

Build a ML-powered product in no time on Google Cloud


Andrea Martelli - Solutions Engineer, Google Cloud
hammers@google.com



My pasta story

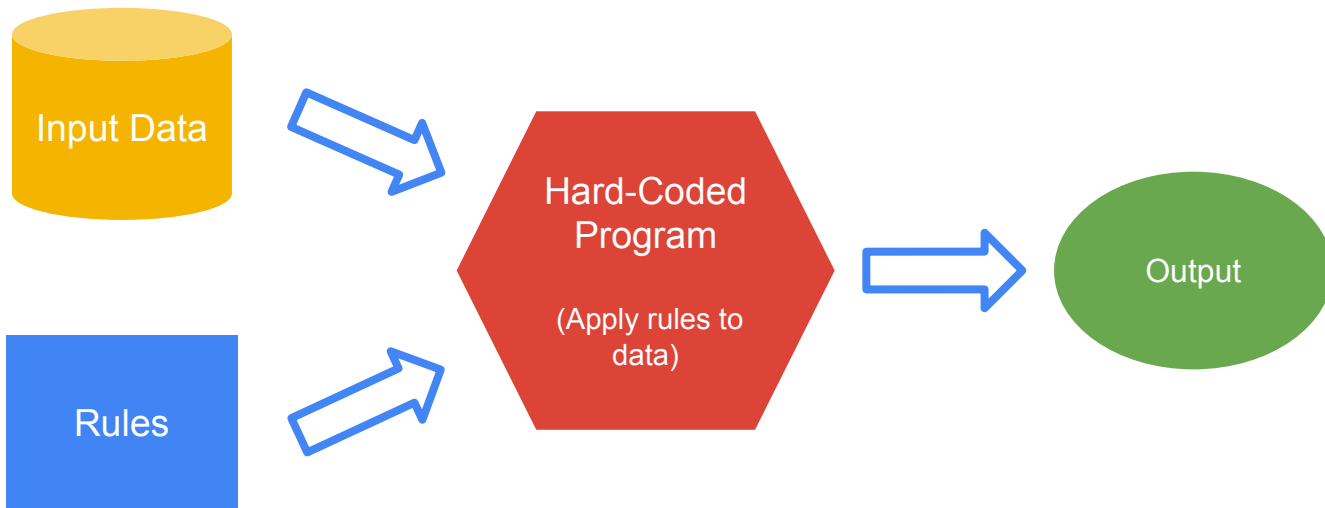


My pasta story



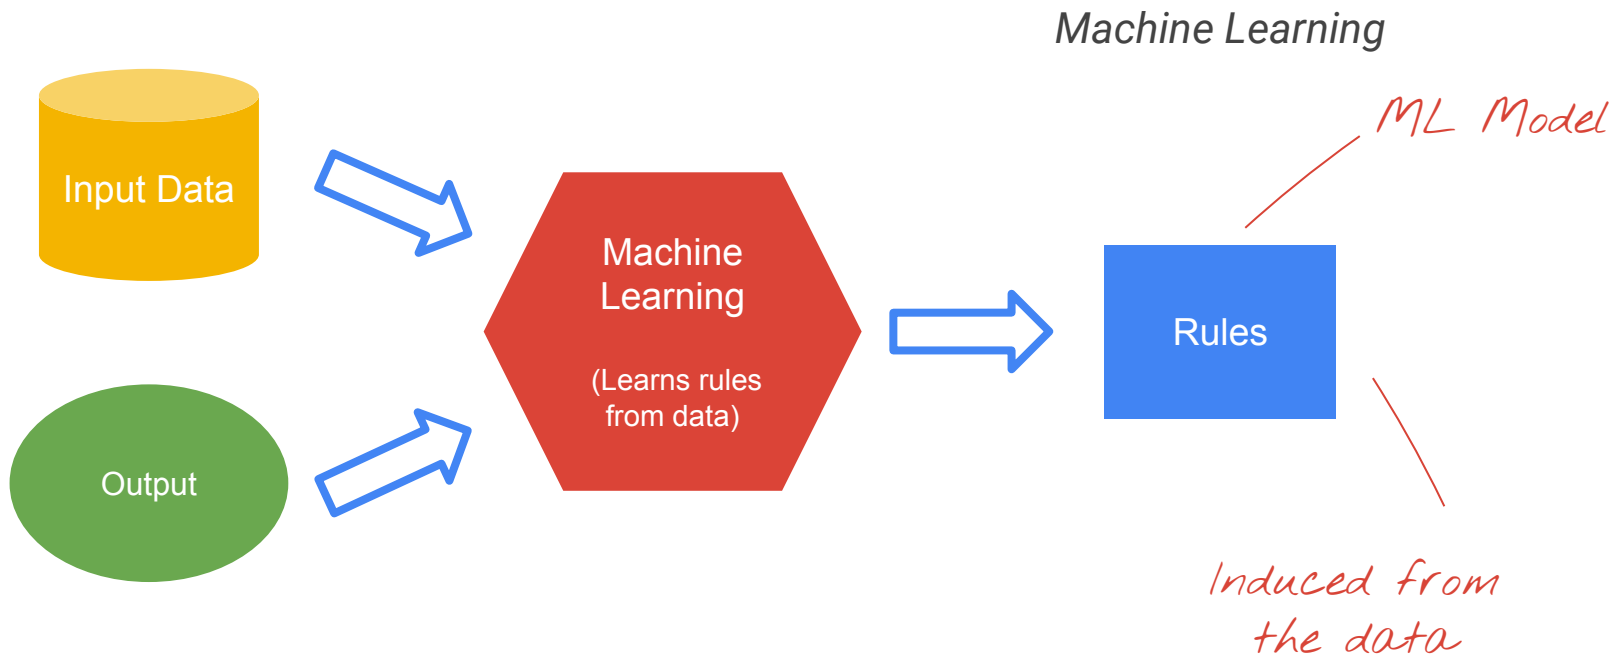
What is Machine Learning?

Normal program



*Written by SMEs,
Statisticians, etc.*

What is Machine Learning?



Gut Feel

AI captures the gut feel in Go, one for how a move feels and another for whether we feel like we're winning or losing



Pixels in, joystick out

Just neurons

At 10 minutes?

it's random luck



030 3 1

At 120 minutes?
it's flawless



022 5 1

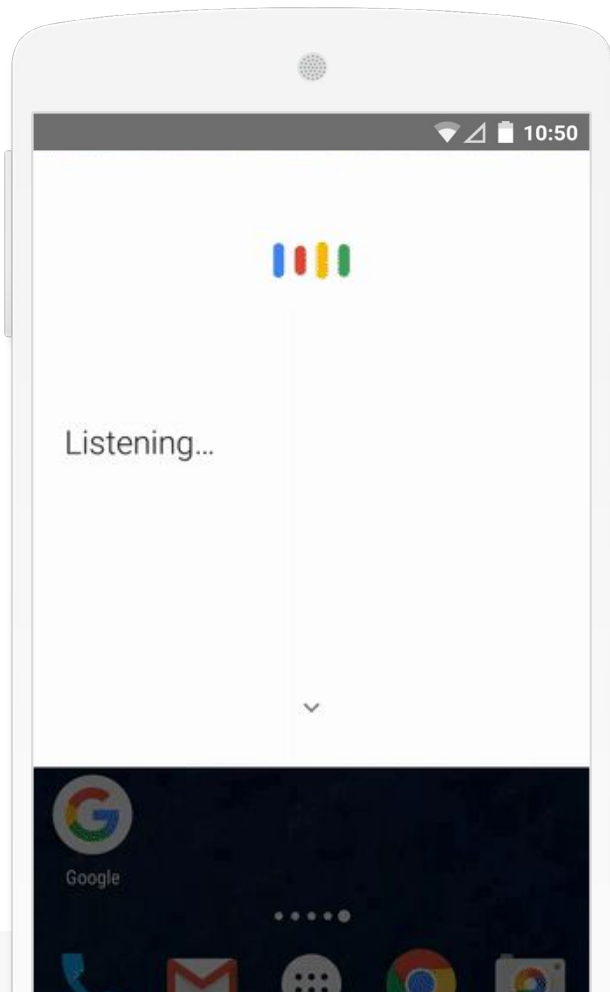
At 240 minutes?
strategy emerges



“Machine learning is a core, transformative way by which we’re rethinking how we’re doing everything.”

– Sundar Pichai





20%

of mobile searches
are via voice




Google app

Google

Proprietary & Confidential

Translate Improvement



old: PBMT

new: GNMT


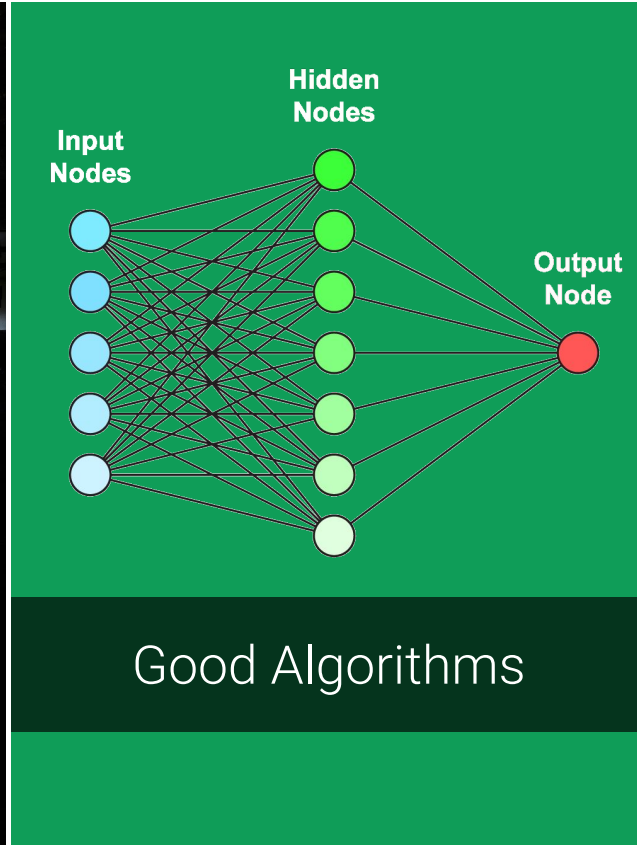



Driving


Five million miles and three years later
the blind are now driving in Waymo
vehicles



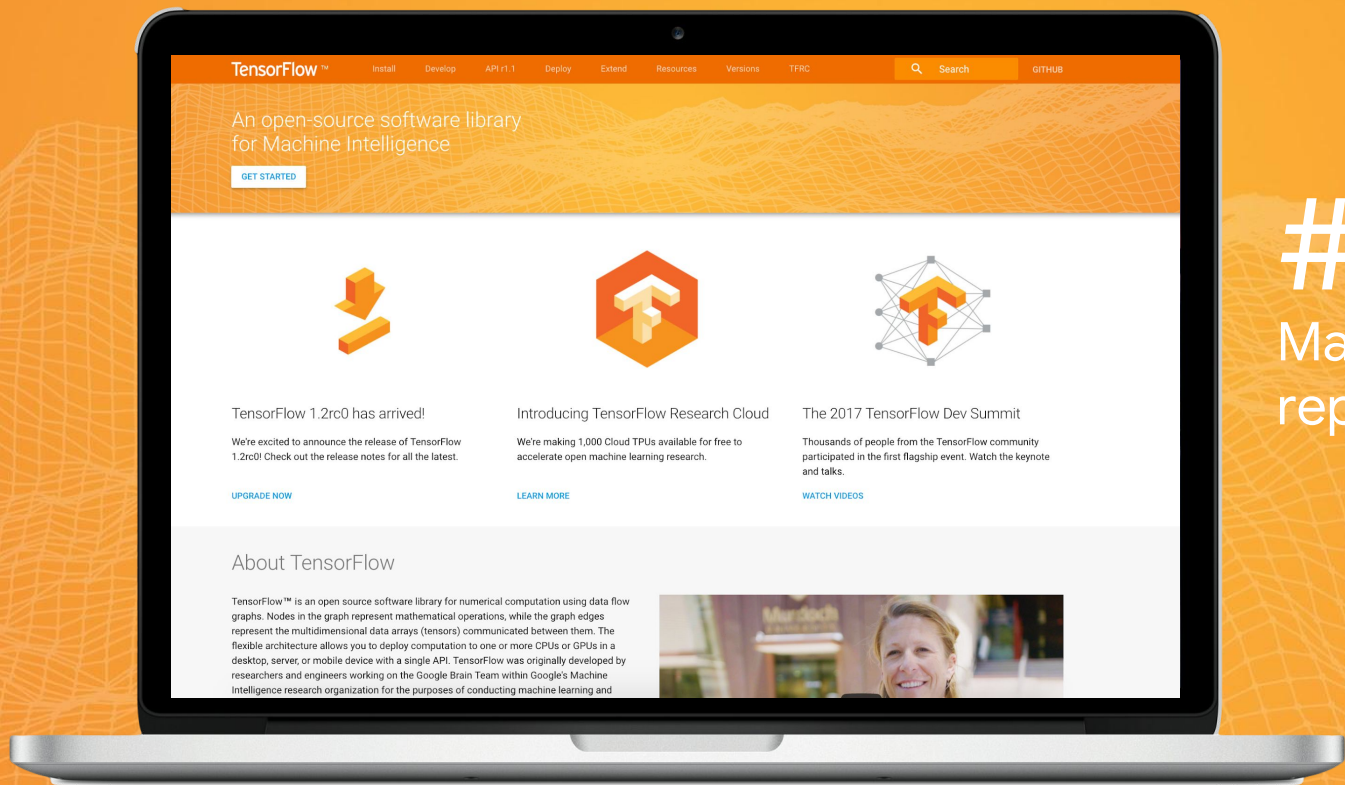
Ingredients for successful ML projects



THE BATTLE FOR TRAINING DATA



CAN BE NASTY



#1


Machine Learning repository on GitHub



“If everyone spoke
to their phone for
three minutes, we’d
exhaust all available
computing resources”

*Jeff Dean, Google Senior Fellow
2014*

Tensor Processing Unit v2



ASIC for TensorFlow
Designed by Google

24 hours on 32 GPUs

6 hours on $\frac{1}{8}$ of a TPU Pod

ML Services spectrum on GCP

Ready to Go

our data + our models

App developers

Total Control

your data + your model

Data Scientists & ML practitioners



Cloud
Vision API



Cloud
Speech API



Cloud
Jobs API



Cloud
Translation API



Cloud Natural
Language API



Cloud Video
Intelligence API



TensorFlow



Cloud Machine
Learning Engine

State of the Industry: Lack of Expertise

Very few users today
can create a custom ML model



State of the Industry: Complex & Time Intensive



Large computational resource

Machine learning expertise

Manual data labeling

Introducing Cloud AutoML

A technology that can automatically create a Machine Learning Model



**DATA
PREPROCESSING**

**ML MODEL
DESIGN**

**TUNE ML MODEL
PARAMETERS**

EVALUATE

DEPLOY

UPDATE

Demo



Demo

10+ kinds of clouds



Let's try the Vision API



| | |
|------------|-----|
| Sky | 99% |
| Cloud | 98% |
| Daytime | 97% |
| Blue | 97% |
| Atmosphere | 93% |

Cloud AutoML Vision

Upload and label images

Train your model

Evaluate



Handbag



Shoe



Hat



Model is now trained and ready to make prediction.
This model can scale as needed to adapt to customer demands.

Dataset: clouds

IMPORT

LABEL

TRAIN

EVALUATE

PREDICT

EXPORT

Import training data

[Ways to format your CSV file](#) ▾

[Known limitations](#) ▾

[Troubleshooting issues](#) ▾

☒ Import your training images from GCS [?]

The CSV file path and all image paths must start with "gs://hammers-it-vcm/".

[Past CSV imports](#)

gs://hammers-it-vcm/path/to/file.csv

☐ Upload training CSV from your computer

Select a CSV file that contains a list of gs:// paths to your images.

All image paths within your CSV must start with "gs://hammers-it-vcm/".

SELECT FILE

Dataset: clouds

IMPORT

LABEL

TRAIN

EVALUATE

PREDICT

EXPORT

Define labels



Apply labels to images



Use human labeling service to label more images




Review image labels



Label statistics



Number of images for each label



Dataset: clouds

IMPORT

LABEL

TRAIN

EVALUATE

PREDICT

EXPORT

Define labels



Apply labels to images



Use human labeling service to label more images



Review image labels



Filter by: --any-label--



Last model trained for the dataset: clouds_advanced

Train a new model for this dataset ?

- ☒ **Base model (free) about 1 hour train time**
- ☐ **Advanced model (\$550) up to 24 hours train time.**

* The minimum number of images required for training the advanced model is higher.

Requirements:



Between 2 and 100 labels.
Currently there are 10 labels.




No more than 100,000 labeled images total.
Currently there are 1,614 labeled images.



At least 10 example images per label.
Currently there are 0 labels with fewer than 10 example images.

[Name your model](#)

clouds_201804110932_base

Previously trained models can be found in  Models section

Confusion matrix



| | Predicted label | | | | | | | | | |
|---------------|-----------------|---------|---------------|---------|--------|-------------|--------------|--------------|--------------|-------------|
| | nimbostratus | stratus | stratocumulus | cumulus | cirrus | altocumulus | cumulonimbus | cirrostratus | cirrocumulus | altostratus |
| nimbostratus | 25% | 37% | 12% | 12% | 0% | 12% | 0% | 0% | 0% | 0% |
| stratus | 50% | 0% | 0% | 0% | 0% | 0% | 50% | 0% | 0% | 0% |
| stratocumulus | 0% | 0% | 42% | 28% | 14% | 14% | 0% | 0% | 0% | 0% |
| cumulus | 0% | 0% | 0% | 85% | 0% | 4% | 9% | 0% | 0% | 0% |
| cirrus | 2% | 0% | 0% | 0% | 97% | 0% | 0% | 0% | 0% | 0% |
| altocumulus | 0% | 0% | 0% | 16% | 16% | 33% | 0% | 0% | 33% | 0% |
| cumulonimbus | 0% | 0% | 0% | 0% | 0% | 0% | 100% | 0% | 0% | 0% |
| cirrostratus | 0% | 0% | 0% | 0% | 50% | 0% | 0% | 50% | 0% | 0% |
| cirrocumulus | 0% | 16% | 0% | 0% | 0% | 0% | 0% | 16% | 50% | 16% |
| altostratus | 0% | 0% | 7% | 0% | 14% | 7% | 7% | 0% | 0% | 64% |

Dataset: clouds

IMPORT

LABEL

TRAIN

EVALUATE

PREDICT

EXPORT

Try Google model serving for free



You can test your model with 100,000 free predictions. To increase number of queries, please contact sales: automl-vision-feedback@google.com

Use model:

Query online



Query Cloud ML Engine from the command line



Query Cloud ML Engine from Python



Query Vision API



1614 out of 1614 images are labeled

Important: During the alpha period, there may be a warmup delay with your model. If your initial requests return an error, please wait a few seconds before trying again.



0.950 cirrus


Select up to 10 images to make predictions on.

Choose Files No file chosen


Cloud AutoML - Best in Class Research

Learning to learn


Transfer Learning




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.

Large-Scale Evolution of Image Classifiers

Esteban Real, Sherry Moore, Andrew Selle, Saurabh Saxena, Yutaka Leon Suematsu, Quoc Le, Alex Kurakin.
International Conference on Machine Learning, 2017.

Bayesian Optimization for a Better Dessert

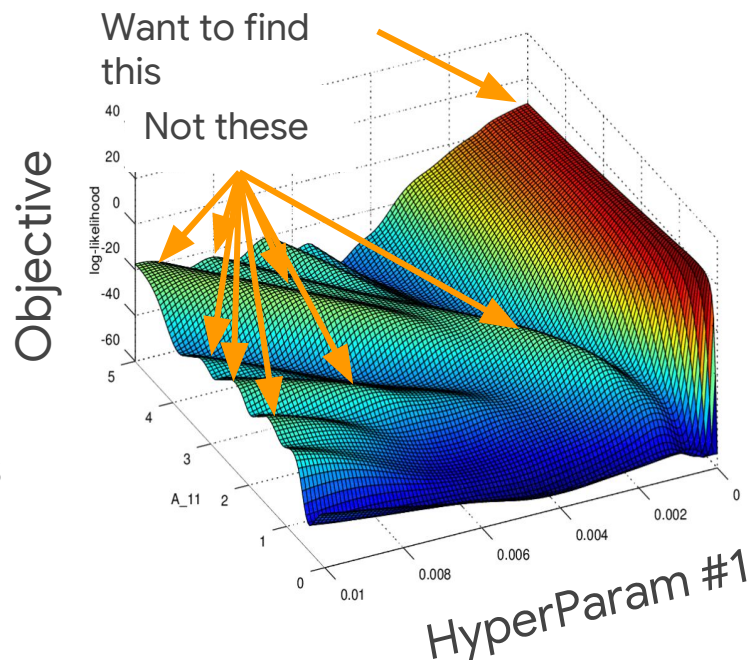
Benjamin Solnik, Daniel Golovin, Greg Kochanski, John Elliot Karr

Hyperparameter tuning with HyperTune

Automatic **hyperparameter tuning**

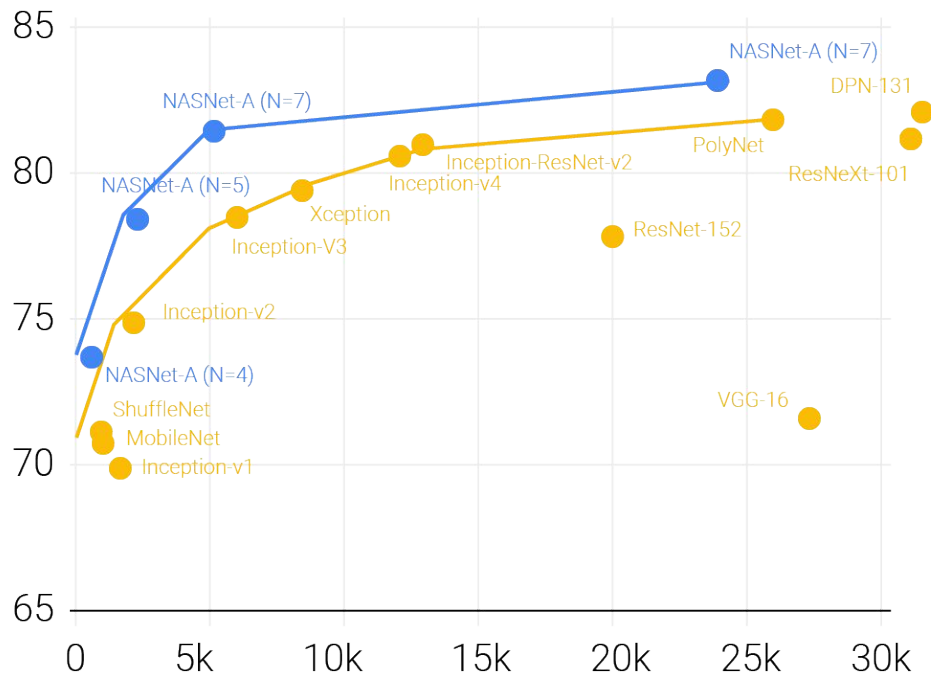
Runs multiple trials in a training with specified HP metrics

Gaussian process to optimize the trials



ML does ML

Systematic exploration of the model space, using the techniques finessed in AlphaGo, yields super-human performance in ML network design



Wrap up

Limited ML expertise needed

Unlike all other ML services, you need limited ML expertise to train a machine learning model.

Your own custom models

Easily create and customize your own state-of-the-art ML models for your unique use case.

Simple

Simple graphical user interface (GUI) where you just have to specify content, and Cloud AutoML turn content into a high quality model from your dataset.

High Quality

State-of-the-art performance on public benchmark CIFAR and ImageNet



Thank you!

Google Cloud