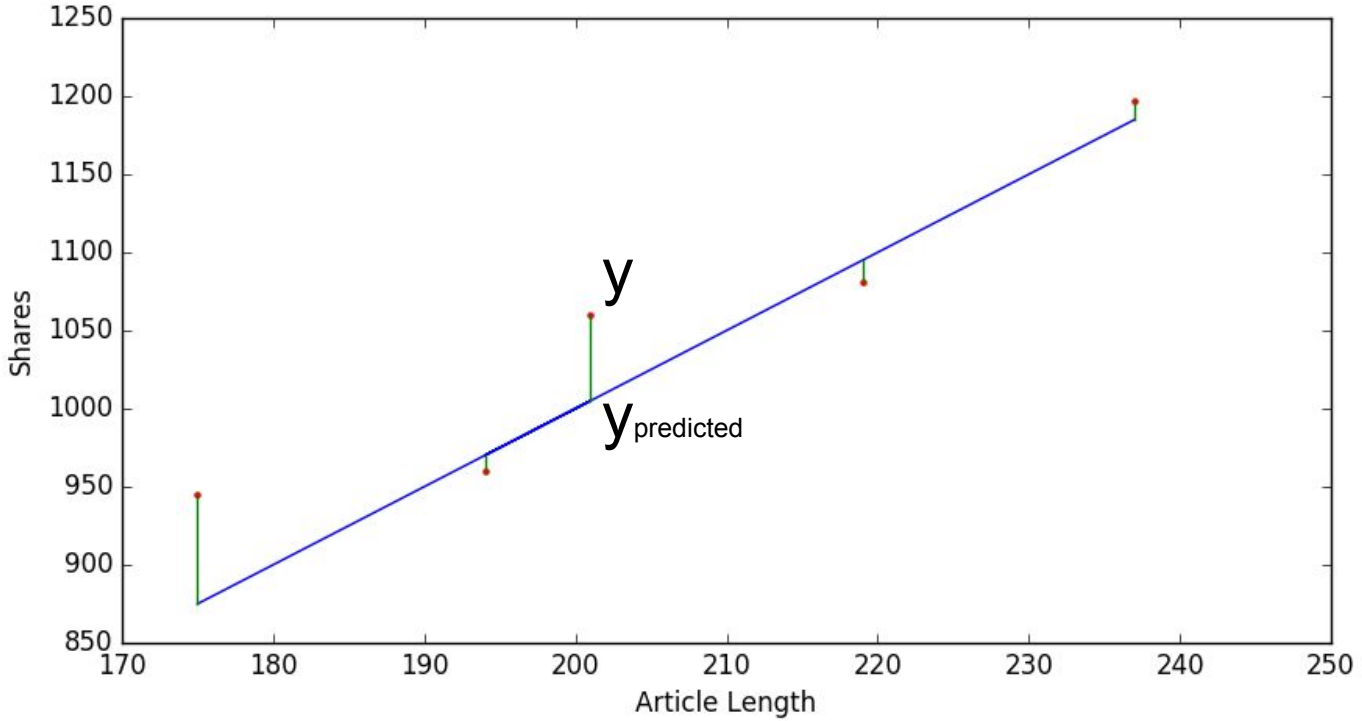

Defining the error

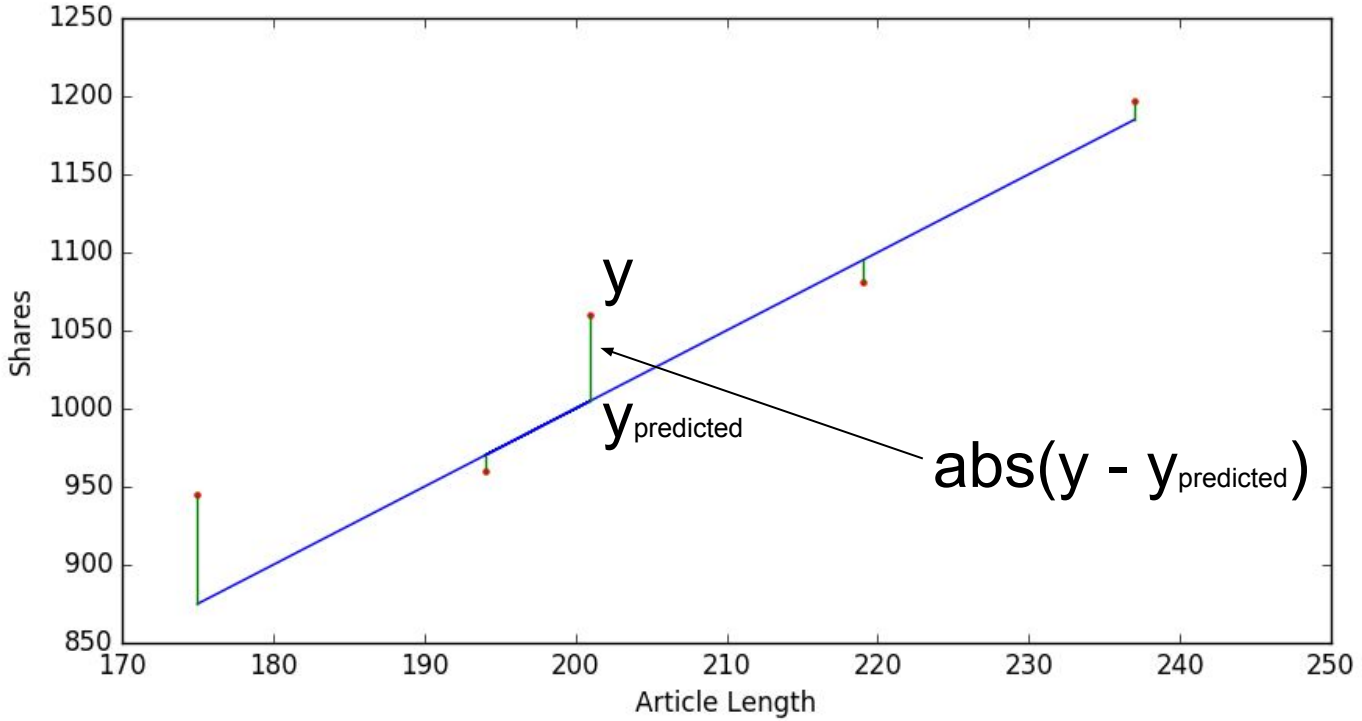

Defining the error

- 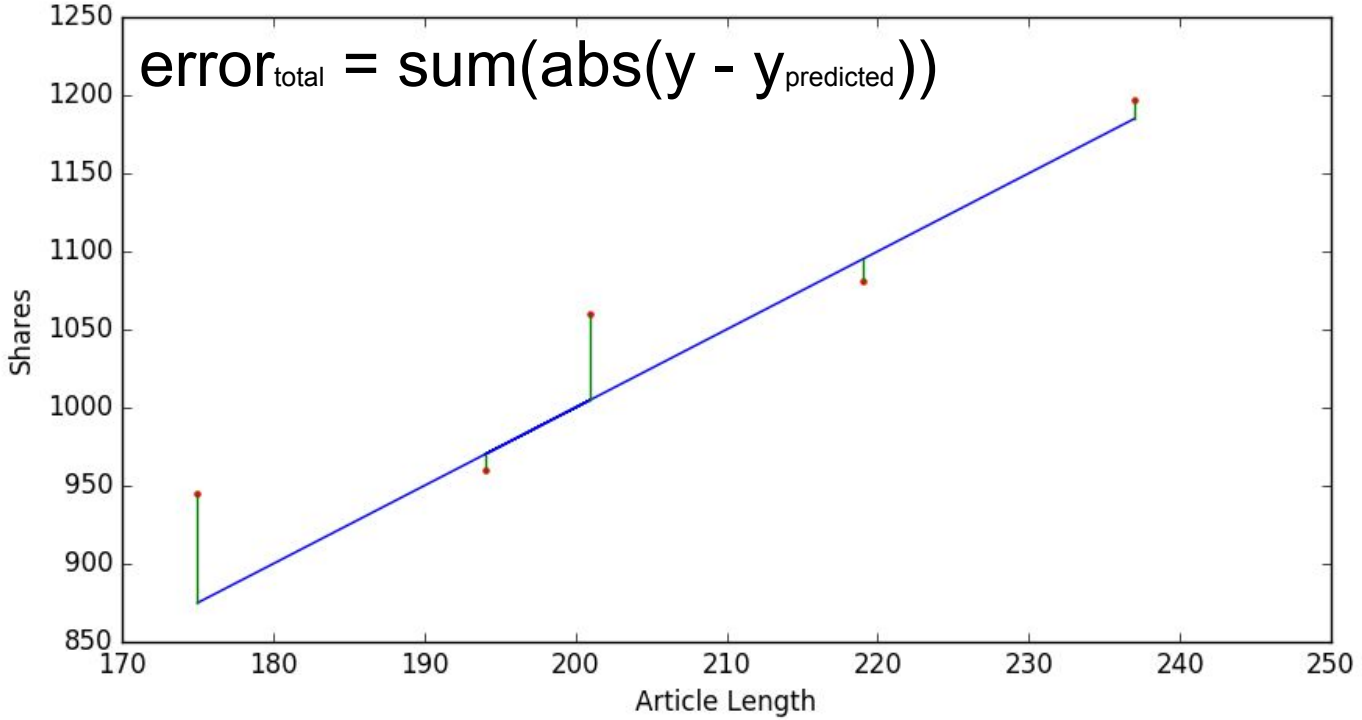

Defining the error

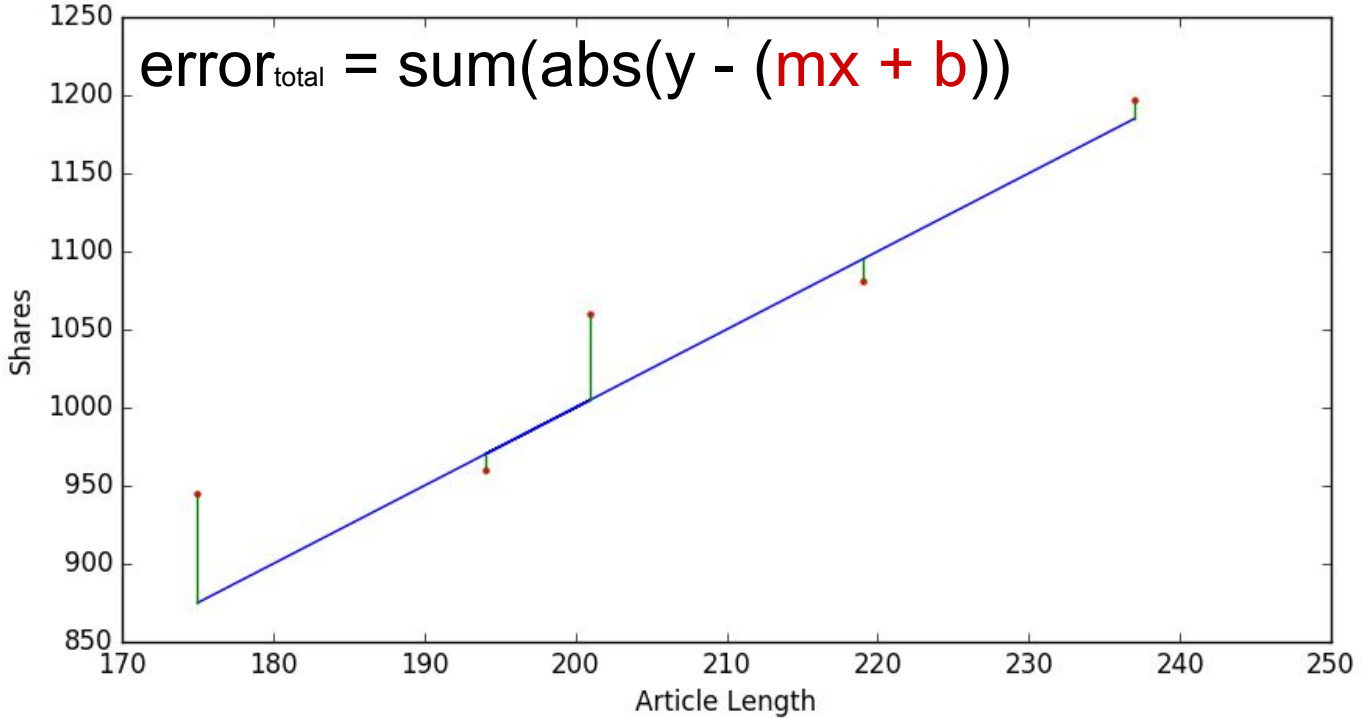

Hot and cold

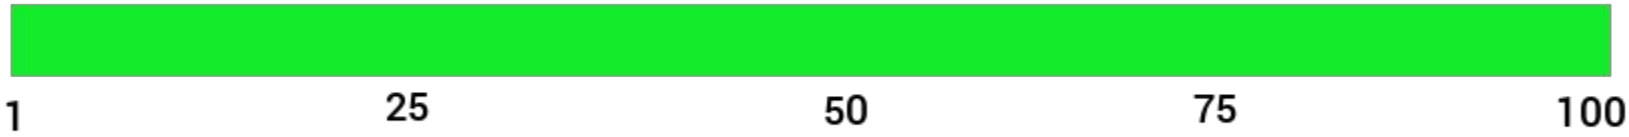

Hot and cold

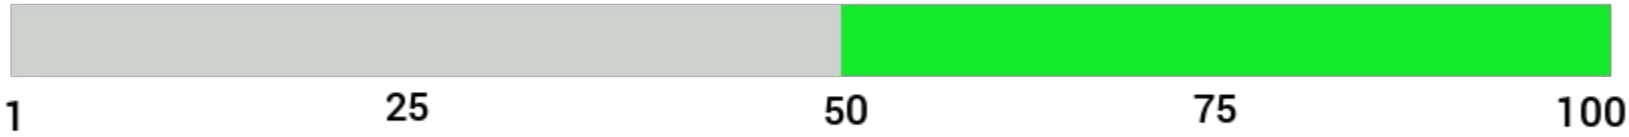

Hot and cold

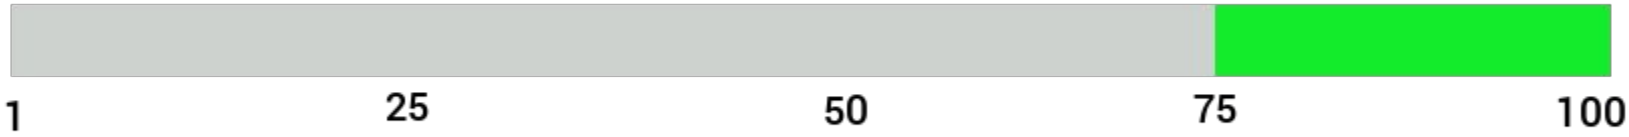

Hot and cold

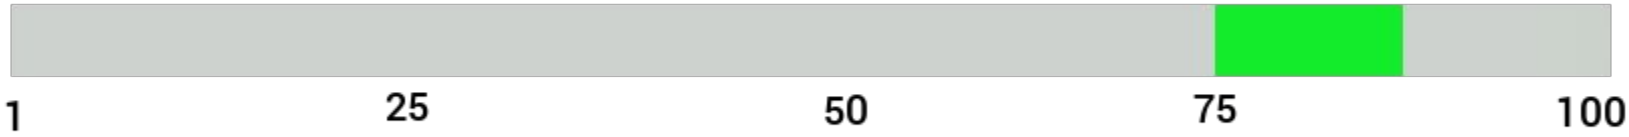

Hot and cold

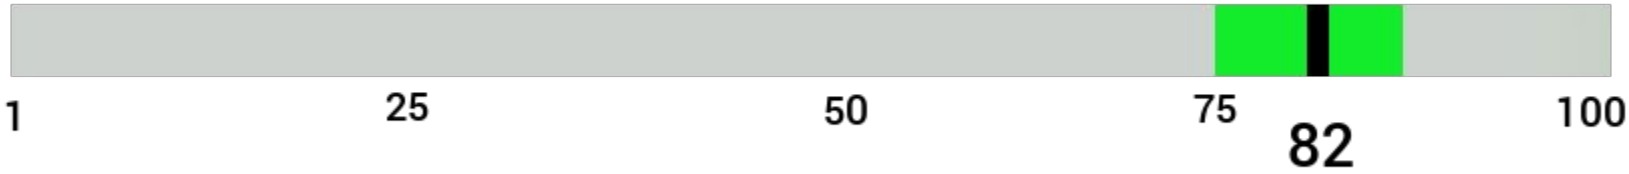

Hot and cold

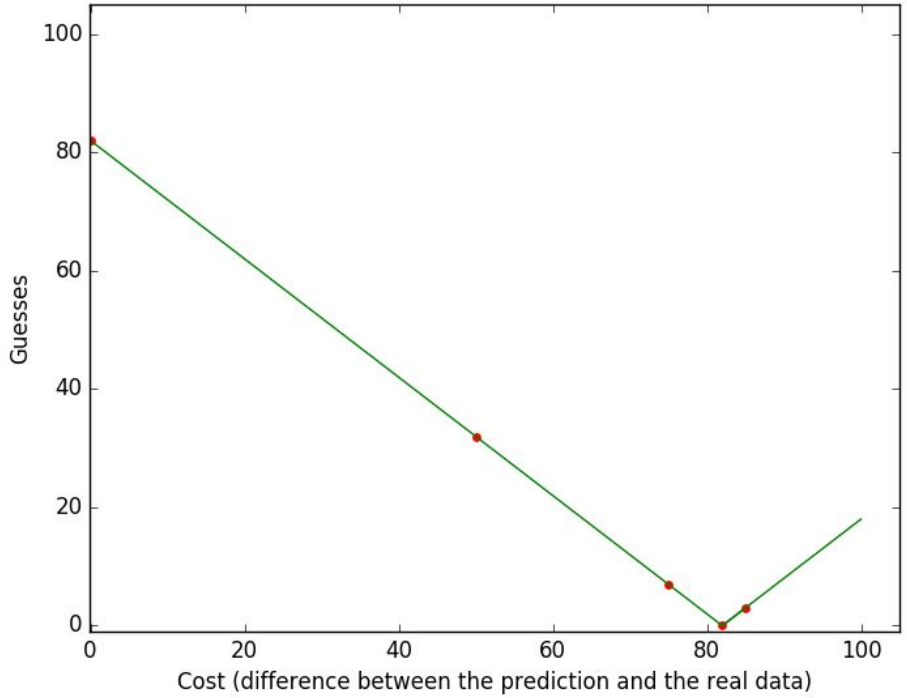

Hot and cold

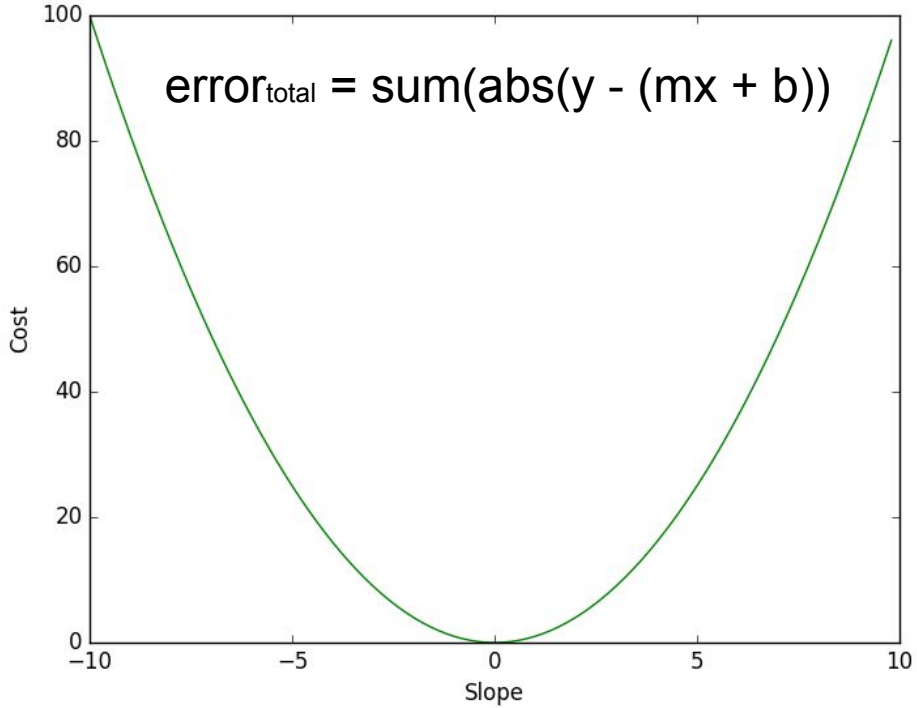

Gradient descent

An iterative algorithm to find the minimum of the cost/error function and that parameters that achieve that minimum.

Putting it together

- Data
- Python3
- Tensorflow
- Matplotlib (optional)
- Ubuntu 16.04

What we'll do

- Read the data
- Create the model
- Define the error
- Define the optimizer
- Repeatedly run the optimizer until we get as close as we can

DEMO

The actual answer

$$y = .37x + 12$$

$$m = .37$$

$$b = 12$$

DEMO

Where do we go from here?

More practically

- Don't expect to go from B&W to VR in one step
  - Smaller steps in between

More practically

- Don't expect to go from B&W to VR in one step
  - Smaller steps in between
- Start with analytics

More practically

- Don't expect to go from B&W to VR in one step
  - Smaller steps in between
- Start with analytics
- Ultimately, know what you're trying to achieve

Q&A

Thank You