

Deep Learning 101

Tristan Attwood, ATPCO

atpco

Things to Know



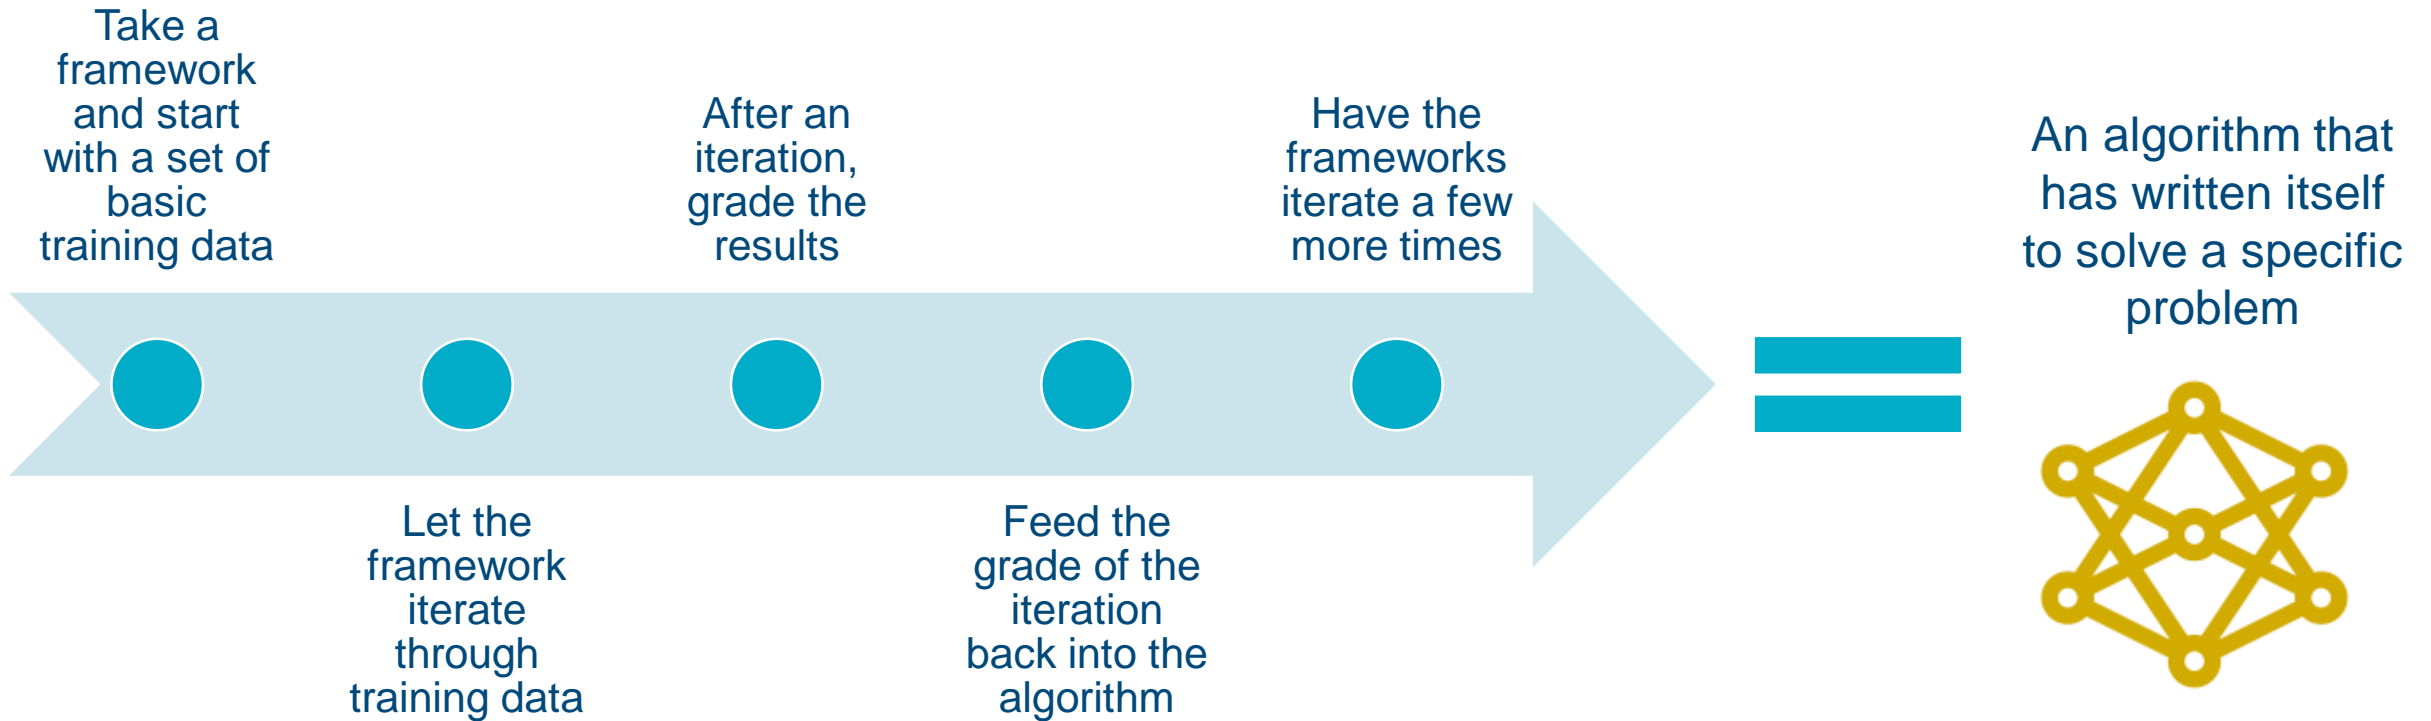
Terminology

- Deep Learning
- Machine Learning
- Artificial Intelligence (AI)

Training: the act of teaching the program to do what you want; uses a vast amount of data and processing power

Deep Learning Algorithm: a specific program created for a specific task

How is it made?



What can it do?

Deep learning lets a computer make qualitative decisions based on complex or ambiguous data.



Voice Recognition
(Siri, Google Assistant,
Cortana, Alexa)



Image Recognition
(Google, Facebook)



Autonomous
Navigation
(Cars, Drones)



Beating humans at our
own games
(Chess, Go)

DeepMind

A subsidiary of
Google AI



May 2017

Defeated world's top ranked
Go player, 3-0

December 5th, 2017

Released paper: *Mastering
Chess and Shogi by Self-Play
with a General reinforcement
Learning Algorithm*

Alpha Zero Program

After given the basic rules of
chess, within 24 hours had
taught itself to be the best
player on the planet.



Using it in the Industry

Scenario: I want to book a non-rev between two cities on your carrier or vacation.

Process: You have to figure out what are the chances I get the seat.

Step 1.

Think about all the factors you would have to consider...

Step 2.


Go through all “historical” data that you have regarding your knowledge of non-rev tickets

Step 3.

After decades of experience, you know which ticket is best for this specific scenario.

Now you try!

- What is this a picture of?
- How would you teach a computer to know that?



But there are downsides...



Requires an **ENORMOUS** and clean dataset to train

- Any errors in the dataset will be taken as correct and replicated

Training is **VERY** computationally expensive


- Google is building proprietary custom chips to train their deep learning algorithms

The result is a **black box**

- In a normal program you can find a bug, go in to the code, and fix it
- That is currently not possible with a deep learning algorithm

This leads to...

Strange Errors!



From MIT Tech Review, "Smart" Software Can Be Tricked into Seeing What Isn't There (December 24, 2014)
[Link](#)


One pixel attack for fooling deep neural networks, Jiawei Su, Danilo Vasconcellos Vargas, Kouichi Sakurai, Kyushu University Japan (October 24, 2017)
[Link](#)



Strange Errors!

Think of these errors as the robot version
of optical illusions...

...humans have errors in the way we
process images too.





So, what does this mean for me?

Deep learning algorithms have been proposed for multiple applications within the air travel industry (e.g, security screenings, managing overbooking).

Here are things to keep in mind:

When they work, it's incredibly well. And when they don't...accept it or start over.

Figuring out why they don't work is an open area of academic research. But if you have the answer, you make all the money in Big Tech.

Think of Deep Learning as hiring an incredibly productive employee...

An employee that never sleeps, needs a break, and works at superhuman speed.

However, that incredibly productive employee can't be coached or given guidance.

If it becomes convinced that having an odd number of magnets on the break room fridge means they need to restart the email server.
Do you accept this? Or do you fire them and hire someone else?

atpcc@