More Low-Hanging Fruit Than You Think

Enabling non-Experts with Google ML Services

Kevin Kissell, Office of the CTO

Google Cloud

Google is an AI first company

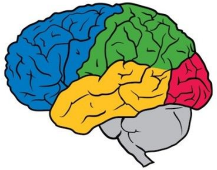

Used across products:

AlphaGo

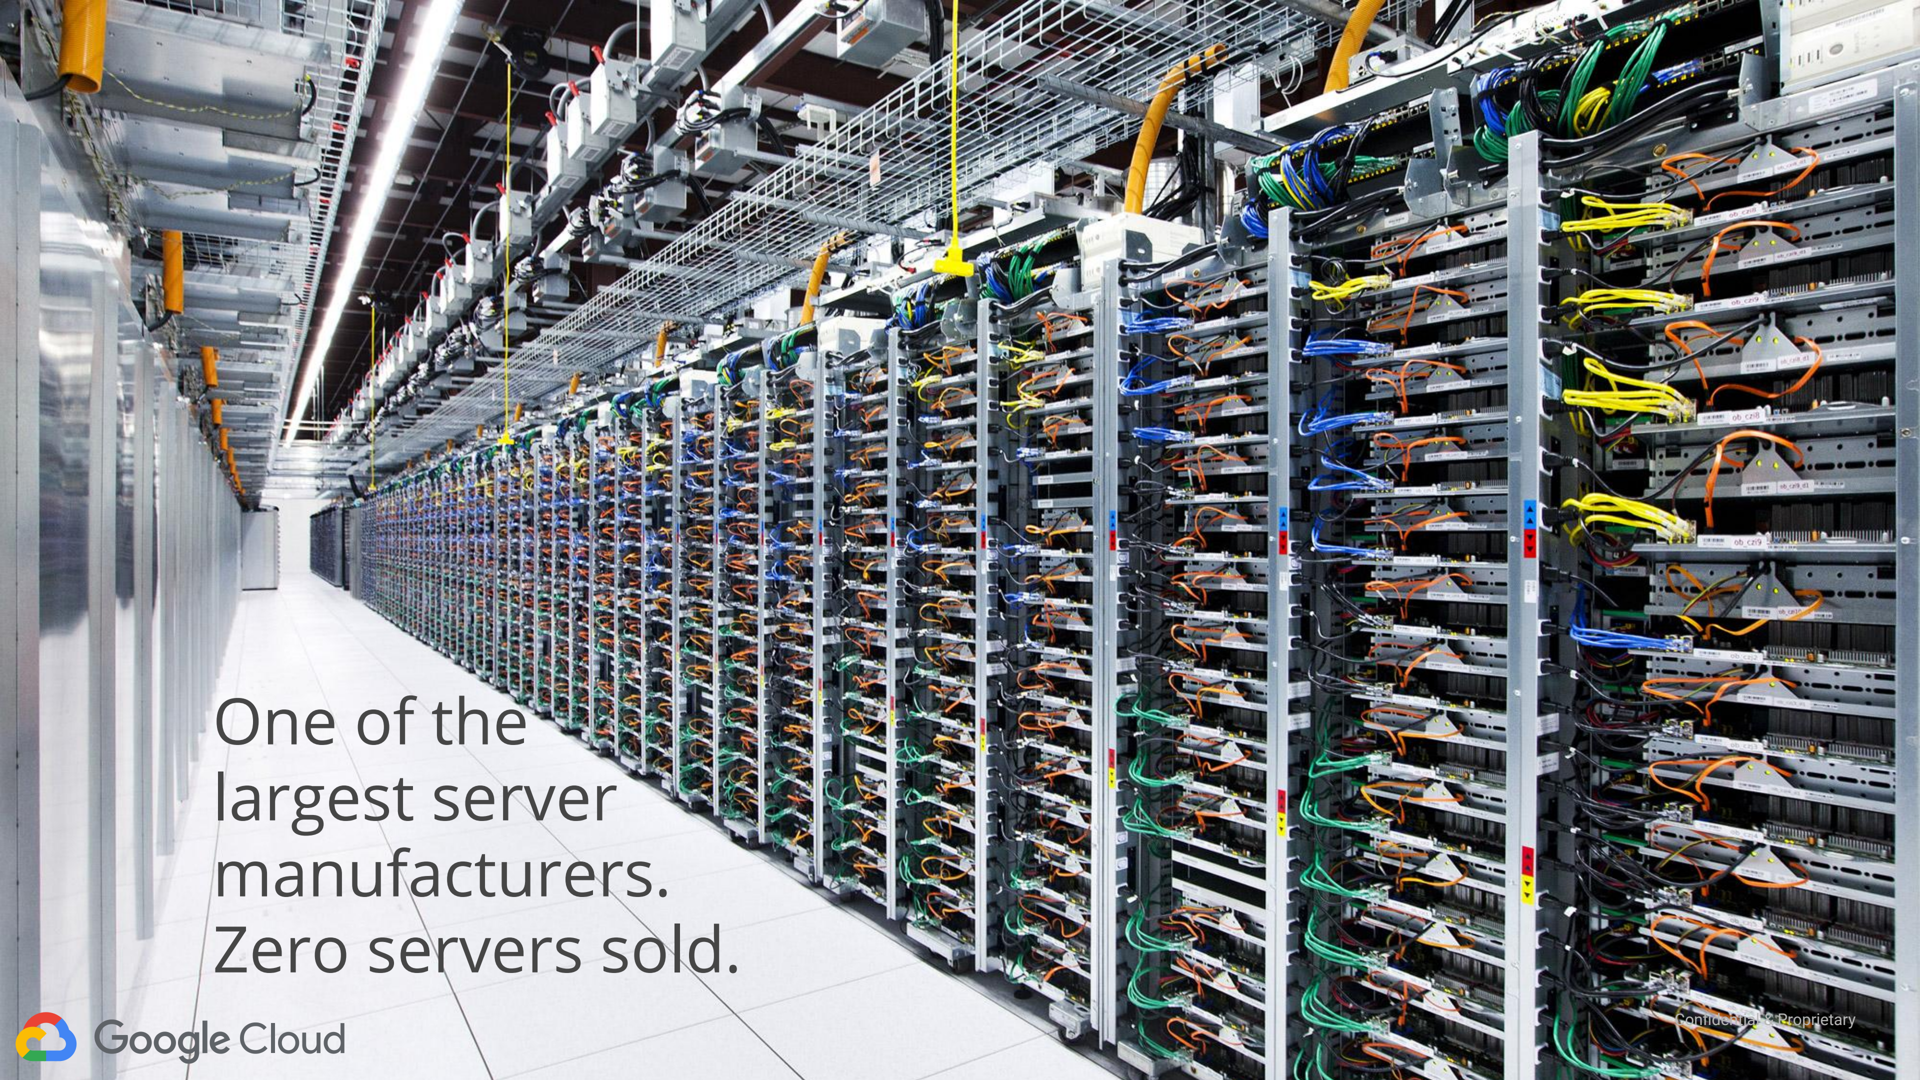
One of the largest server manufacturers.
Zero servers sold.
Google Cloud
Confidential & Proprietary

“Eating Our Own Dog Food”

Machine Learning to drive down PUE

~40% Reduction in power use for cooling

TPUs

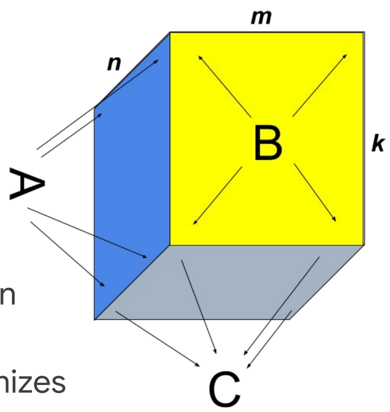

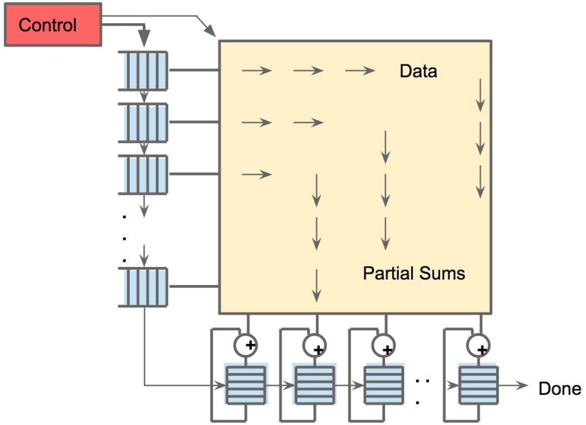

Matrix Multiplies Dominate ML Computation

Systolic Array Multiplier Architecture Maximizes Parallelism with Minimal Architectural State

> 100 TOPs per chip
> 100 POPs per pod in a toroidal mesh

Support processors handle communication, perform JIT compilation on XLA from RPCs

Only programmable in TensorFlow and PyTorch

ML-enabled Astronomy

NASA/Kepler Mission Data

UT Austin and Google, trained deep learning network with known exoplanet transit data.

Discovered 2 New Exoplanets, Kepler 90g and Kepler 90i

ML Predicting Weather

Google Brain Research for US NOAA

MetNet Neural Weather Model

Runs on 256 Google TPUs

Outperforms current physics based models for speed and accuracy out to 8 days

Parallel scaling allows prediction for entire US in seconds.

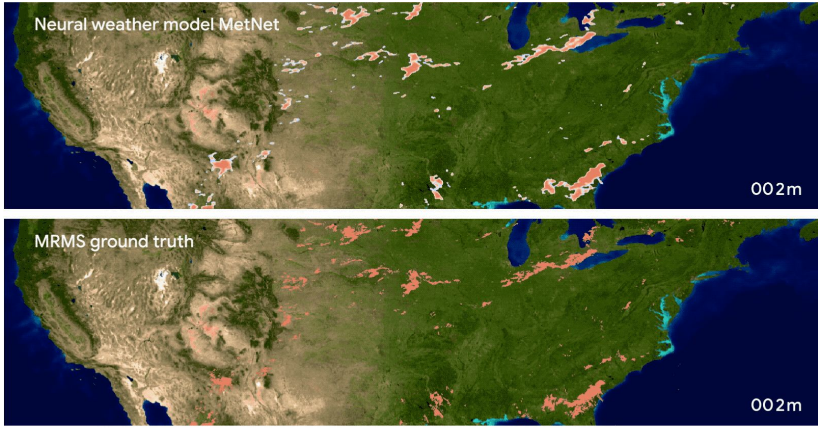

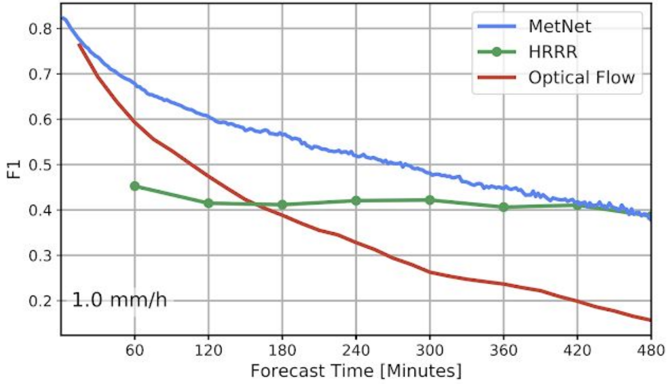

https://arxiv.org/abs/2003.12140

AlphaFold: ML for a Scientific Grand Challenge

Prediction of 3D Protein Structure, based on “1D” molecular sequence data

DeepMind used deep neural networks to predict Structure after training on ~170,000 known proteins

RMS Error of 0.1nm, scale of single atom

Comparable to results from experimental methods, but can be used for proteins not easily analysed.

This computational work represents a stunning advance on the protein-folding problem, a 50-year-old grand challenge in biology. It has occurred decades before many people in the field would have predicted. It will be exciting to see the many ways in which it will fundamentally change biological research.

PROFESSOR VENKI RAMAKRISHNAN
NOBEL LAUREATE AND PRESIDENT OF THE ROYAL SOCIETY

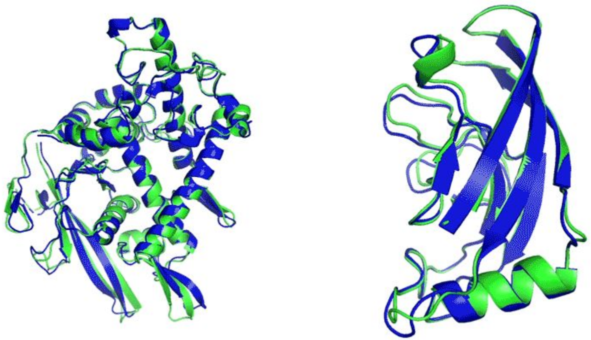

T1037 / 6vr4
90.7 GDT
(RNA polymerase domain)

T1049 / 6y4f
93.3 GDT
(adhesin tip)

Experimental result
Computational prediction

ML **must** become easier and more accessible

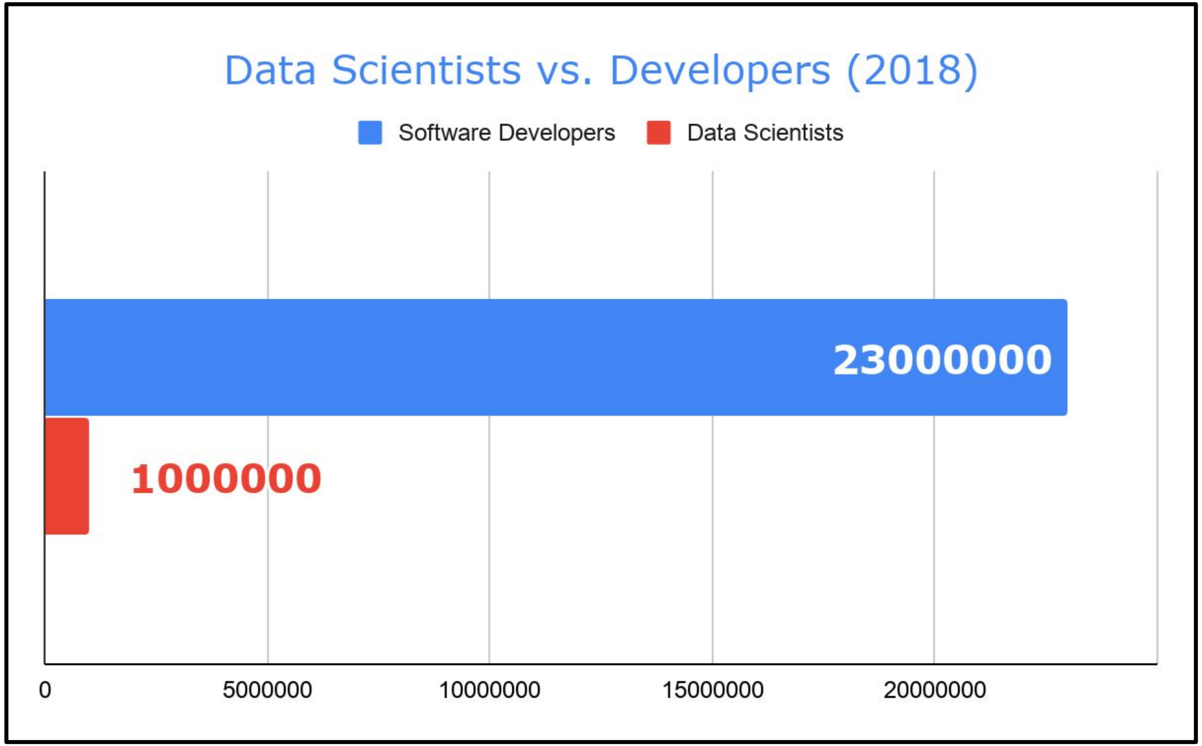

https://www.kdnuggets.com/2018/09/how-many-data-scientists-are-there.html
https://www.quora.com/How-many-software-developers-are-there-in-the-world

The Vertex AI Platform

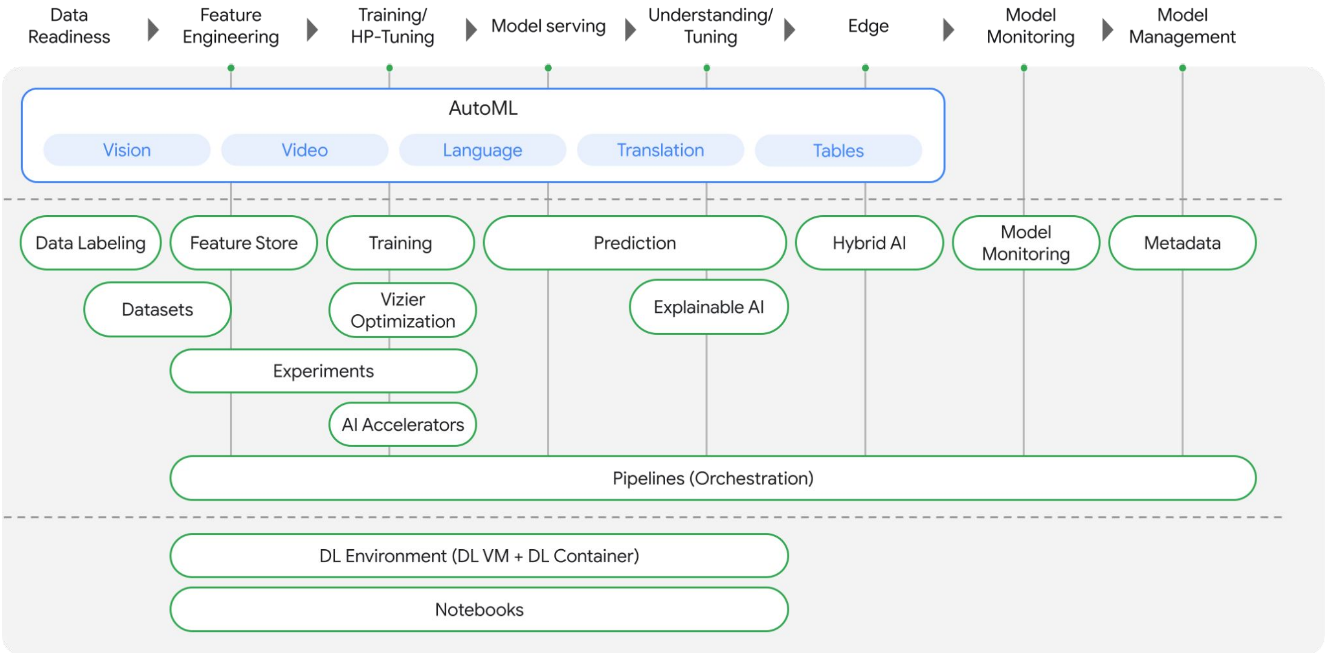

APIs and AutoML

Making ML accessible to all developers

Sight

- Cloud Vision
- Cloud Video Intelligence
- AutoML Vision
- AutoML Video Intelligence (new)

Language

- Cloud Translation
- Cloud Natural Language
- AutoML Translation
- AutoML Natural Language

Conversation

- Dialogflow Enterprise Edition
- Cloud Text-to-Speech
- Cloud Speech-to-Text

Structured Data (new)

- AutoML Tables (new)
- Recommendation AI (new)

What's the right tool for you?

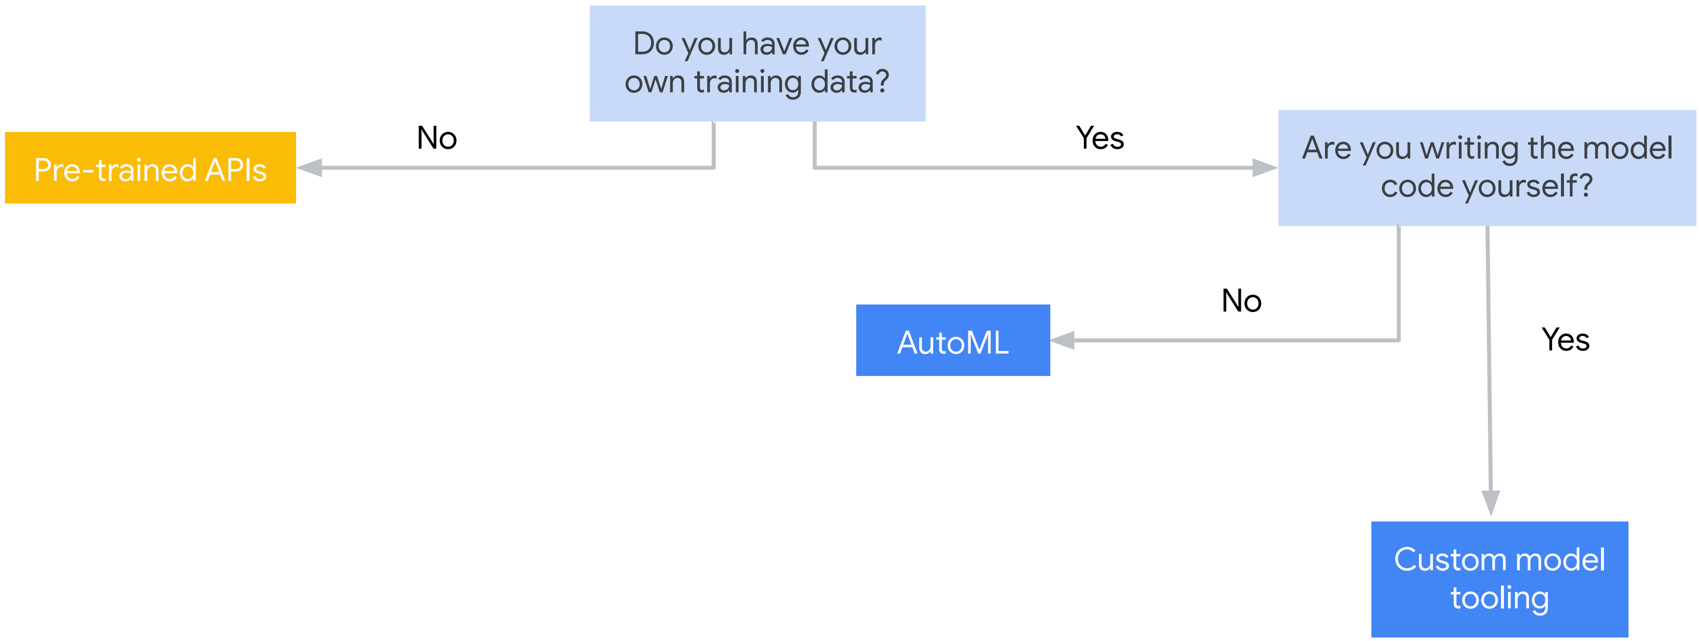

Use a **pre-trained model** to accomplish common ML tasks

Cloud Vision

Cloud Video Intelligence

Cloud Speech-to-Text & Text-to-Speech

Cloud Natural Language

Cloud Translation

What’s the right tool for you?

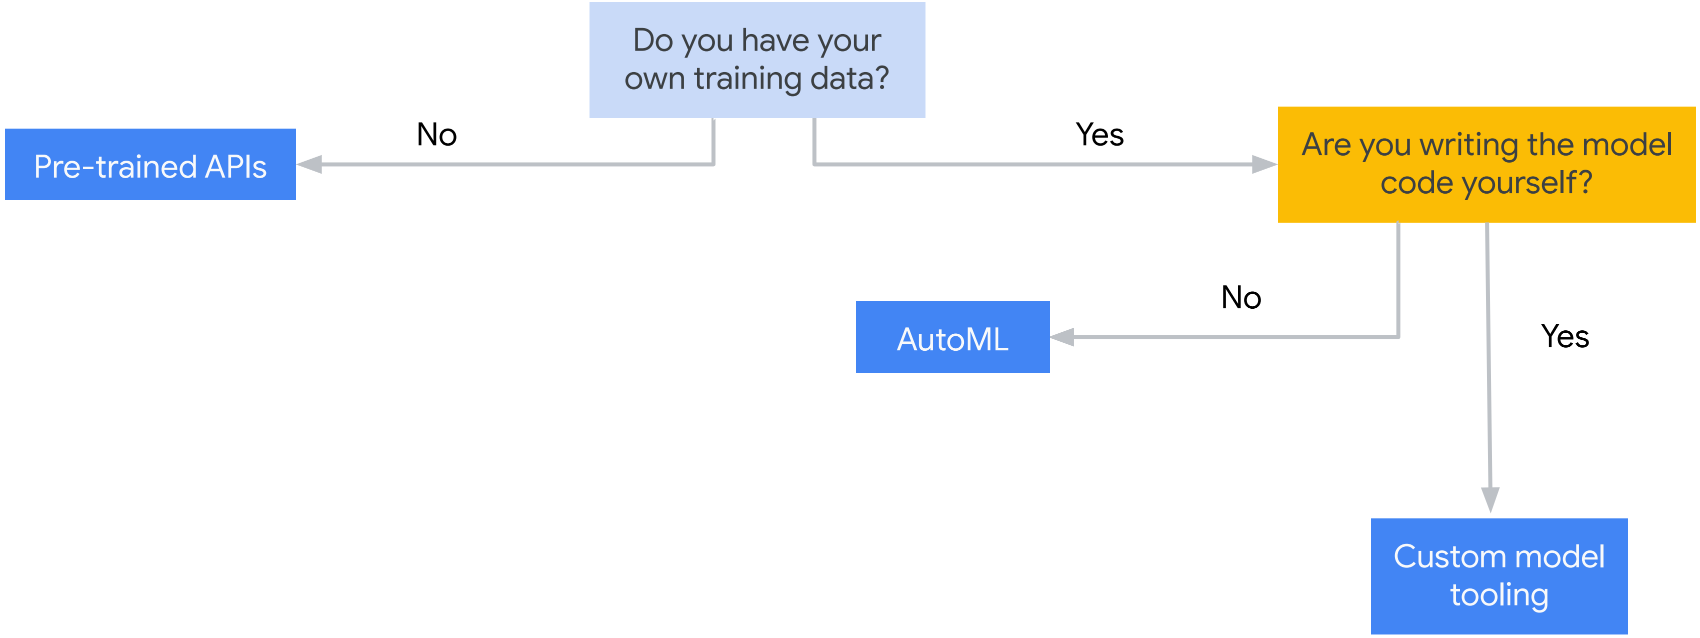

Bringing your OWN Data? First choose your data type..

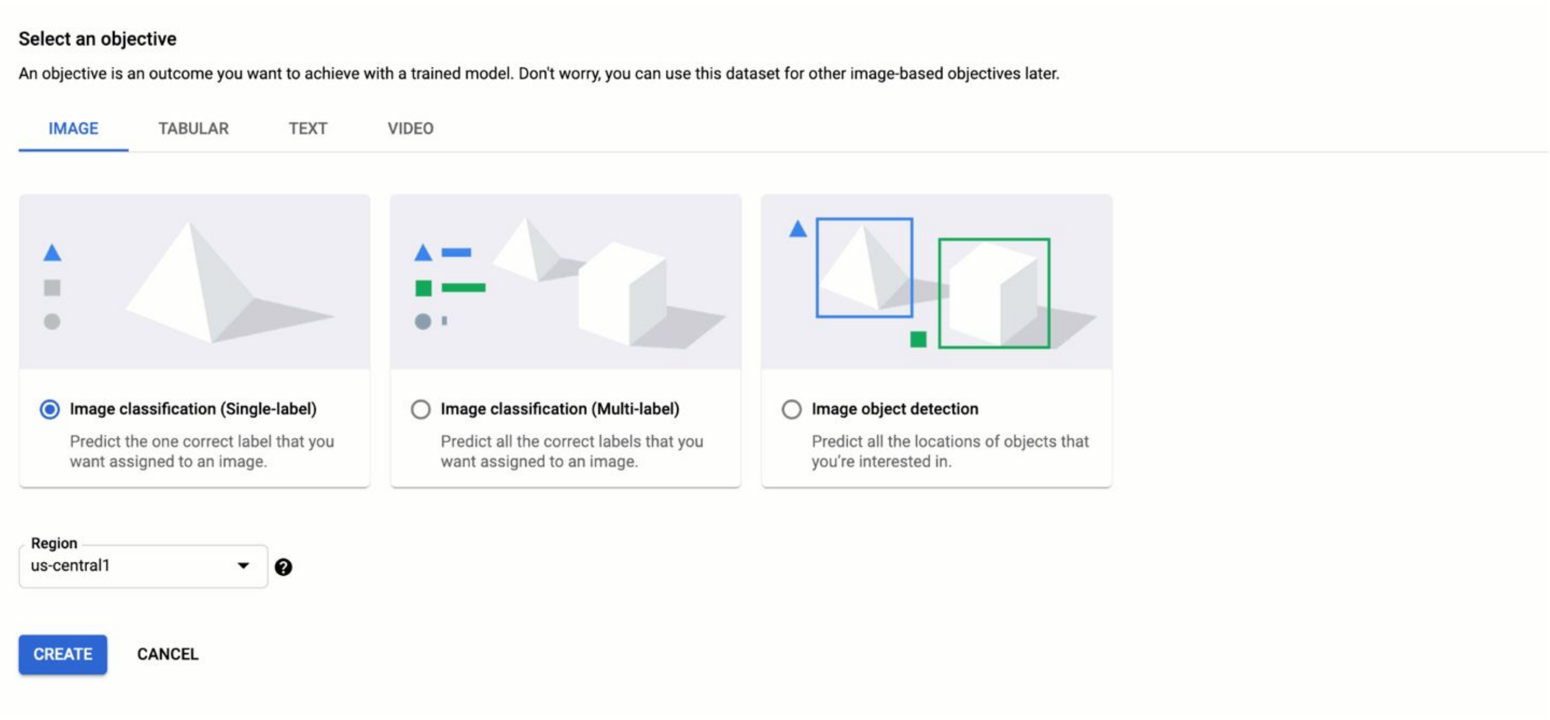

Image datasets

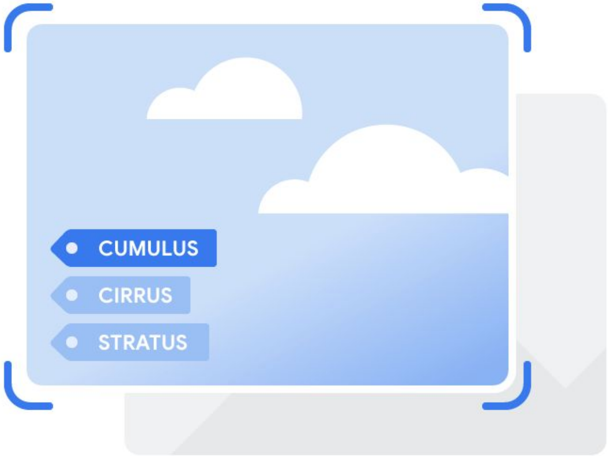

Image classification models predict one (or many) labels for an image. For example, identifying types of clouds from images of the sky.

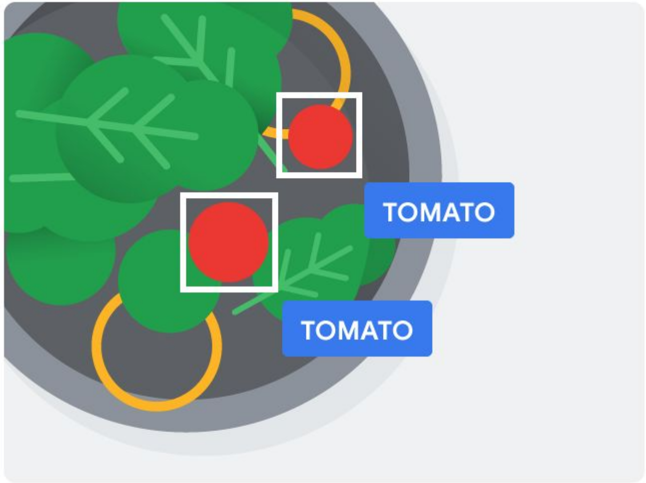

Object detection models draw bounding boxes around items of interest. For example, identifying vegetables from images of food.

Should you write your own model code? It depends.

Use AutoML if...	Build a custom model if...
Your use case fits into our supported AutoML offerings: • Image & video classification or object detection • Text classification, entity extraction, sentiment • Tabular regression, classification, or forecasting	Your use case doesn’t fit into any AutoML use cases OR Your model takes mixed input types, like images + tabular metadata
You don’t need to know specifics about the underlying model	You want control over your model’s architecture, framework , or exported model assets. For example, maybe your model needs to be built with TensorFlow
You want to develop a quick initial model to use as a baseline (which could end up being your production model)	You already have a baseline or heuristic and you want to see if you can improve upon it

What’s the right tool for you?

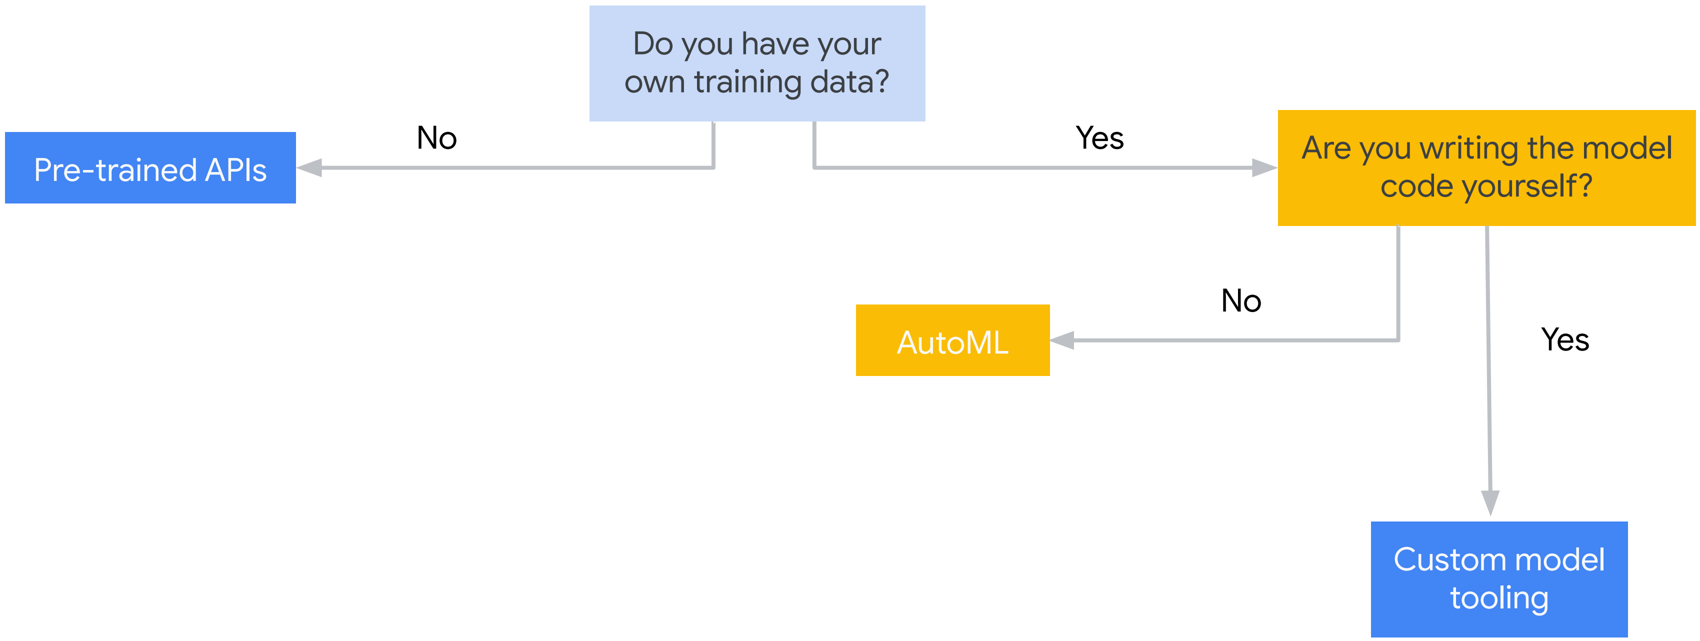

AutoML

ML that creates ML

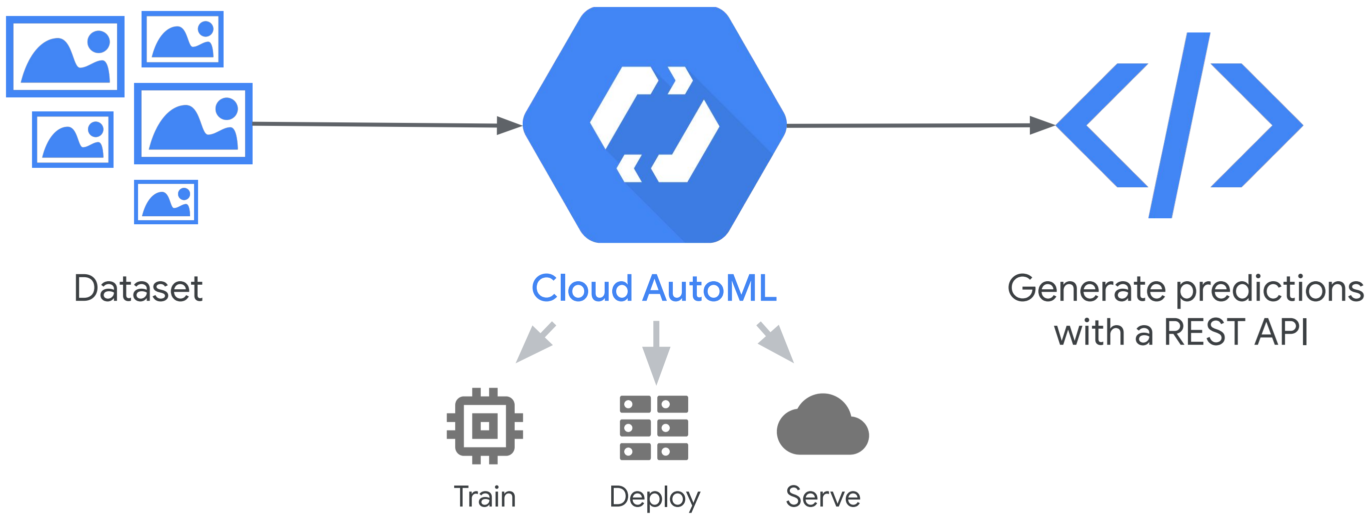

Cloud AutoML - Best in Class Research

Learning to learn

Transfer Learning

Hyperparameter Tuning

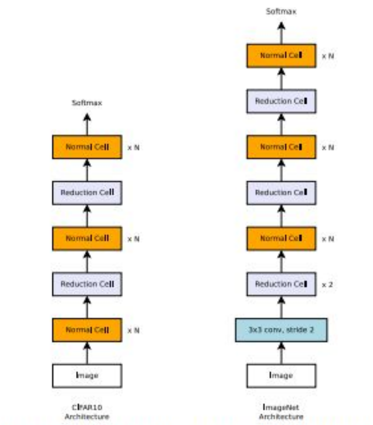

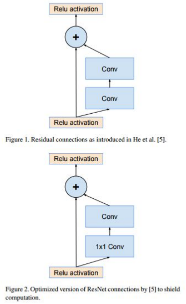

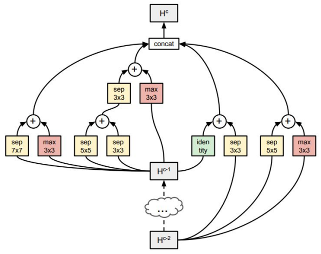

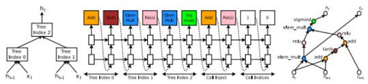

Learning Transferable Architectures for Scalable Image Recognition, Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V. Le. Arxiv, 2017.

Inception-v4, Inception-ResNet and the Impact of Residual Connections on Learning Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alex Alemi. AAAI, 2017.

Progressive Neural Architecture Search Chenxi Liu, Barret Zoph, Jonathon Shlens, Wei Hua, Li-Jia Li, Li Fei-Fei, Alan Yuille, Jonathan Huang, Kevin Murphy, Arxiv, 2017

Neural Architecture Search with Reinforcement Learning Barret Zoph, Quoc V. Le. ICLR 2017.

How does AutoML work?

Controller: proposes ML models

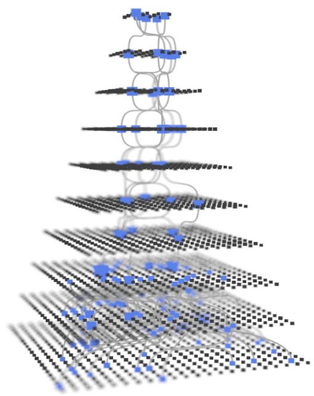

20K times

Iterate to find the most accurate model

Train & evaluate models

AutoML outperforms handcrafted models

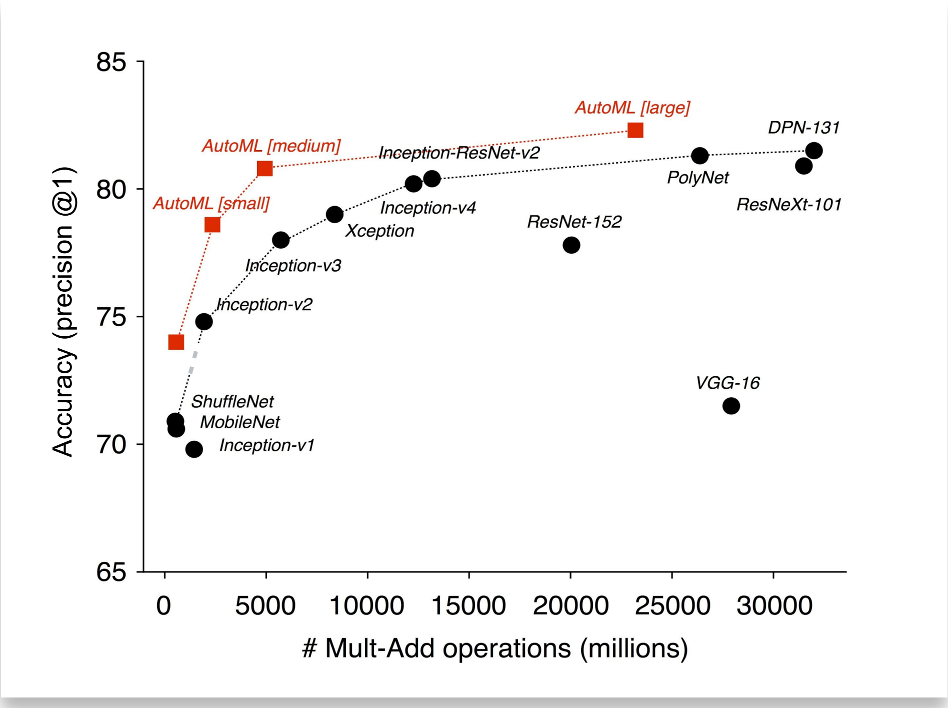

accuracy

computational cost

Google Cloud

Learning Transferable Architectures for Scalable Image Recognition, Zoph et al. 2017, https://arxiv.org/abs/1707.07012

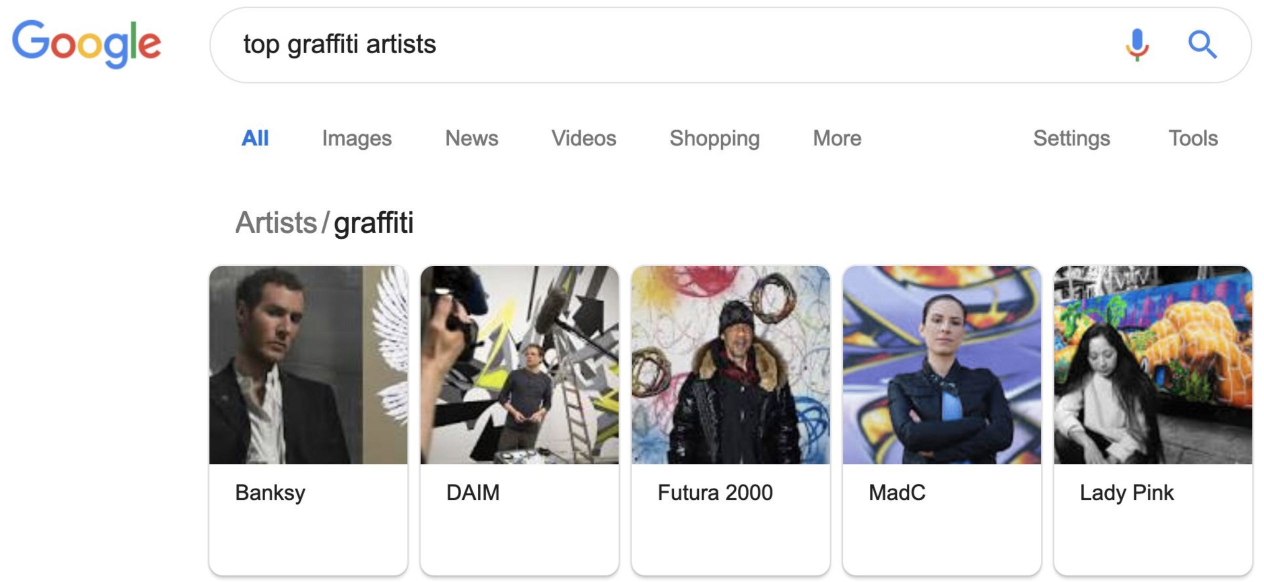

Google
top graffiti artists
All
Images
News
Videos
Shopping
More
Settings
Tools
Artists / graffiti
Banksy
DAIM
Futura 2000
MadC
Lady Pink

AutoML Step 1: Show me examples

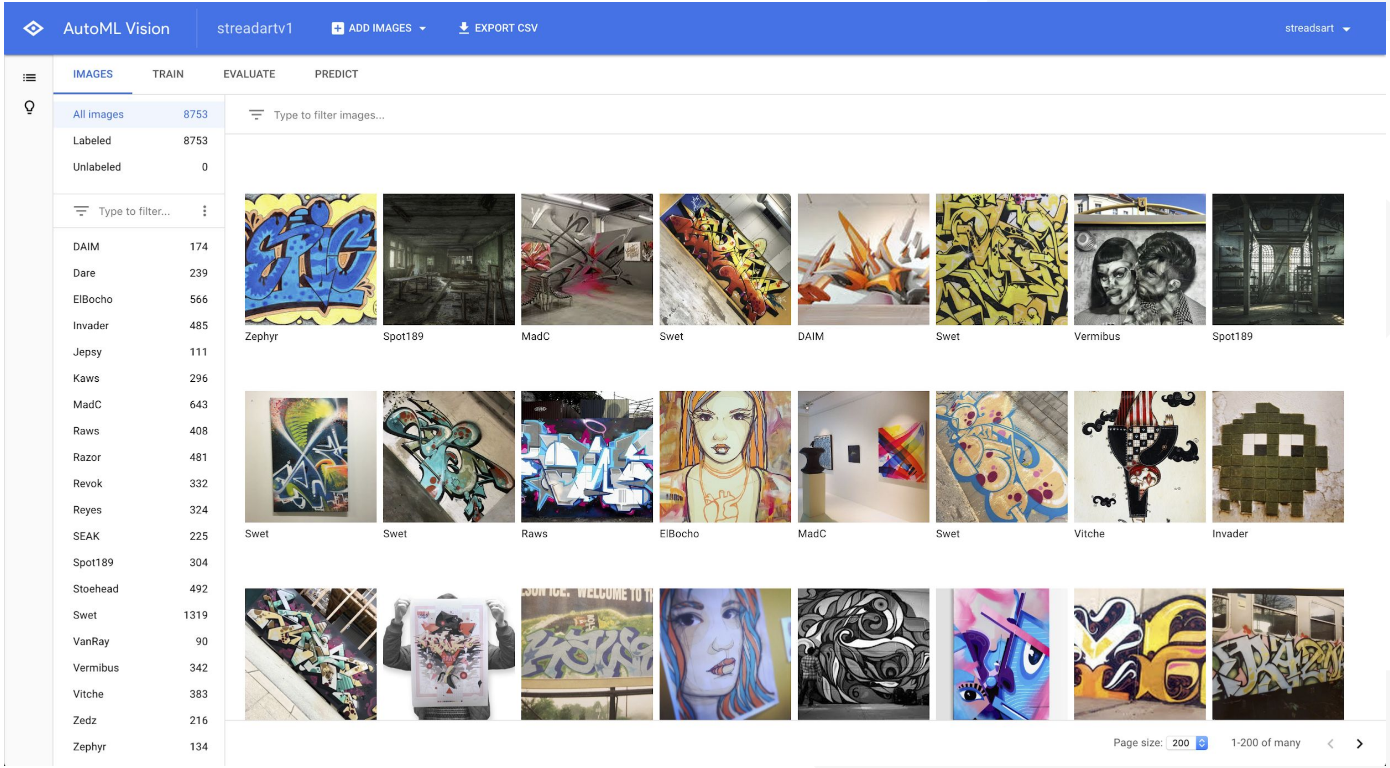

Google Cloud

AutoML Step 2: Click a button and wait

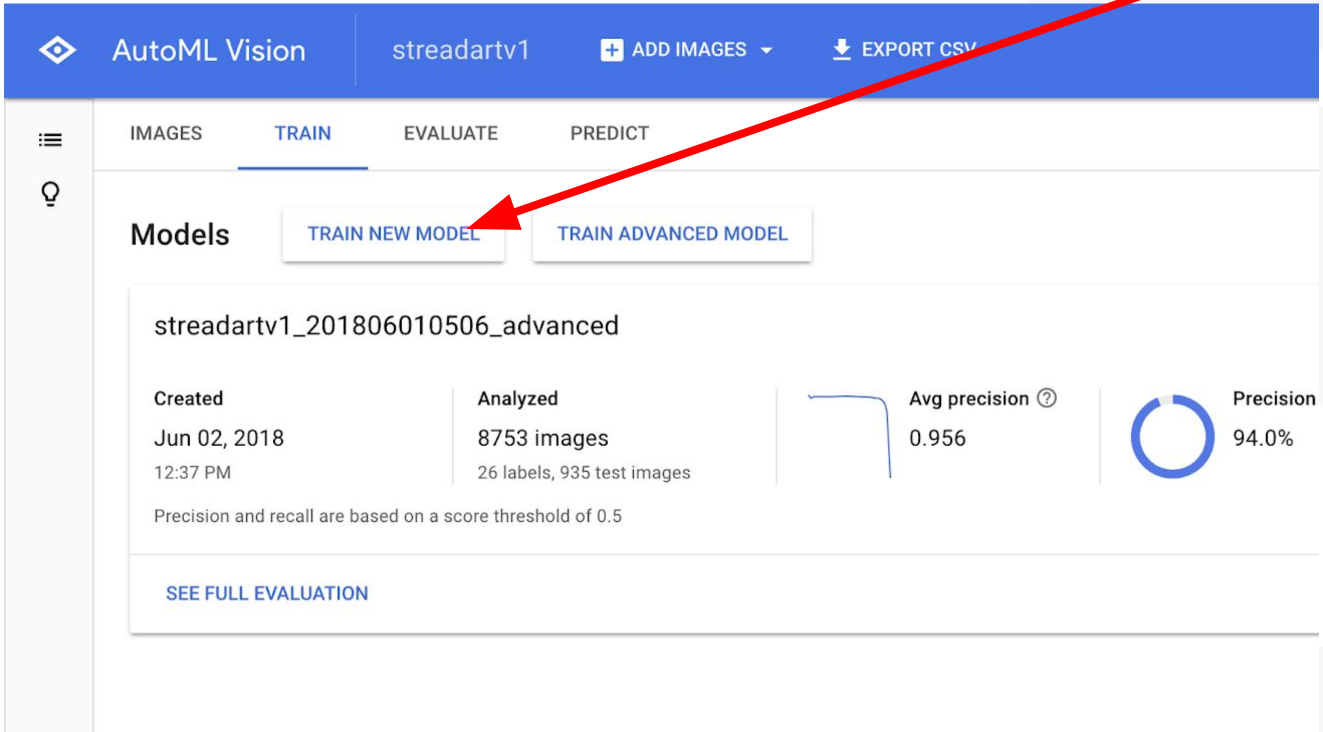

AutoML Step 3: Evaluate

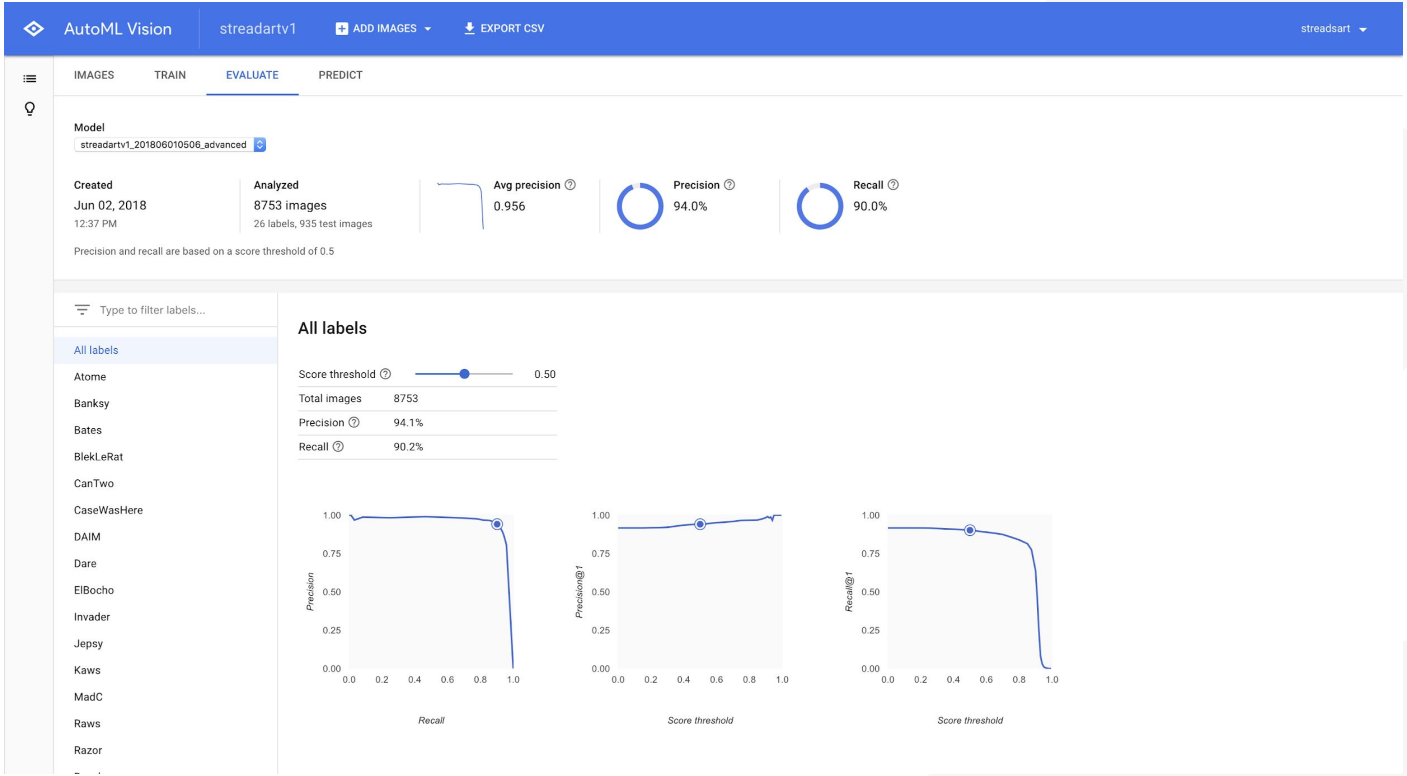

AutoML Step 3: Evaluate

AutoML Vision > Datasets > streadartv1

streadsart

Confusion matrix

Predicted label

True Label	Banksy	BlekLeRat	Vitche	Swet	Zephyr	VanRay	Reyes	MadC	Invader	Dare
Banksy	82%	0%	5%	0%	0%	2%	2%	2%	2%	0%
BlekLeRat	11%	88%	0%	0%	0%	0%	0%	0%	0%	0%
Vitche	0%	0%	100%	0%	0%	0%	0%	0%	0%	0%
Swet	1%	0%	0%	98%	0%	0%	0%	0%	0%	0%
Zephyr	7%	0%	0%	0%	85%	0%	0%	0%	0%	7%
VanRay	0%	0%	0%	0%	0%	100%	0%	0%	0%	0%
Reyes	0%	0%	0%	0%	0%	0%	86%	8%	0%	4%
MadC	0%	0%	0%	0%	0%	0%	0%	94%	0%	5%
Invader	0%	2%	0%	0%	0%	0%	0%	0%	97%	0%
Dare	3%	0%	0%	3%	3%	3%	0%	10%	0%	75%

AutoML Step 3: Evaluate

Predicted label

True Label	Banksy	BlekLeRat	Vitche	Swet	Zephyr	VanRay	Reyes	MadC	Invader	Dare
Banksy	82%	0%	5%	0%	0%	2%	2%	2%	2%	0%
BlekLeRat	11%	88%	0%	0%	0%	0%	0%	0%	0%	0%
Vitche	0%	0%	100%	0%	0%	0%	0%	0%	0%	0%
Swet	1%	0%	0%	98%	0%	0%	0%	0%	0%	0%
Zephyr	7%	0%	0%	0%	85%	0%	0%	0%	0%	7%
VanRay	0%	0%	0%	0%	0%	100%	0%	0%	0%	0%
Reyes	0%	0%	0%	0%	0%	0%	86%	8%	0%	4%
MadC	0%	0%	0%	0%	0%	0%	0%	94%	0%	5%
Invader	0%	2%	0%	0%	0%	0%	0%	0%	97%	0%
Dare	3%	0%	0%	3%	3%	3%	0%	10%	0%	75%

AutoML Step 4: Use the Model (streetartist.app)

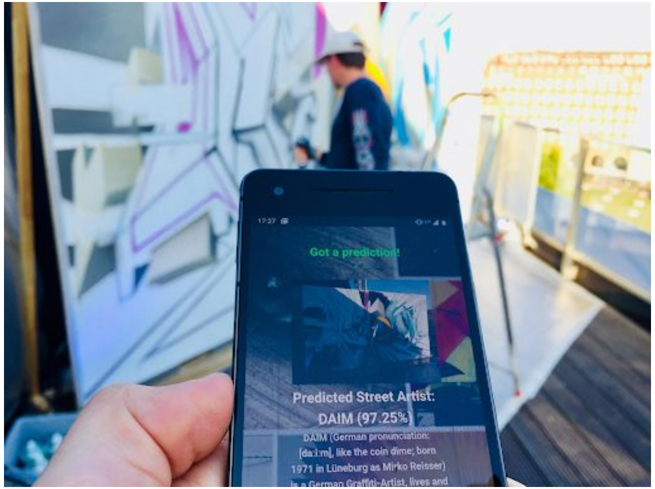

AutoML Tables

Go from raw data to models in days instead of months

Produce state-of-the art models with one click

Guides users through the full machine learning lifecycle without code

Uses Google's serving infrastructure making deployment fast and easy

Handle data as found in the wild

Automated feature engineering for:

- Numbers
- Timestamps
- Classes
- Lists
- Strings
- Nested fields

Resilient to + guardrails for:

- Imbalanced data
- Highly correlated features
- Missing values
- High cardinality features (like IDs)
- Outliers

Back to Proteins: Imagination Before Brute Force

Reduce to a Solved Problem

Estrada et. al. at UNM wanted to identify proteins automatically

Purpose-built NNs slow and expensive to develop

So...encoded 3D structure and properties in synthetic color 2D images

Used off-the-shelf image classifiers

https://par.nsf.gov/servlets/purl/10109139

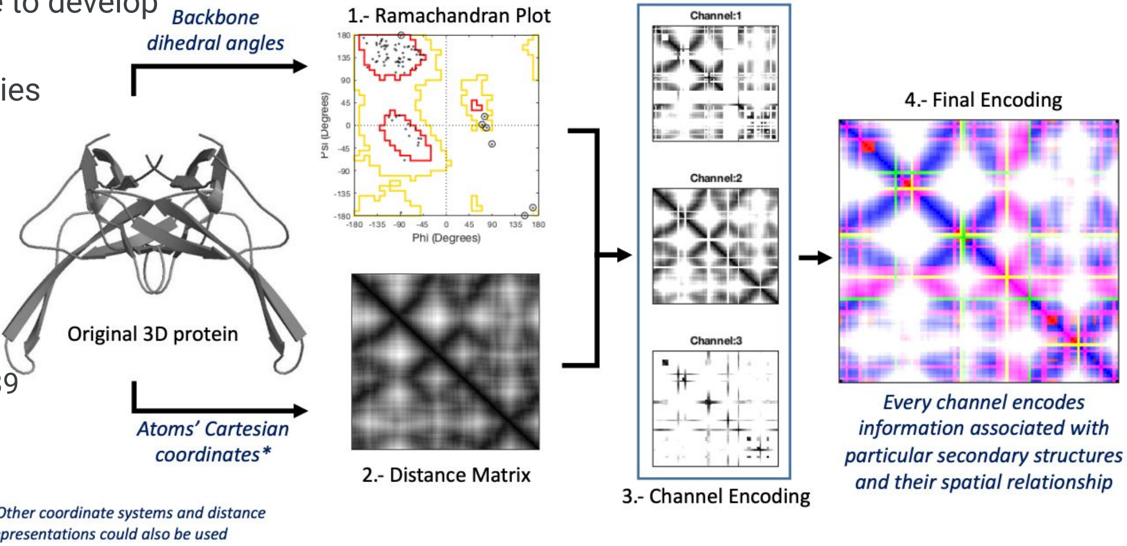

The Machine Learning Use Cases are there (so is the data)

Manufacturing

- Predictive maintenance
- Energy efficiency
- Warranty reserve estimation
- Propensity to buy
- Demand forecasting
- Process optimization
- Telematics

Retail

- Predictive inventory planning
- Recommendation engines
- Upsell and cross-channel marketing
- Market segmentation and targeting
- Customer ROI and lifetime value

Healthcare and Life Sciences

- Alerts and diagnostics from real-time patient data
- Disease identification and risk stratification
- Patient triage optimization
- Proactive health management
- Healthcare provider sentiment analysis

Travel and Hospitality

- Aircraft scheduling
- Dynamic pricing
- Social media – consumer feedback and interaction analysis
- Customer complaint resolution
- Traffic patterns and congestion management

Financial Services

- Risk analytics and regulation
- Customer Segmentation
- Cross-selling and up-selling
- Sales and marketing campaign management
- Credit worthiness evaluation

Energy, Utilities and Raw Materials

- Power usage analytics
- Seismic data processing
- Carbon emissions and trading
- Customer-specific pricing
- Smart grid management
- Energy demand and supply optimization

Custom ML models

For advanced ML use cases

AI Platform

A new middle pathway

Majority of enterprises

Cloud Auto ML

Pre-packaged AI solutions

No ML skills required

Pre-built, pre-trained APIs

TensorFlow
Open Source Machine Learning Model

Cloud TPUs
Hardware optimised for machine learning

Google Cloud

Thanks for Your Attention!

And think about the ways
that Google Cloud
could become part
of YOUR lab equipment...