Proprietary + Confidential

Google

Google Cloud HPC & AI Capabilities

Overview

Leonid Kuligin
Karan Bhatia

Dec 6 2019

Proprietary + Confidential

Doubling LHC CMS compute capacity.

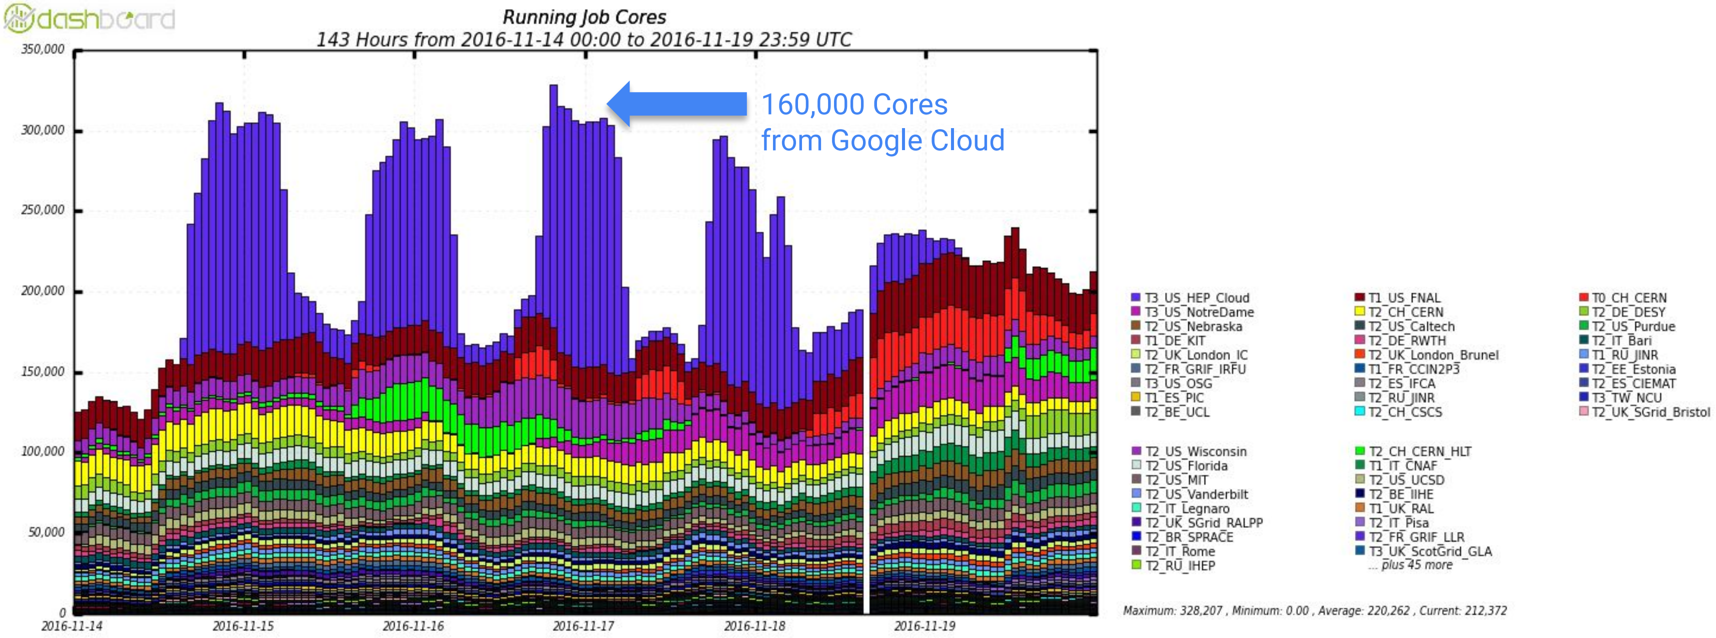

Google Cloud

Combating Traffic Congestion using Massive HPC Analytics in the Google Cloud Platform

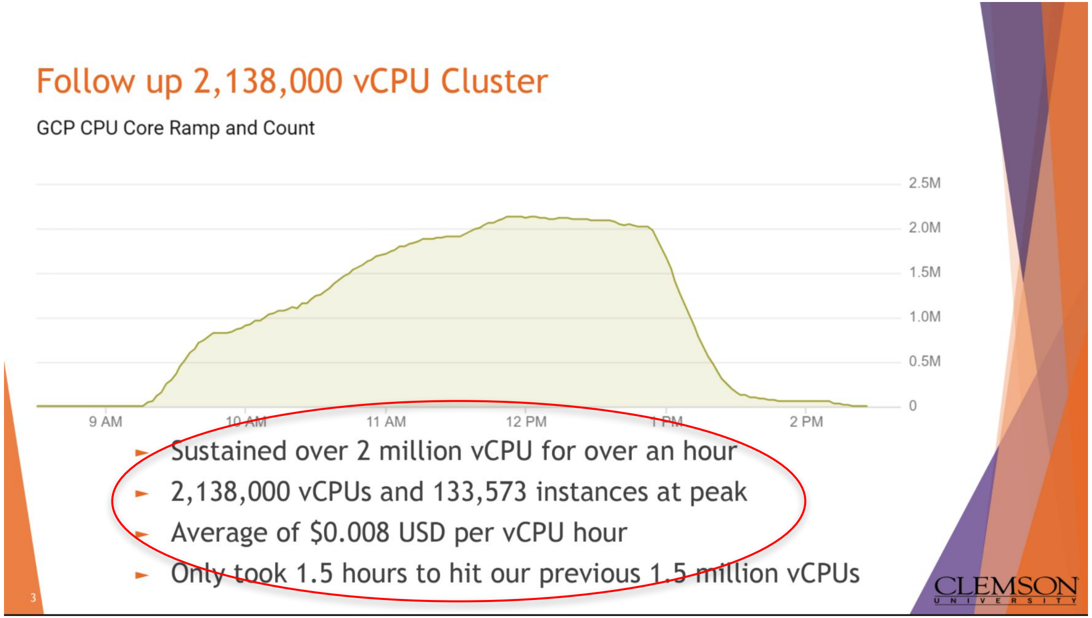

> 1Tbps throughput

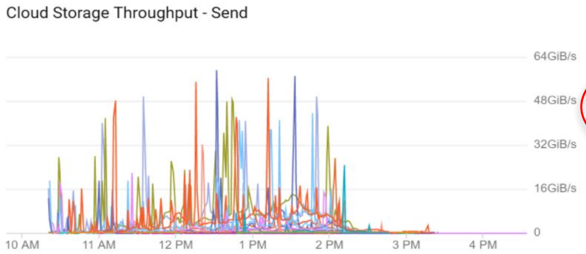

- 128 GiB/s peak in Cloud Storage
- 59 GiB/s peak in one bucket

Data Intensive Computing Ecosystems (DICE)
School of Computing, Clemson University

Proprietary + Confidential

Rediscovering Higgs boson on GCP

- https://github.com/mmm/higgs-tutorial
- Kubecon Barcelona 2019 keynote https://www.youtube.com/watch?v=CTfp2woVEkA

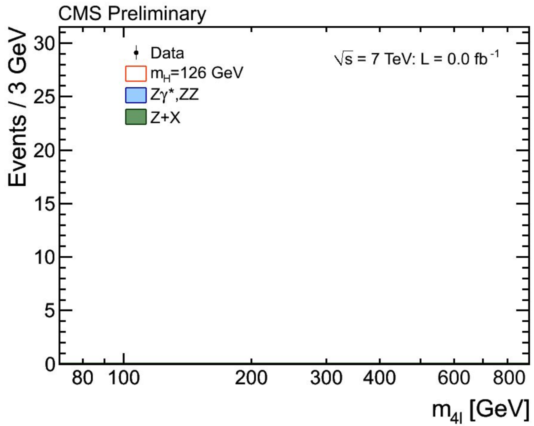

Google Cloud

Proprietary + Confidential

Google Cloud Background

Proprietary + Confidential

HPC Infrastructure

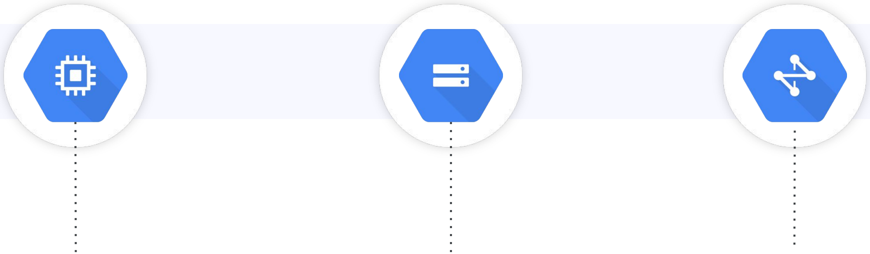

Compute

Storage

Networking

Google Cloud

Google confidential & proprietary

TPUs are ASICs focused on Machine Learning

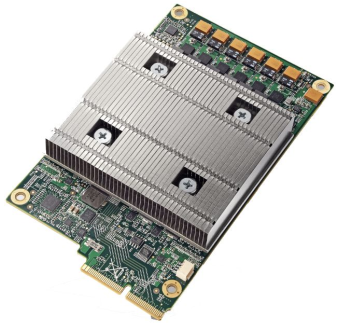

TPU v1
(2015)
92 teraops
First Generation

TPU v2
(2017)
180 teraflops
Available via Google Cloud

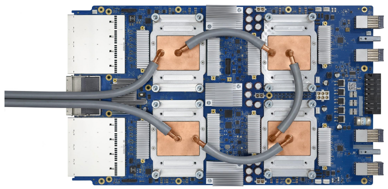

TPU v3
(2018)
420 teraflops
Available via Google Cloud
~2.3x the power of v2

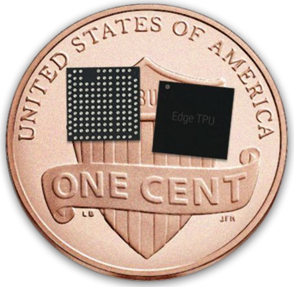

Edge TPU
(2018 EAP)
Inference Accelerator

Cloud TPU Pods - Product Offerings

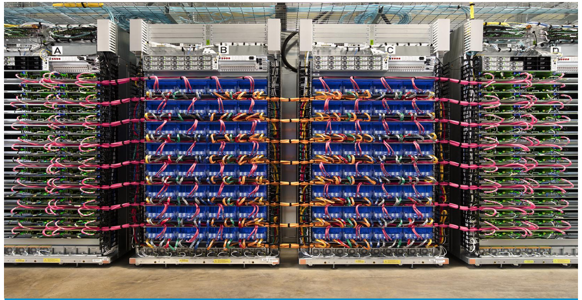

Cloud TPU v2 Pod Beta

11,500 teraflops

Up to 256 chips

4,000 GB HBM

2-D toroidal mesh network

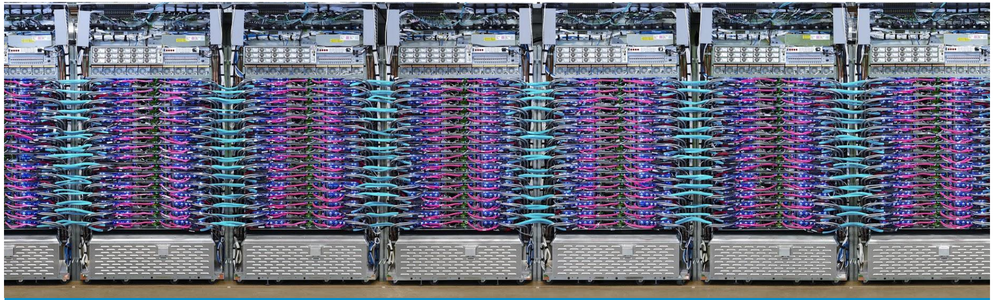

Cloud TPU v3 Pod Beta

100,000+ teraflops

Up to 1,024 chips

32,000 GB HBM

2-D toroidal mesh network

Cloud TPU Configurations

- TPU: The Tensor Processing Unit (TPU) is a custom-design chip, built from the ground up by Google for machine learning workloads.
- Cloud TPU: a device containing four TPU chips along with a fraction of a CPU host.
- Cloud TPU pods: Cloud TPUs are connected via a high-speed 2D toroidal mesh network to form Cloud TPU Pods.
- Cloud TPU slices: Slices, or smaller sections of pods, are scalable to address as much performance is needed for the workload. Slices are internal allocations consisting of different numbers of TPU cores. Pod slices come in 32, 128, 256, 512, 1024, and 2048 core-count configurations.

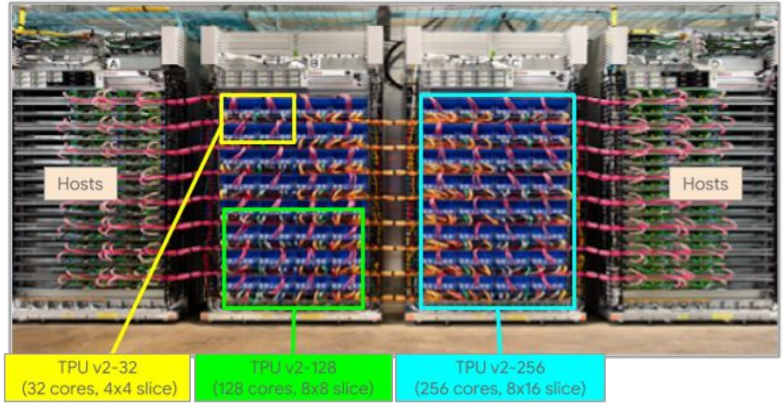

The network matters

134 points of presence and 13 subsea cable investments around the globe

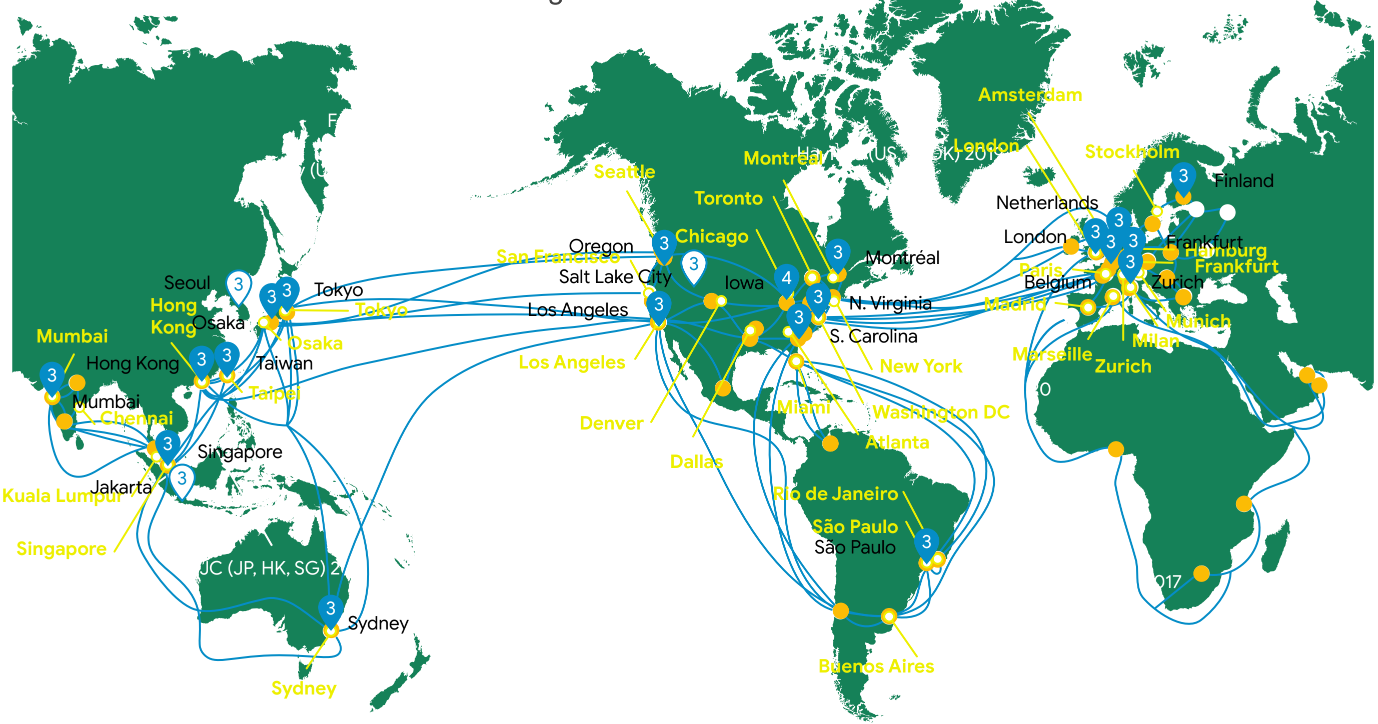

Current regions and number of zones

Future regions and number of zones

Edge points of presence

CDN nodes

Network

Dedicated Interconnect

Proprietary + Confidential

HPC → Machine Learning

ML improves with data size

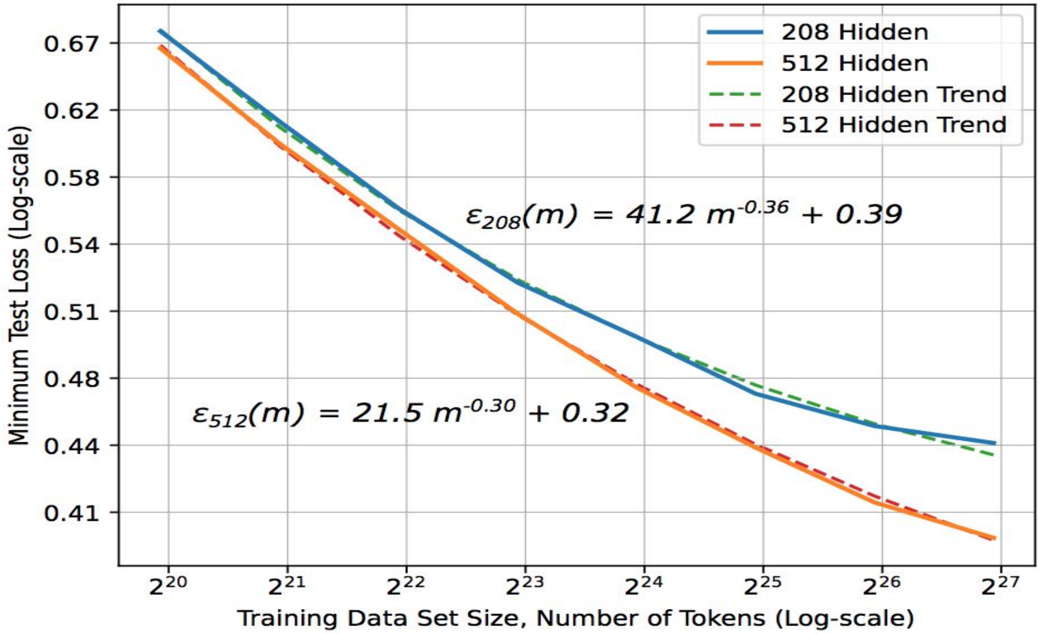

The unreasonable effectiveness of data
https://static.googleusercontent.com/media/research.google.com/en//pubs/archive/35179.pdf

Deep Learning scaling is predictable, empirically
https://arxiv.org/abs/1712.00409

Proprietary + Confidential

Increases in accuracy require much more compute.

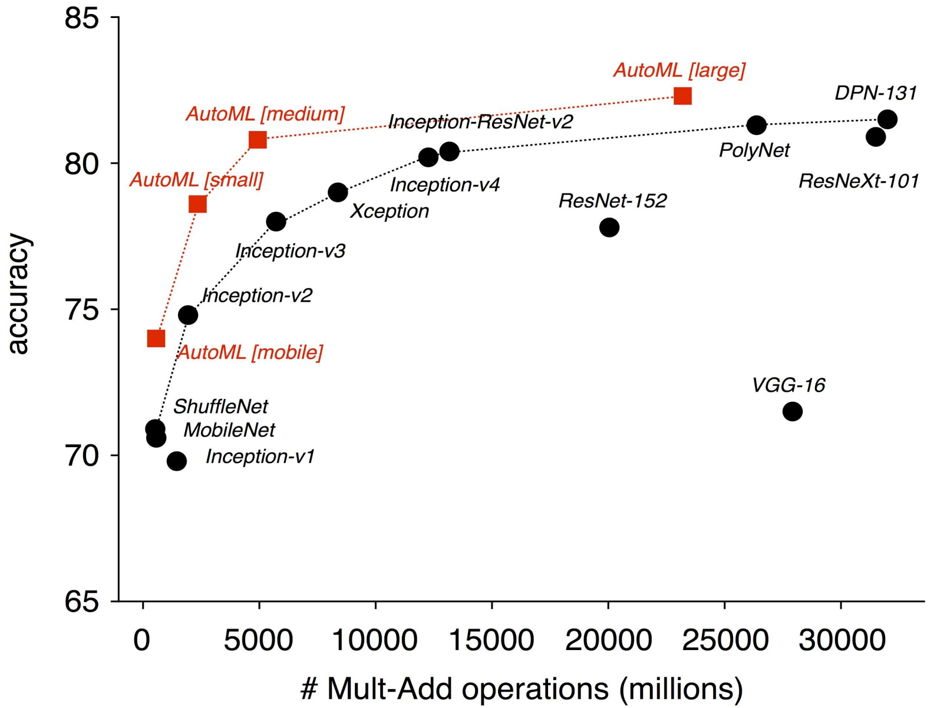

Learning Transferable Architectures for Scalable Image Recognition
Barret Zoph, Vijay Vasudevan, Jonathon Shlens, Quoc V. Le
https://arxiv.org/abs/1707.07012

Google Cloud

Proprietary + Confidential

OpenAI

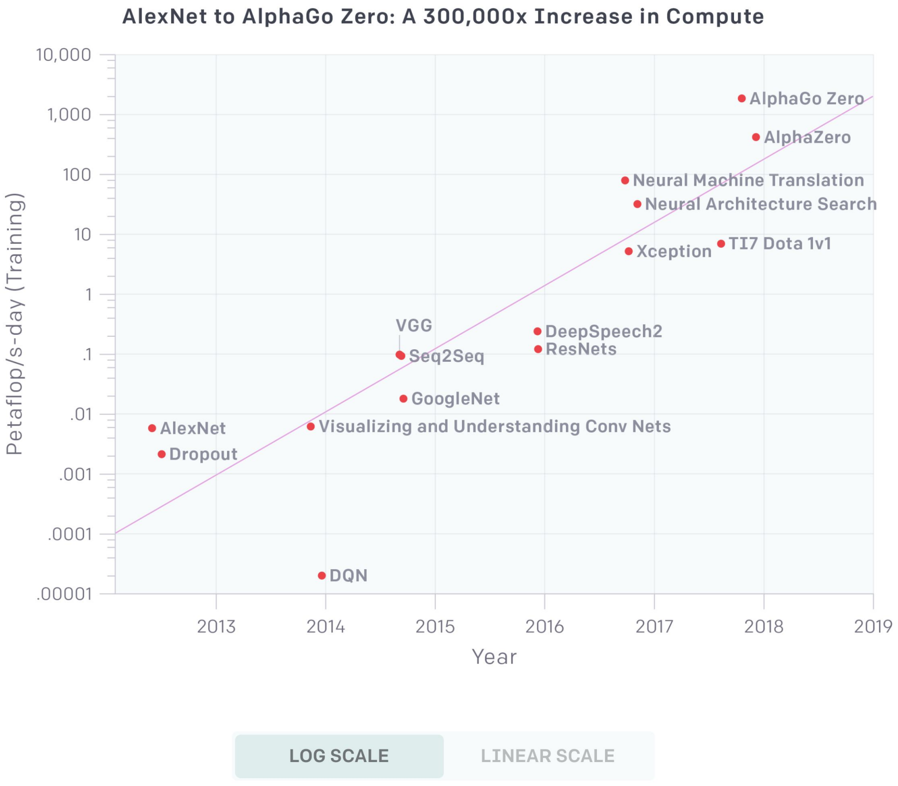

https://blog.openai.com/ai-and-compute/

Google Cloud

Proprietary + Confidential

OpenAI

https://blog.openai.com/ai-and-compute/

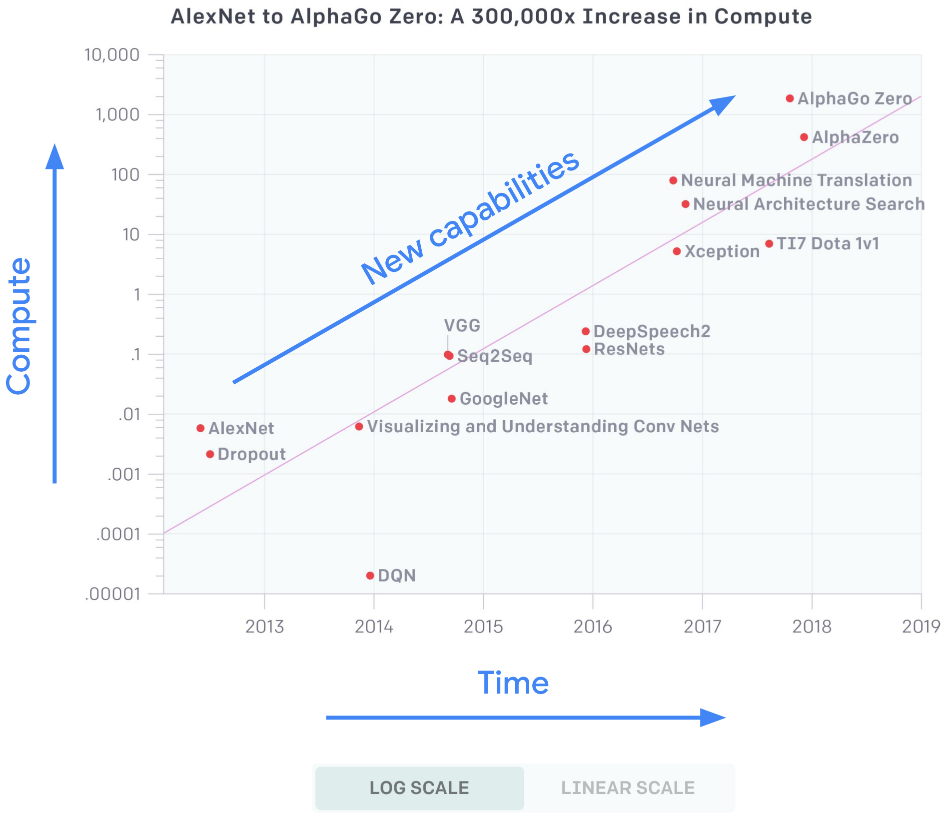

Google Cloud

OpenAI

... since 2012, the amount of compute used in the largest AI training runs **has been increasing exponentially with a 3.5-month doubling time** (by comparison, Moore's Law had an 18-month doubling period).

Since 2012, **this metric has grown by more than 300,000x** (an 18-month doubling period would yield only a 12x increase).

... within many current domains, **more compute seems to lead predictably to better performance**, and is often complementary to algorithmic advances.

... we believe **the relevant number** is not the speed of a single GPU, nor the capacity of the biggest datacenter, but **the amount of compute that is used to train a single model** – this is the number most likely to correlate to how powerful our best models are.

https://blog.openai.com/ai-and-compute/

Google Cloud

Leading to increased model quality

% ranking on Kaggle private leaderboard

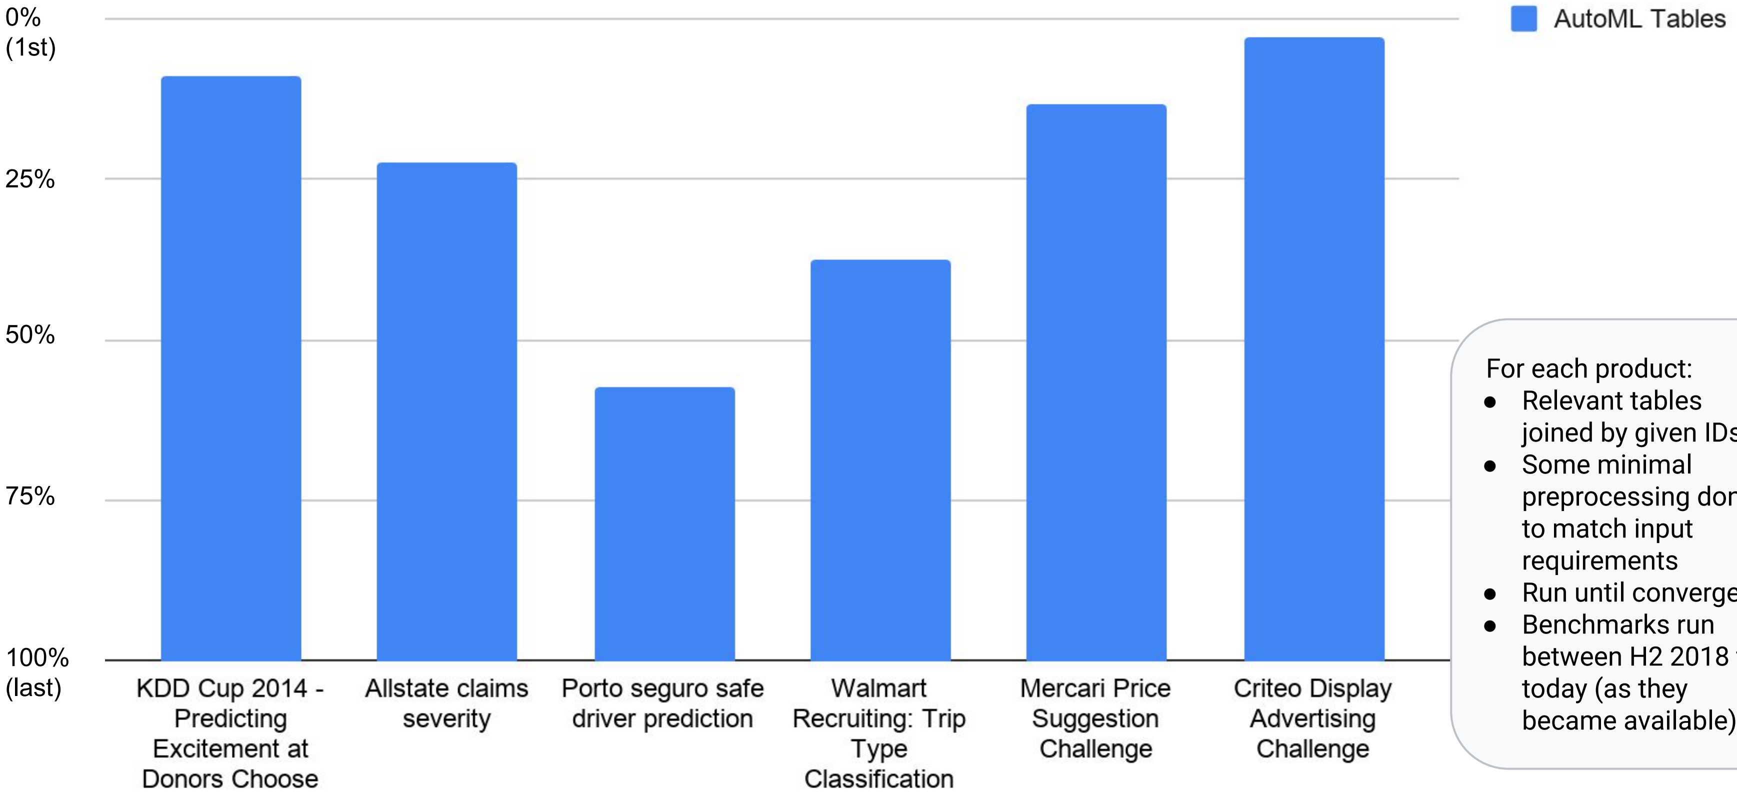

For each product:

- Relevant tables joined by given IDs
- Some minimal preprocessing done to match input requirements
- Run until converge
- Benchmarks run between H2 2018 to today (as they became available)

AutoML + experts = even better!

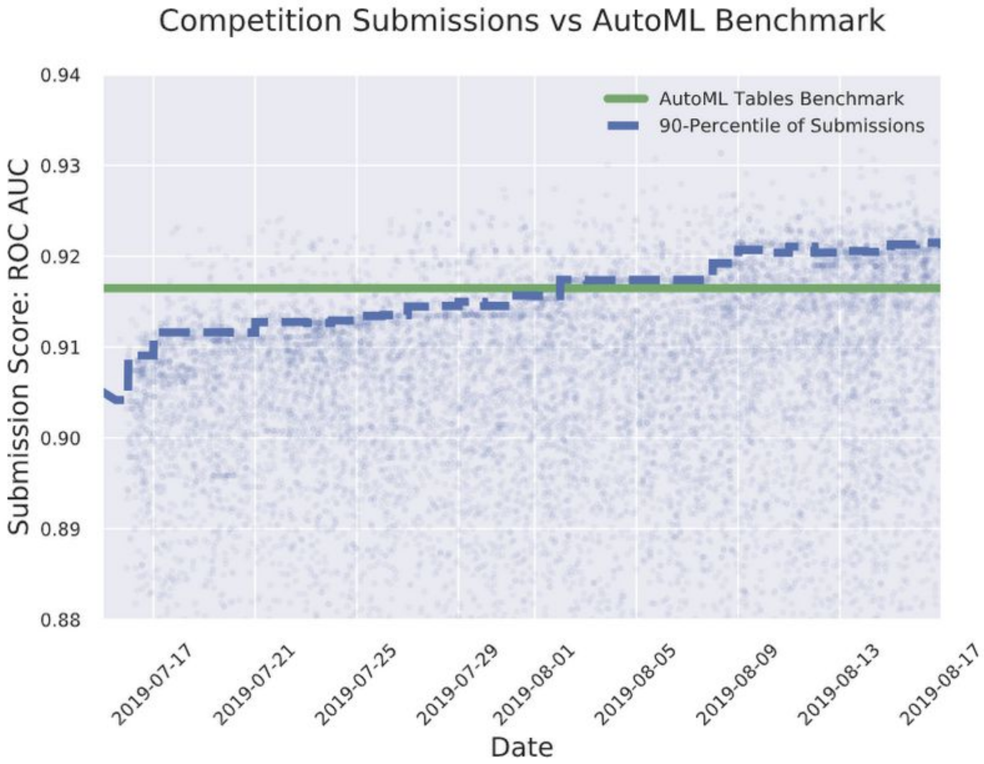

	score
team	
Erkut & Mark,Google AutoML	0.618492
Erkut & Mark	0.616913
Google AutoML	0.615982
Erkut & Mark,Google AutoML,Sweet Deal	0.615858
Sweet Deal	0.615766
Arno Candel @ H2O.ai	0.615492
ALDAPOP	0.615040
9hr Overfitness	0.614371
Shlandryn	0.614132
Erin (H2O AutoML 100 mins)	0.612657

To do ML in production, in addition to the actual ML...

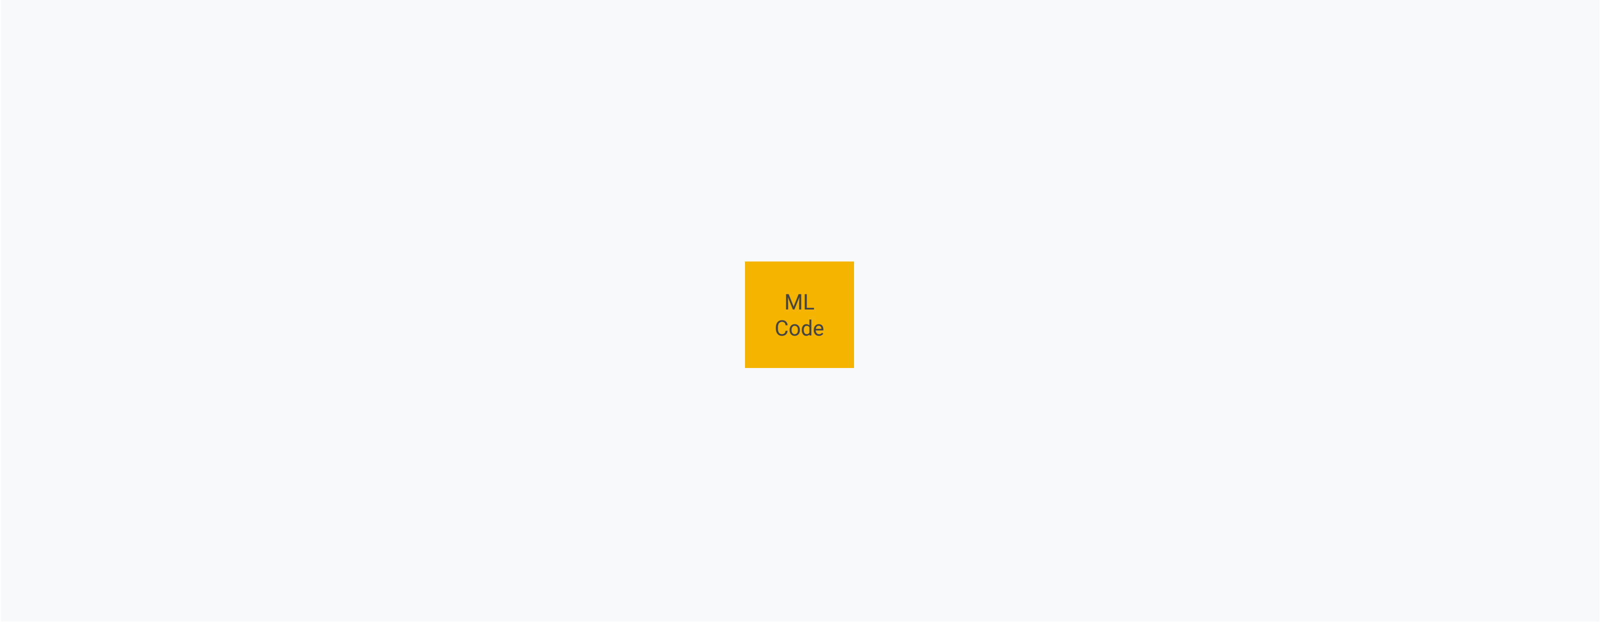

...you have to worry about so much more.

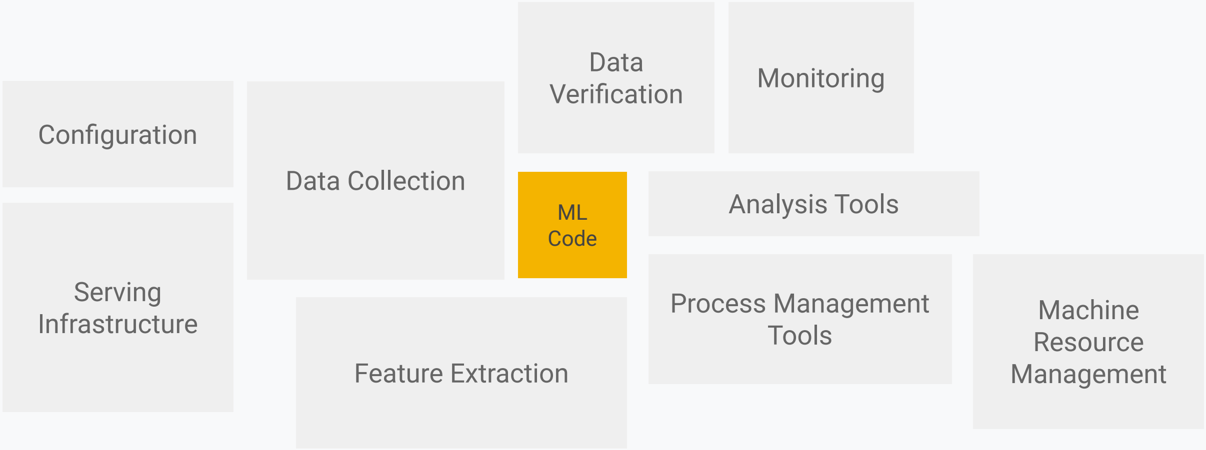

Source: Sculley et al.: Hidden Technical Debt in Machine Learning Systems

Operatize ML pipelines, not ML models

- Input validation
- Model retraining
- Reusable and shareable components
- ML microservices
- Serverless

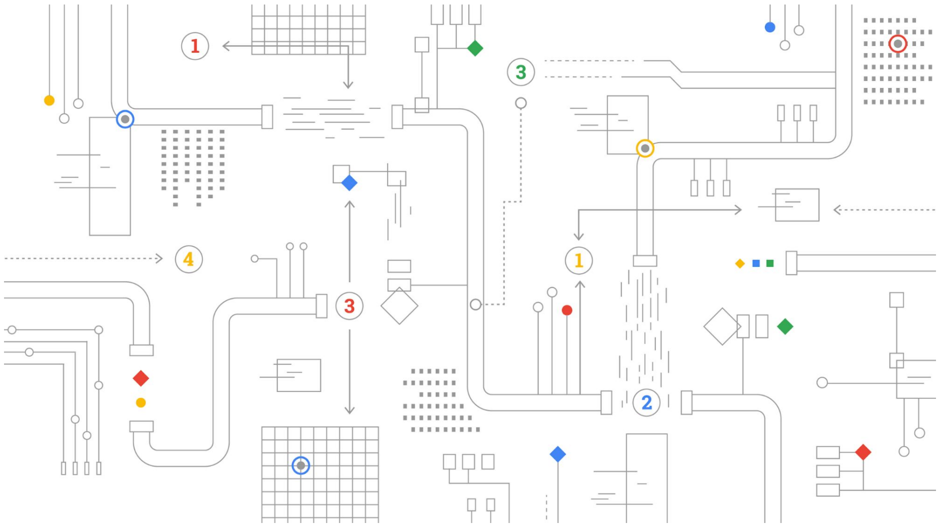

AI Hub

Public Content		+ Private Content
By Google	**By Partners**	**By Customers**
Unique AI assets by Google	Created, shared & monetized by anyone.	Content shared Securely within and with other organizations.

AutoML, TPUs, kaggle
Cloud AI Platform, etc.

Research at Google

DeepMind

Proprietary + Confidential

AI Platform Overview

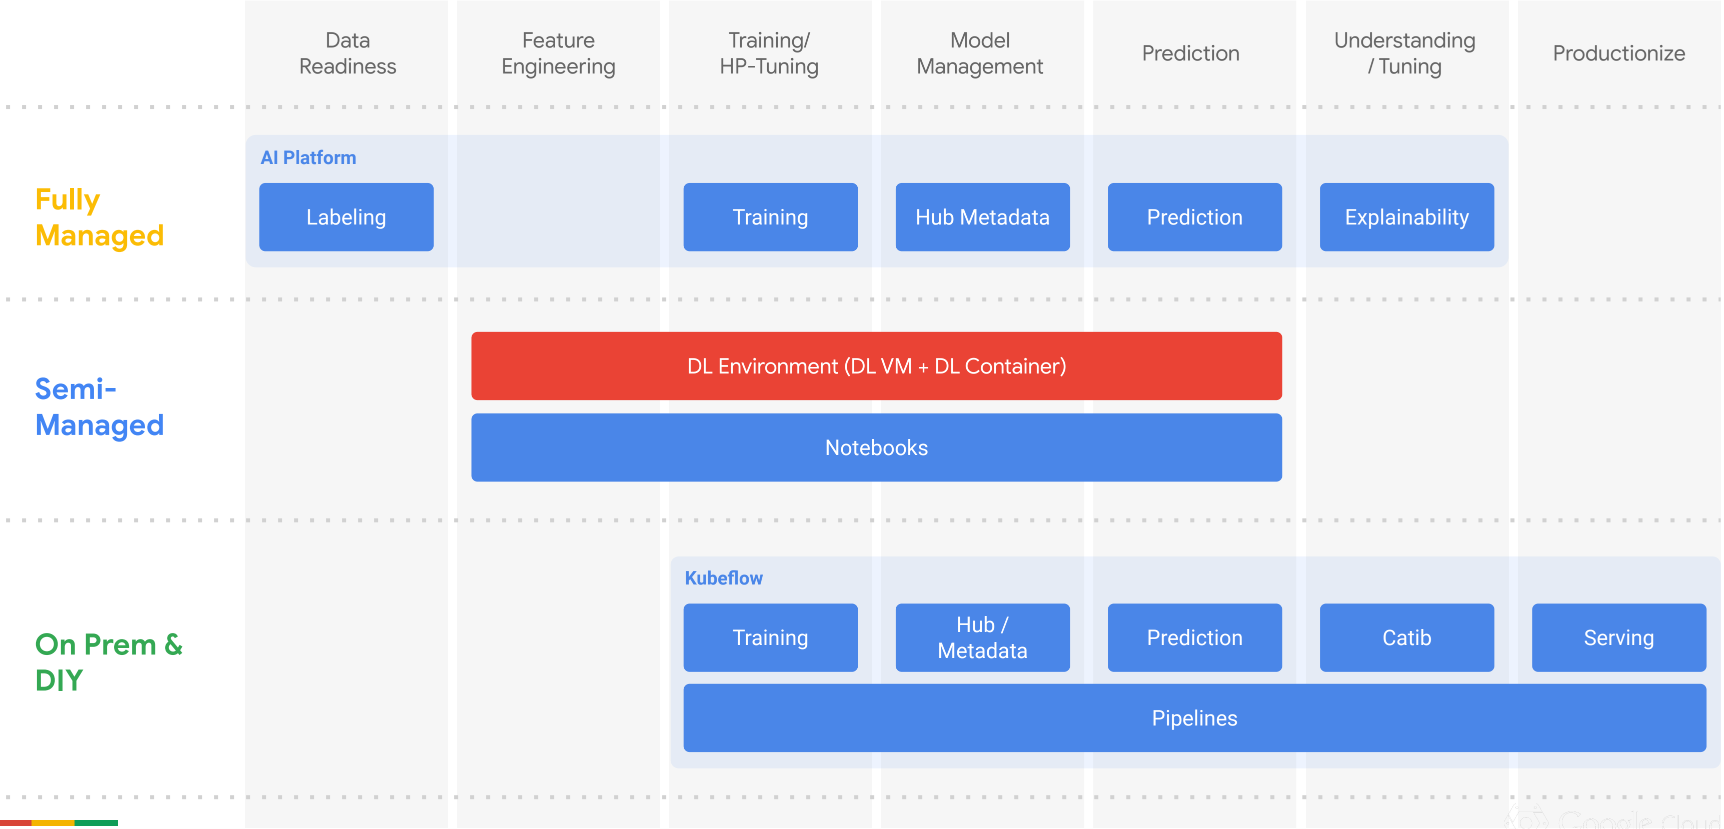

Google Cloud Platform

Choose operational complexity based on level of control needed

Cloud AI Platform

Turnkey serverless training and batch/online predictions

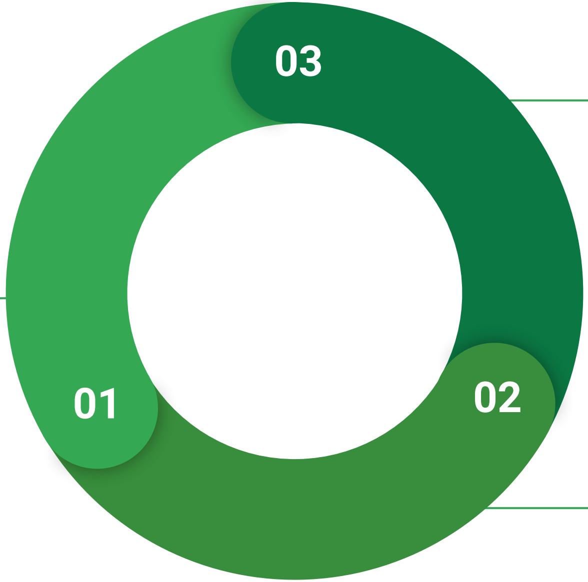

Kubeflow pipelines on Google Kubernetes Engine

Fully customizable

Deep Learning VM
Cloud Functions
Cloud Dataflow

Customized training, online/batch predictions

Proprietary + Confidential

Takeaways

GCP for HTC and AI

GCP well suited for high throughput computing with many partners, schedulers and cost effective solutions.

AI capabilities both for quick prototyping as well as scaled training

Cloud is ideal both for quick prototyping, sharing and reusing code and ML models, as well as for reproducible workloads (both large-scale training and high-throughput inference)

Scale and hybrid approach

You can use GCP capabilities to achieve enormous scale to add them to your existing on-premise resources

Thank you

Google

Karan Bhatia
karanbhatia@google.com
Leonid Kuligin
kuligin@google.com

Google