Harvard

Yá'át'ééh 👋

EASI-22

Edge AI Summer Institute 2022

with Navajo Tech

IEEE COMPUTER SOCIETY

Hi! I'm Brian!

I'm an Assistant Professor of Computer Science
at **Barnard College, Columbia University**

The Barnard Accessible and
Accelerated Robotics Lab

Our **team!**






with help from **many more**

VE RI TAS

B

NAVAJO TECHNICAL UNIVERSITY

TM

Our **website**!

tinyMLedu.org/EASI-22

home base for **all information**!

This workshop is based on materials from the TinyMLedu initiative. To learn more about TinyMLedu check out **their website**!

EASI-22

Edge AI Summer Institute
July 19-21 2022

Home
Schedule and Materials
Student Application
Educator Application
Apply by July 1st
Team
Workshop Flyer

Updated: 7/22
by @plancherb1

Schedule and Materials

The workshop will be held on Zoom.

The workshop will run each day from **12:00 PM to 3:00 PM MDT (New Mexico Time) which is 2:00 PM to 5:00 PM in your local timezone** (according to your computer system time). Times below adjusted to that time zone. Exact timing and topics subject to change.

Day	Date	Topics	Speakers and Materials
Day 1	Tuesday	**Introduction to Artificial Intelligence and (Tiny)ML** *2:00 PM* Conference Opening and Schedule *2:15 PM* Buy2Pay Overview *2:25 PM* Introduction to Artifical Intelligence and Machine Learning *4:45 PM* Day Closing	Brian Plancher of Barnard College, Columbia University and of Harvard University Slides as PDF | As Google Slides Molly Marshall of Harvard University Slides | Buy2Pay Login
Day 2	Wednesday	**Keyword Spotting for the Navajo Language** *2:00 PM* Day Opening *2:10 PM* Keyword Spotting with Convolutional Neural Networks *3:10 PM* Hands-On Lab *4:45 PM* Day Closing	Brian Plancher of Barnard College, Columbia University and of Harvard University
Day 3	Thursday	**Bringing AI/ML from the Cloud to the Edge** *2:00 PM* Day Opening *2:10 PM* Introduction to the Arduino Tiny Machine Learning Kit *3:10 PM* Hands-On Lab *4:15 PM* Roundtable Discussion: Next Steps *4:45 PM* Workshop Closing	Dhilan Ramaprasad of Harvard University

Questions?

Contact easi-staff@googlegroups.com with any questions regarding this workshop.

Supporters

We would like to thank the IEEE CS for their generous support of this program through the IEEE CS Diversity & Inclusion Fund.

Make Sure to Pick Up an Arduino Kit!

Question? Contact:
Monsuru Ramoni
mramoni@navajotech.edu

Teachers Sign up for Buy2Pay

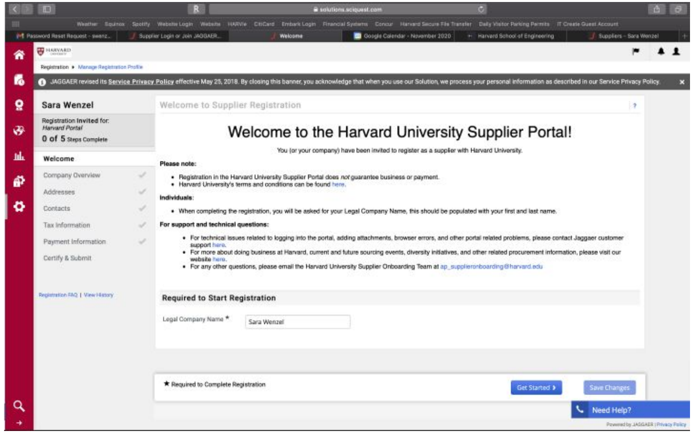

Question? Contact:
Molly Marshall
mmarshall@seas.harvard.edu

Workshop **Agenda**

Day 1	Introduction to AI and (Tiny)ML	Cloud ML
Day 2	Keyword Spotting for the Navajo Language	Mobile ML
Day 3	Bringing AI/ML from the Cloud to the Edge	Embedded ML

Today's Agenda

- What is Artificial Intelligence?
- Hands-on: AutoDraw
- What is (Deep) Machine Learning?
- Hands-on: ThingTranslator
- What is Responsible TinyML?
- Summary

Today's Agenda

- **What is Artificial Intelligence?**
- Hands-on: AutoDraw
- What is (Deep) Machine Learning?
- Hands-on: ThingTranslator
- What is Responsible TinyML?
- Summary

Artificial Intelligence (AI) is when a computer can...

Think Like A Human	**Think Rationally**
Act Like a Human	**Act Rationally**

The Turing Test

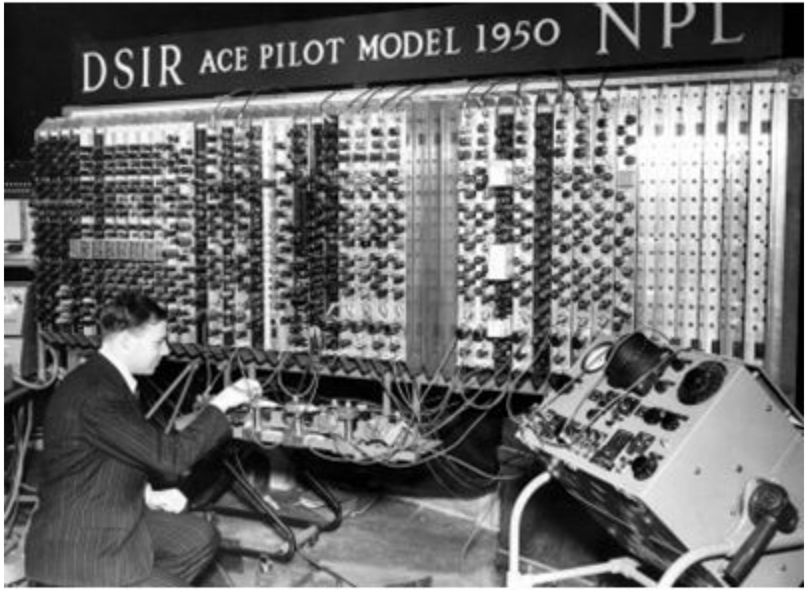

The Turing Test

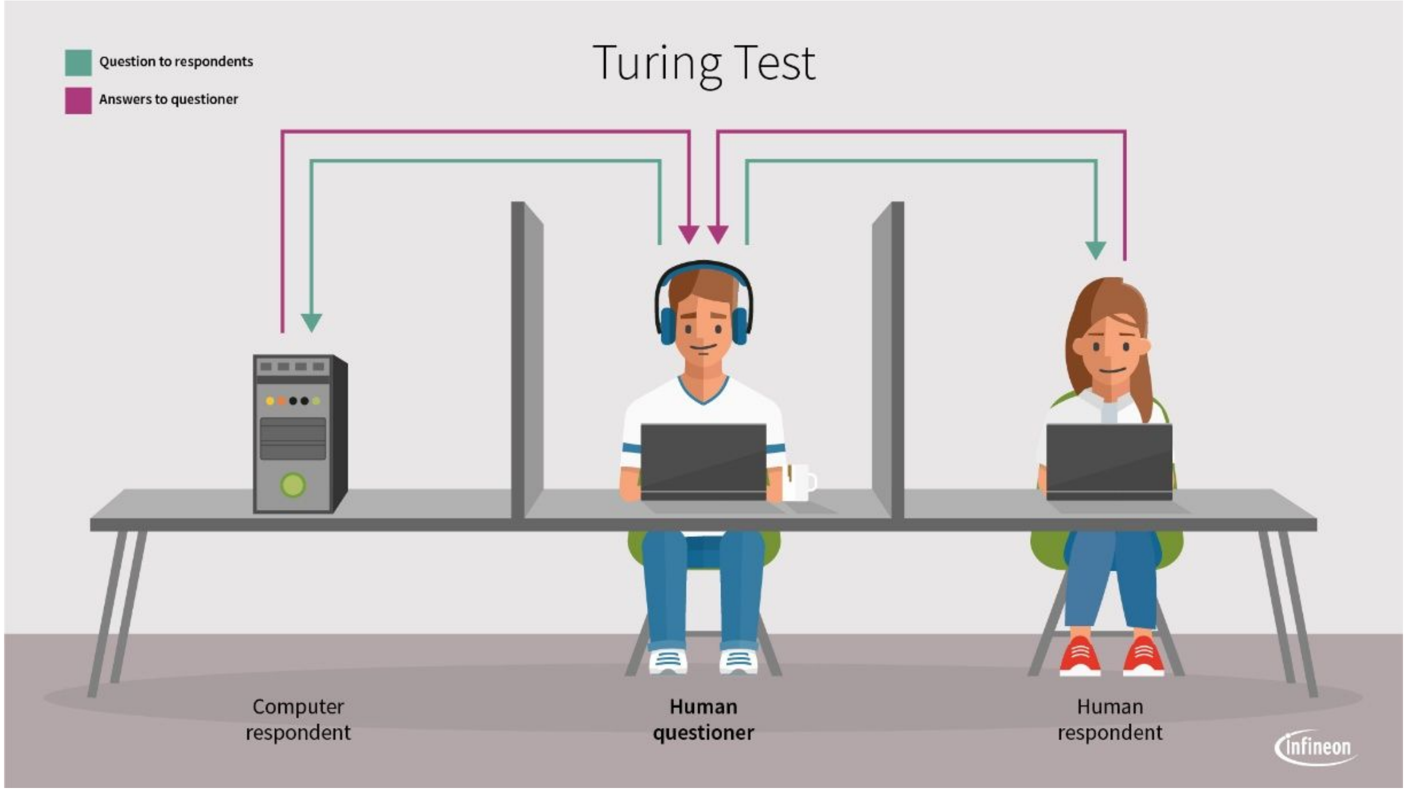

The Turing Test

Chatbots can be convincing

```
Q: What is your favorite animal?
A: My favorite animal is a dog.

Q: Why?
A: Because dogs are loyal and friendly.

Q: What are two reasons that a dog might be in a bad mood?
A: Two reasons that a dog might be in a bad mood are if it is
hungry or if it is hot.

Q: How many eyes does a giraffe have?
A: A giraffe has two eyes.

Q: How many legs does a frog have?
A: A frog has four legs.

Q: Are there any animals with three legs?
A: No, there are no animals with three legs.

Q: Why don't animals have three legs?
A: Animals don't have three legs because they would fall over.
```

https://lacker.io/ai/2020/07/06/giving-gpt-3-a-turing-test.html

The Turing Test

Chatbots can be convincing

```
Q: What is your favorite animal?
A: My favorite animal is a dog.

Q: Why?
A: Because dogs are loyal and friendly.

Q: What are two reasons that a dog might be in a bad mood?
A: Two reasons that a dog might be in a bad mood are if it is
hungry or if it is hot.

Q: How many eyes does a giraffe have?
A: A giraffe has two eyes.

Q: How many legs does a frog have?
A: A frog has four legs.

Q: Are there any animals with three legs?
A: No, there are no animals with three legs.

Q: Why don't animals have three legs?
A: Animals don't have three legs because they would fall over.
```

https://lacker.io/ai/2020/07/06/giving-gpt-3-a-turing-test.html

The Turing Test

But they only know "what they know"

```
Q: How many eyes does a giraffe have?
A: A giraffe has two eyes.

Q: How many eyes does my foot have?
A: Your foot has two eyes.

Q: How many eyes does a spider have?
A: A spider has eight eyes.

Q: How many eyes does the sun have?
A: The sun has one eye.

Q: How many eyes does a blade of grass have?
A: A blade of grass has one eye.
```

https://lacker.io/ai/2020/07/06/giving-gpt-3-a-turing-test.html

Traditional AI Programming

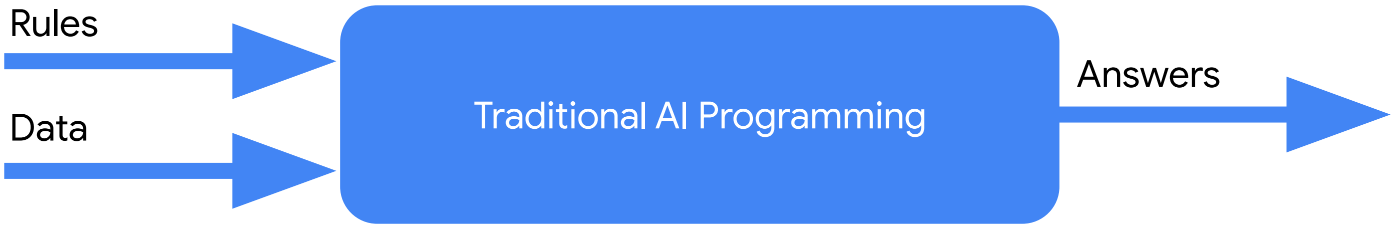

Traditional AI Programming

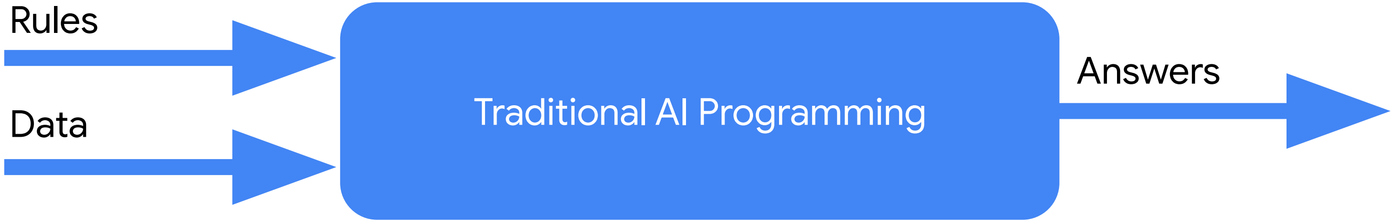

AI is the study of **algorithms** that can give computers the rules they need to be **“intelligent”!**

Traditional AI Programming

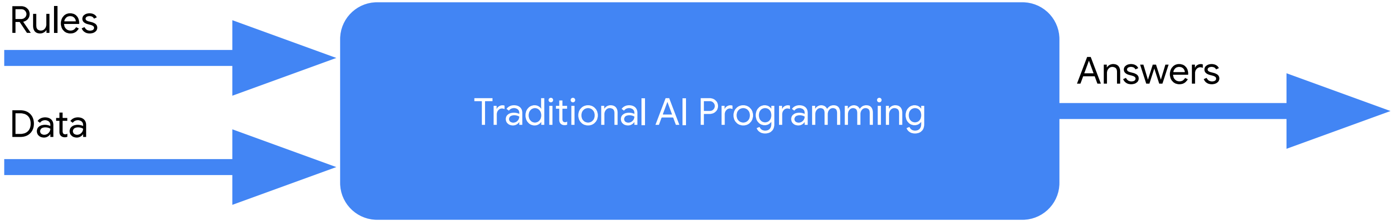

```
if(speed<4){
    status=WALKING;
}
```

Traditional AI Programming

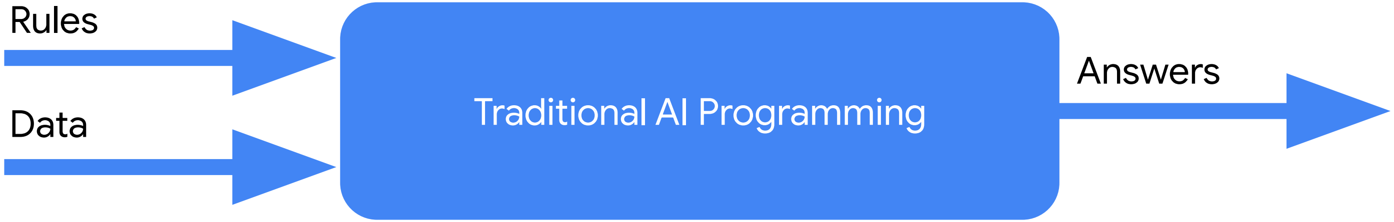

```
if(speed<4){
     status=WALKING;
}
```

```
if(speed<4){
     status=WALKING;
} else {
     status=RUNNING;
}
```

Traditional AI Programming

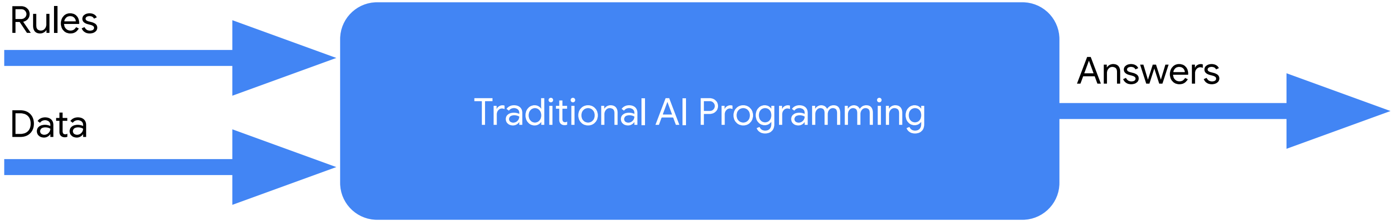

```
if(speed<4){
    status=WALKING;
}
```

```
if(speed<4){
    status=WALKING;
} else {
    status=RUNNING;
}
```

```
if(speed<4){
    status=WALKING;
} else if(speed<12){
    status=RUNNING;
} else {
    status=BIKING;
}
```

Deep Blue

Deep Blue
IBM chess computer

Garry Kasparov
World Chess Champion

Deep Blue

On average in any board configuration there are **35** possible moves in chess.

Deep Blue

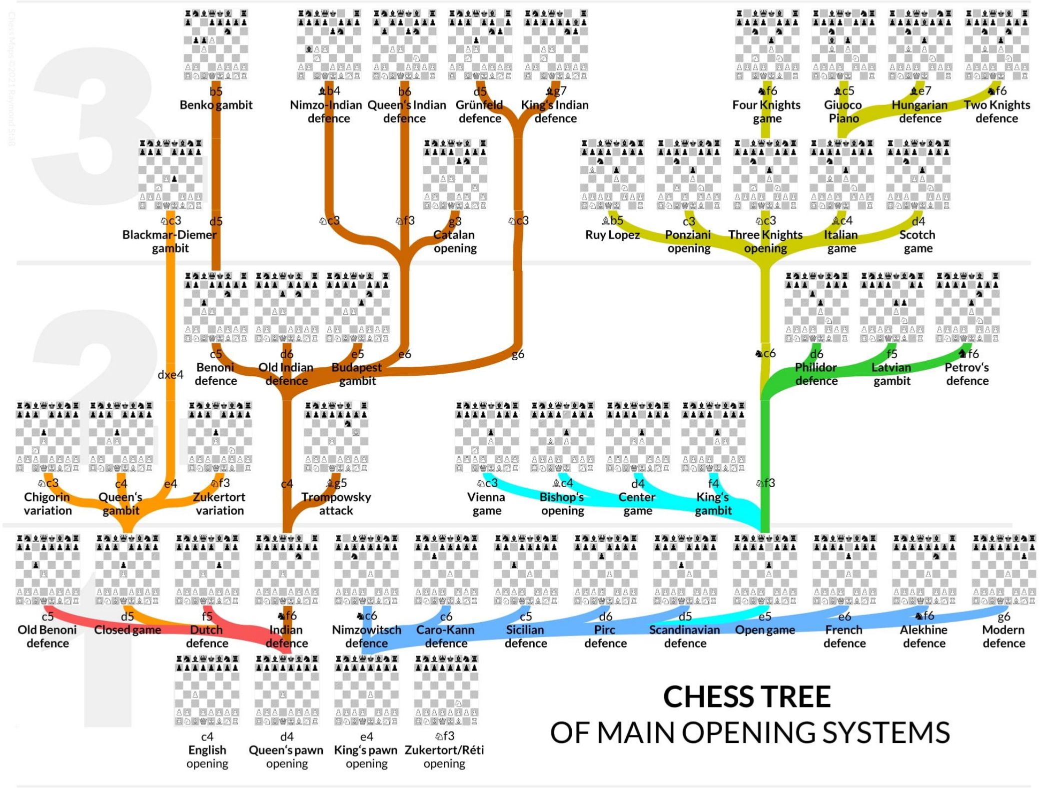

On average in any board configuration there are **35** possible moves in chess.

That means that the computer can search for the move that eventually leads to success with a **35^{turns}** guesses

Deep Blue

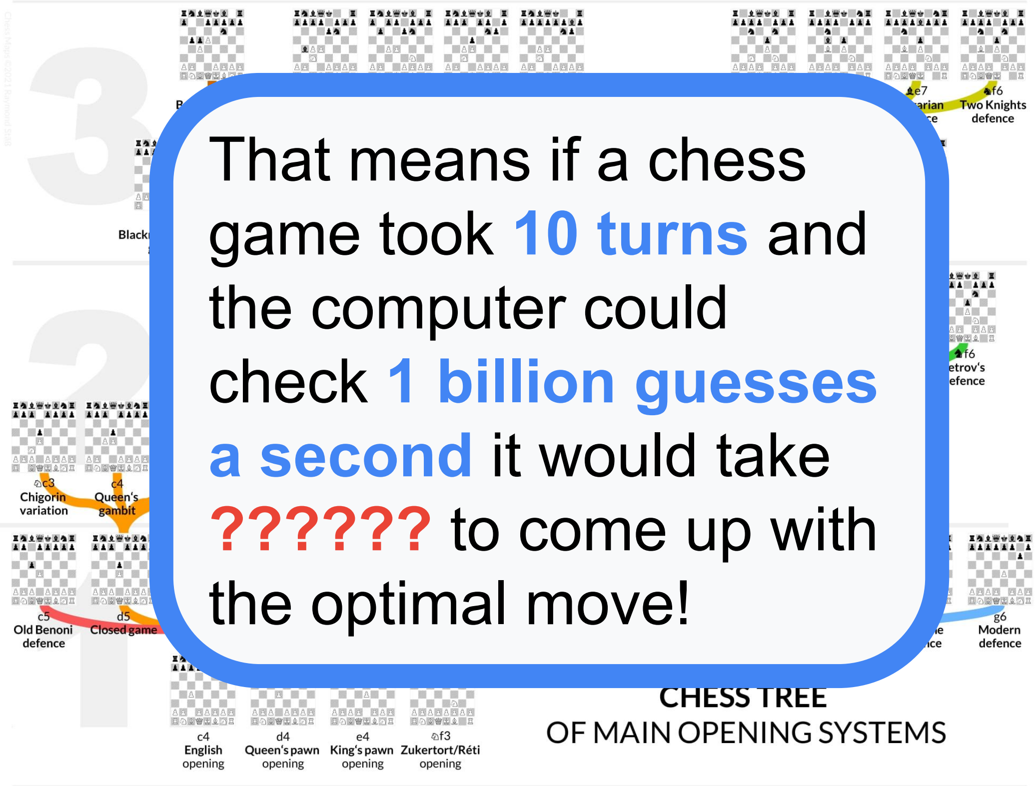

That means if a chess game took **10 turns** and the computer could check **1 billion guesses a second** it would take **??????** to come up with the optimal move!

On average in any board configuration there are **35** possible moves in chess.

That means that the computer can search for the move that eventually leads to success with a 35^{turns} guesses

Deep Blue

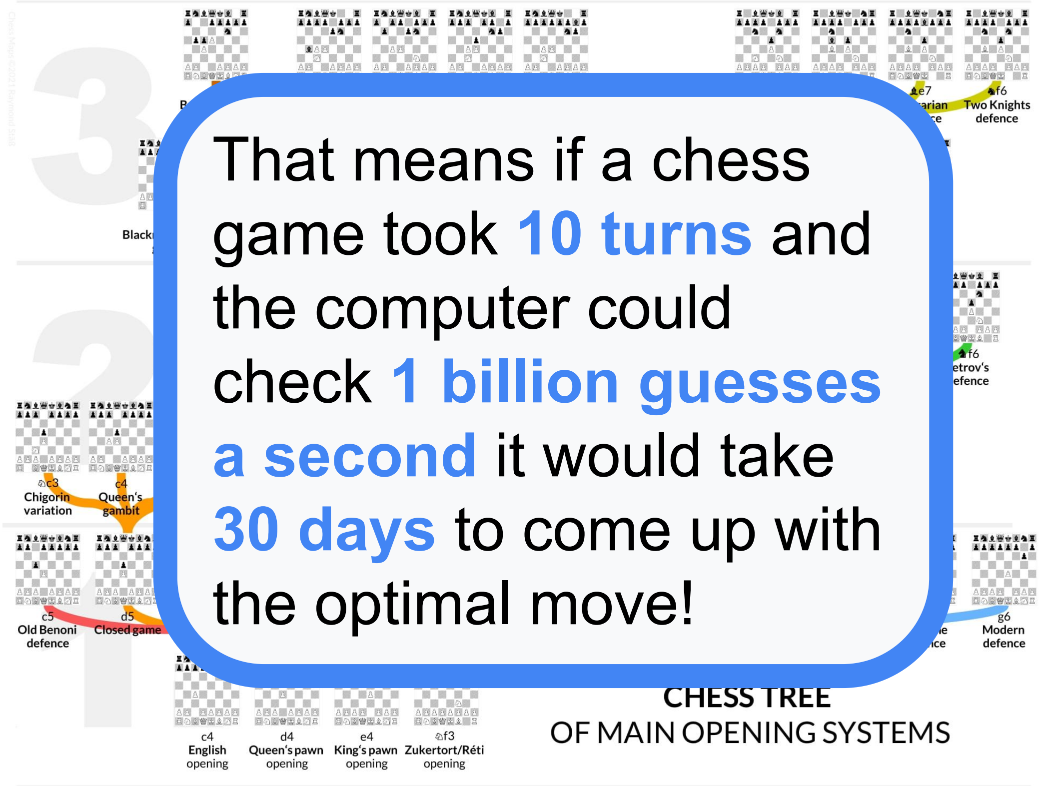

That means if a chess game took **10 turns** and the computer could check **1 billion guesses a second** it would take **30 days** to come up with the optimal move!

On average in any board configuration there are **35** possible moves in chess.

That means that the computer can search for the move that eventually leads to success with a 35^{turns} guesses

Deep Blue

Deep Blue
IBM chess computer

Garry Kasparov
World Chess Champion

So deep blue searched ~7 turns ahead and relied on a **board scoring rule** created by the programmers!

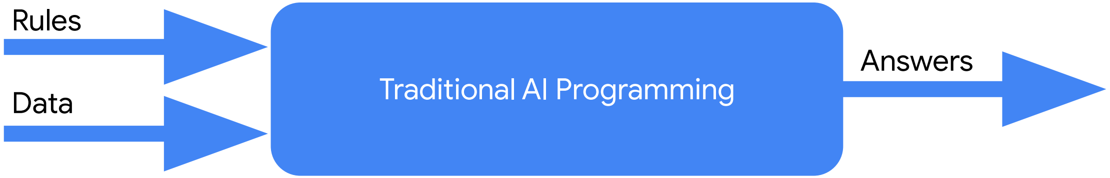

Today's Agenda

- **What is Artificial Intelligence?**
- Hands-on: AutoDraw
- What is (Deep) Machine Learning?
- Hands-on: ThingTranslator
- What is Responsible TinyML?
- Summary

Today's Agenda

- What is Artificial Intelligence?
- **Hands-on: AutoDraw**
- What is (Deep) Machine Learning?
- Hands-on: ThingTranslator
- What is Responsible TinyML?
- Summary

Experiments with **Google**

This is an

A.I.
Experiment

Fill in the blank

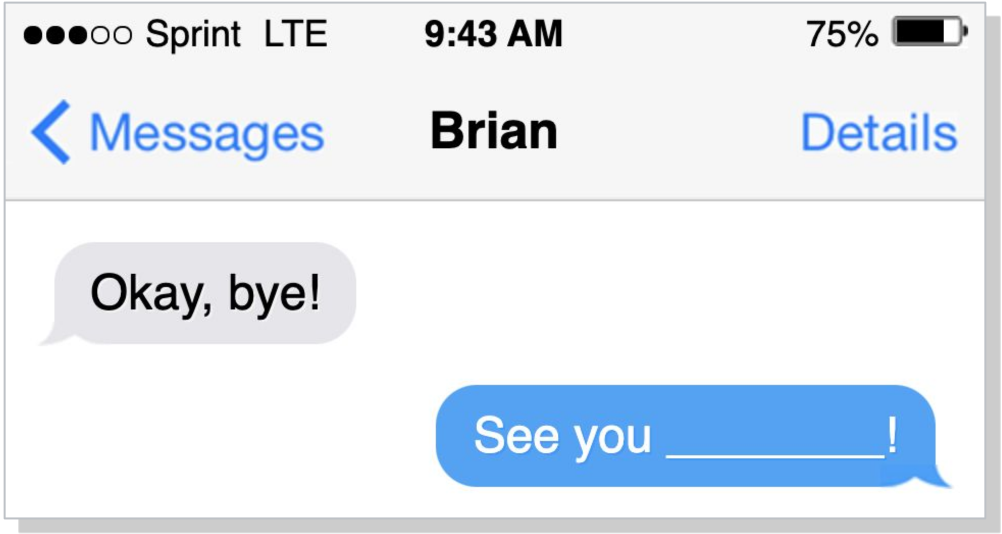

Fill in the blank

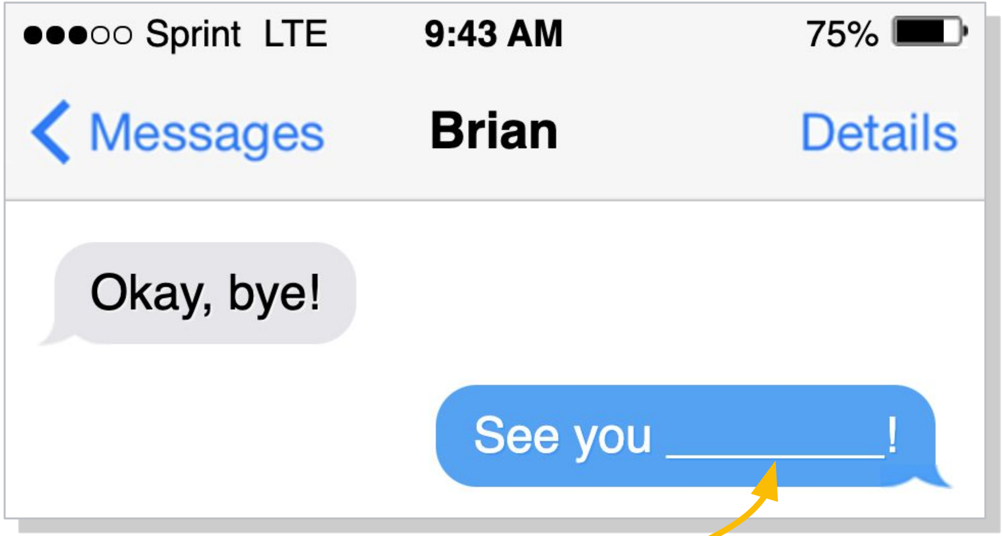

tomorrow
later

Prediction: autocomplete

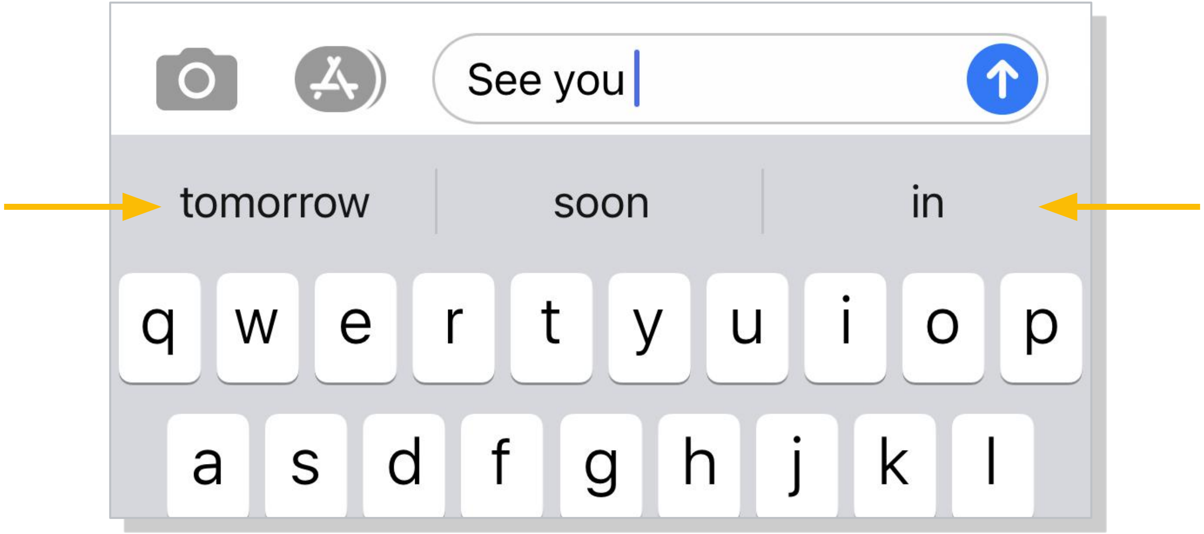

AutoDraw

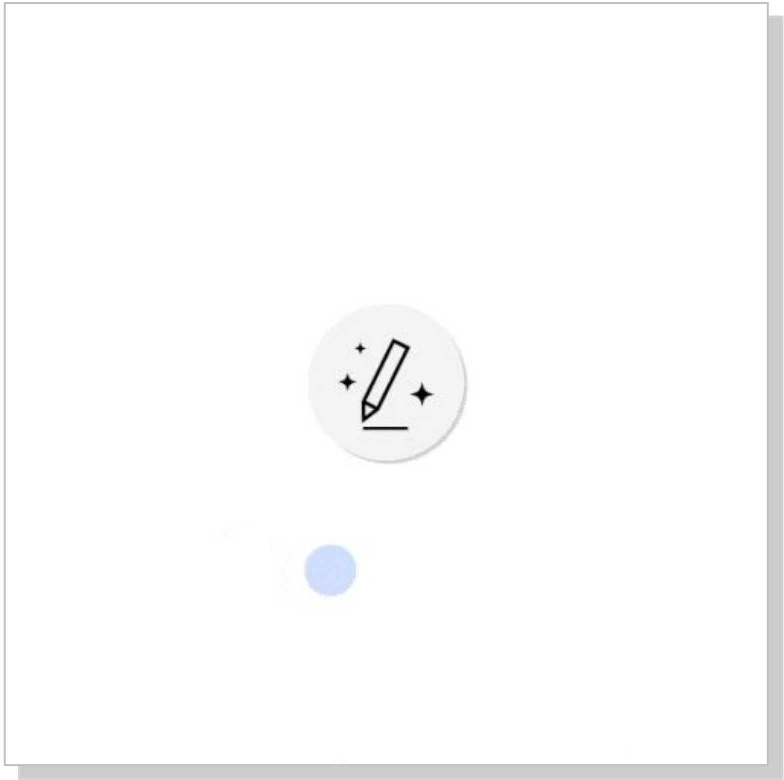

AutoDraw

autodraw.com

Discussion **Groups**

1. Do you think the AI did a **good job**? 👍 / 👎
2. **Why** do you think the AI did (or did not) **work well**?
3. **How** do you think the AI is working to solve this task? 🤔
4. What types of things were **particularly hard or easy** for the AI?

anything **else?**

Today's Agenda

- What is Artificial Intelligence?
- **Hands-on: AutoDraw**
- What is (Deep) Machine Learning?
- Hands-on: ThingTranslator
- What is Responsible TinyML?
- Summary

Today’s Agenda

- What is Artificial Intelligence?
- Hands-on: AutoDraw
- **What is (Deep) Machine Learning?**
- Hands-on: ThingTranslator
- What is Responsible TinyML?
- Summary

Artificial Intelligence

Machine Learning

What's the difference?

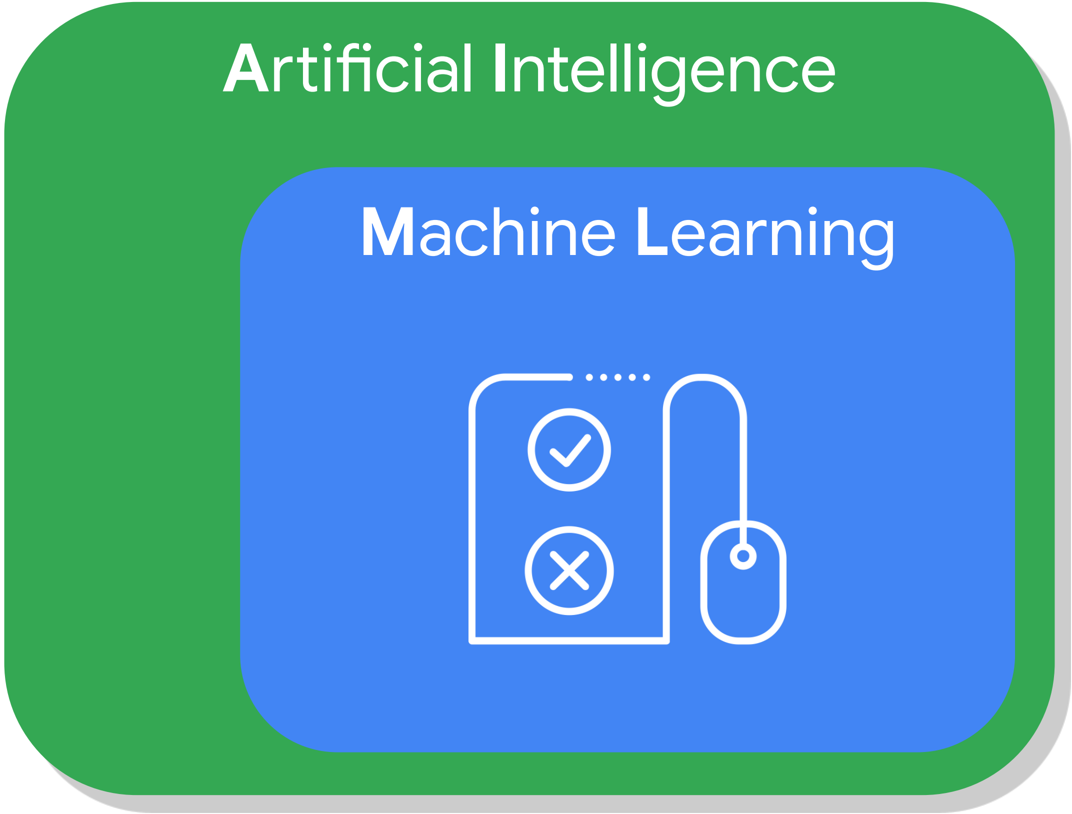
Artificial Intelligence
Machine Learning

Traditional AI Programming

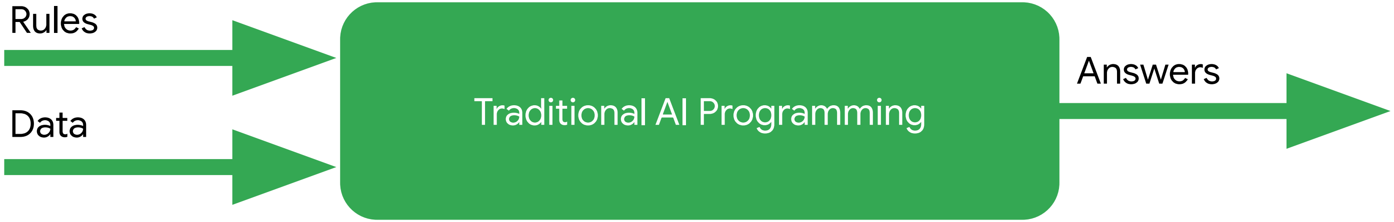

```
if(speed<4){
    status=WALKING;
}
```

```
if(speed<4){
    status=WALKING;
} else {
    status=RUNNING;
}
```

```
if(speed<4){
    status=WALKING;
} else if(speed<12){
    status=RUNNING;
} else {
    status=BIKING;
}
```

Traditional AI Programming

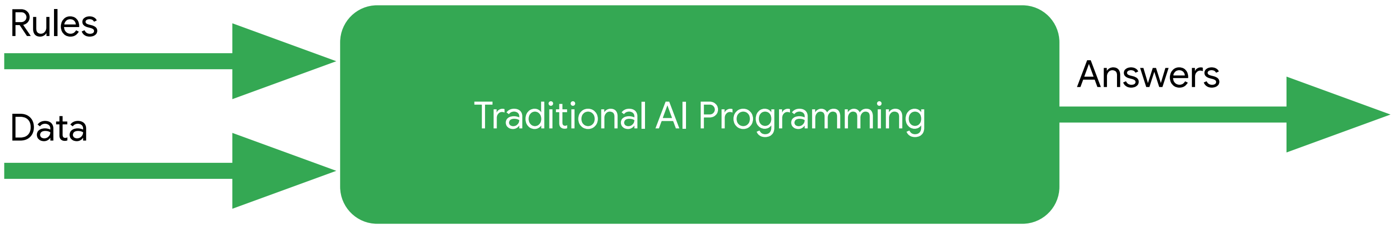

```
if(speed<4){
    status=WALKING;
}
```

```
if(speed<4){
    status=WALKING;
} else {
    status=RUNNING;
}
```

```
if(speed<4){
    status=WALKING;
} else if(speed<12){
    status=RUNNING;
} else {
    status=BIKING;
}
```

```
// ???
```

Categorize

1

2

3

4

5

6

7

8

The Machine Learning Paradigm

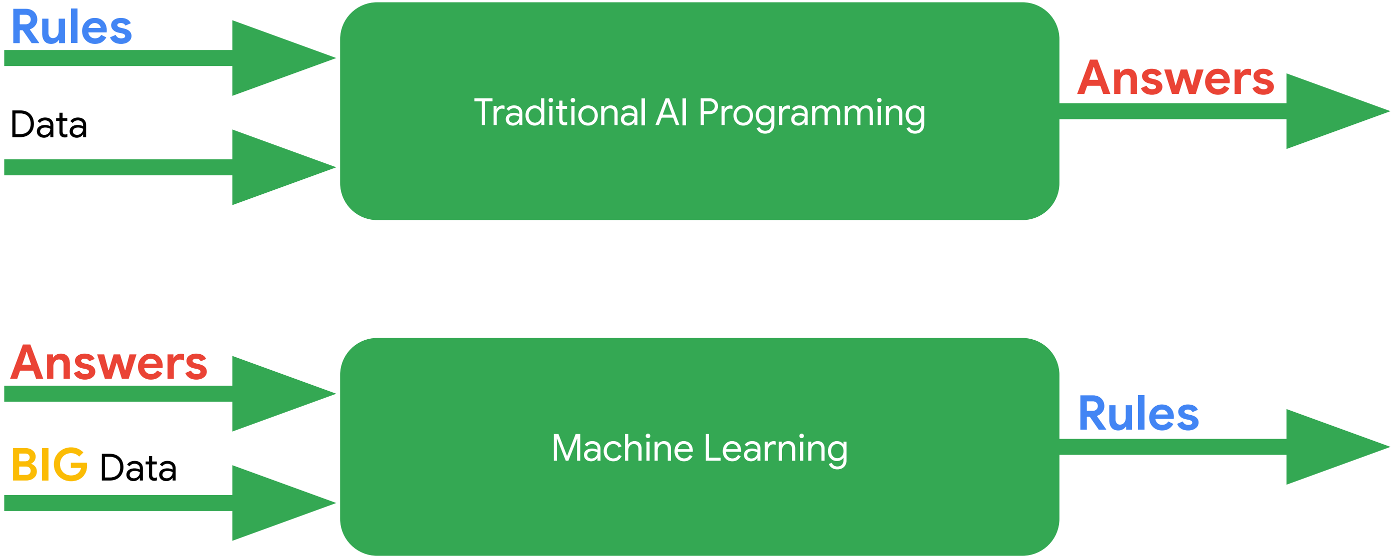

Activity Detection with Machine Learning

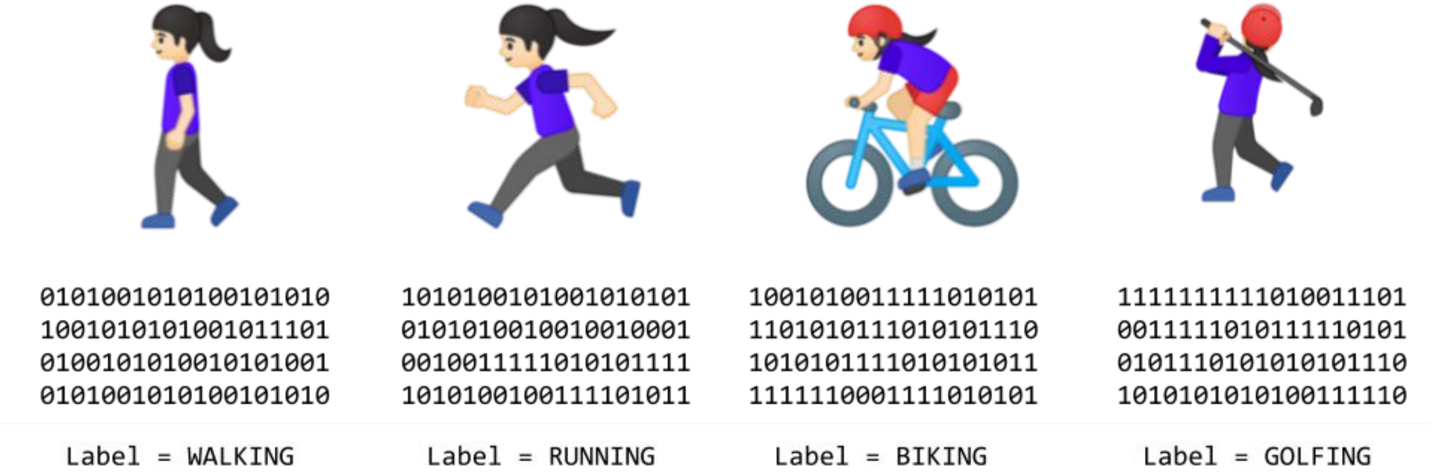

Activity Detection with Machine Learning

0101001010100101010	1010100101001010101	1001010011111010101	1111111111010011101
1001010101001011101	0101010010010010001	1101010111010101110	0011111010111110101
0100101010010101001	0010011111010101111	1010101111010101011	0101110101010101110
0101001010100101010	1010100100111101011	1111110001111010101	1010101010100111110
Label = WALKING	Label = RUNNING	Label = BIKING	Label = GOLFING

Activity Detection with Machine Learning

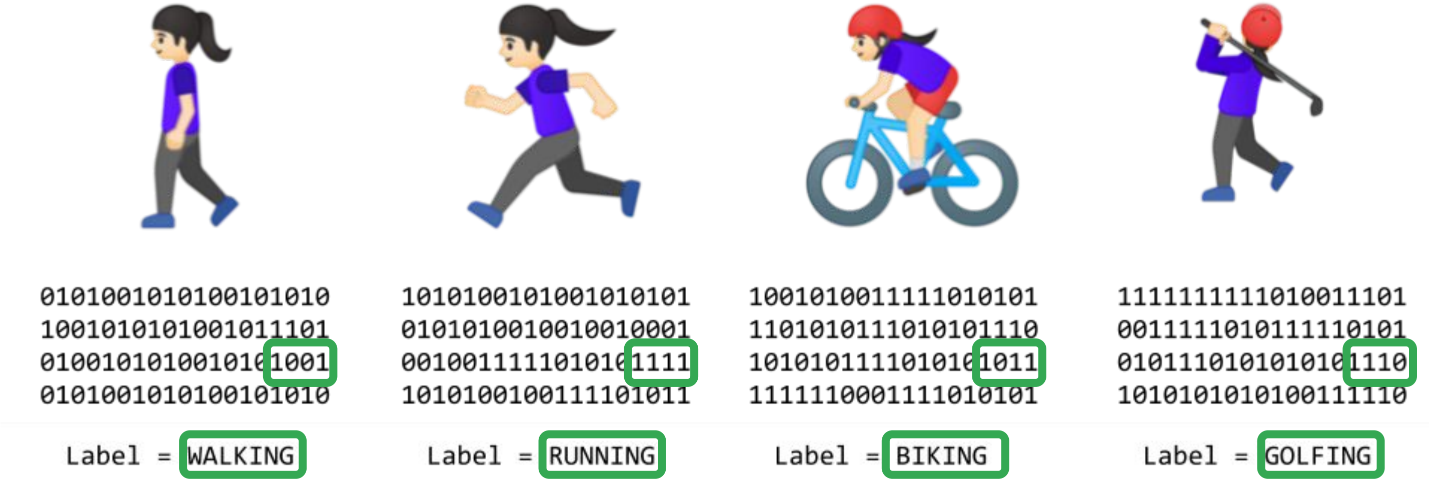

Review what we've learned

Machine learning provides a computer with data, **rather than explicit instructions.** Using these data, the computer learns to **recognize patterns** and becomes able to execute tasks on its own.

Object Detection

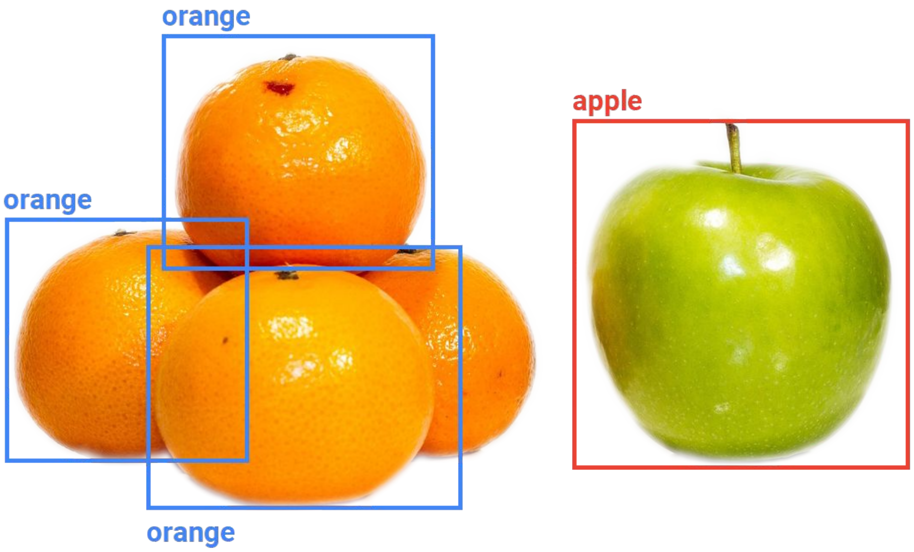

Segmentation

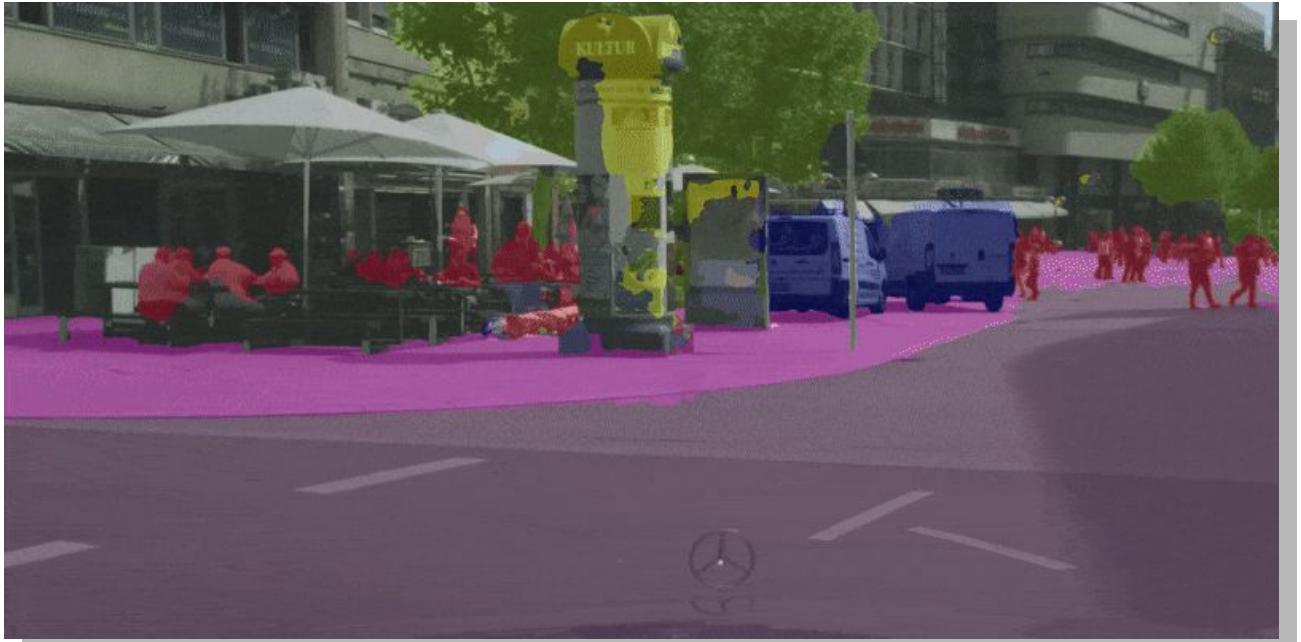

Training the machine

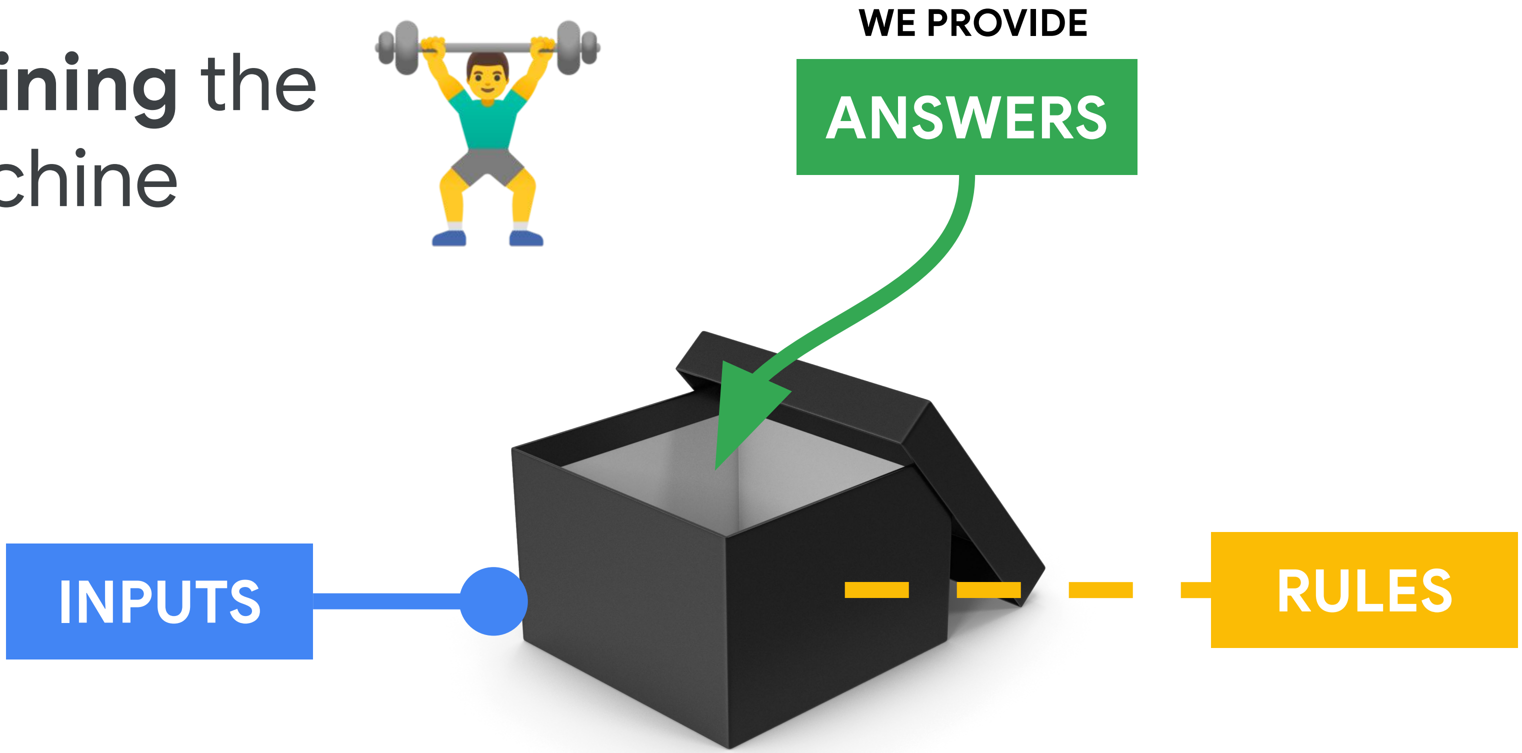

Training the machine

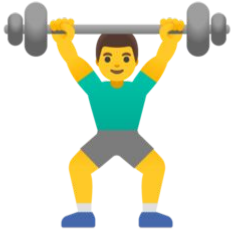

For a set of Input Data

Input, Label

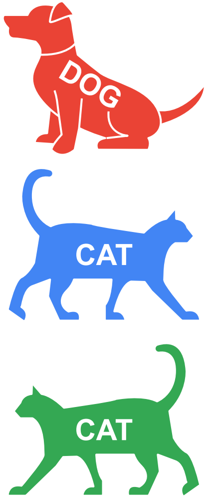

Training the machine

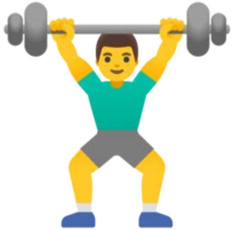

For a set of Input Data

Guess the Answer
and count mistakes

Input, Label	Result
	Dog ✓
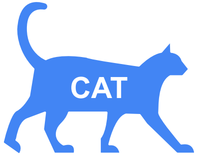	Dog ✗
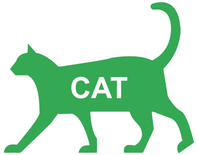	Cat ✓

Training the machine

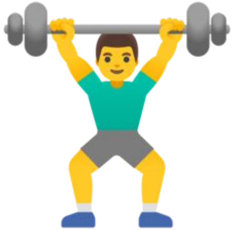

Loss
function of mistakes

Guess the Answer
and count mistakes

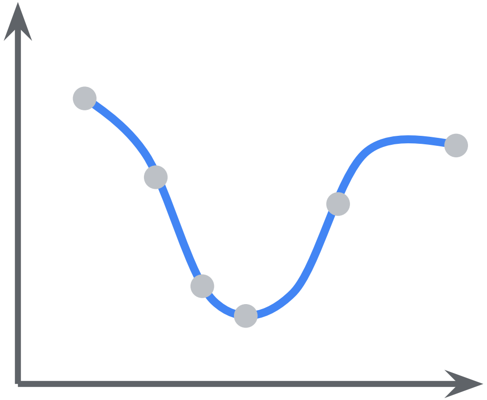

Training the machine

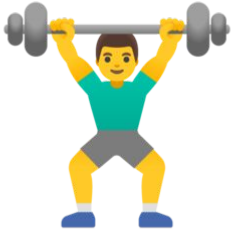

For a set of Input Data

Guess the Answer
and count mistakes

Improve the model to be more correct

Training the machine

Loss

number of mistakes

Improve the model to be more correct

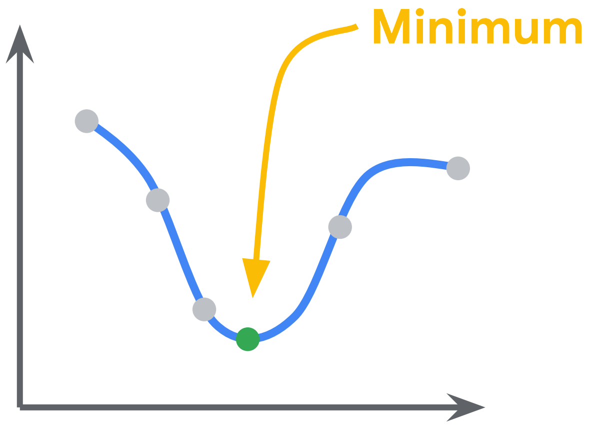

Training the machine

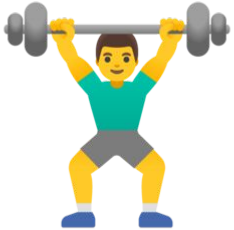

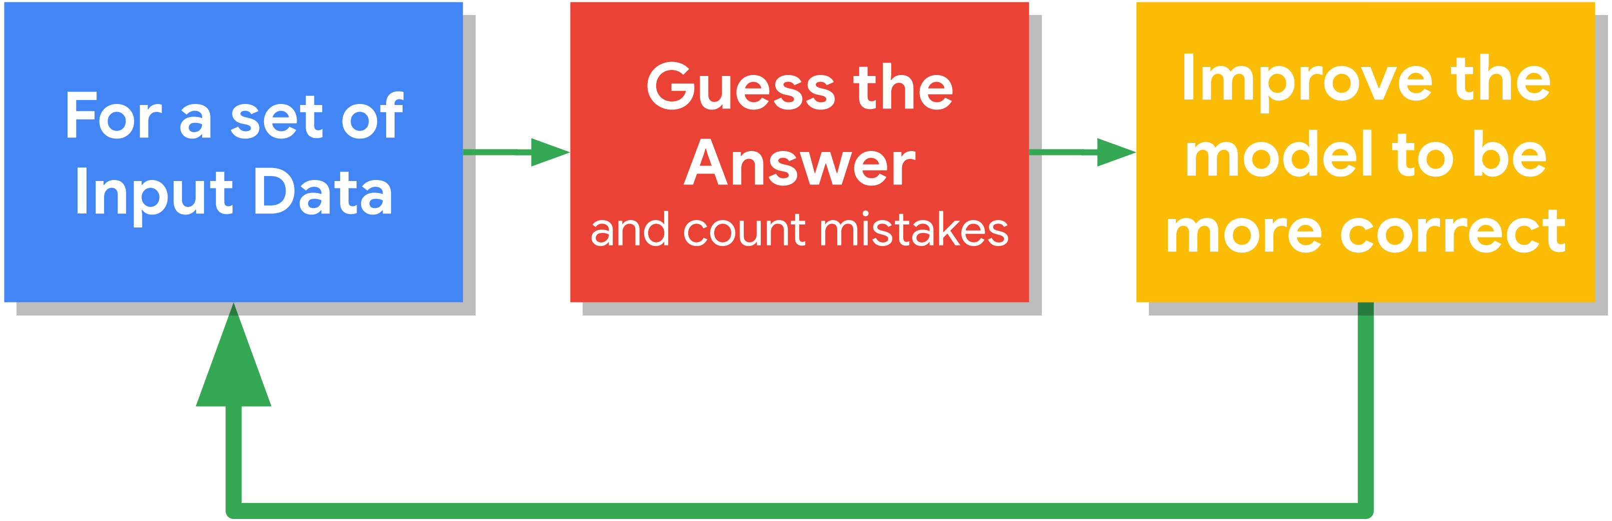

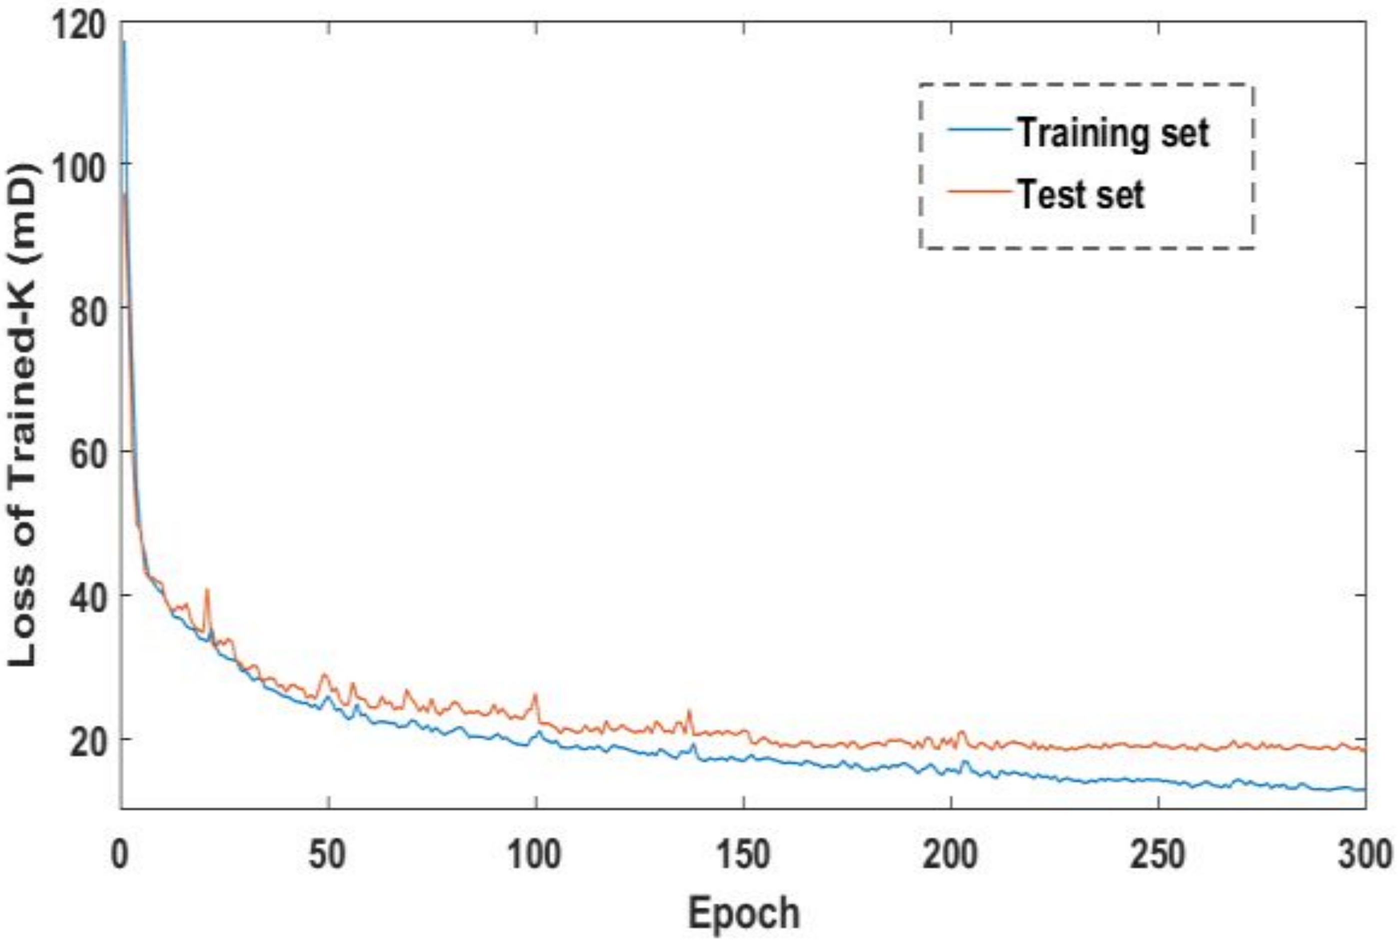

https://www.researchgate.net/publication/339545333_The_Connectivity_Evaluation_Among_Wells_in_Reservoir_Utilizing_Machine_Learning_Methods/figures?lo=1

Training the machine

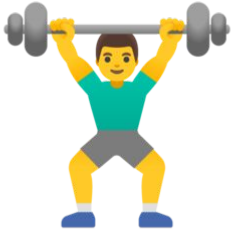

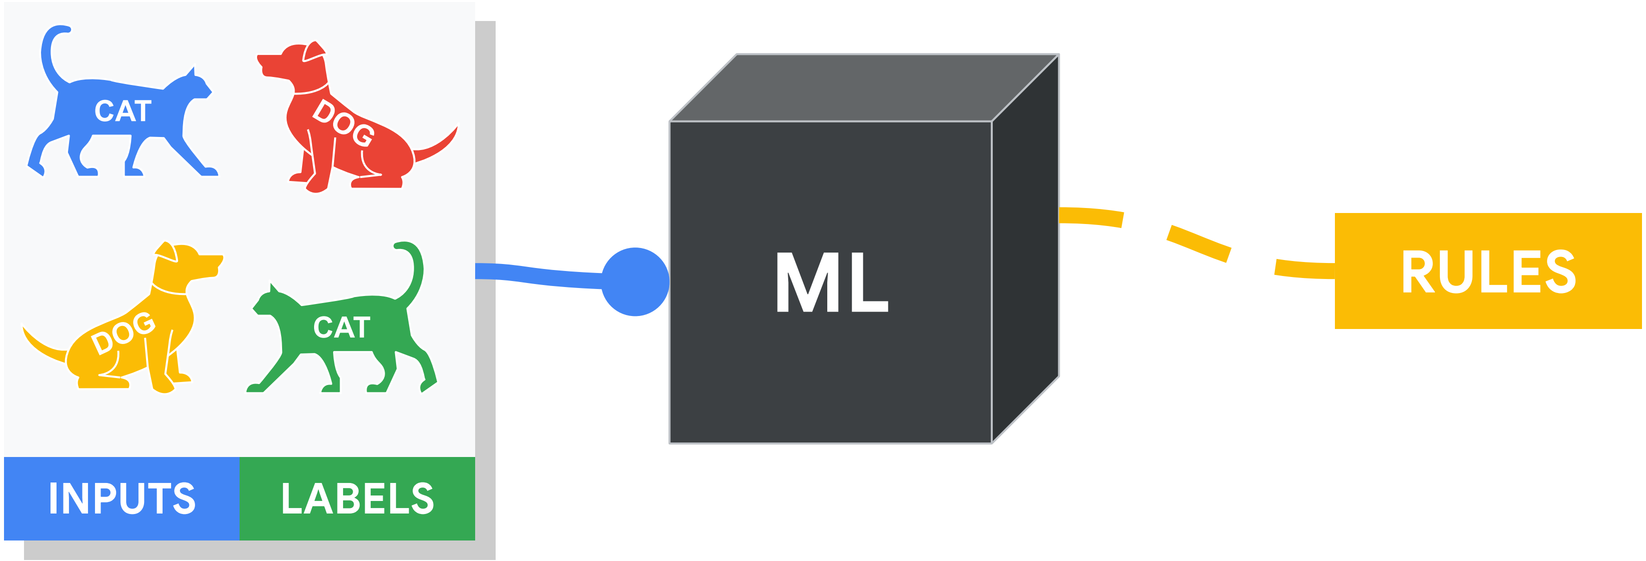

After it's **learned:**

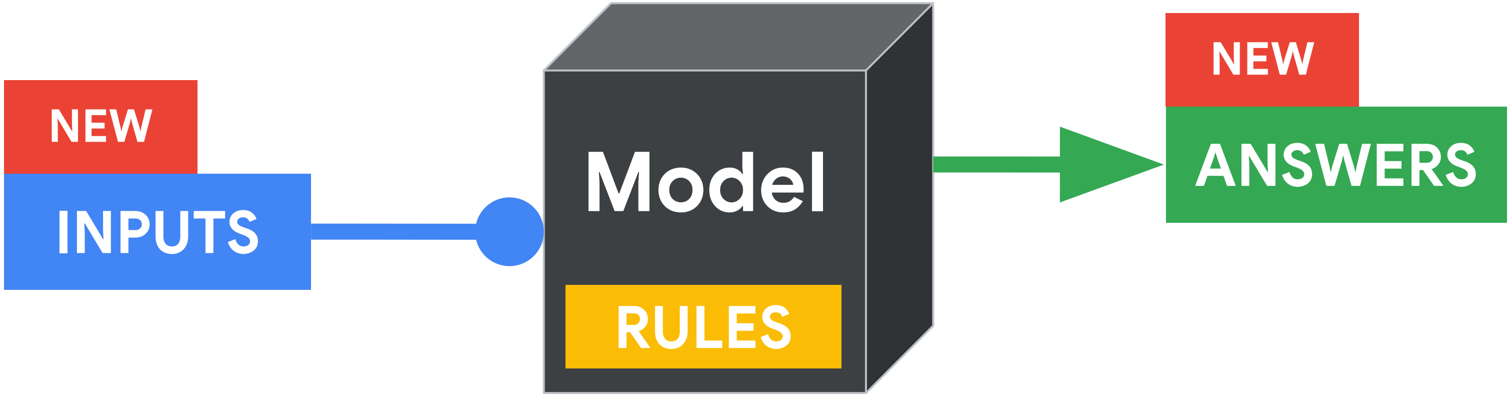

After it's **learned:**

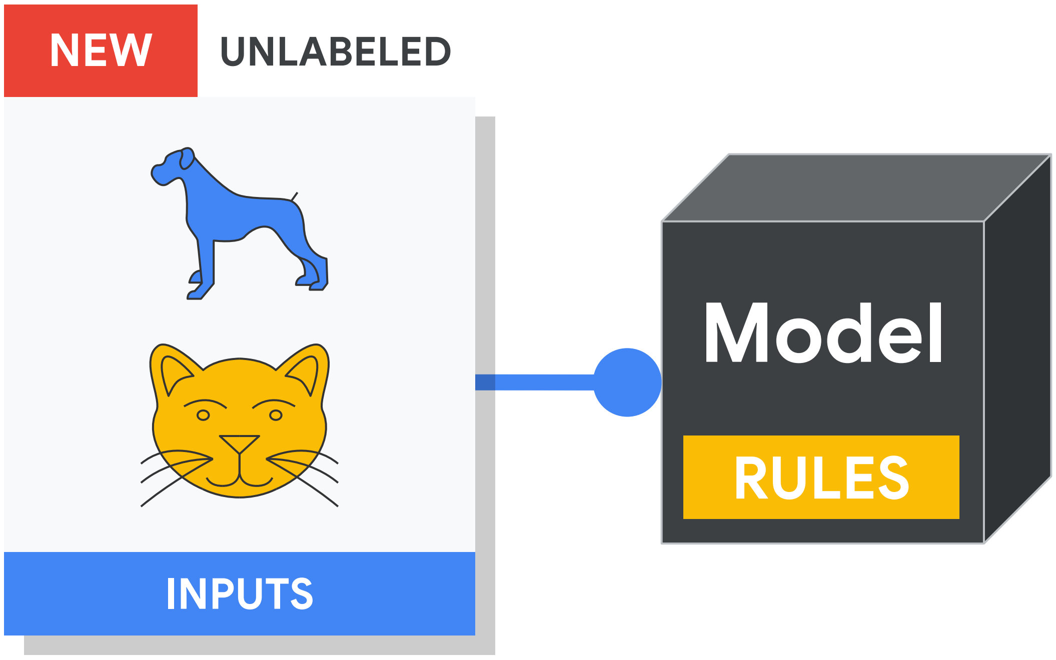

Making **predictions:**

This is often called **INFERENCE**!

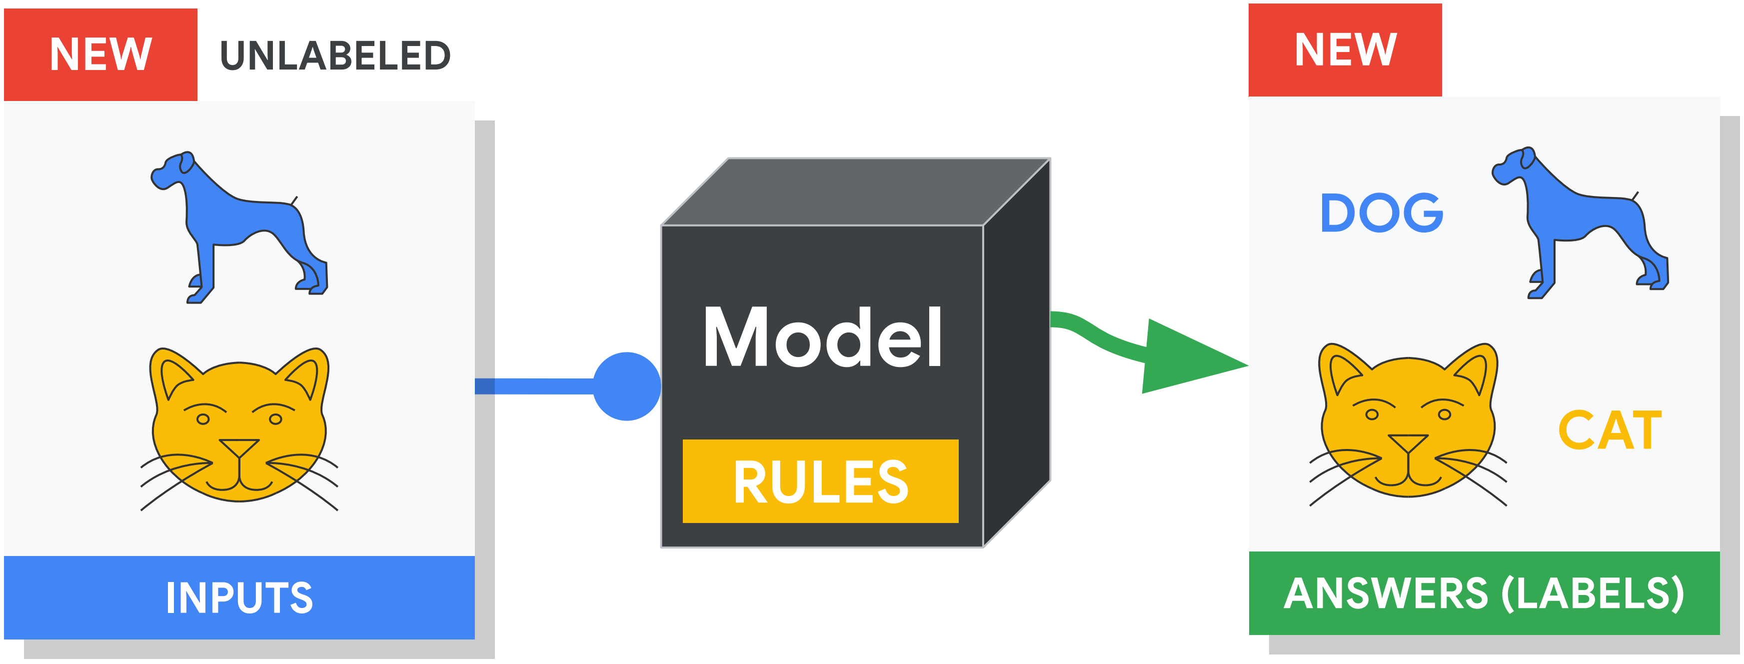

Deep Learning

Machine Learning

Ok so what about ***Deep*** *Learning?*

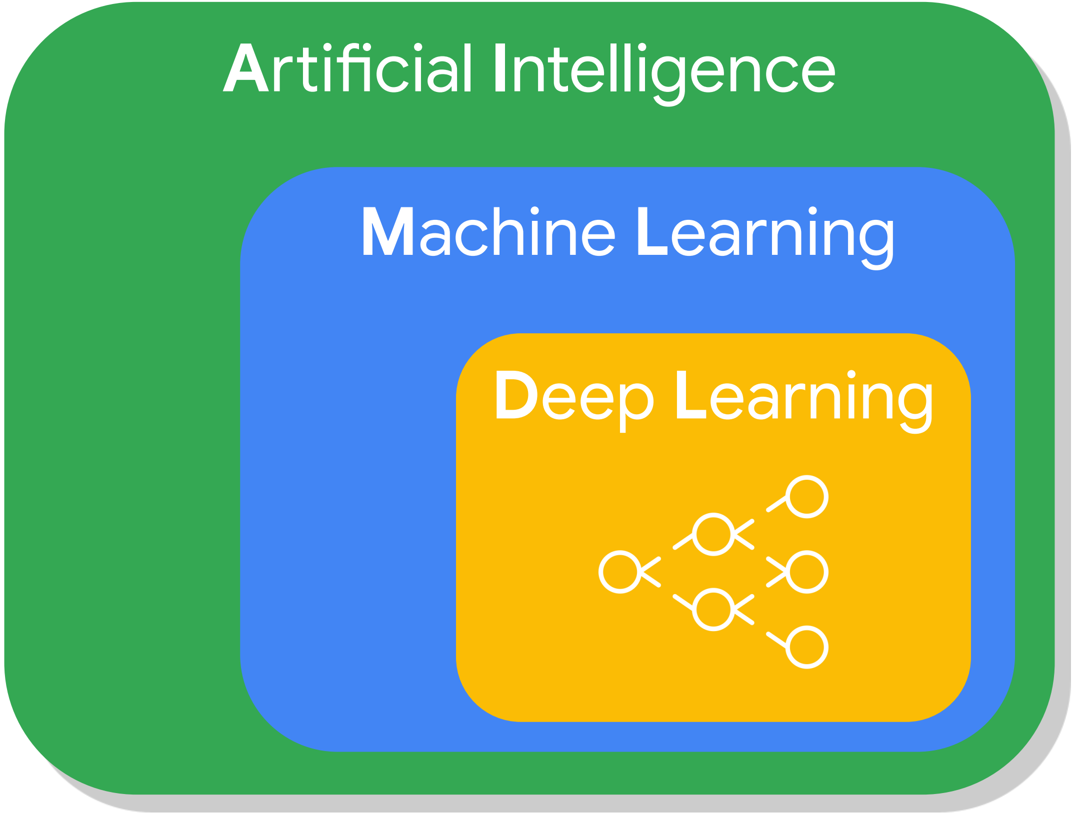
Artificial Intelligence
Machine Learning
Deep Learning

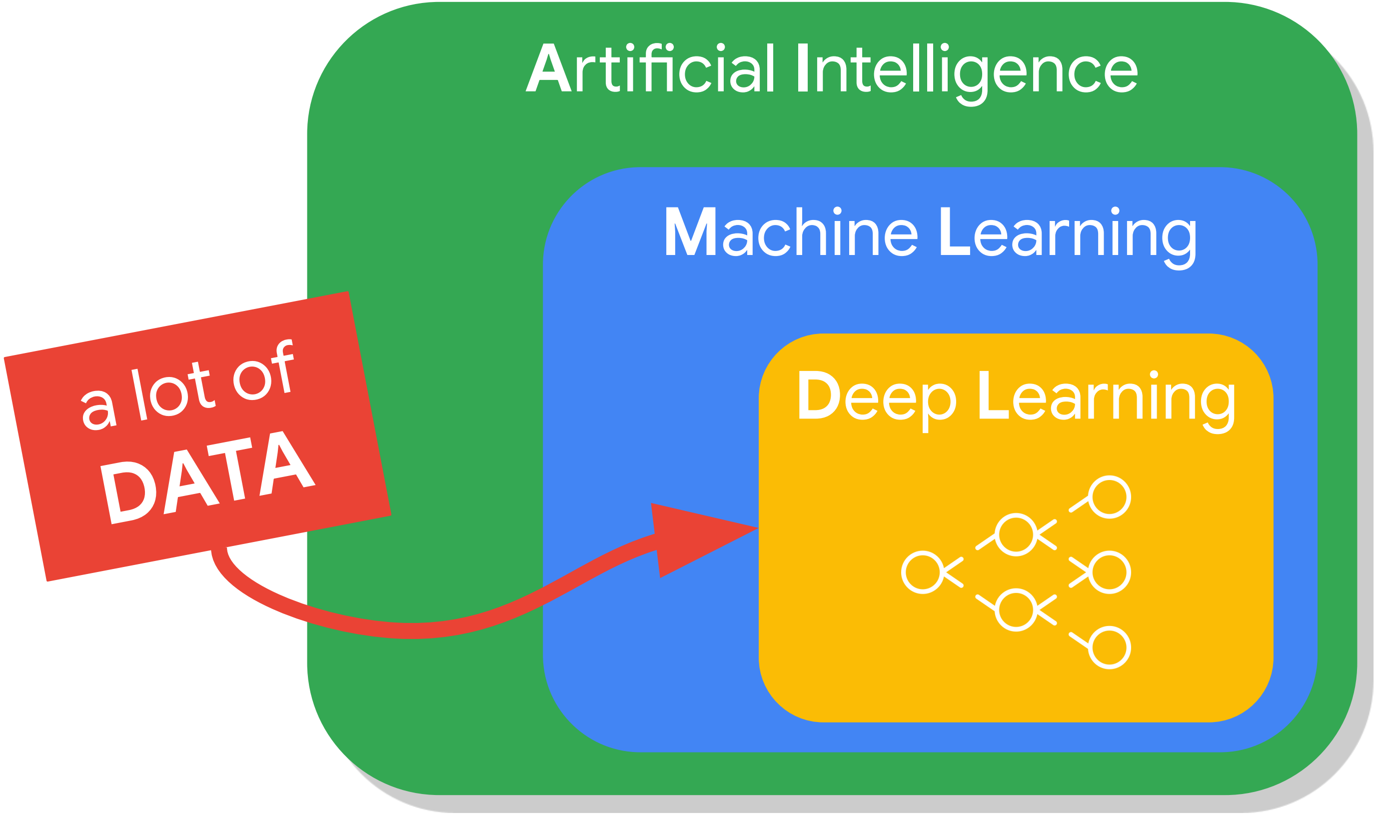
Artificial Intelligence
Machine Learning
Deep Learning
a lot of
DATA

Training the machine

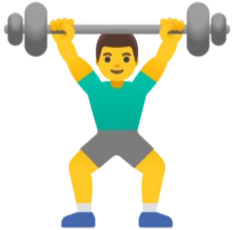

Make a Guess!

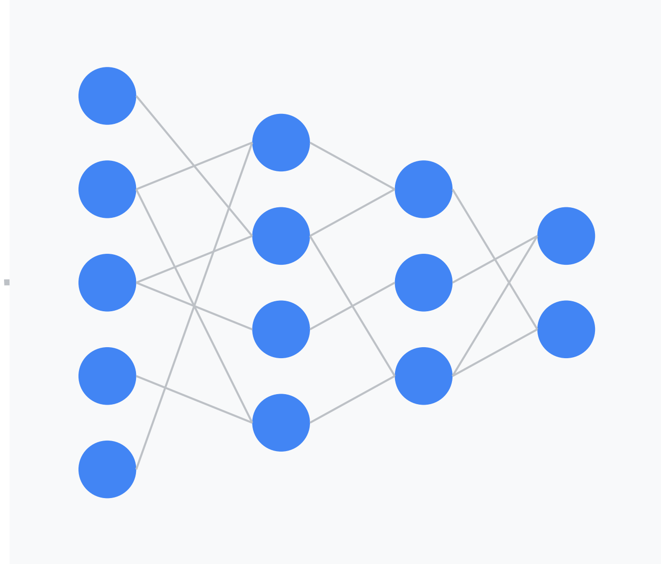

Neural Network

Neural network

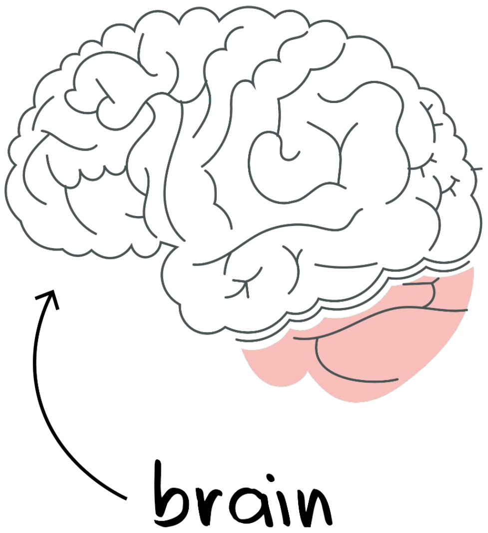

Neural network

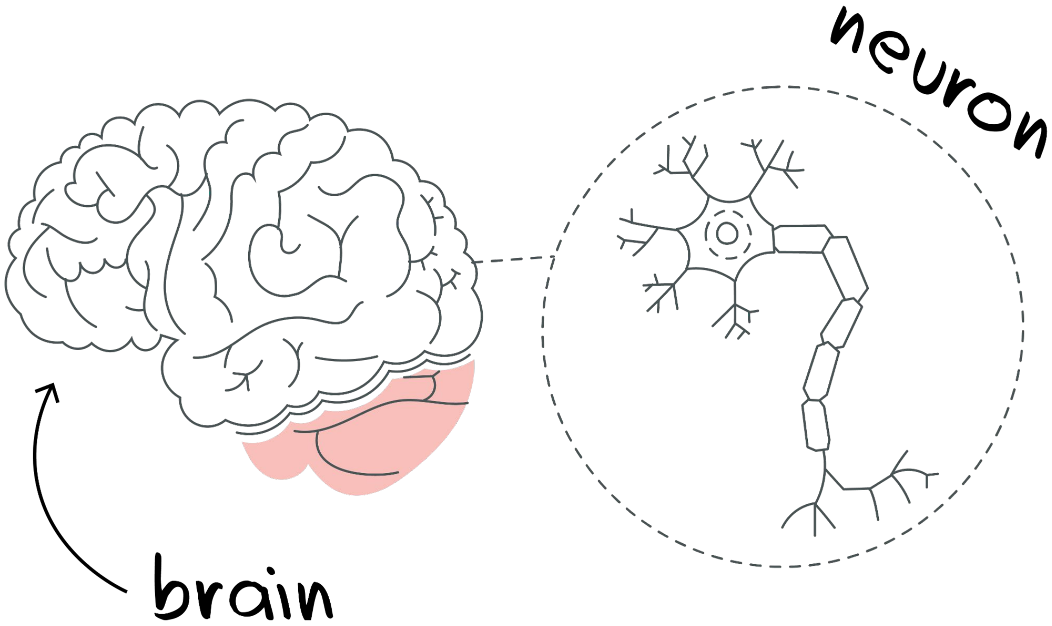

Neural network

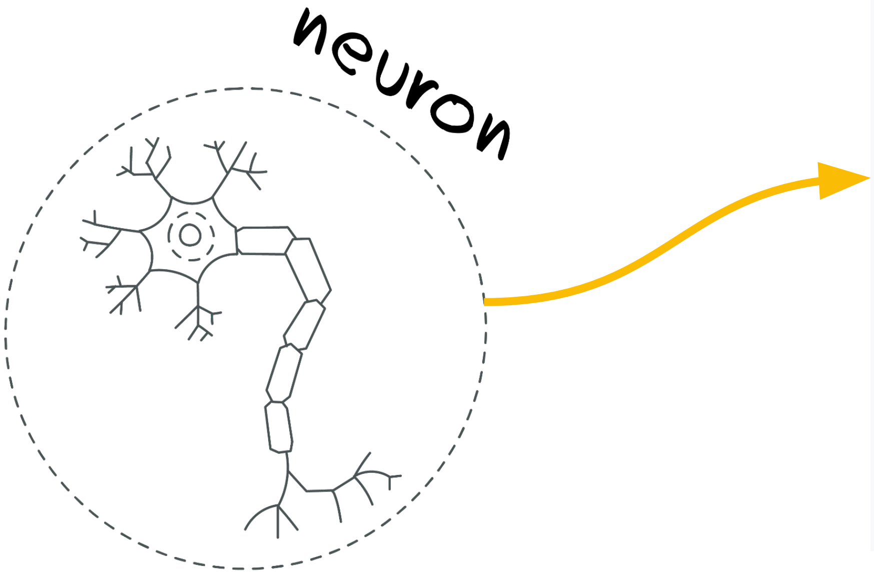

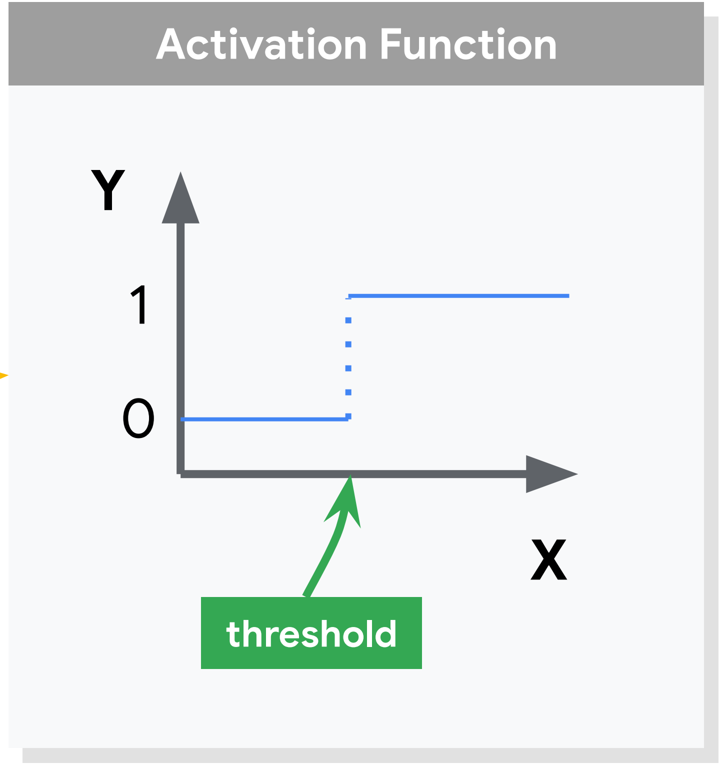

Neural network

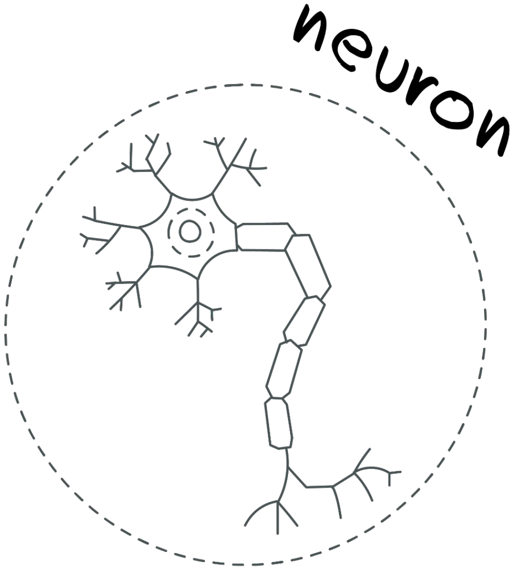

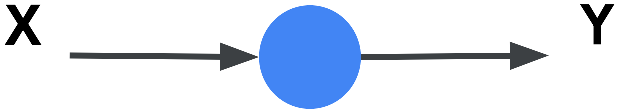

artificial

Neural network

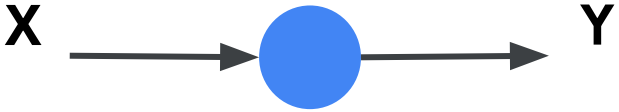

Neural network

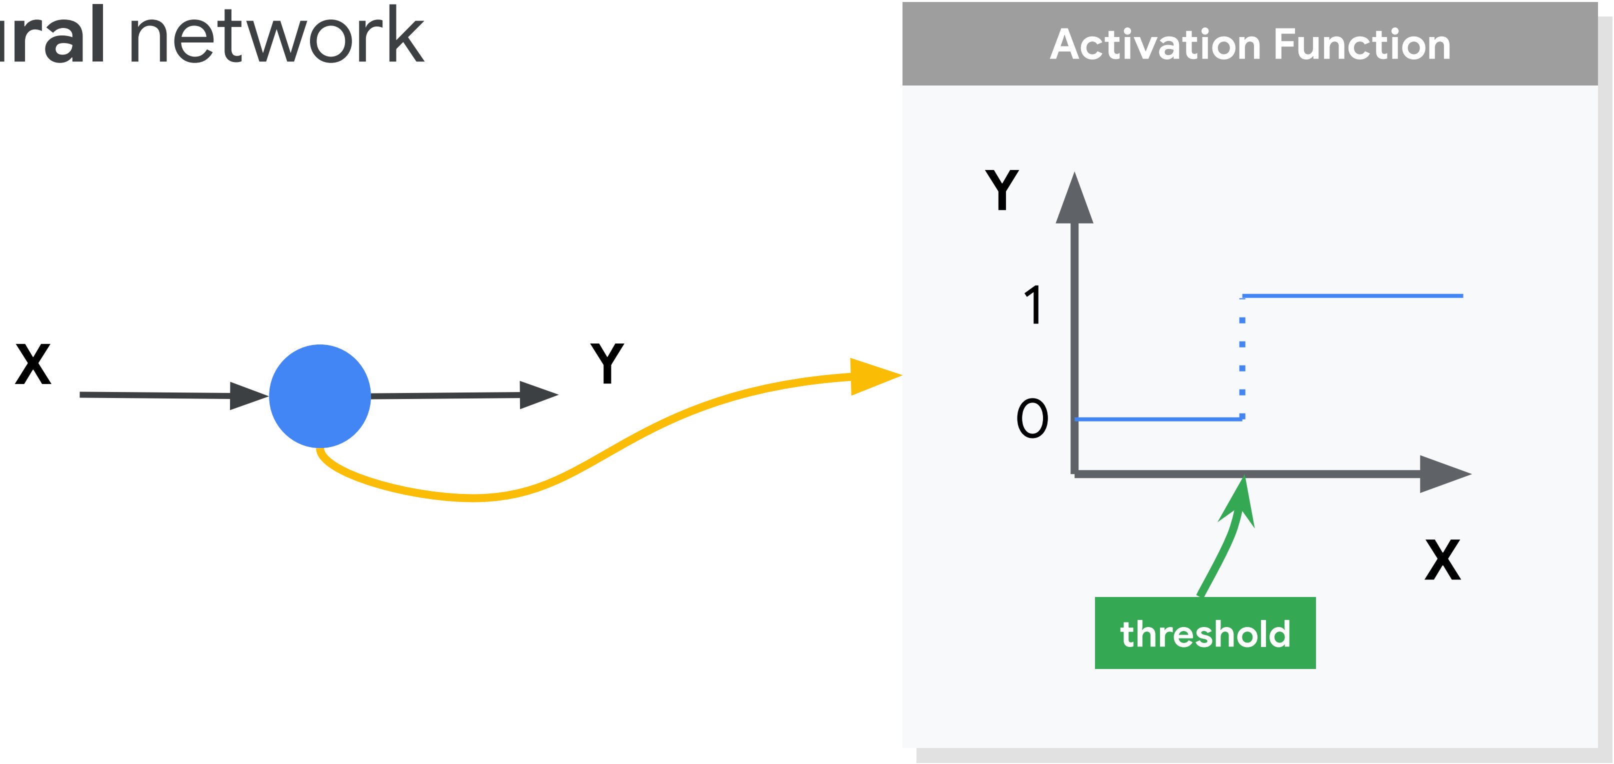

Multi-layer
neural network

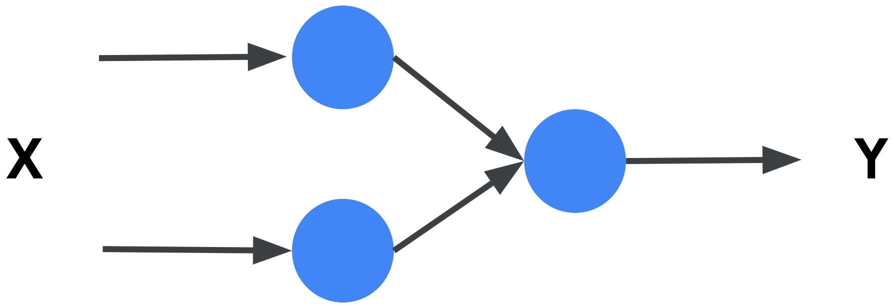

Deep Learning with **Neural Networks**

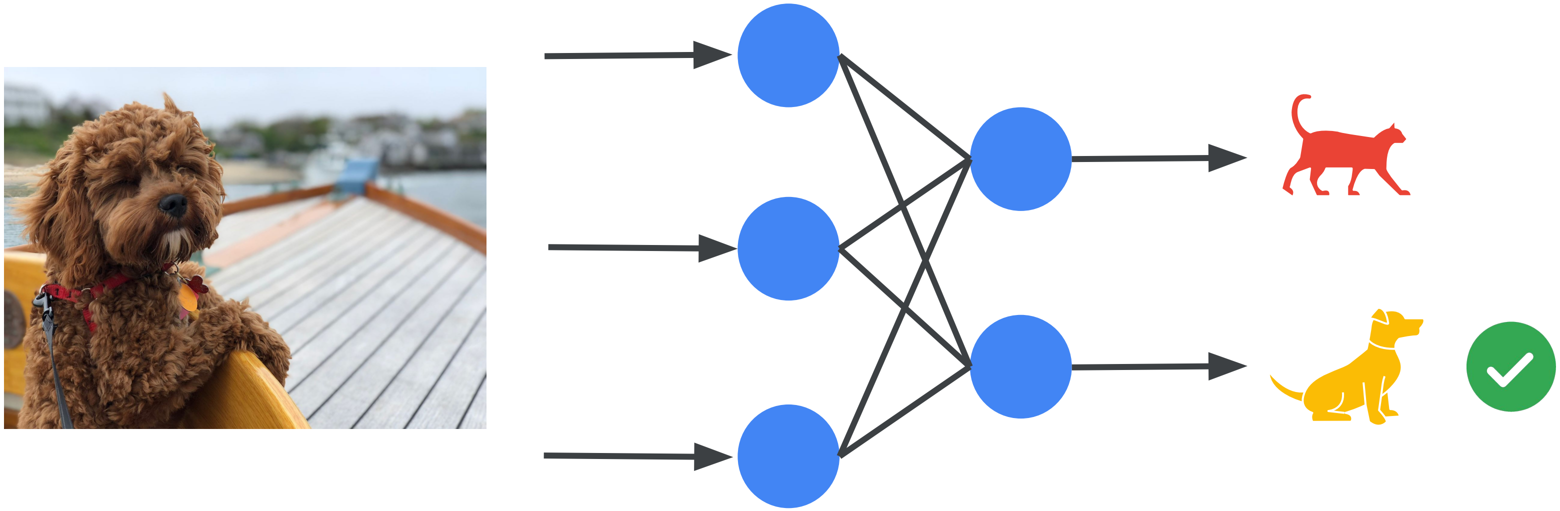

Training the machine

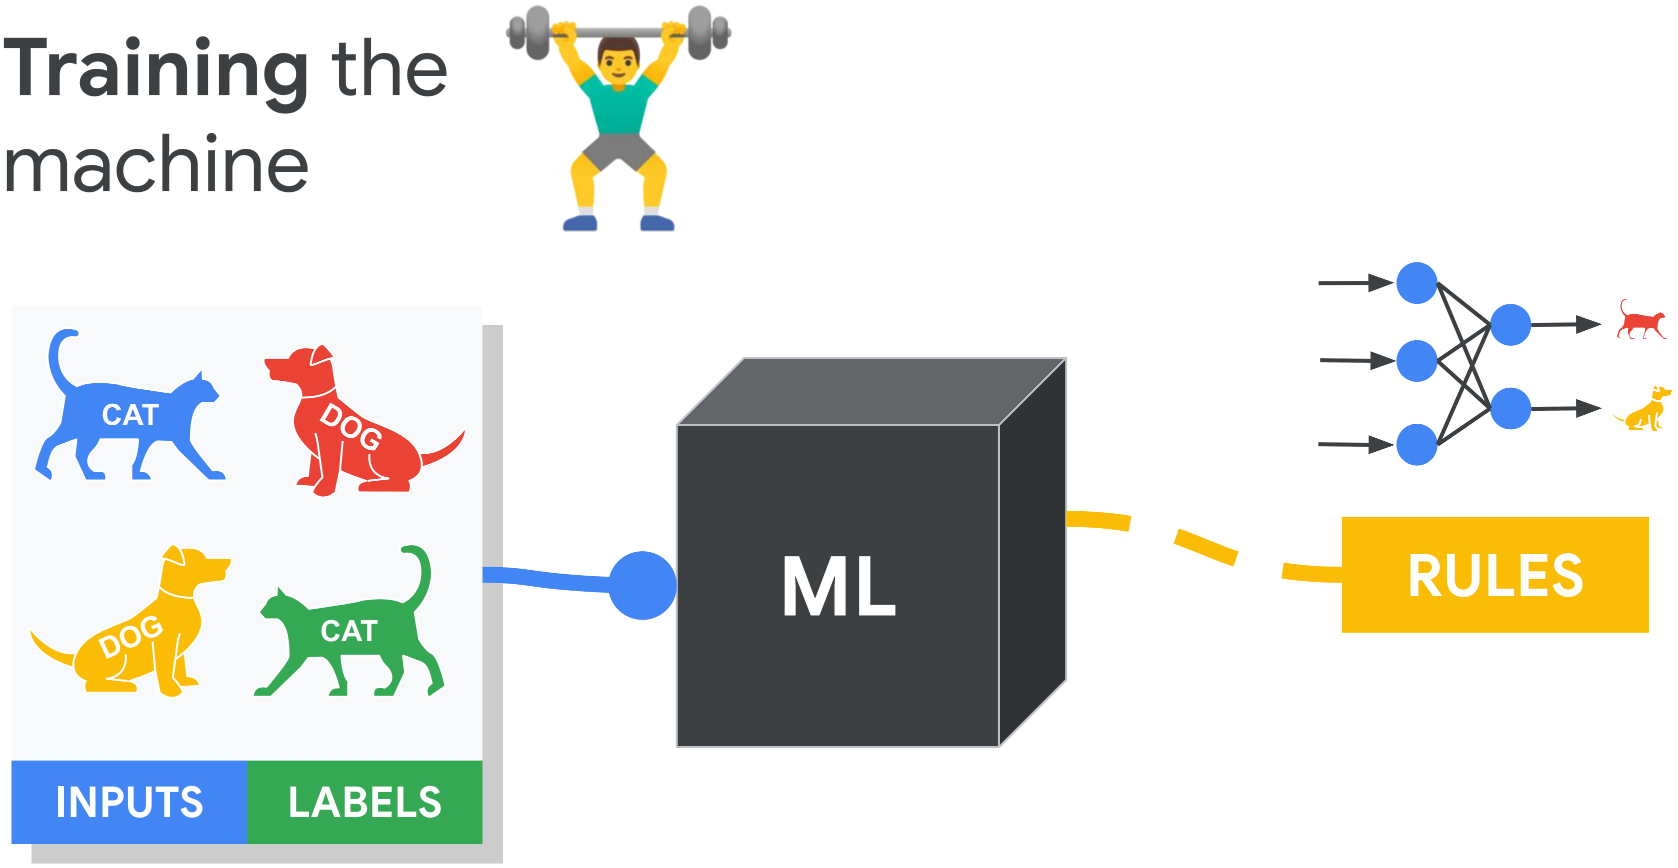

Case Study:
Handwriting

practice
DATA

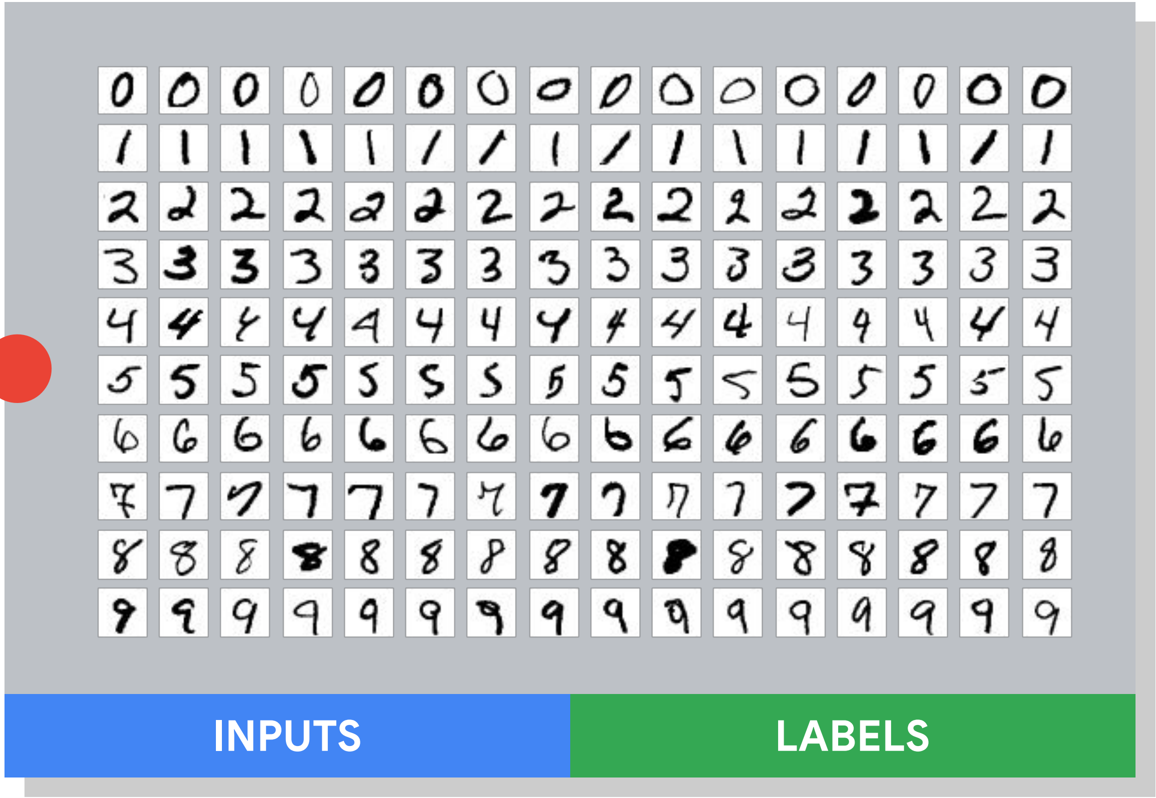

Case Study: **Handwriting**

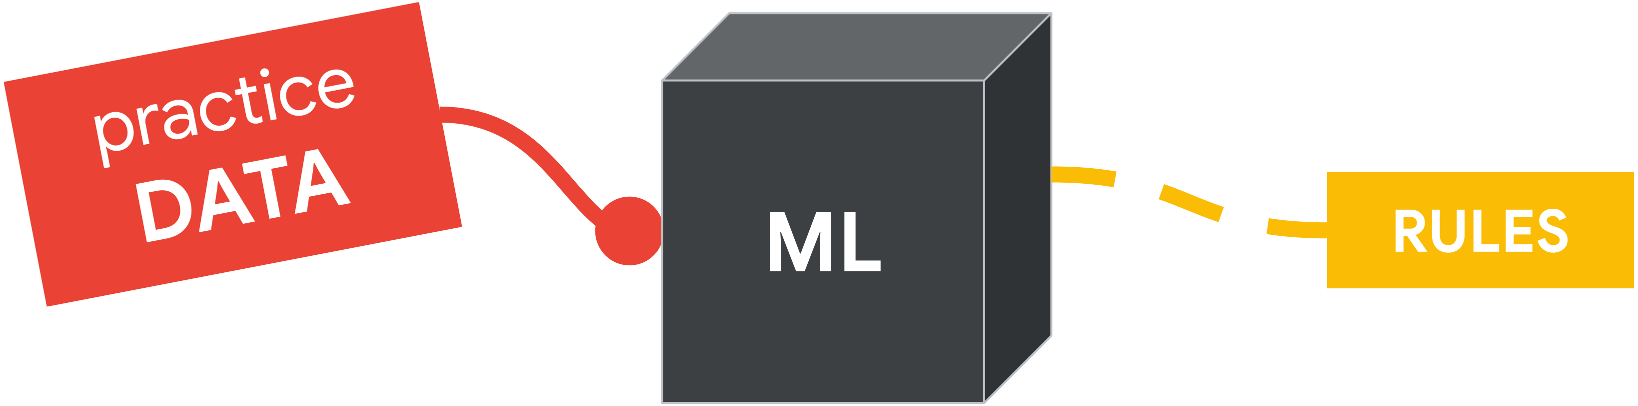

Case Study: **Handwriting**

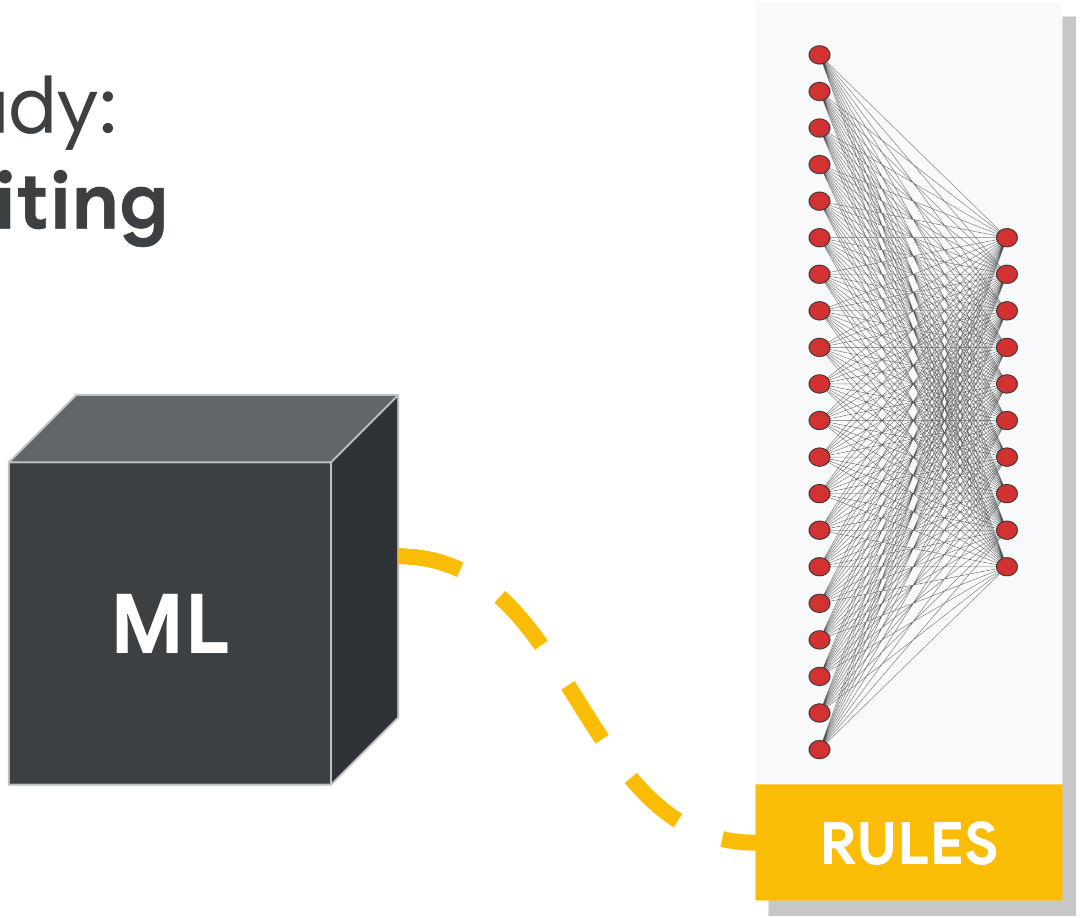

What **number?**

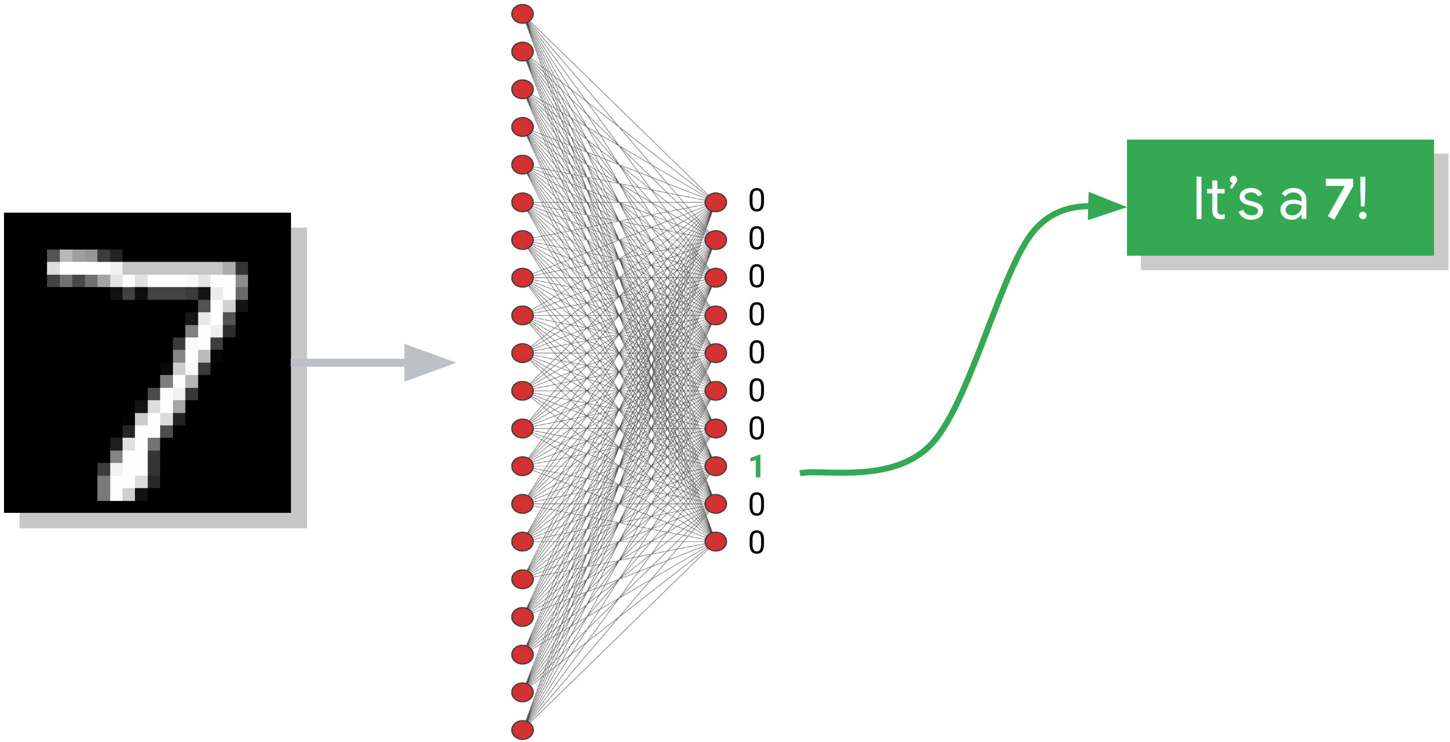

What **number?**

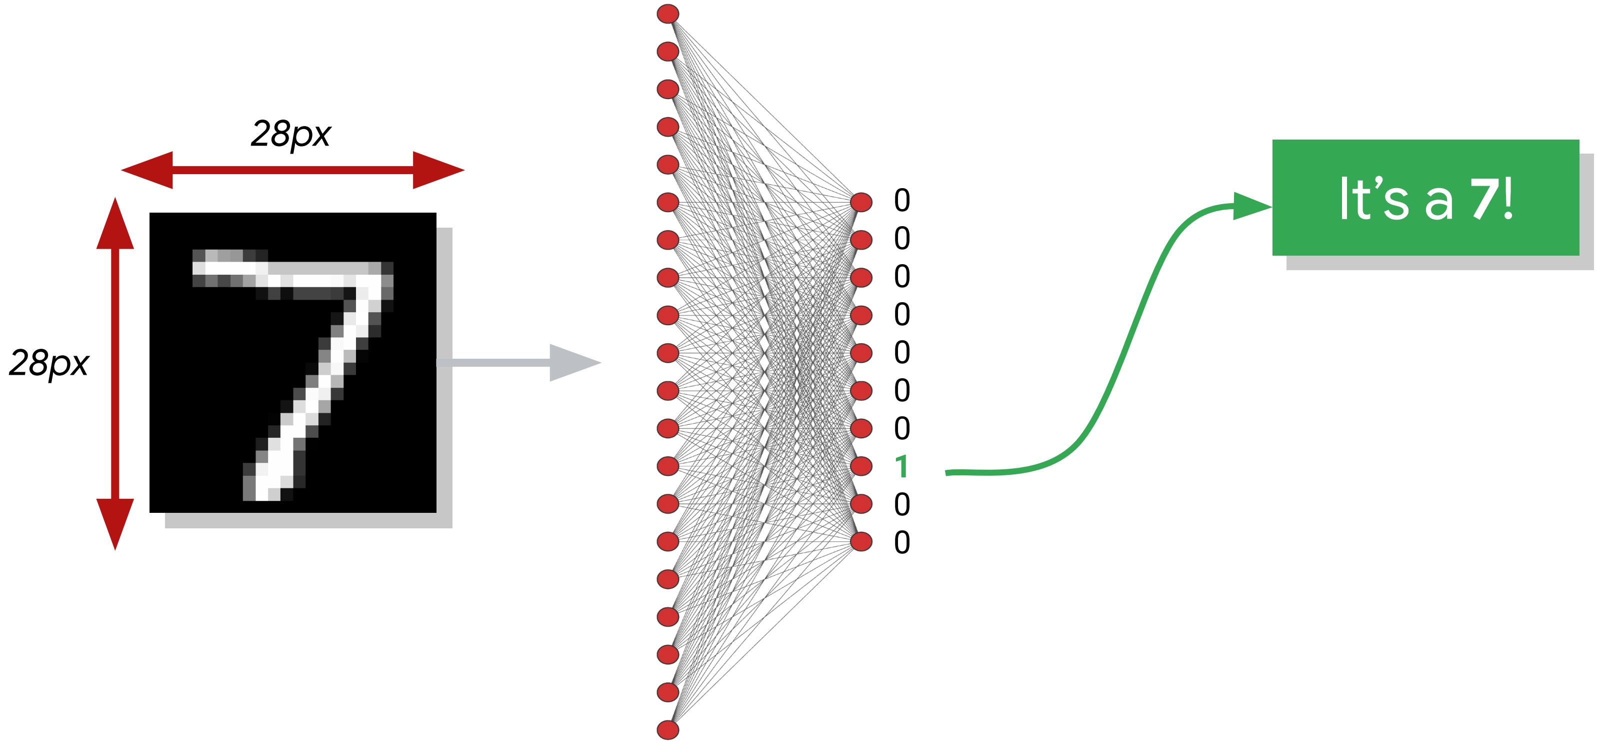

What **number**?

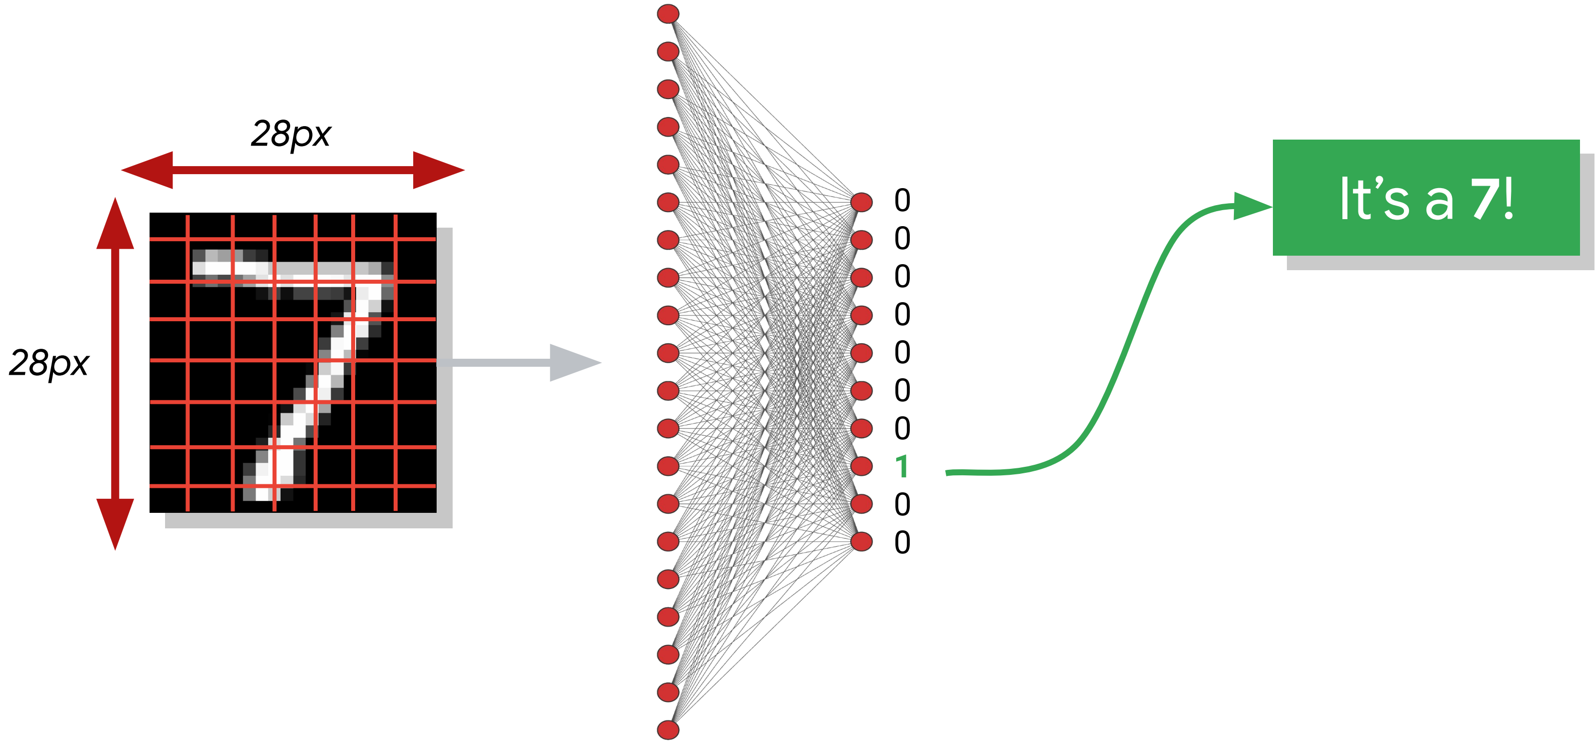

Transform: flatten

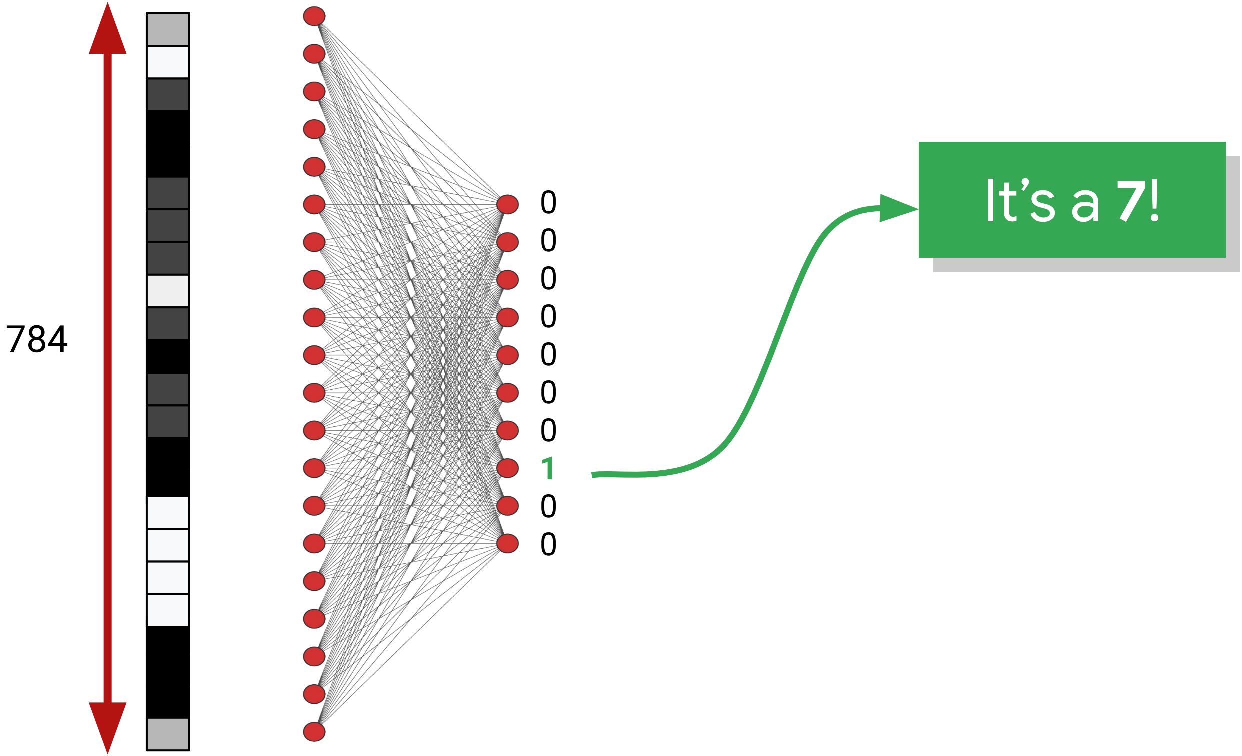

[Google Colab](#)

After Training the Model is VERY good!

Rank	Model	Percentage error	Accuracy	Trainable Parameters	Error rate	Percentage correct	Extra Training Data	Paper	Code	Result	Year	Tags
1	**Однородный ансамбль с простым CNN**	0.09	99.91				×	An Ensemble of Simple Convolutional Neural Network Models for MNIST Digit Recognition			2020	
2	**Branching/Merging CNN + Homogeneous Vector Capsules**	0.13	99.87	1,514,187			×	No Routing Needed Between Capsules			2020	
3	**EnsNet** (Ensemble learning in CNN augmented with fully connected subnetworks)	0.16	99.84				×	Ensemble learning in CNN augmented with fully connected subnetworks			2020	
4	**Efficient-CapsNet**	0.16	99.84	161,824			×	Efficient-CapsNet: Capsule Network with Self-Attention Routing			2021	
5	**SOPCNN** (Only a single Model)	0.17	99.83	1,400,000			×	Stochastic Optimization of Plain Convolutional Neural Networks with Simple methods			2020	
6	**RMDL** (30 RDLs)	0.18	99.82				×	RMDL: Random Multimodel Deep Learning for Classification			2018	
7	**DropConnect**	0.21	99.79				×	Regularization of Neural Networks using DropConnect			2013	

https://paperswithcode.com/sota/image-classification-on-mnist

And it can solve problems we couldn’t solve without ML!

DeepBlue

On average in any board configuration there are **35** possible moves in chess.

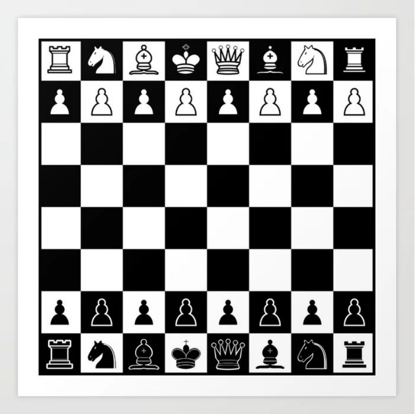

And it can solve problems we couldn't solve without ML!

AlphaGo

"There are an astonishing **10 to the power of 170 possible board configurations** - more than the number of atoms in the known universe. This makes the game of Go a **googol times more complex than chess**."

DeepBlue

On average in any board configuration there are **35** possible moves in chess.

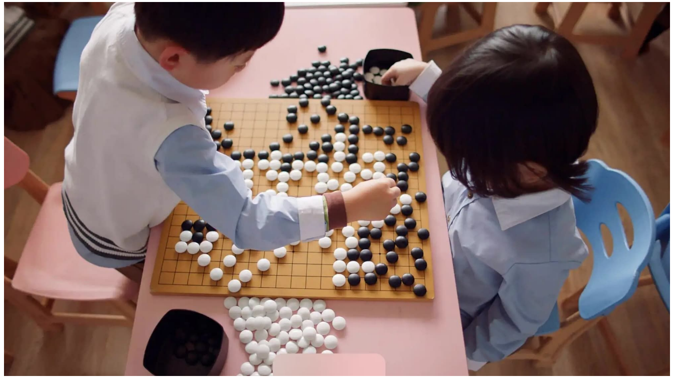

https://www.deepmind.com/research/highlighted-research/alphago

But It Need Lots of Data

This is considered a SMALL and simple dataset (~45MB)

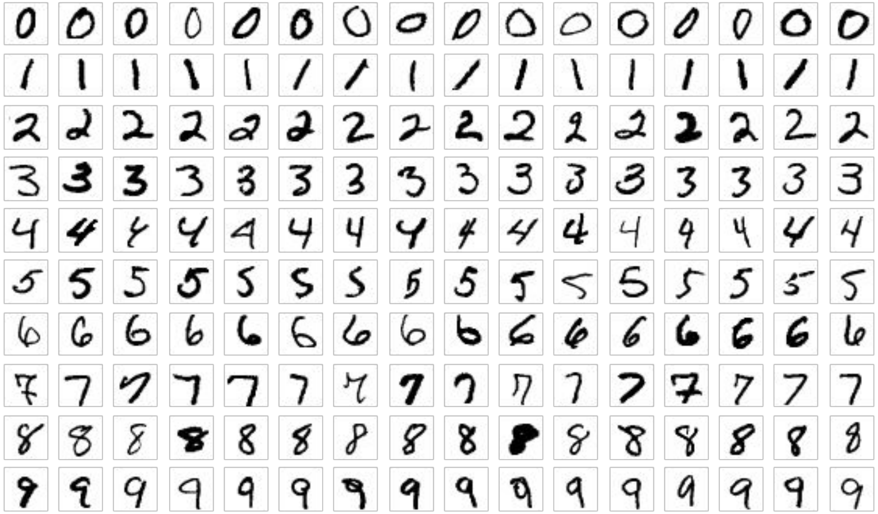

10 Classes

6000 Images / Class

But It Need Lots of Data

GPT-3 Used ~45TB of data that's ~1,000,000 times more data than MNIST!

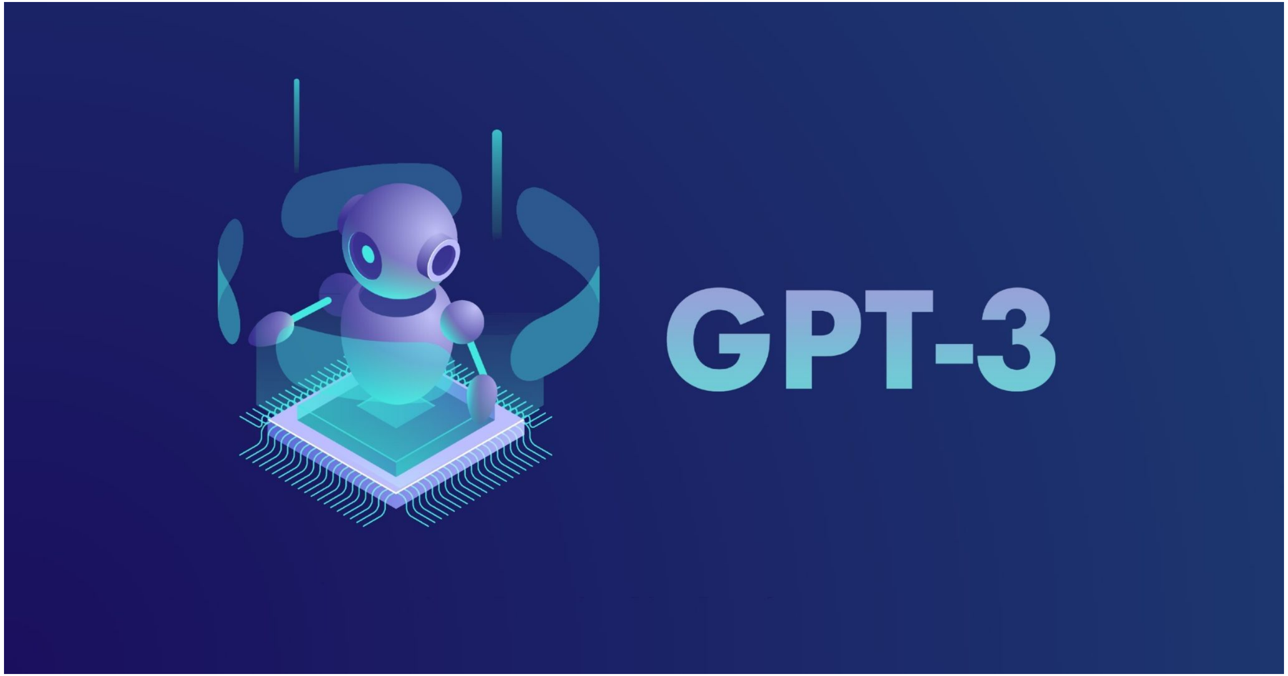

Today's Agenda

- What is Artificial Intelligence?
- Hands-on: AutoDraw
- **What is (Deep) Machine Learning?**
- Hands-on: ThingTranslator
- What is Responsible TinyML?
- Summary

Today's Agenda

- What is Artificial Intelligence?
- Hands-on: AutoDraw
- What is (Deep) Machine Learning?
- **Hands-on: ThingTranslator**
- What is Responsible TinyML?
- Summary

Thing **Translator**

This is an
A.I.
Experiment

Thing Translator

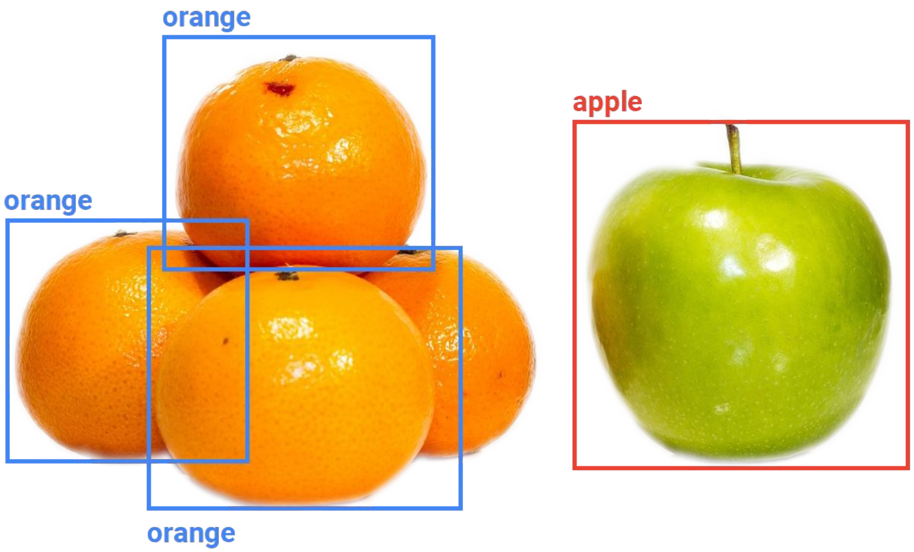

Thing Translator

Thing Translator

thing-translator.appspot.com

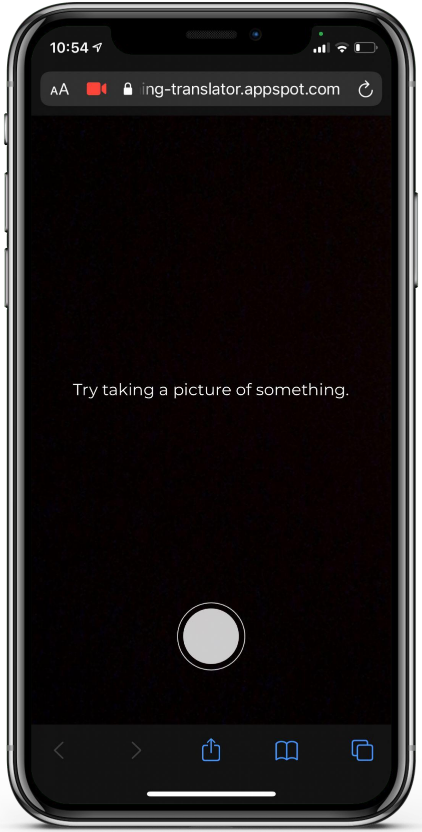

Discussion **Groups**

1. Do you think the AI did a **good job**? 👍 / 👎
2. **Why** do you think the AI **worked well**?
3. **How** did the AI solve this task? 🤔
4. What types of things were **particularly hard or easy** for the AI?
5. Was the AI **better or worse** in this experiment? **Why** do you think?

anything **else**?

Today's Agenda

- What is Artificial Intelligence?
- Hands-on: AutoDraw
- What is (Deep) Machine Learning?
- **Hands-on: ThingTranslator**
- What is Responsible TinyML?
- Summary

Today's Agenda

- What is Artificial Intelligence?
- Hands-on: AutoDraw
- What is (Deep) Machine Learning?
- Hands-on: ThingTranslator
- **What is Responsible TinyML?**
- Summary

TinyML

ML

What’s the difference?

Cloud / Server

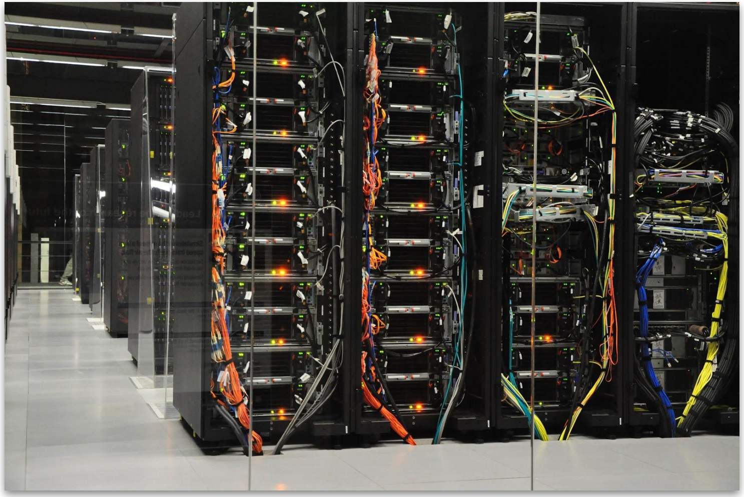

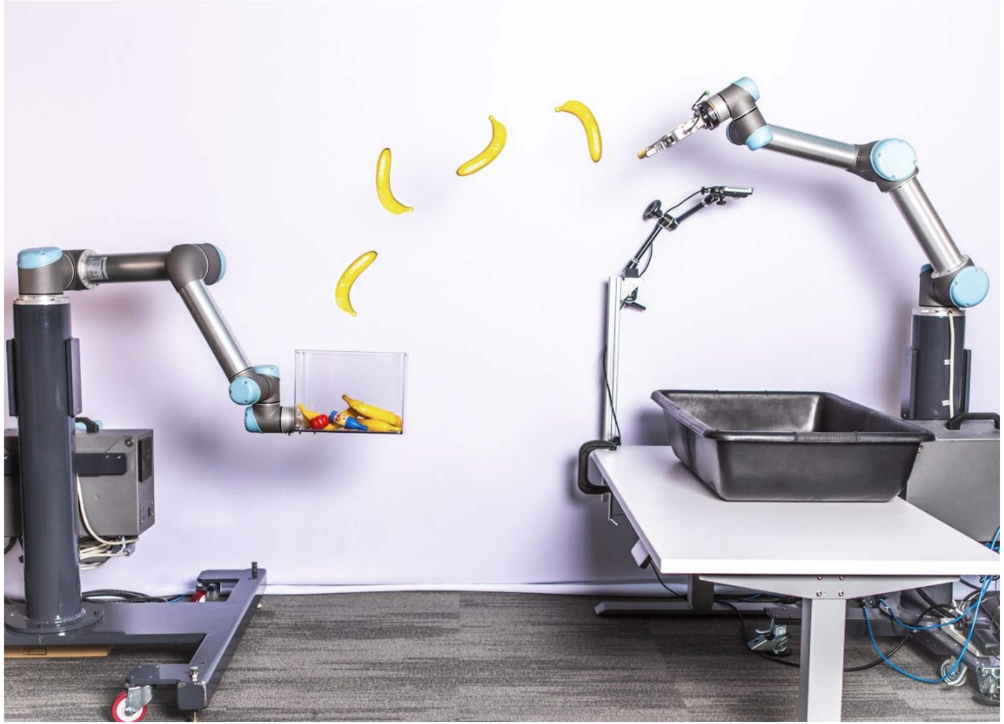

https://tossingbot.cs.princeton.edu/

Mobile

https://plantvillage.psu.edu/

ML on Embedded Devices

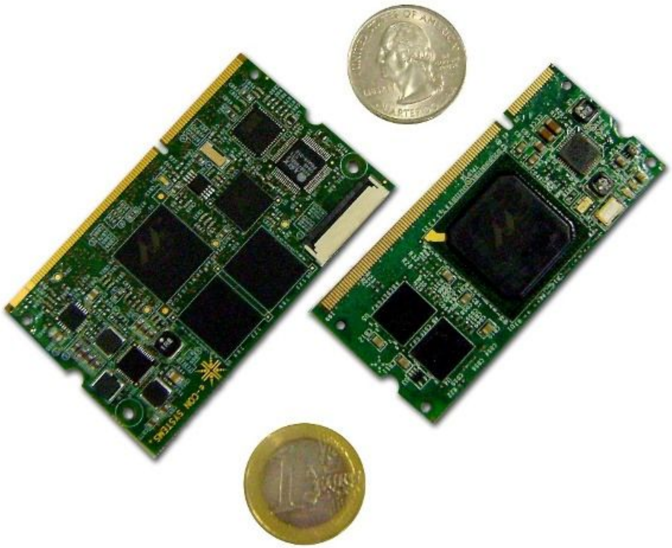

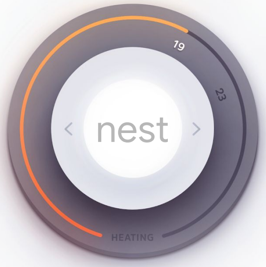

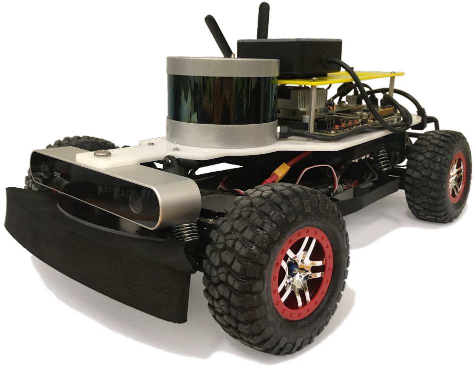

Google Assistant

IoT 1.0:
Internet
of Things

IoT 2.0:
Intelligence
on Things

No Good Data Left Behind

5 Quintillion

bytes of data produced every day by IoT

<1%

of unstructured data is analyzed or used at all

Source: Harvard Business Review, What's Your Data Strategy?, April 18, 2017
Cisco, Internet of Things (IoT) Data Continues to Explode Exponentially. Who Is Using That Data and How?, Feb 5, 2018

IoT 1.0:
Internet
of Things

IoT 2.0:
Intelligence
on Things

B
L
E
R
P

IoT 1.0:
Internet
of Things

IoT 2.0:
Intelligence
on Things

Bandwidth

Latency

Energy

Reliability

Privacy

IoT 1.0:
Internet
of Things

IoT 2.0:
Intelligence
on Things

Bandwidth

Latency

Energy

Battery Life is only O(months) and only sends GPS signal

IoT 1.0:
Internet
of Things

↓

IoT 2.0:
Intelligence
on Things

Bandwidth

Latency

Energy

IoT 1.0:
Internet
of Things

IoT 2.0:
Intelligence
on Things

Reliability

Privacy

Google Assistant

IoT 1.0:
Internet
of Things

↓

IoT 2.0:
Intelligence
on Things

Bandwidth
Latency
Energy
Reliability
Privacy

TinyML to the rescue!

What is Tiny Machine Learning (**TinyML**)?

TinyML

Fastest-growing field of **ML**

What is Tiny Machine Learning (**TinyML**)?

TinyML

Fastest-growing field of **ML**

Algorithms, hardware, software

What is Tiny Machine Learning (**TinyML**)?

TinyML

Fastest-growing field of **ML**

Algorithms, hardware, software

On-device sensor analytics

What is Tiny Machine Learning (**TinyML**)?

TinyML

Fastest-growing field of **ML**

Algorithms, hardware, software

On-device sensor analytics

Low power consumption

What is Tiny Machine Learning (TinyML)?

TinyML

Fastest-growing field of **ML**

Algorithms, hardware, software

On-device sensor analytics

Low power consumption

Battery-operated

What is Tiny Machine Learning (**TinyML**)?

TinyML

- Fastest-growing field of **ML**
- Algorithms, hardware, software
- **On-device** sensor analytics
- **Low power** consumption
- **Battery-operated**
- **Always-on ML**

What is Tiny Machine Learning (TinyML)?

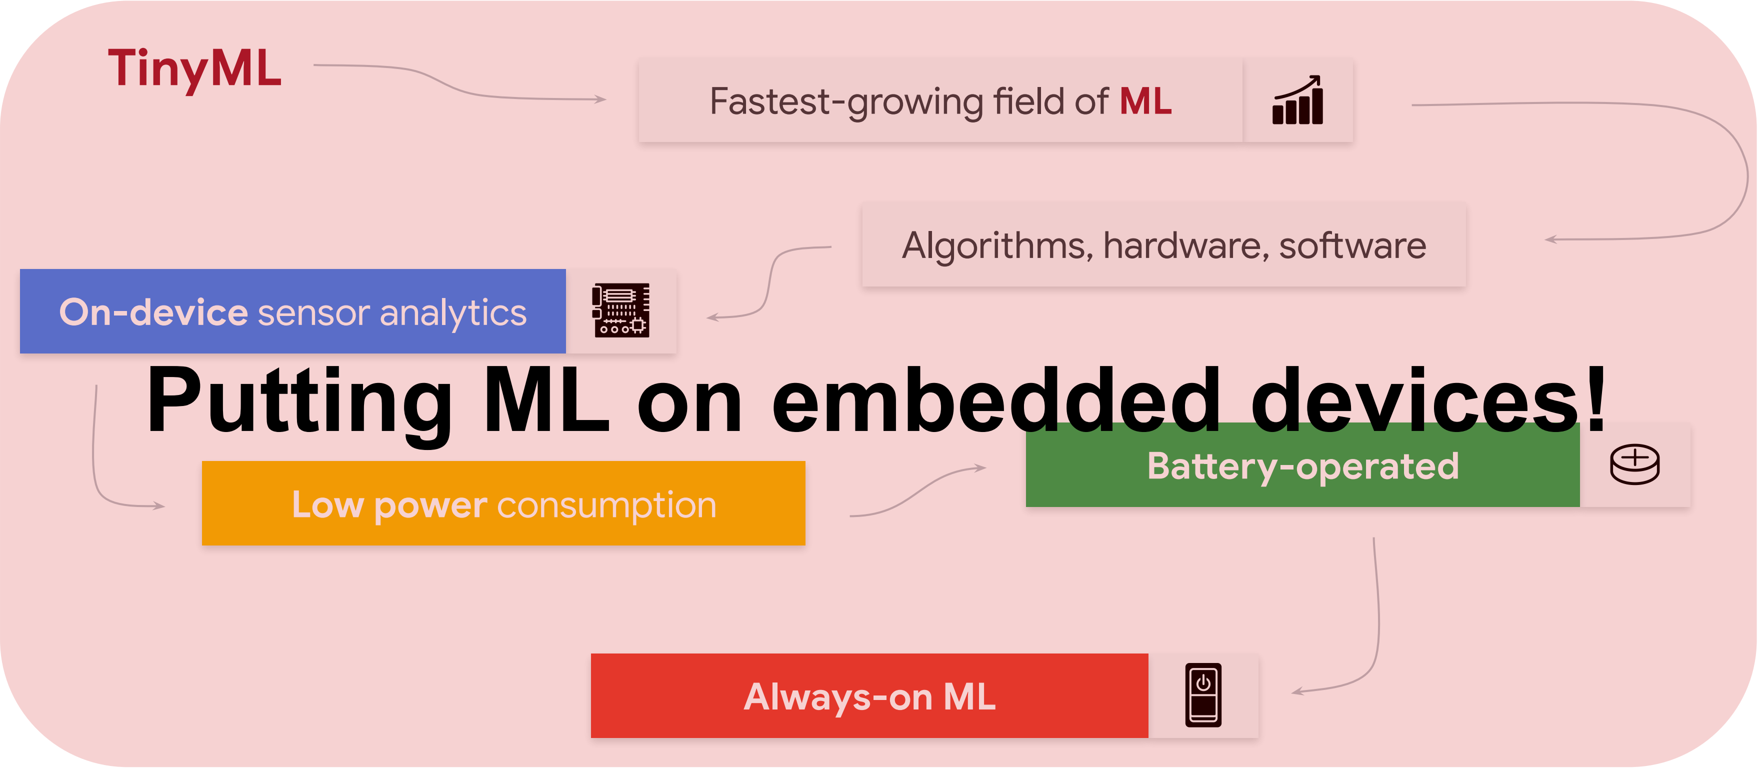

Promising **Social Applications** of TinyML

Wildlife conservation

ElephantEdge

Building The World's Most Advanced
Wildlife Tracker.

Agriculture

May be able to reduce agrichemical use to 0.1% of conventional blanket spraying

Technology: The Future of Agriculture

Anthony King

Nature **544**, S21–S23 (2017) | Cite this article

161k Accesses | **132** Citations | **209** Altmetric | Metrics

TinyRL: Autonomous Navigation on Nano Drone

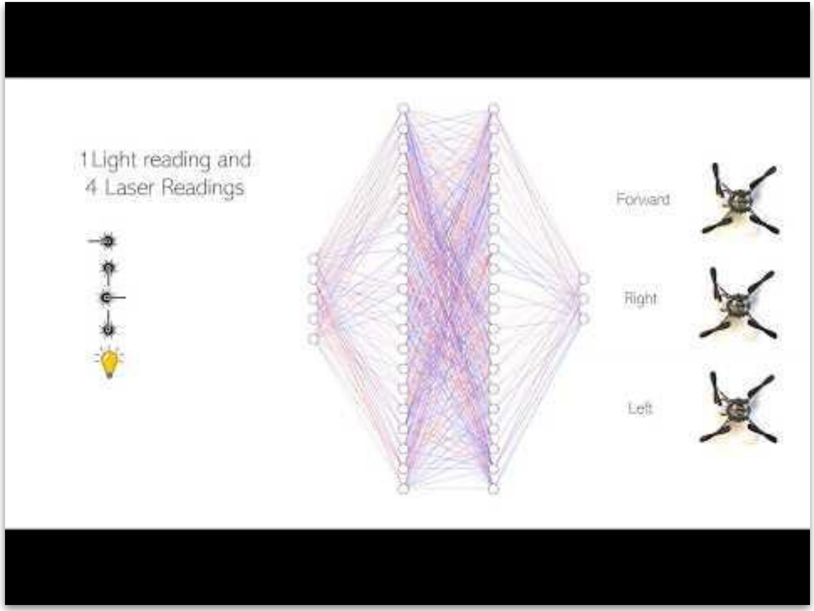

[ICRA'21]

[IROS'21]

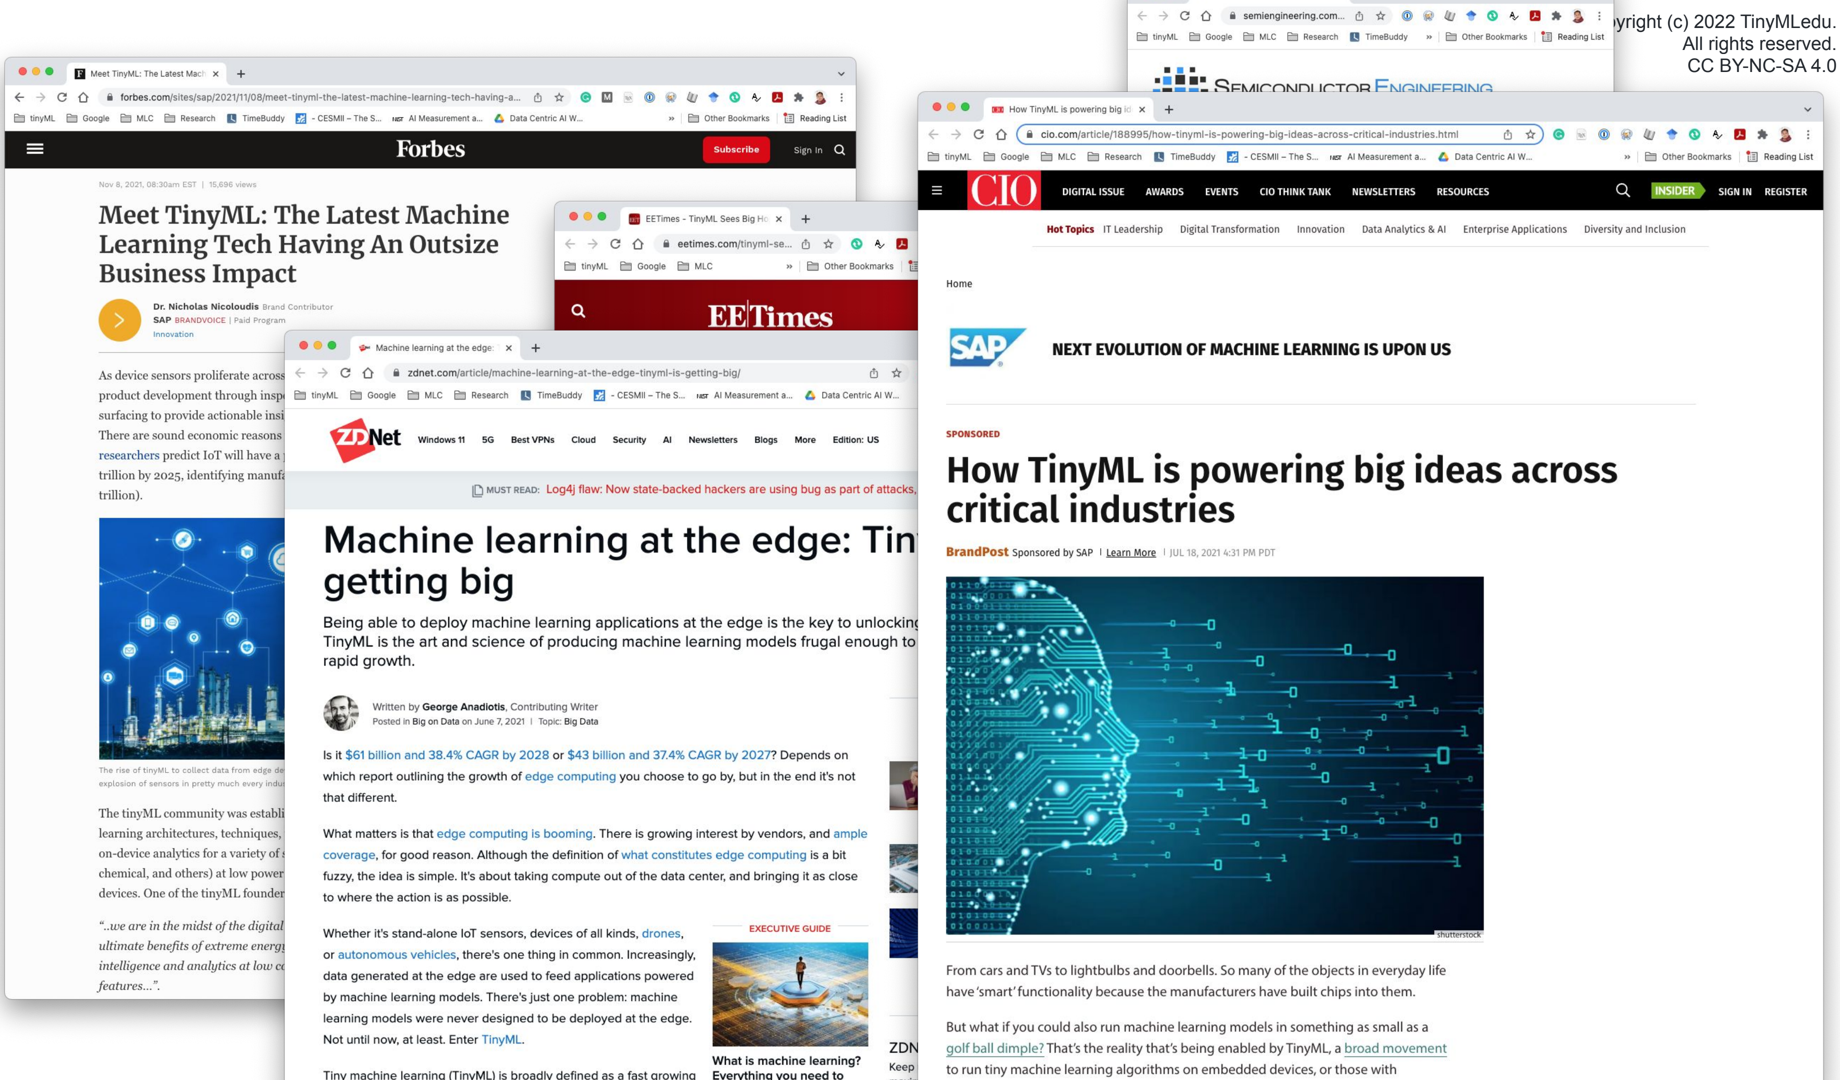

Meet TinyML: The Latest Machine Learning Tech Having An Outsize Business Impact

Nov 8, 2021, 08:30am EST | 15,696 views

Dr. Nicholas Nicoloudis Brand Contributor
SAP BRANDVOICE | Paid Program
Innovation

As device sensors proliferate across
product development through insp
surfacing to provide actionable insi
There are sound economic reasons
researchers predict IoT will have a
trillion by 2025, identifying manufa
trillion).

The rise of tinyML to collect data from edge de
explosion of sensors in pretty much every indu

The tinyML community was establi
learning architectures, techniques,
on-device analytics for a variety of
chemical, and others) at low power
devices. One of the tinyML founder

"..we are in the midst of the digital
ultimate benefits of extreme energ
intelligence and analytics at low c
features...".

EETimes - TinyML Sees Big Ho

Machine learning at the edge: Tin
getting big

Being able to deploy machine learning applications at the edge is the key to unlocking
TinyML is the art and science of producing machine learning models frugal enough to
rapid growth.

Written by George Anadiotis, Contributing Writer
Posted in Big on Data on June 7, 2021 | Topic: Big Data

Is it $61 billion and 38.4% CAGR by 2028 or $43 billion and 37.4% CAGR by 2027? Depends on which report outlining the growth of edge computing you choose to go by, but in the end it's not that different.

What matters is that edge computing is booming. There is growing interest by vendors, and ample coverage, for good reason. Although the definition of what constitutes edge computing is a bit fuzzy, the idea is simple. It's about taking compute out of the data center, and bringing it as close to where the action is as possible.

Whether it's stand-alone IoT sensors, devices of all kinds, drones, or autonomous vehicles, there's one thing in common. Increasingly, data generated at the edge are used to feed applications powered by machine learning models. There's just one problem: machine learning models were never designed to be deployed at the edge. Not until now, at least. Enter TinyML.

EXECUTIVE GUIDE
What is machine learning? Everything you need to

SEMICONDUCTOR ENGINEERING

SAP — NEXT EVOLUTION OF MACHINE LEARNING IS UPON US

SPONSORED

How TinyML is powering big ideas across critical industries

BrandPost Sponsored by SAP | Learn More | JUL 18, 2021 4:31 PM PDT

From cars and TVs to lightbulbs and doorbells. So many of the objects in everyday life have 'smart' functionality because the manufacturers have built chips into them.

But what if you could also run machine learning models in something as small as a golf ball dimple? That's the reality that's being enabled by TinyML, a broad movement to run tiny machine learning algorithms on embedded devices, or those with

The (Tiny) Machine Learning **Workflow**

The (Tiny) Machine Learning **Workflow**

Collect & Transform Data → Design & Train a Model → Deploy Model → Make Inferences

If ML is going to be everywhere we need to consider how to best collect GOOD data RESPONSIBLY

Good Data is Necessary for Accuracy

What problem are you trying to *solve*?

- Your data must contain useful features
- Can a human (expert) distinguish between examples of each class?
- How will you measure performance?

Good Data is Necessary for Accuracy

What problem are you trying to solve?

- Your data must contain useful features
- Can a human (expert) distinguish between examples of each class?
- How will you measure performance?

Both *quantity* and *quality* will influence your model's performance

- Wide variety of training examples
- Correct labels (answers)
- Good Balance (e.g., dog, cat, random)

Potential **Bias** in Speech Recognition

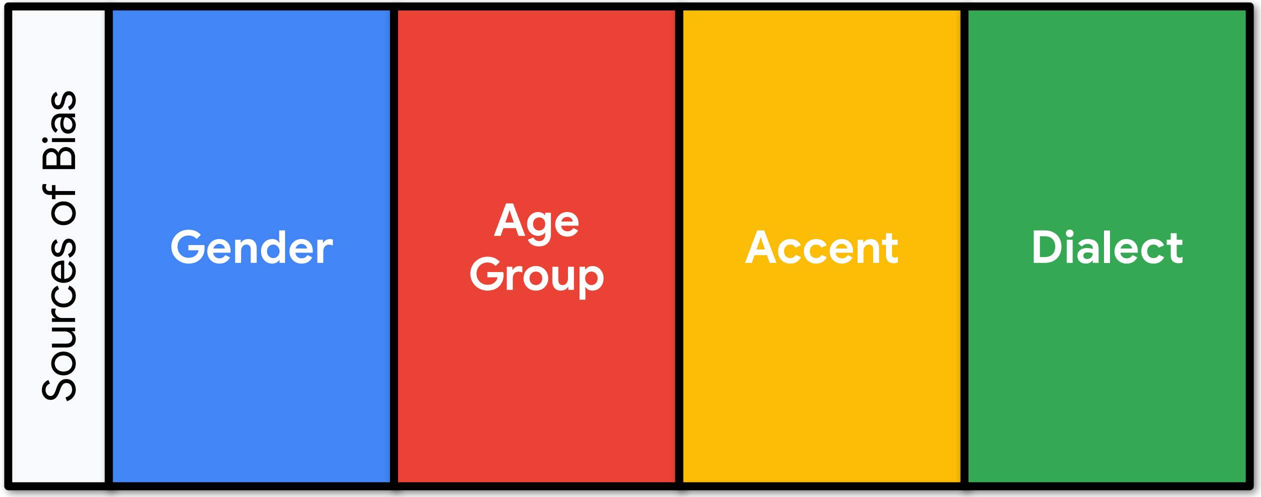


MIT Technology Review
Featured
Topics
Newsletters
Events
Podcasts
Sign in
Subs
ARTIFICIAL INTELL
Predictive
They need
Lack of transparency and b
purpose. If we can't fix then
By Will
TOM SIMONITE
BUSINE
The Bes
US government t
NEWSLETTERS
Sign up to read our regular email newsletters
NewScientist
News
Podcasts
Video
Technology
Space
Physics
Health
More
Shop
Courses
Events
Voice assistants seem to be worse at understanding commands from women
TECHNOLOGY 9 May 2019
By Nicole Kobie
Color Matters in Computer Vision
Facial recognition algorithms made by Microsoft, IBM and Face++ were more likely to misidentify the gender of black women than white men.
1 percent of lighter-skinned males in a set of
7 percent of lighter-skinned females in a set of
12 percent of darker-skinned males in a set of
Gender was misidentified in 35 percent of darker-skinned females in a set of 271
es.

Today's Agenda

- What is Artificial Intelligence?
- Hands-on: AutoDraw
- What is (Deep) Machine Learning?
- Hands-on: ThingTranslator
- **What is Responsible TinyML?**
- Summary

Today's Agenda

- What is Artificial Intelligence?
- Hands-on: AutoDraw
- What is (Deep) Machine Learning?
- Hands-on: ThingTranslator
- What is Responsible TinyML?
- **Summary**

Machine Learning **Workflow**

Machine Learning **Workflow**

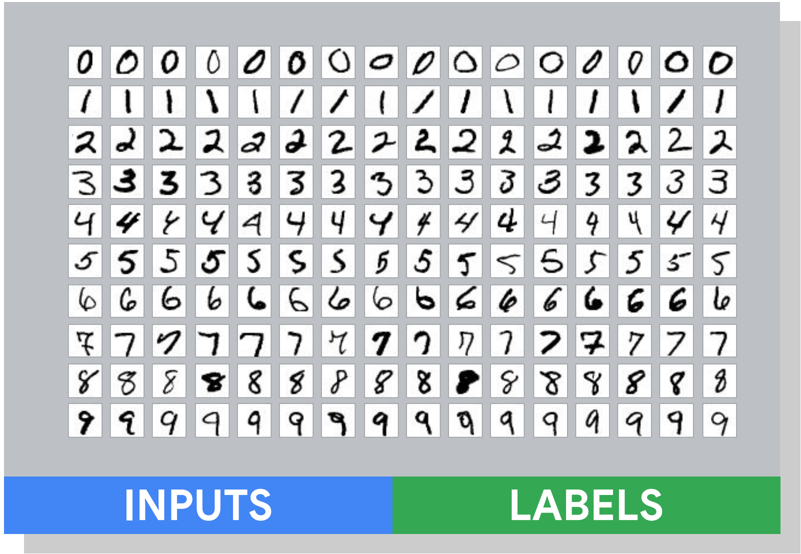

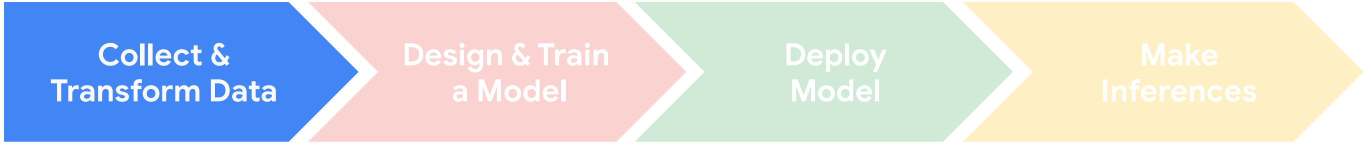

Machine Learning **Workflow**

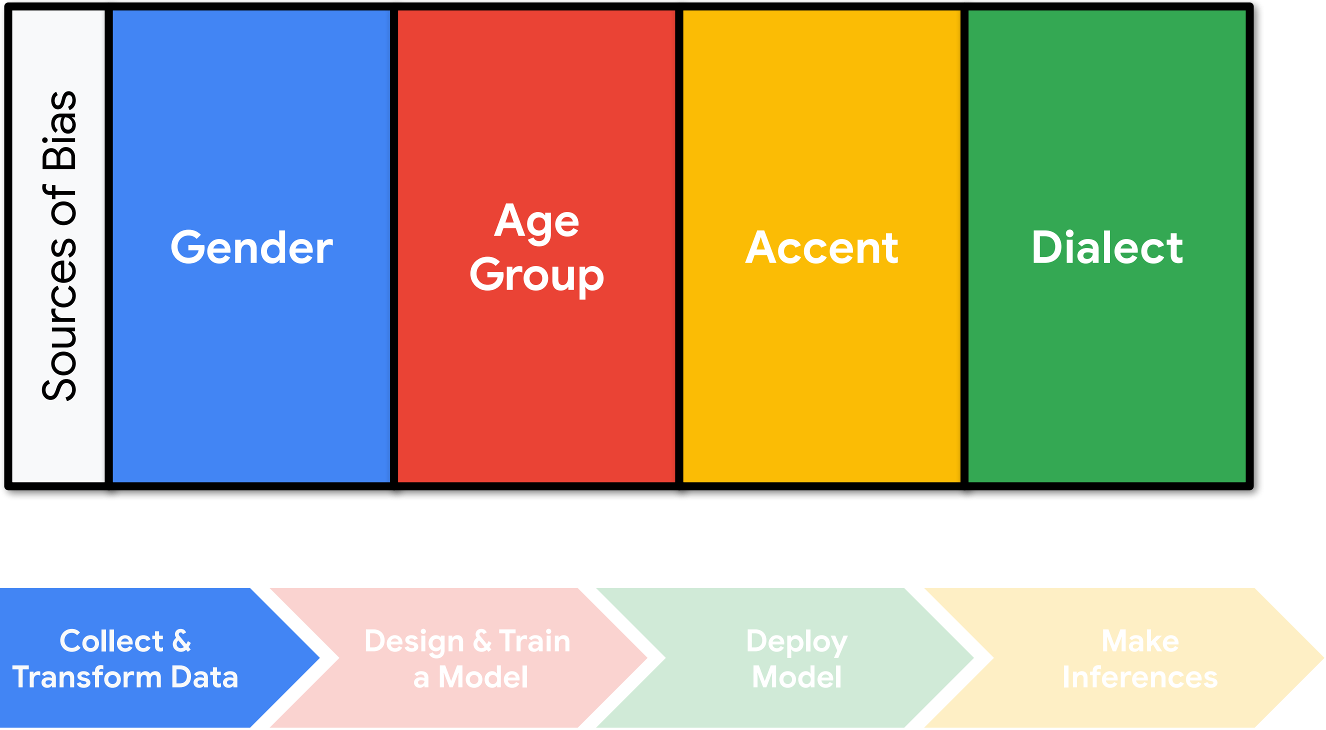

Machine Learning **Workflow**

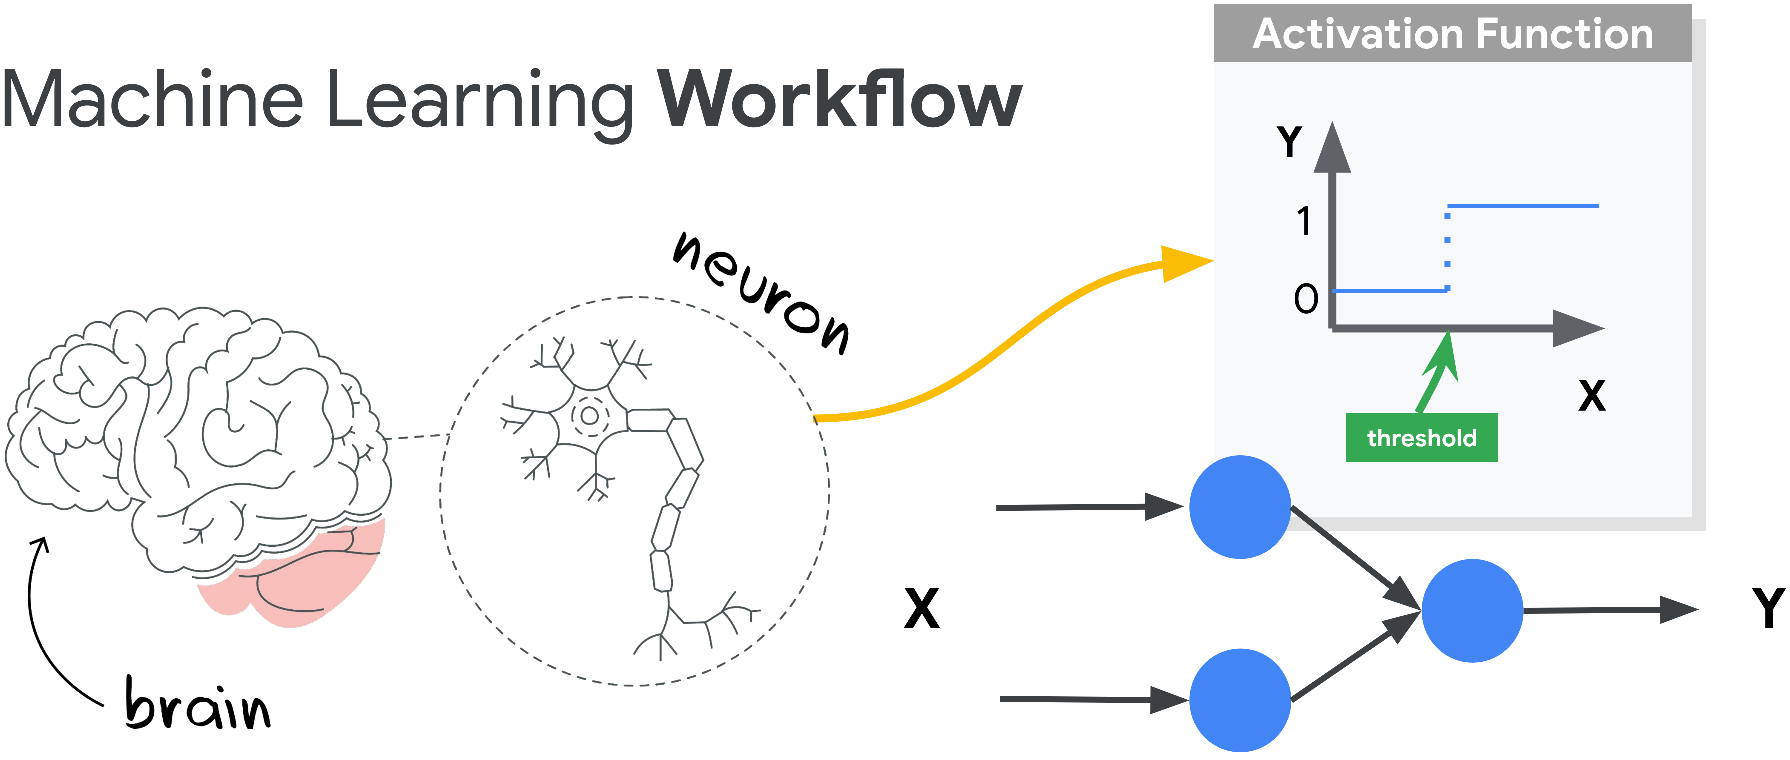

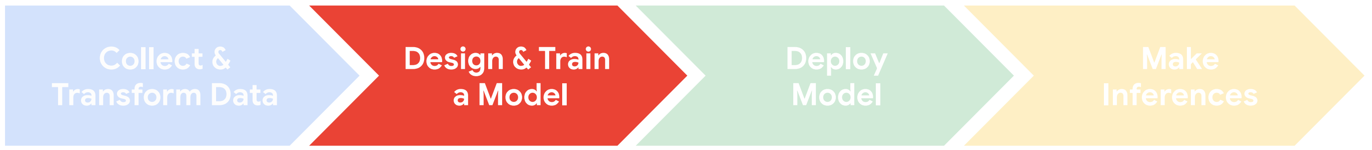

Machine Learning **Workflow**

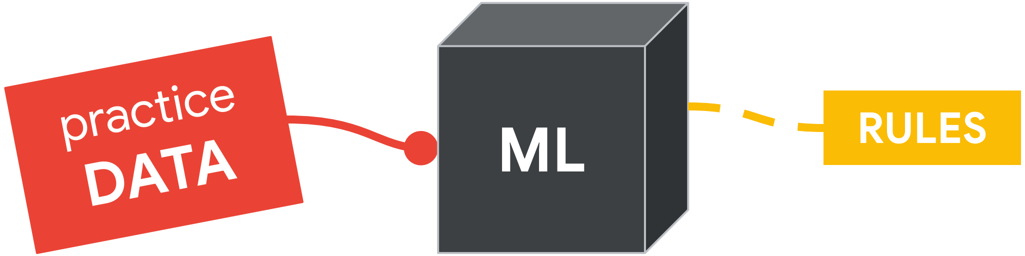

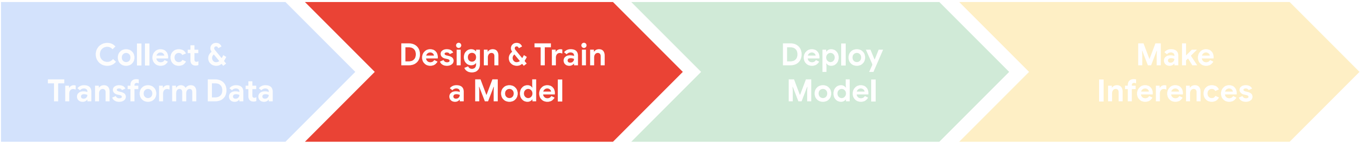

Machine Learning **Workflow**

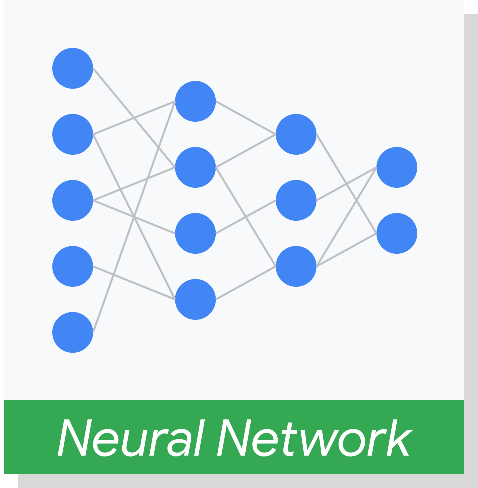

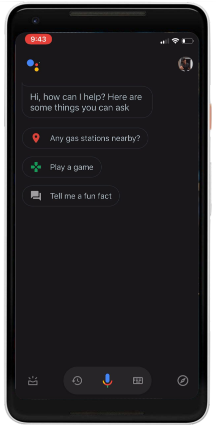

Collect & Transform Data → Design & Train a Model → **Deploy Model** → Make Inferences

Machine Learning **Workflow**

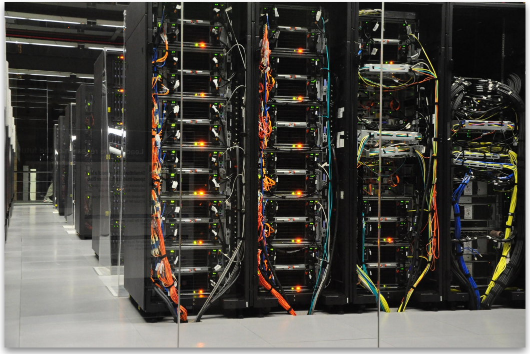

Google
Assistant

Collect & Transform Data → Design & Train a Model → **Deploy Model** → Make Inferences

Machine Learning **Workflow**

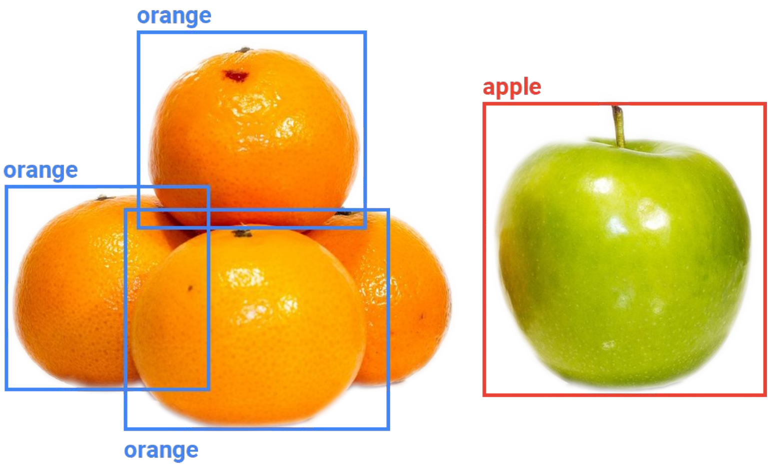

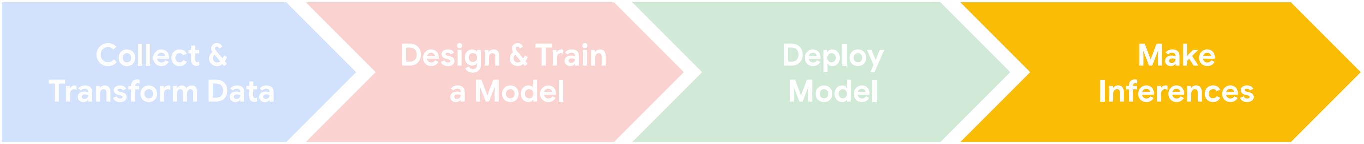

Bonus Content: Scaling TinyML

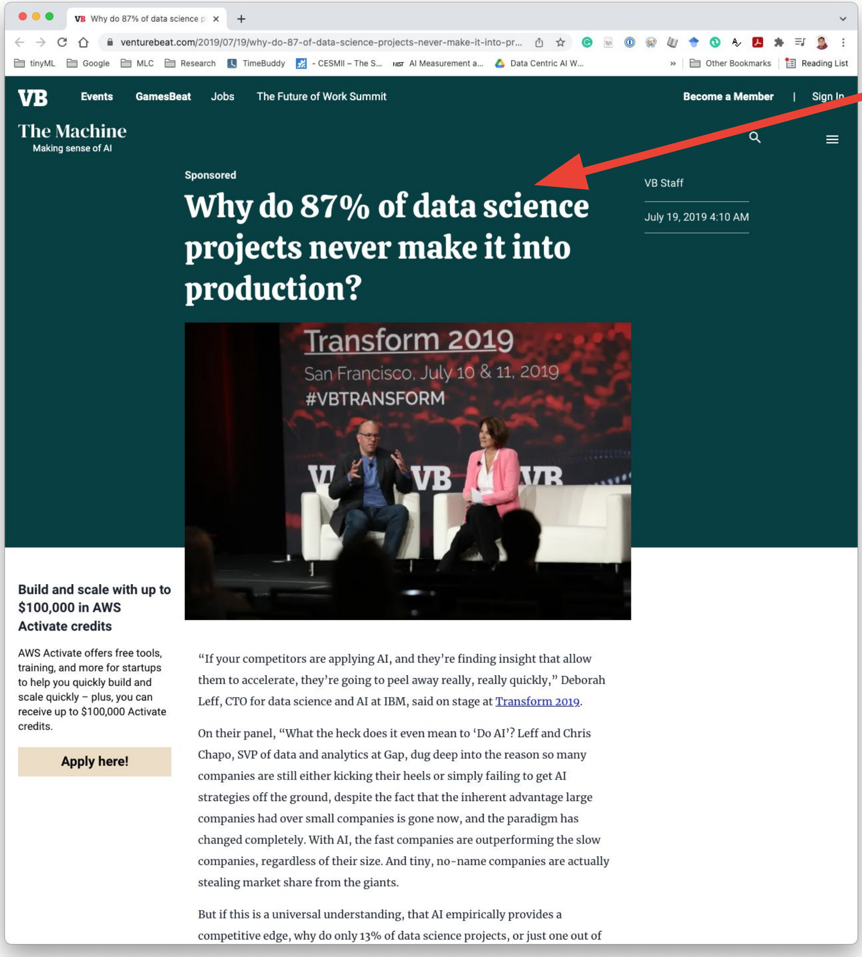

VB Events GamesBeat Jobs The Future of Work Summit Become a Member | Sign In

The Machine
Making sense of AI

Sponsored

Why do 87% of data science projects never make it into production?

VB Staff

July 19, 2019 4:10 AM

Transform 2019
San Francisco, July 10 & 11, 2019
#VBTRANSFORM

Build and scale with up to $100,000 in AWS Activate credits

AWS Activate offers free tools, training, and more for startups to help you quickly build and scale quickly – plus, you can receive up to $100,000 Activate credits.

Apply here!

"If your competitors are applying AI, and they're finding insight that allow them to accelerate, they're going to peel away really, really quickly," Deborah Leff, CTO for data science and AI at IBM, said on stage at Transform 2019.

On their panel, "What the heck does it even mean to 'Do AI'? Leff and Chris Chapo, SVP of data and analytics at Gap, dug deep into the reason so many companies are still either kicking their heels or simply failing to get AI strategies off the ground, despite the fact that the inherent advantage large companies had over small companies is gone now, and the paradigm has changed completely. With AI, the fast companies are outperforming the slow companies, regardless of their size. And tiny, no-name companies are actually stealing market share from the giants.

But if this is a universal understanding, that AI empirically provides a competitive edge, why do only 13% of data science projects, or just one out of

Let's quantify this a bit. In 2019 alone, approximately **USD 40 billions** were invested into privately held AI companies. If we extrapolate this and throw the approximated success rate of AI projects into these figures (and completely exclude intracompany ML investments), we reach the conclusion that in 2019, around **USD 38 billions were wasted due to unsuccessful Machine Learning projects.**

?

Predicts 2019: Analytics and BI Solutions

- Through 2020, 80% of AI projects will remain alchemy, run by wizards whose talents will not scale in the organization.
- Through 2022, only 20% of analytic insights will deliver business outcomes.
- By 2021, proof-of-concept analytic projects using quantum computing infrastructure will have outperformed traditional analytic approaches in multiple domains by at least a factor of 10

Source: https://blogs.gartner.com/andrew_white/2019/01/03/our-top-data-and-analytics-predicts-for-2019/


ML
Code

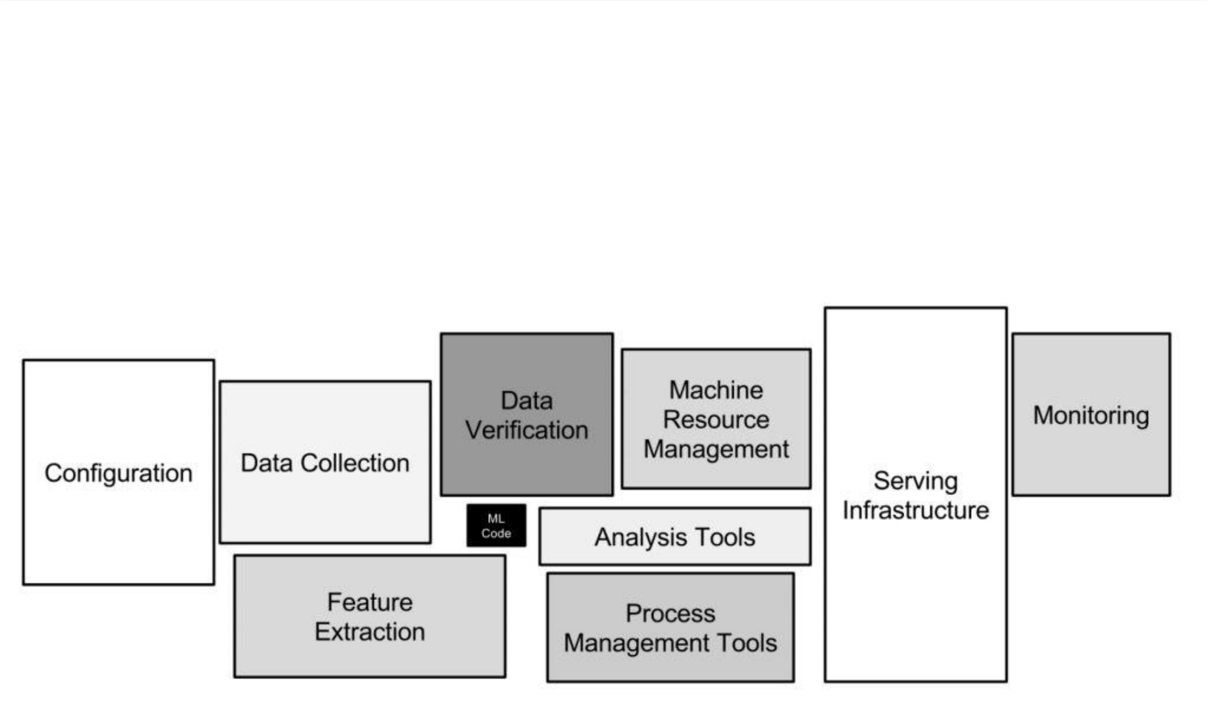
Configuration
Data Collection
Data Verification
Machine Resource Management
Serving Infrastructure
Monitoring
ML Code
Analysis Tools
Feature Extraction
Process Management Tools

Hidden Technical Debt in Machine Learning Systems

D. Sculley, Gary Holt, Daniel Golovin, Eugene Davydov, Todd Phillips
{dsculley,gholt,dgg,edavydov,toddphillips}@google.com
Google, Inc.

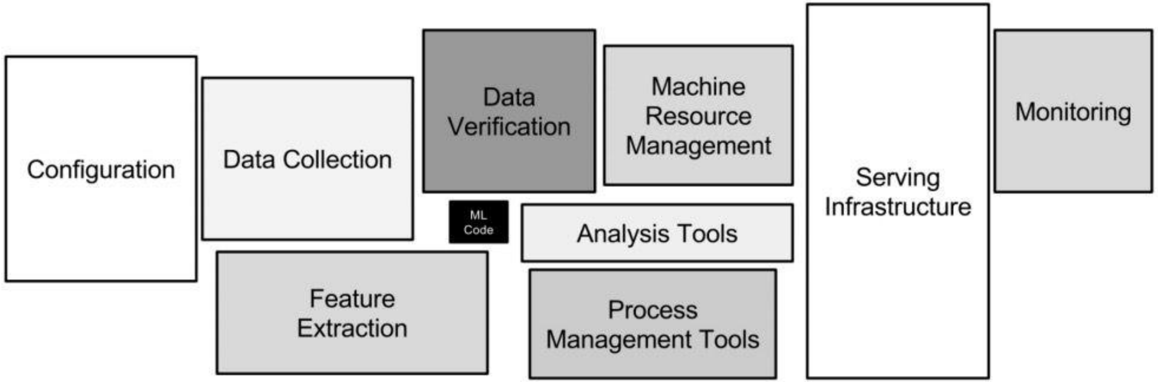

MLCode

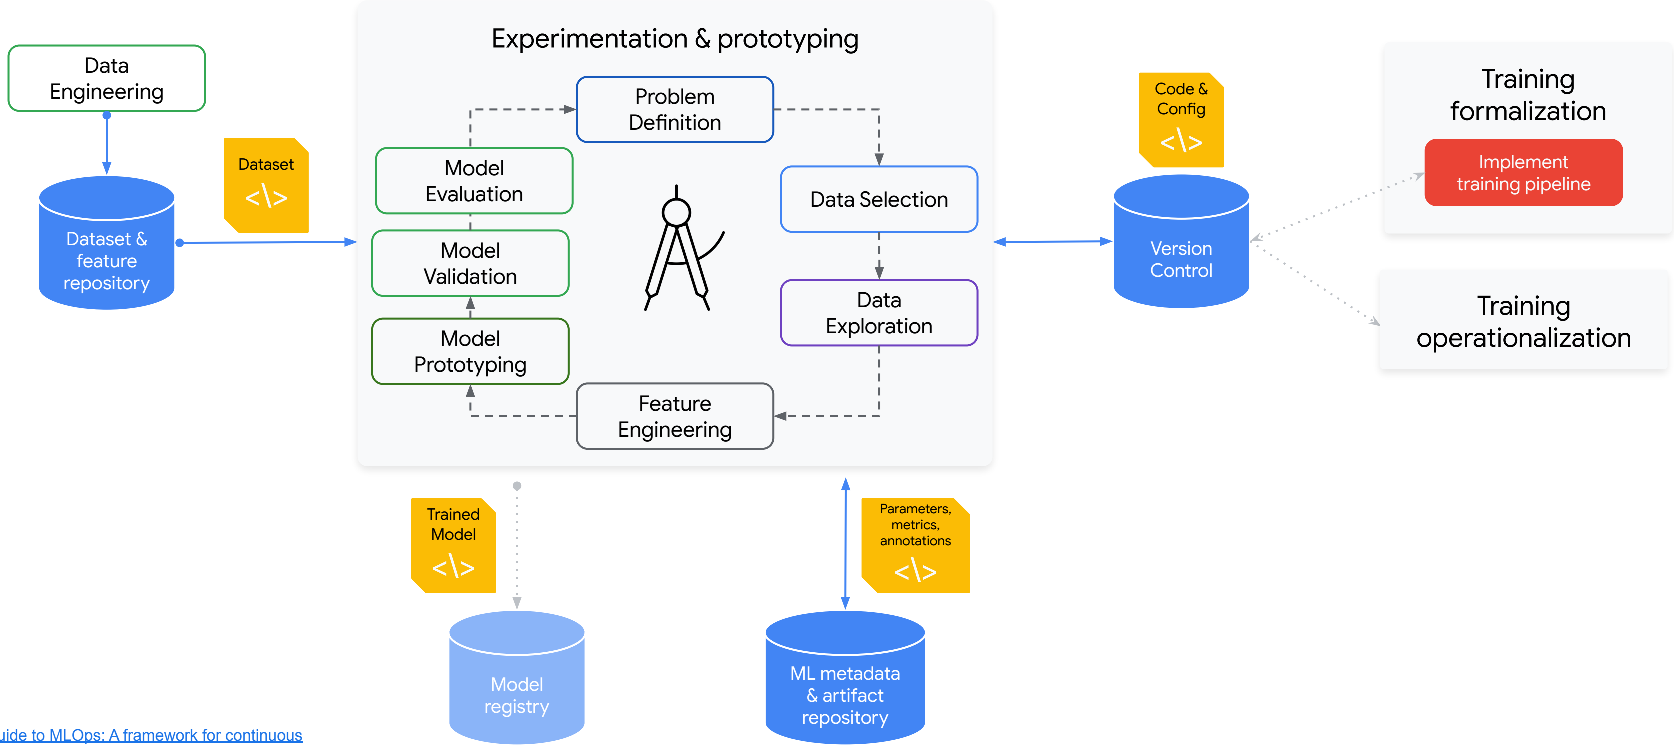

[Practitioners guide to MLOps: A framework for continuous delivery and automation of machine learning]()

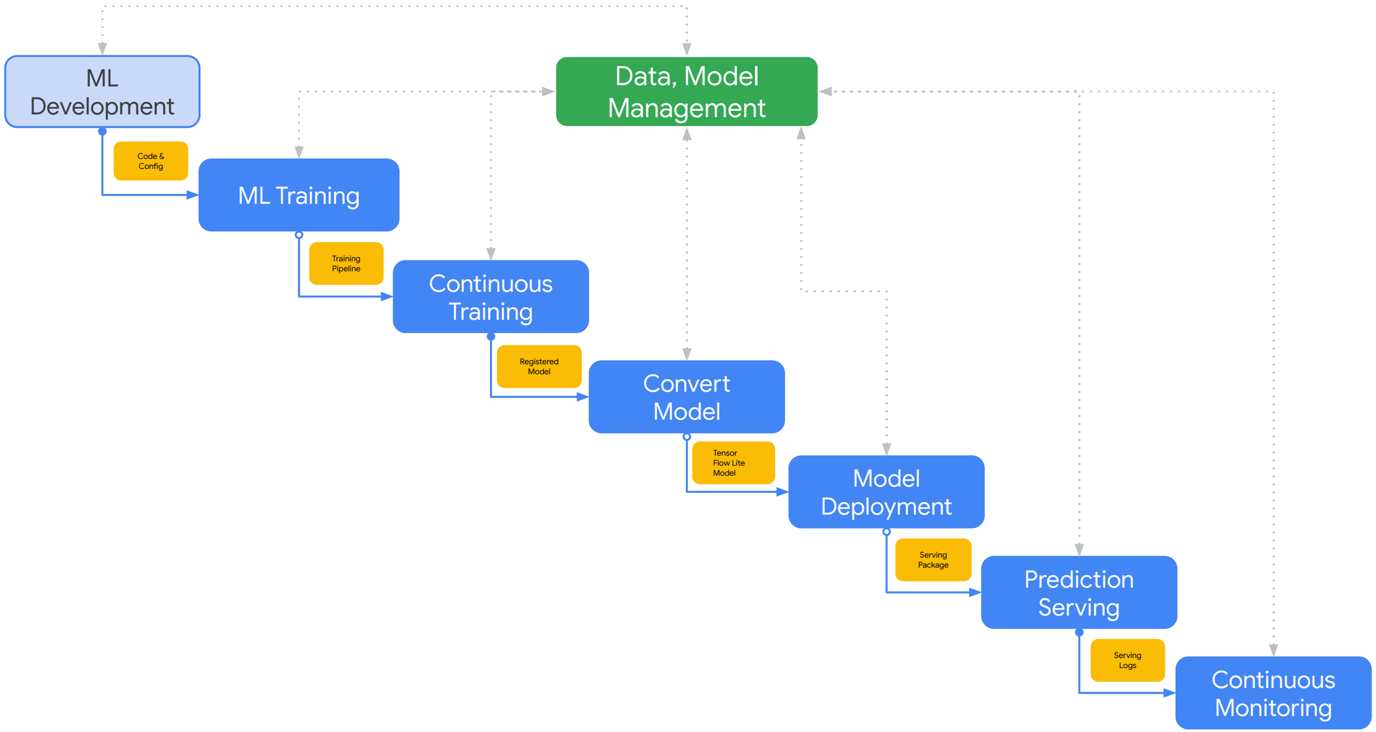
ML Development
Code & Config
ML Training
Training Pipeline
Continuous Training
Registered Model
Convert Model
Tensor Flow Lite Model
Model Deployment
Serving Package
Prediction Serving
Serving Logs
Continuous Monitoring
Data, Model Management

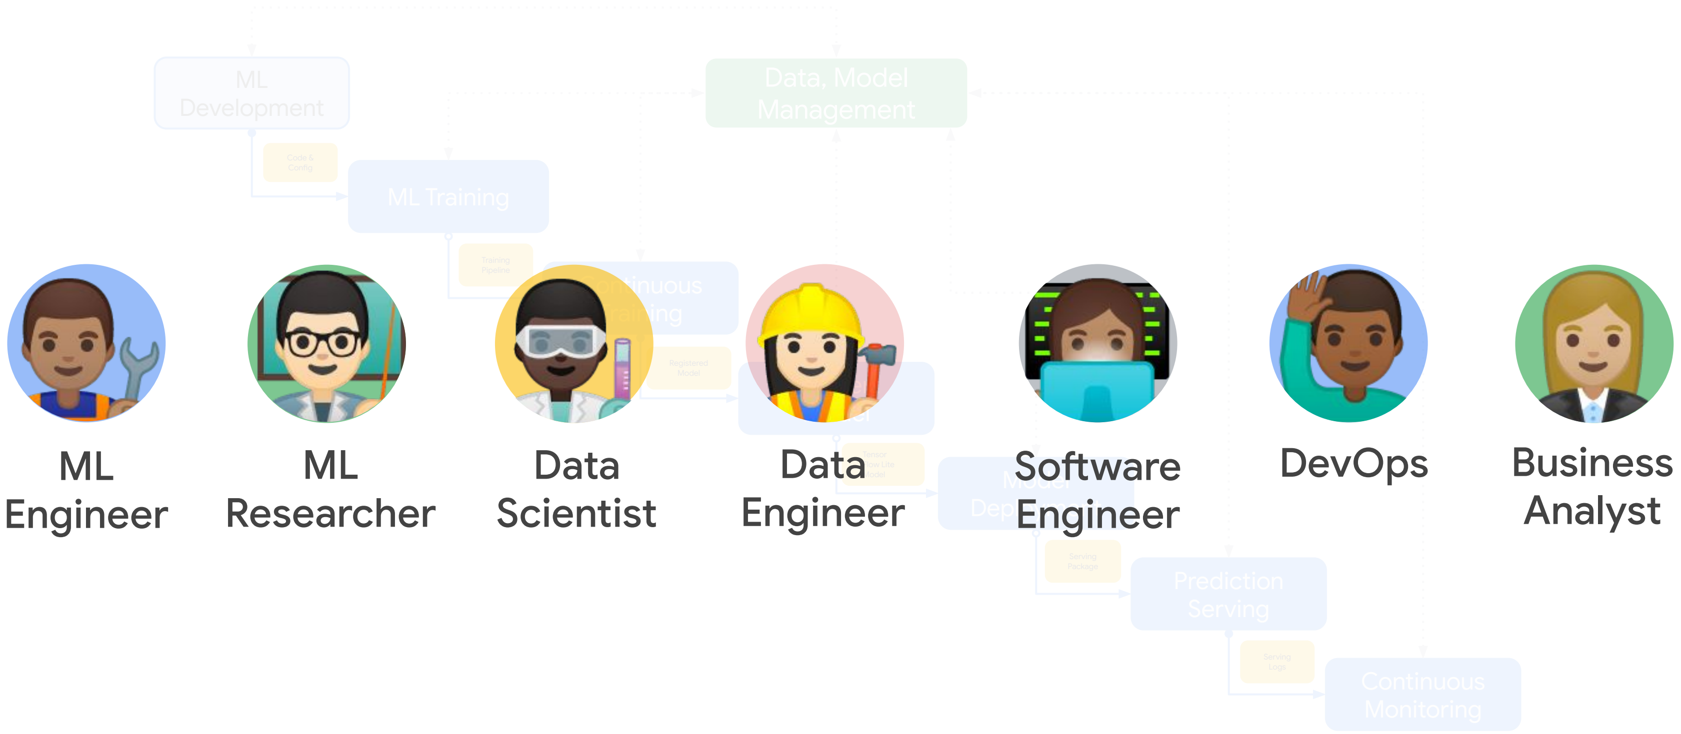
ML Development
Data, Model Management
ML Training
ML Engineer
ML Researcher
Data Scientist
Data Engineer
Software Engineer
DevOps
Business Analyst
Prediction Serving
Continuous Monitoring

ML Expertise

Deployment Expertise

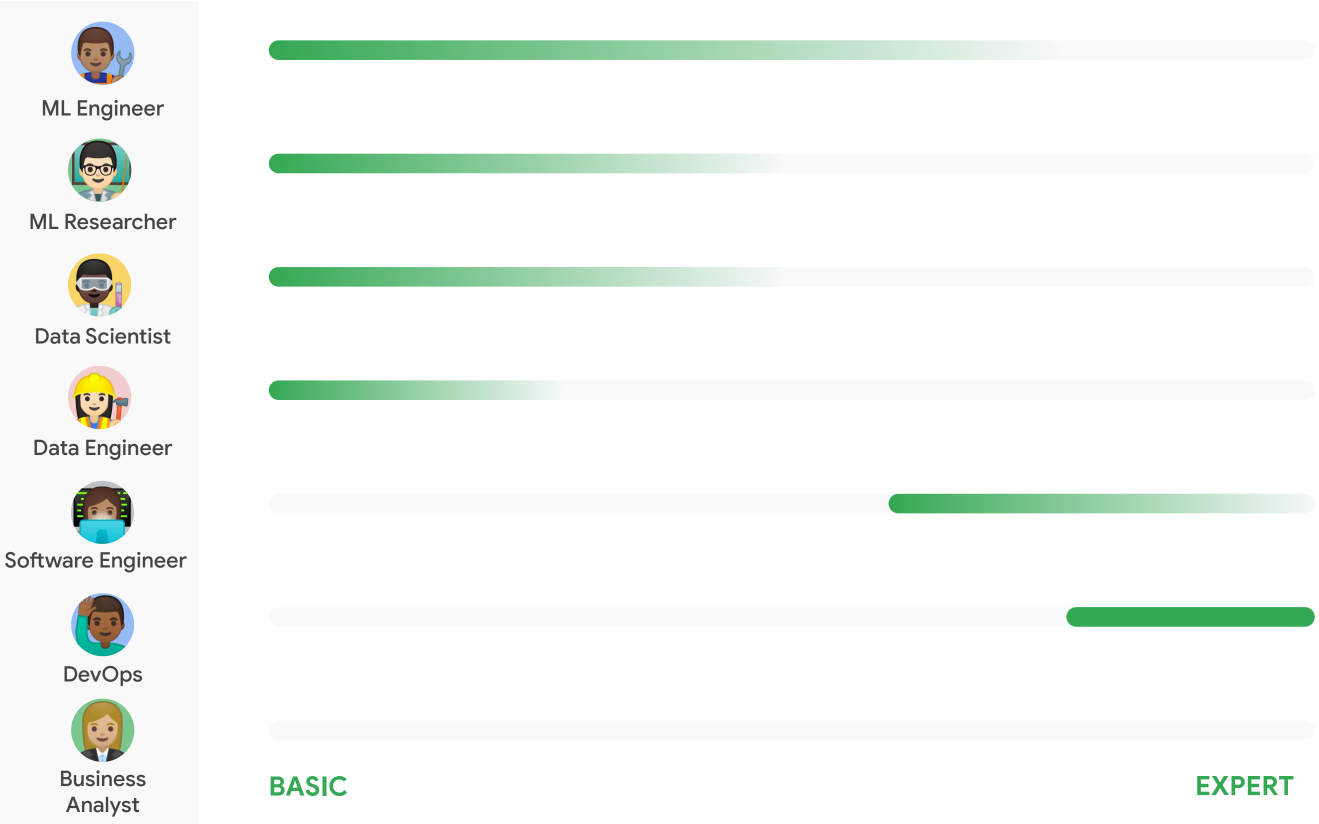

BREADTH

of experience, knowledge, & sectors

DEPTH

high skills in one discipline

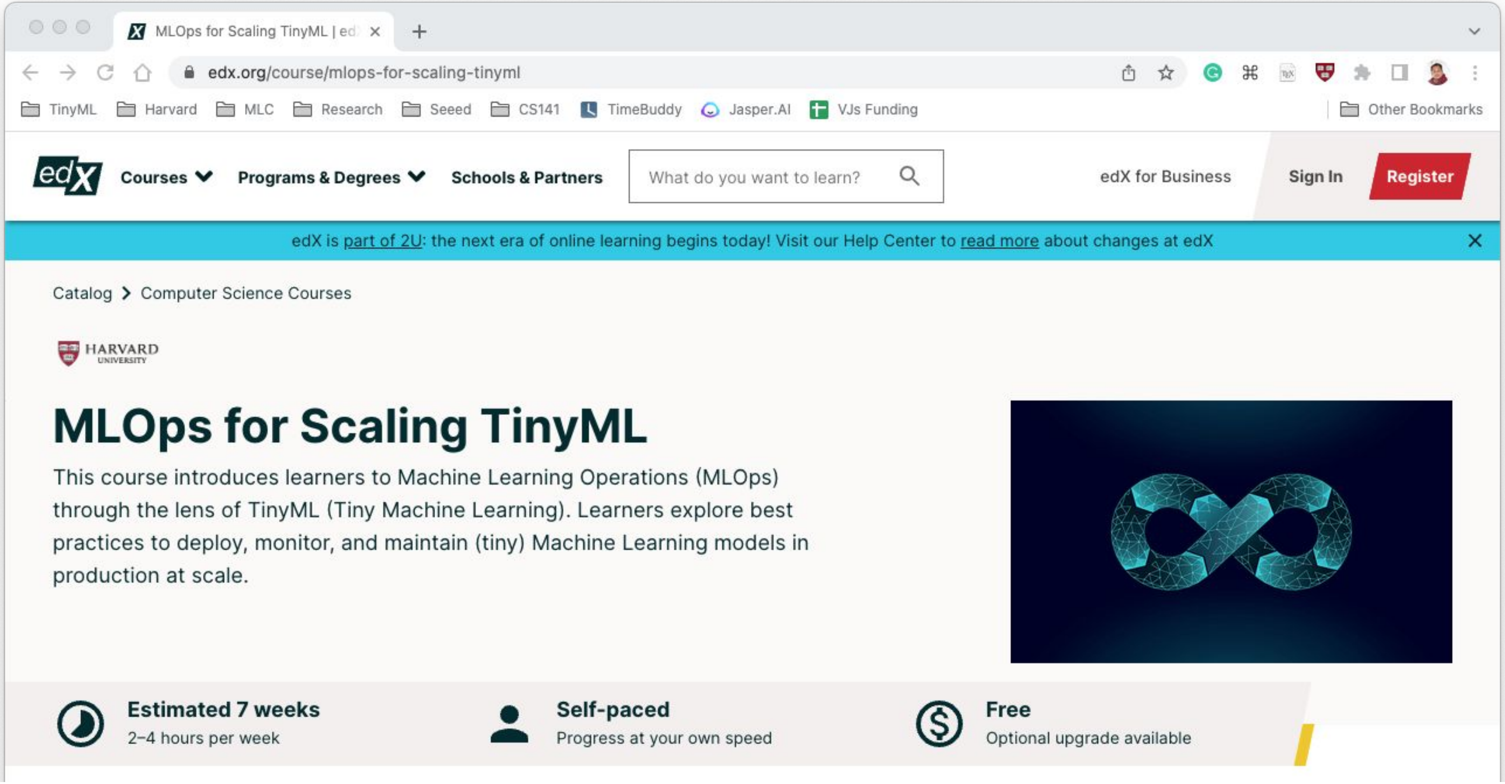
MLOps for Scaling TinyML | ed
edx.org/course/mlops-for-scaling-tinyml
TinyML
Harvard
MLC
Research
Seeed
CS141
TimeBuddy
Jasper.AI
VJs Funding
Other Bookmarks
edX
Courses
Programs & Degrees
Schools & Partners
What do you want to learn?
edX for Business
Sign In
Register
edX is part of 2U: the next era of online learning begins today! Visit our Help Center to read more about changes at edX
Catalog
Computer Science Courses
HARVARD
UNIVERSITY
MLOps for Scaling TinyML
This course introduces learners to Machine Learning Operations (MLOps) through the lens of TinyML (Tiny Machine Learning). Learners explore best practices to deploy, monitor, and maintain (tiny) Machine Learning models in production at scale.
Estimated 7 weeks
2–4 hours per week
Self-paced
Progress at your own speed
Free
Optional upgrade available

Our **website!**

tinyMLedu.org/EASI-22

home base for **all information**!

Our **team!**


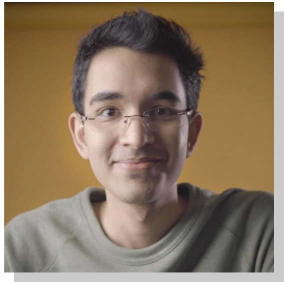



with help from **many more**

hágoónee’ 👋

see you **again at 12pm (Mountain Time)**