CS294:
Scheduling Deep Learning Workloads

Joey Gonzalez and Ion Stoica

April 24, 2019

What is the problem?

You have a bunch of training jobs sharing the cluster

You want to optimze for a scheduling metric:

- Makespan / average completion time
- Throughput
- Fairness
- A combination of the above

Why care about this?

Training a single model very expensive

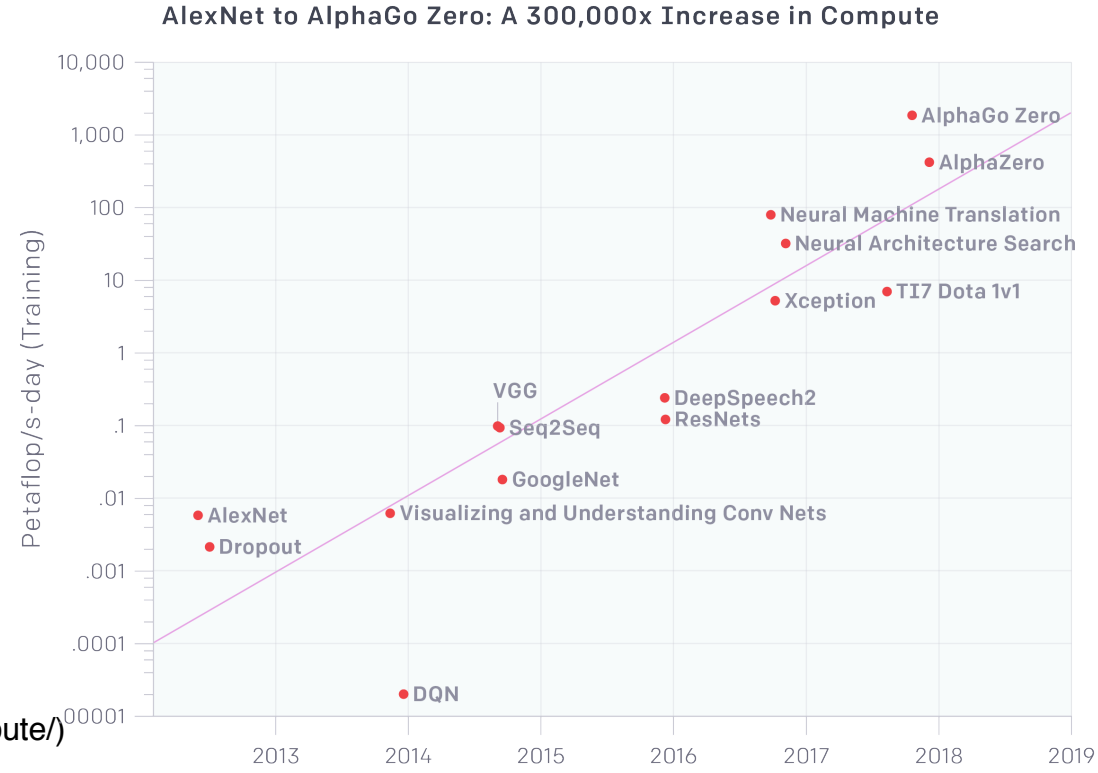

AI and Compute (https://blog.openai.com/ai-and-compute/)

Why care about this?

Training a single model very expensive

Models brittle so might want to train multiple times

And, we want to do hyperparameter search!

- Network architecture design
- AutoML
- …

Minimize job average completion time

For one processor: shortest job first is optimal

For parallel systems is an approximation but need to bin the jobs carefully

Assumption: you need to know completion time

- Otherwise need to predict it
- Complex because it can depend on parallelism, time when you get a particular resource, etc

Maximize throughput

This is what cluster operators want!

- Justify their investment

It can be gamed!

Need to be careful about how we pack jobs

- Can lead to starvation

Fairness

A trip in the history lane…

- see Dominant Resource Fairness (DRF) presentation