What is Deep Learning?

Deep Learning

Brad Quinton, Scott Chin

What is Deep Learning?

- Deep Learning is a specific type of **Machine Learning** using (normally large, or "Deep") neural networks

- Ok, then what is **Machine Learning**?

- **Machine Learning** is simply a technique to create new functions using <u>**example behavior**</u> rather than explicit instructions…

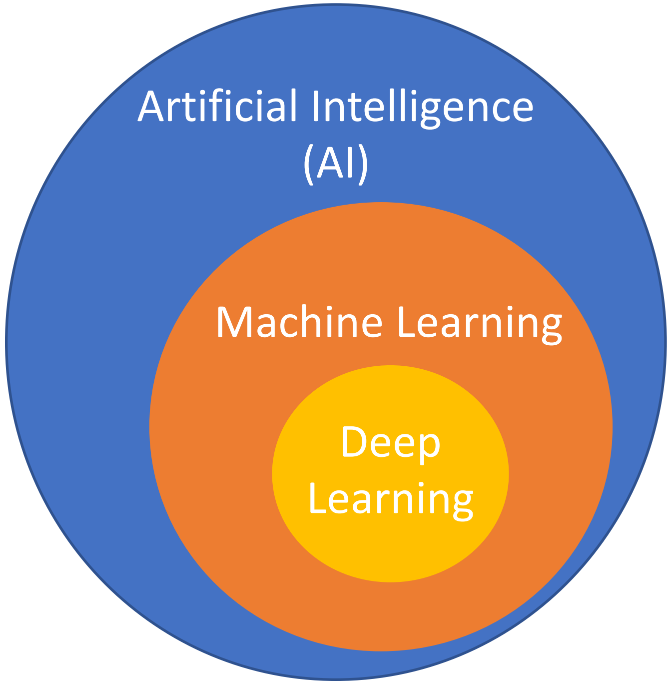

Why Learn From Examples?

- There are classes of problems for which it is difficult (or potentially impossible) to articulate as an explicit set of instructions

- Try to think of an explicit algorithm to sort pictures of cats from dogs…

Why Learn From Examples?

- Dogs :
  - 2 ears,
  - 2 eyes,
  - 4 legs,
  - 1 nose
  - whiskers
  - Colors: brown, white, black
- Cats :
  - 2 ears,
  - 2 eyes,
  - 4 legs,
  - 1 nose
  - whiskers
  - Colors: brown, white, black

- And, yet, I could show a child cat and dog pictures all day and they would be able to classify them with very high accuracy….

Why Learn From Examples?

- Some problems are just much easier to both articulate and solve using examples rather than creating explicit instructions

- In machine learning, these examples are normally called data, and when the example comes with the expected answer it is call **labelled data**

- **Labelled data** provides an example of an input to output mapping (x=>y) from which we would like the ML system to **generalize** for other similar inputs

ML creates *approximate* functions

- Machine Learning algorithms allow you to create **approximate** functions using a set of example data

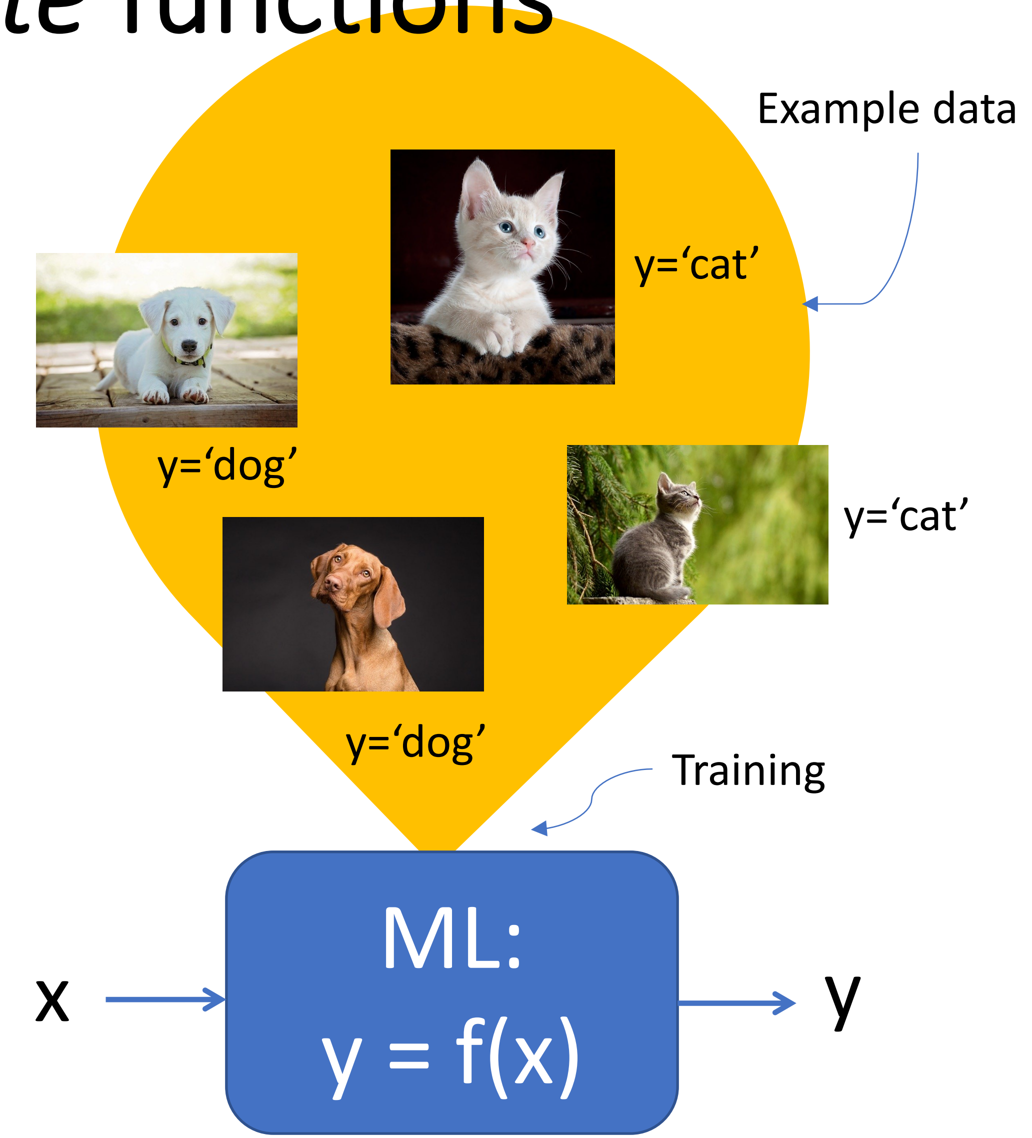

Brad Quinton, Scott Chin

ML creates *approximate* functions

- Machine Learning algorithms allow you to create **approximate** functions using a set of example data

- The resulting approximate function can be applied to **new data**

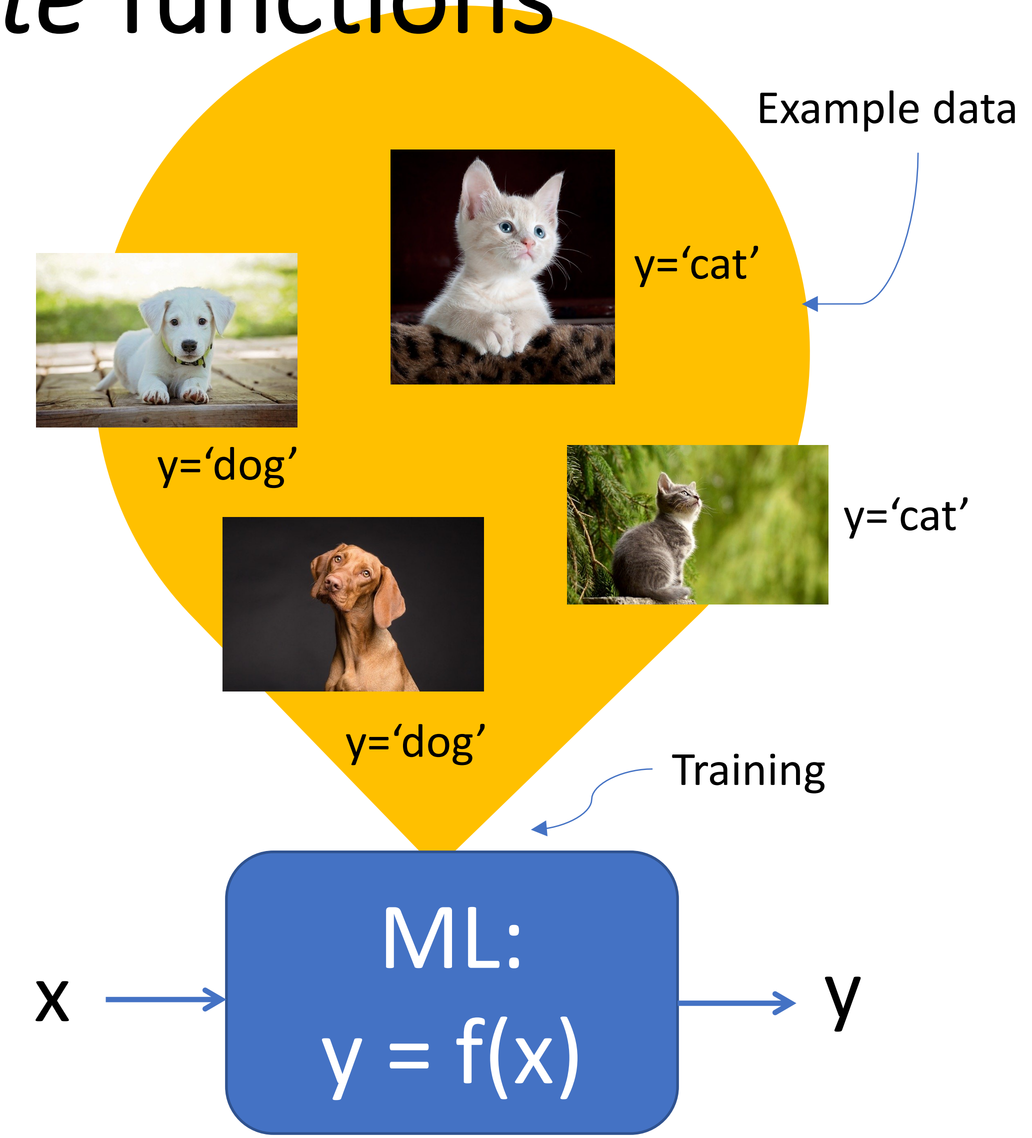

Brad Quinton, Scott Chin

ML creates *approximate* functions

- Machine Learning algorithms allow you to create **approximate** functions using a set of example data

- The resulting approximate function can be applied to **new data**

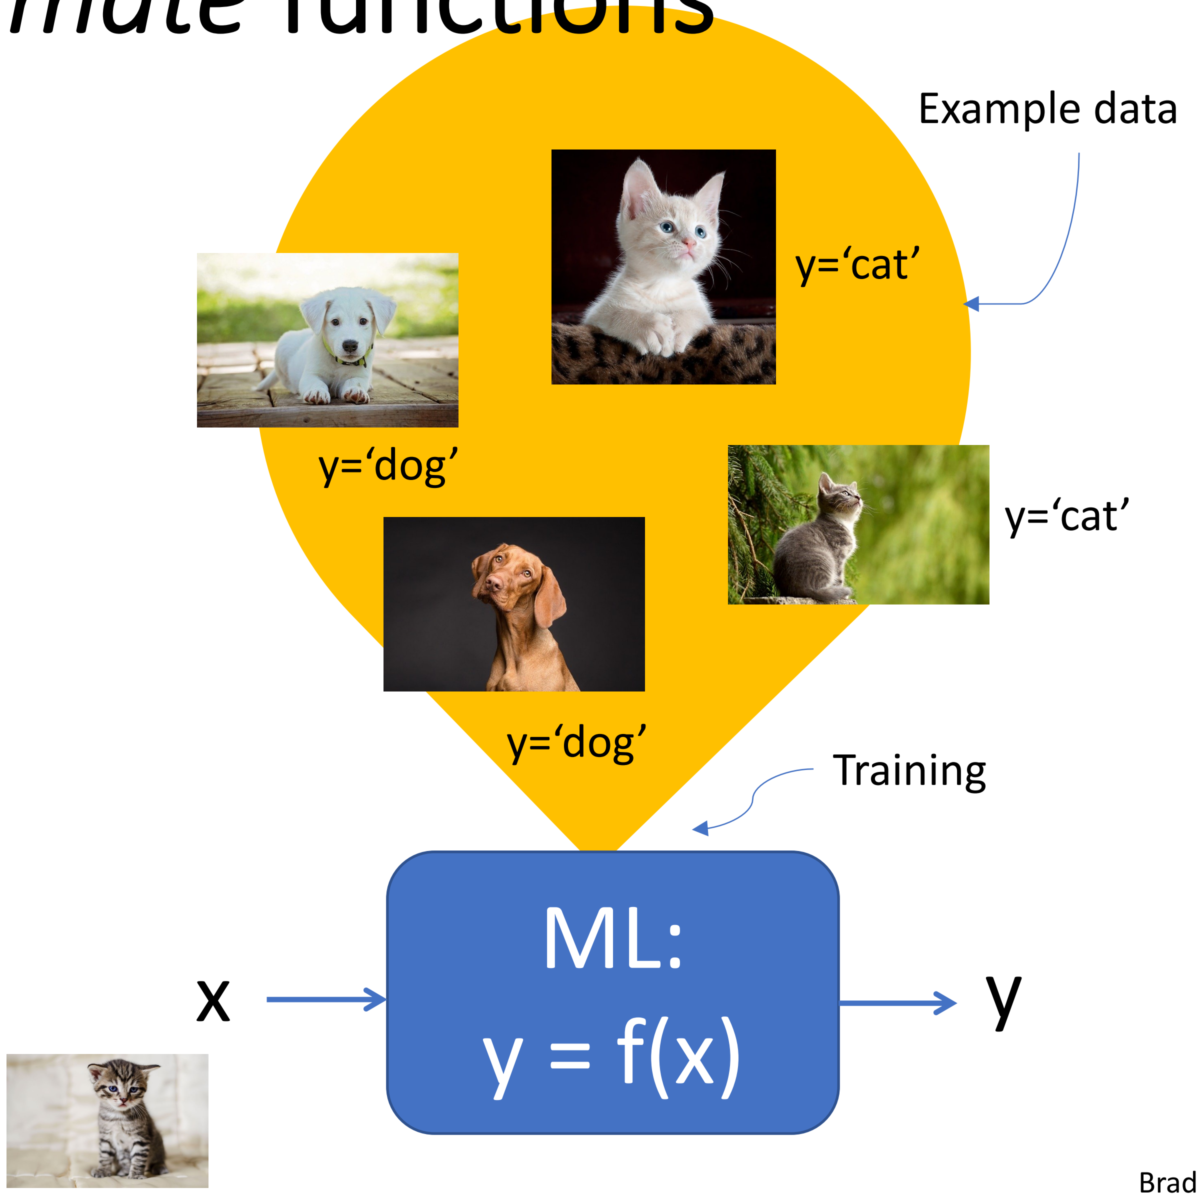

Brad Quinton, Scott Chin

ML creates *approximate* functions

- Machine Learning algorithms allow you to create **approximate** functions using a set of example data

- The resulting approximate function can be applied to **new data**

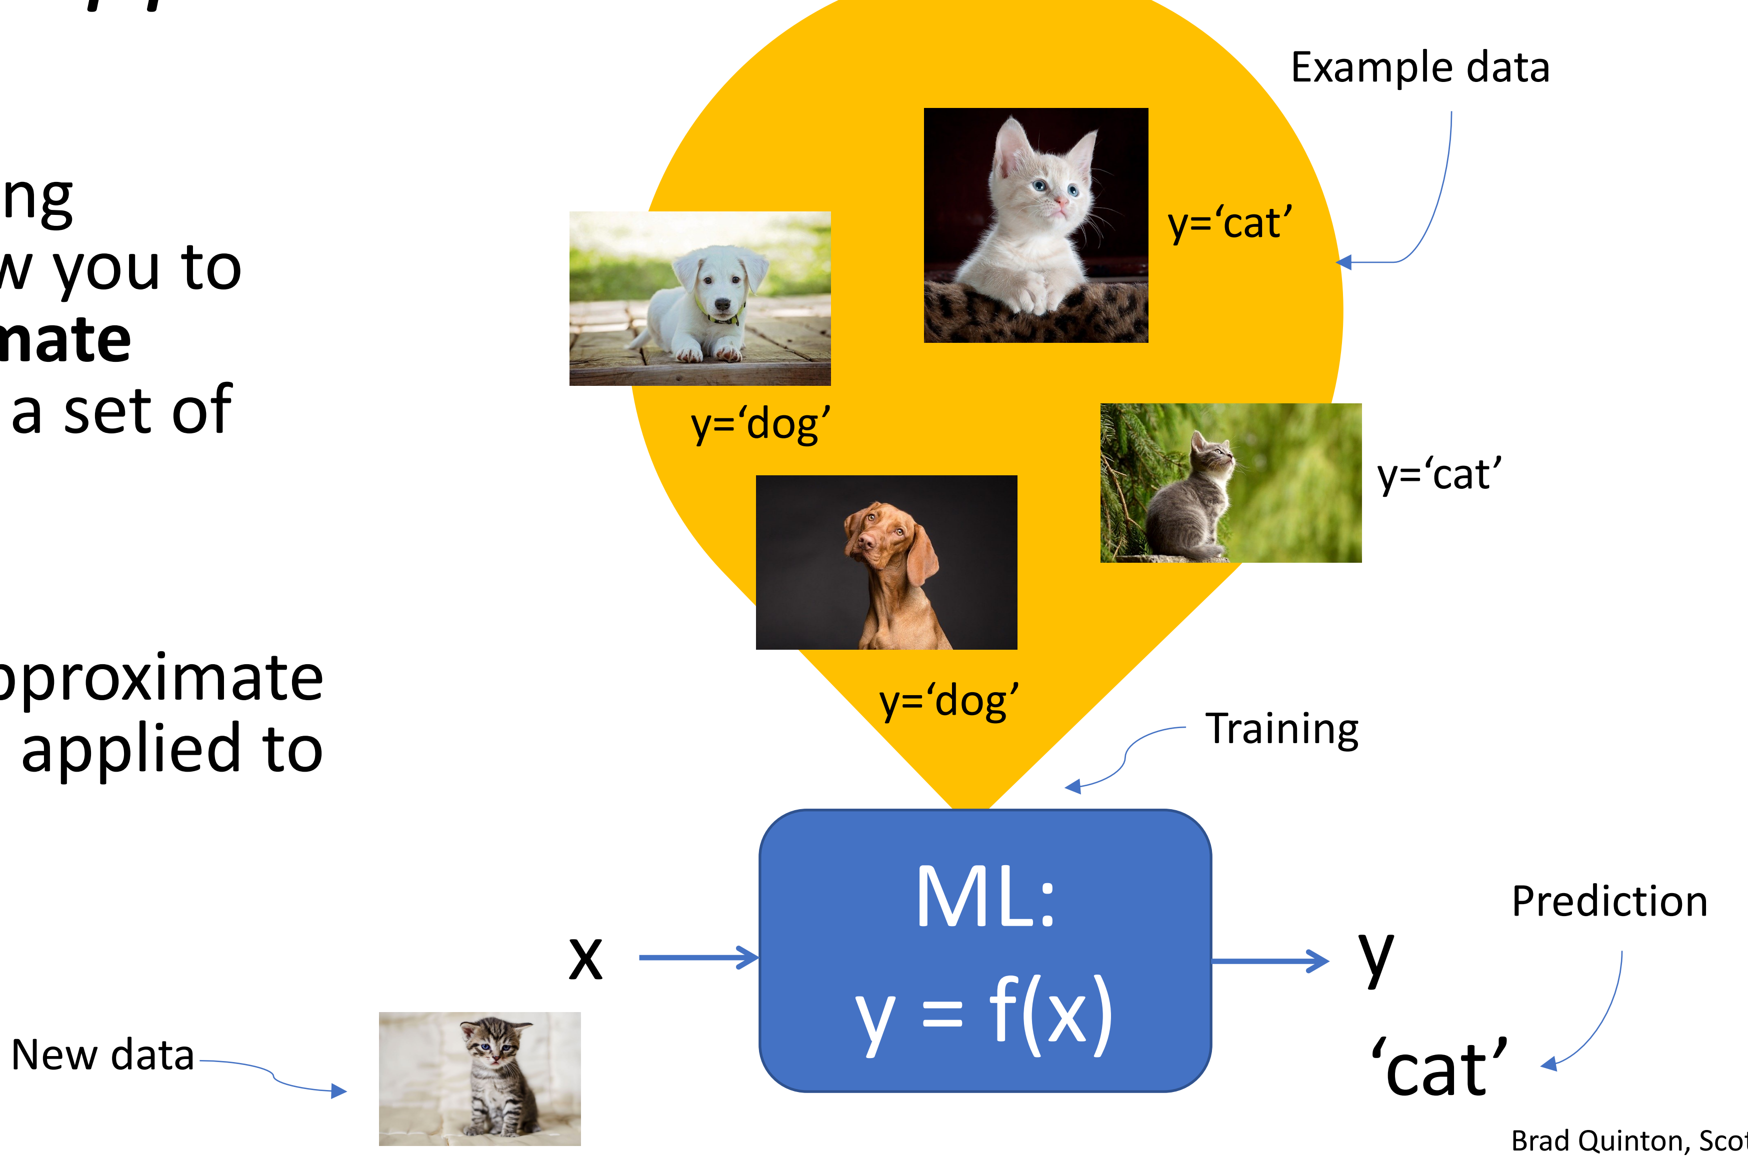

Brad Quinton, Scott Chin

The Prediction Can be Wrong!

- Because this is an *approximate* function, it is **NOT** always correct

- After training, if resulting approximate function is '**accurate enough'** on **new data** then it is useful, if not…back to the drawing board!

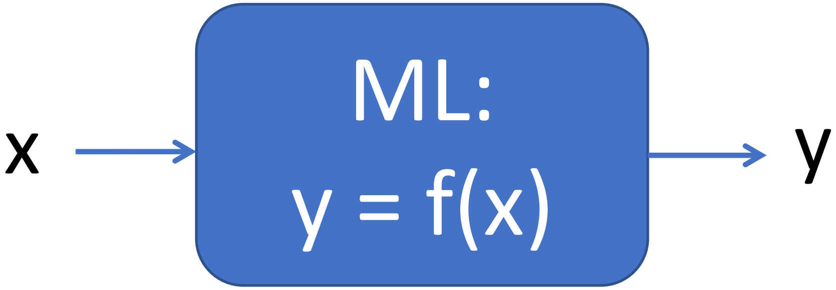

The Prediction Can be Wrong!

- Because this is an *approximate* function, it is **NOT** always correct

- After training, if resulting approximate function is ‘**accurate enough’** on **new data** then it is useful, if not…back to the drawing board!

x → ML: $y = f(x)$ → y ‘dog’

Brad Quinton, Scott Chin

The Prediction Can be Wrong!

- Because this is an *approximate* function, it is **NOT** always correct

- After training, if resulting approximate function is ‘**accurate enough’** on **new data** then it is useful, if not…back to the drawing board!

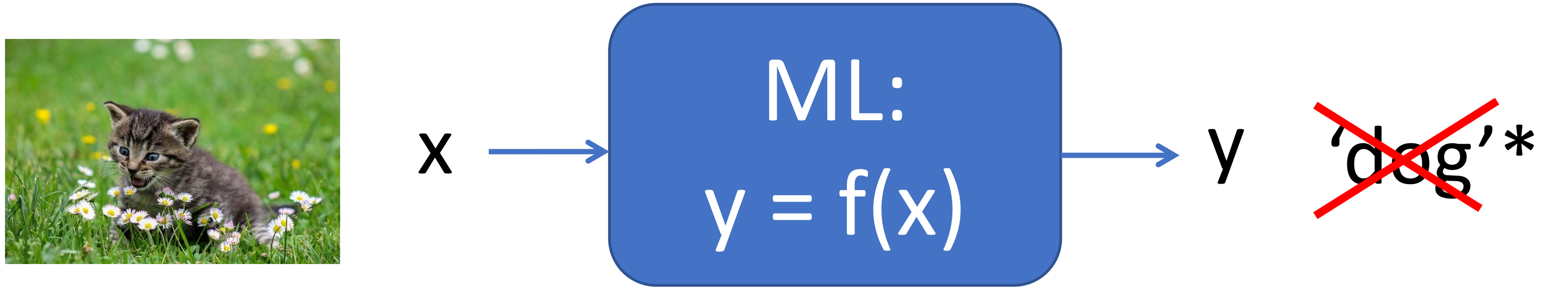

* As we will see in this course, it can be extremely difficult to understand why it is wrong!

What Problems Are Best Solved with ML?

Lots of **high-quality** data is available (or can be created).

AND

The desired output is clear, unambiguous and **testable**.

AND

The input to output relationship is **not already well understood**.

Things that DON'T need ML:

1. There is a clear and well understood **mathematical** relationship between input and output

```
>> y = 4 + 7
>> y = sqrt(8)
>> y = 5 * x^4 + 7
>> y = sort(3, 18, 9, 101, 3, 5, 10, 19, 884)
>> …
```

Things that DON'T need ML:

2. There is a clear and well understood **<u>physical</u>** relationship between the input and output.

There is no need to "learn" this function…

Things that DON'T need ML:

2. There is a clear and well understood **physical** relationship between the input and output.

There is no need to "learn" this function… $F = \frac{G \cdot M \cdot m}{r^2}$

Brad Quinton, Scott Chin

Things that DON'T need ML:

3. There is a clear and well understood **algorithmic** relationship between the input and output.

```
if ((door_open == FALSE) AND (key_valid == TRUE)) then
  set enable_engine TRUE
else
  set enable_engine FALSE
```

Things that can’t be solved with ML (yet, anyway):

1. Write a report on **why** people love pictures of cats.
2. **Explain** the causes of the French Revolution.
3. Create a successful business **plan** for a new product.
4. Solve crimes that require **deductive** reasoning
5. **Design** better ML systems (or many other complex things…)

Deep Learning Doesn’t know the ‘why’…

- The training process is essentially very complex ”curve fitting”, which provides answers, but not **justifications**…

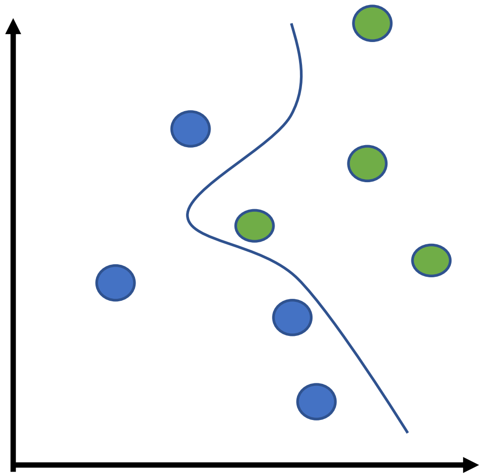

Deep Learning Doesn’t know the ‘why’…

- The training process is essentially very complex ”curve fitting”, which provides answers, but not **justifications**…

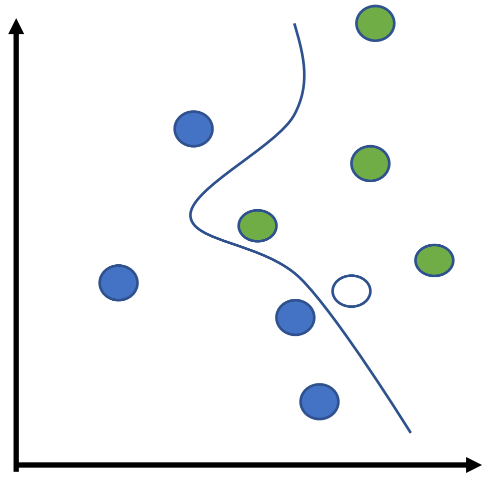

Deep Learning Doesn’t know the ’why’…

- The training process is essentially very complex ”curve fitting”, which provides answers, but not **justifications**…

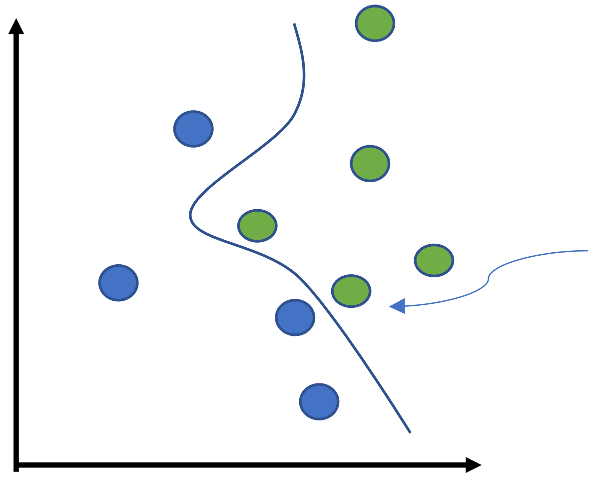

We predict this dot is green, but how do we explain our choice?*

*this is a very active, but largely unsolved research area

Brad Quinton, Scott Chin

Deep Learning Doesn't know the 'why'…

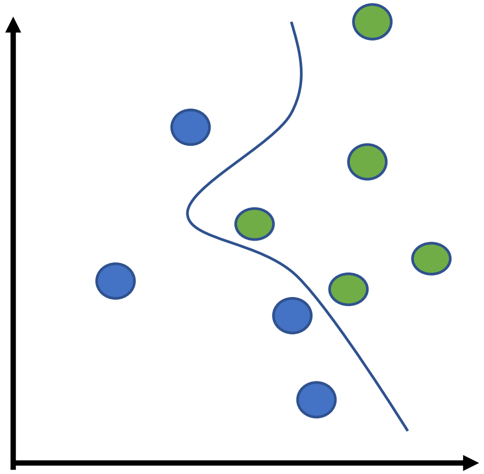

ML Trained Classifier

Vs.

```
function f(x)
  if (x > 0.5 AND y > 0.5)
    f = green
  else
    f = blue
```

"Hand-coded" Explicit Function

Deep Learning Doesn't know the 'why'…

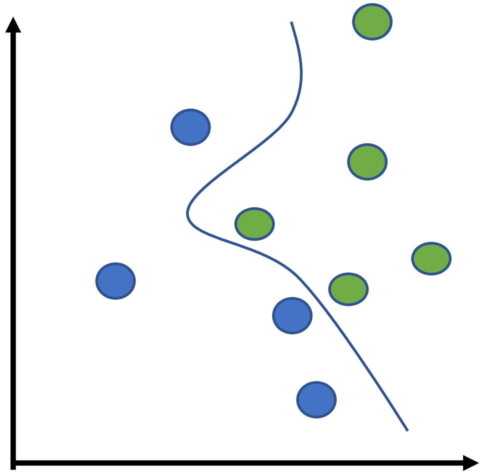

ML Trained Classifier

Vs.

```
function f(x)
  if (x > 0.5 AND y > 0.5)
    f = green
  else
    f = blue
```

"Hand-coded" Explicit Function

Potentially More Accurate

Deep Learning Doesn't know the 'why'…

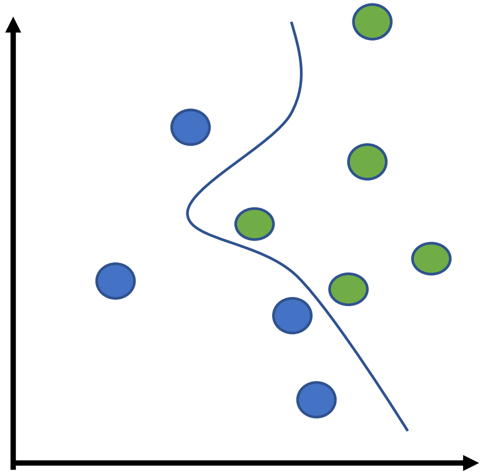

ML Trained Classifier

Vs.

```
function f(x)
  if (x > 0.5 AND y > 0.5)
    f = green
  else
    f = blue
```

"Hand-coded" Explicit Function

Potentially More Accurate

Normally Easier to Understand

Why Neural Networks?

- Neural networks are not the only way to learn from examples. What makes Neural Networks special is that we know how to **train them efficiently**

- As we will see, an algorithm called **backpropagation** will very often quickly and efficiently find a high quality approximate function

- The effectiveness of backpropogration in many layered neural networks is **basis of most of the success** of Deep Learning today…

Why "Deep" Neural Networks?

- Deep Neural Networks are just networks with many trainable layers, which allows them to express very complex (approx.) functions

- They are generally effective when we have a **very large set of training data.**

Small Training Set + Deep NN => likely **poor** results

Large Training Set + "shallow" NN => likely **poor** results

Large Training Set + Deep NN => potential for **high quality** results

Where does Deep Learning work well?

1. Problems where the input is *unstructured data*
    - Images/video
    - Radar/lidar/sonar
    - X-ray/MRI/ultrasound/EEG/EKG
    - Audio/voice
    - Natural Language
    - ‘mixed’ data sources (radar and image data, 3D MRI, etc.)

2. Problems with complex relationships but clear goals
    - Classifying images
    - Identifying objects
    - Winning chess
    - Predicting consumer behavior

Rules of Thumb?

“*Pretty much anything you [or an expert] could do with a second of thought, we can probably now or soon automate using supervised learning”* - **Andrew Ng** (Stanford Professor)

My thoughts: *“If you believe there is truly an underlying, consistent mapping from input to output (i.e. a function does exist) AND that you can get enough examples to ‘sample’ the input space, ML is likely to work well”*

Is ML and Deep Learning the same as AI?

- No, ML and Deep Learning are a type of AI

- There are many types of AI that don't use ML:
  - The IBM DeepBlue systems that famously beat chess champion Gary Kasprov used a "brute-force" search based AI (although, interestingly Google's Alpha Zero does use Deep Learning and is now much better at chess than DeepBlue)
  - Google Maps does not use Deep Learning, at least for route finding
  - The "wizard" in TurboTax is AI not based on ML (it is an "expert system")

- You will often hear people (even me) refer to a Deep NN, as "the AI", this is to mean "this instance of AI", not the only AI!

Why are people so excited by Deep Learning?

1. ML and Deep Learning can address many problem areas **previously** considered **intractable**.
2. People are making **real money** today with ML and Deep Learning.
3. Deep Learning works best with **lots of data** and **lots of computing power**, both are now widely available.
4. Mobile computing is pushing interaction with a lot of **unstructured data** (audio, images, videos, etc.)
5. Confusion between AI and AGI (especially in the media).
6. It just “sounds cool”.

Recent Successes in Deep Learning

- Image classification (AlexNet ImageNet)
- Gaming (Google's DeepMind AlphaZero)
- Natural language processing (OpenAI GPT-2)
- Voice Recognition (Amazon Echo)
- Audio/Video Generation (DeepFake Videos, DeepFake Detection Challenge)
- Face recognition (Apple Face ID)
- Many, many more….

AlexNet Success in ImageNet 2012

- AlexNet was developed at the University of Toronto for an image classification competition called **ImageNet**
- 1 million images, each labeled “cat”, “dog”, “apple”, etc.
- Competitors train their machine learning software on some of the images and then try to predict labels for never seen before images
- In 2012, the AlexNet team **dramatically** won, with a “top-5 error rate” of 16%. The nearest rival that year had a 26% error rate.
- AlexNet was the first to demonstrate the **effectiveness of Deep NNs** for problems when large datasets where available.

AlexNet Success in ImageNet 2012

- What do these images look like?
- Here are some examples…

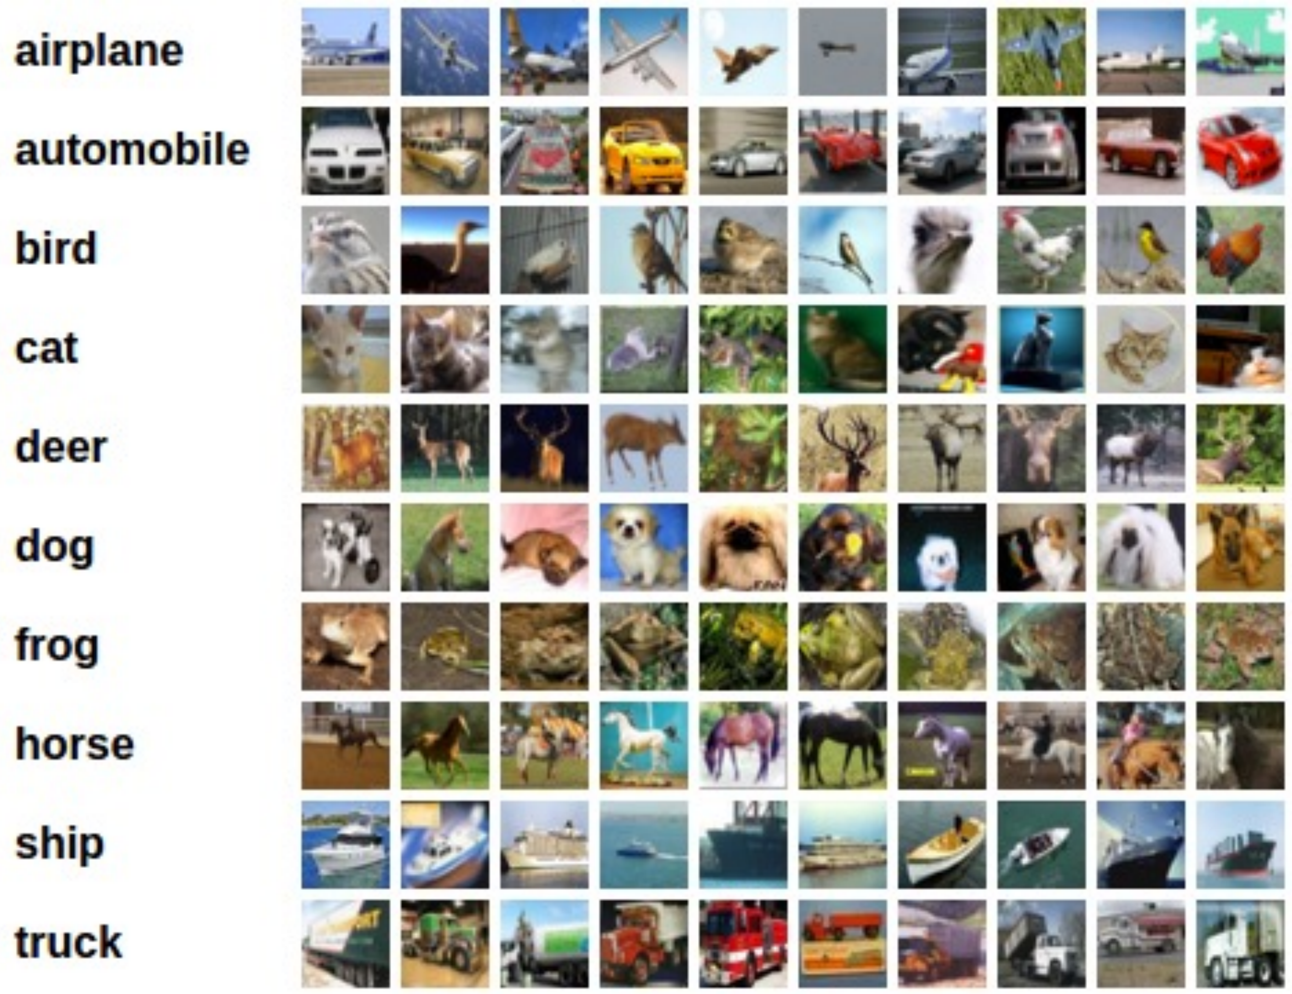

https://pytorch.org/tutorials/beginner/blitz/cifar10_tutorial.html

DeepMind AlphaZero

1. DeepMind's AlphaZero Clobbered The Top AI Champions In Go, Shogi, And Chess

Following its stunning win over the best human Go player in 2016, AlphaGo was upgraded a year later into a generalized and more powerful incarnation, AlphaZero. Free of any human guidance except the basic game rules, AlphaZero learned how to play master-level chess by itself in just four hours. It then proceeded to trounce Stockfish (the top AI chess player) in a 100-game match — without losing a single game.

https://www.forbes.com/sites/mariyayao/2018/02/05/12-amazing-deep-learning-breakthroughs-of-2017/#4986f7cc65db

OpenAI GPT-2 Natural Language Processing

Artificial intelligence (AI)

New AI fake text generator may be too dangerous to release, say creators

The Elon Musk-backed nonprofit company OpenAI declines to release research publicly for fear of misuse

Alex Hern

@alexhern

Thu 14 Feb 2019
12.00 EST

https://www.theguardian.com/technology/2019/feb/14/elon-musk-backed-ai-writes-convincing-news-fiction

OpenAI GPT-2 Natural Language Processing

SYSTEM PROMPT (HUMAN-WRITTEN)

In a shocking finding, scientist discovered a herd of unicorns living in a remote, previously unexplored valley, in the Andes Mountains. Even more surprising to the researchers was the fact that the unicorns spoke perfect English.

MODEL COMPLETION (MACHINE-WRITTEN, 10 TRIES)

The scientist named the population, after their distinctive horn, Ovid's Unicorn. These four-horned, silver-white unicorns were previously unknown to science.

Now, after almost two centuries, the mystery of what sparked this odd phenomenon is finally solved.

Dr. Jorge Pérez, an evolutionary biologist from the University of La Paz, and several companions, were exploring the Andes Mountains when they found a small valley, with no other animals or humans. Pérez noticed that the valley had what appeared to be a natural fountain, surrounded by two peaks of rock and silver snow.

https://openai.com/blog/better-language-models/

Face ID Apple iPhone X

"With the advent of deep learning, and its application to computer vision problems, the state-of-the-art in face detection accuracy took an enormous leap forward…

Compared to traditional computer vision, the learned models in deep learning require orders of magnitude more memory, much more disk storage, and more computational resources."

Figure 1. A revised DCN architecture for face detection

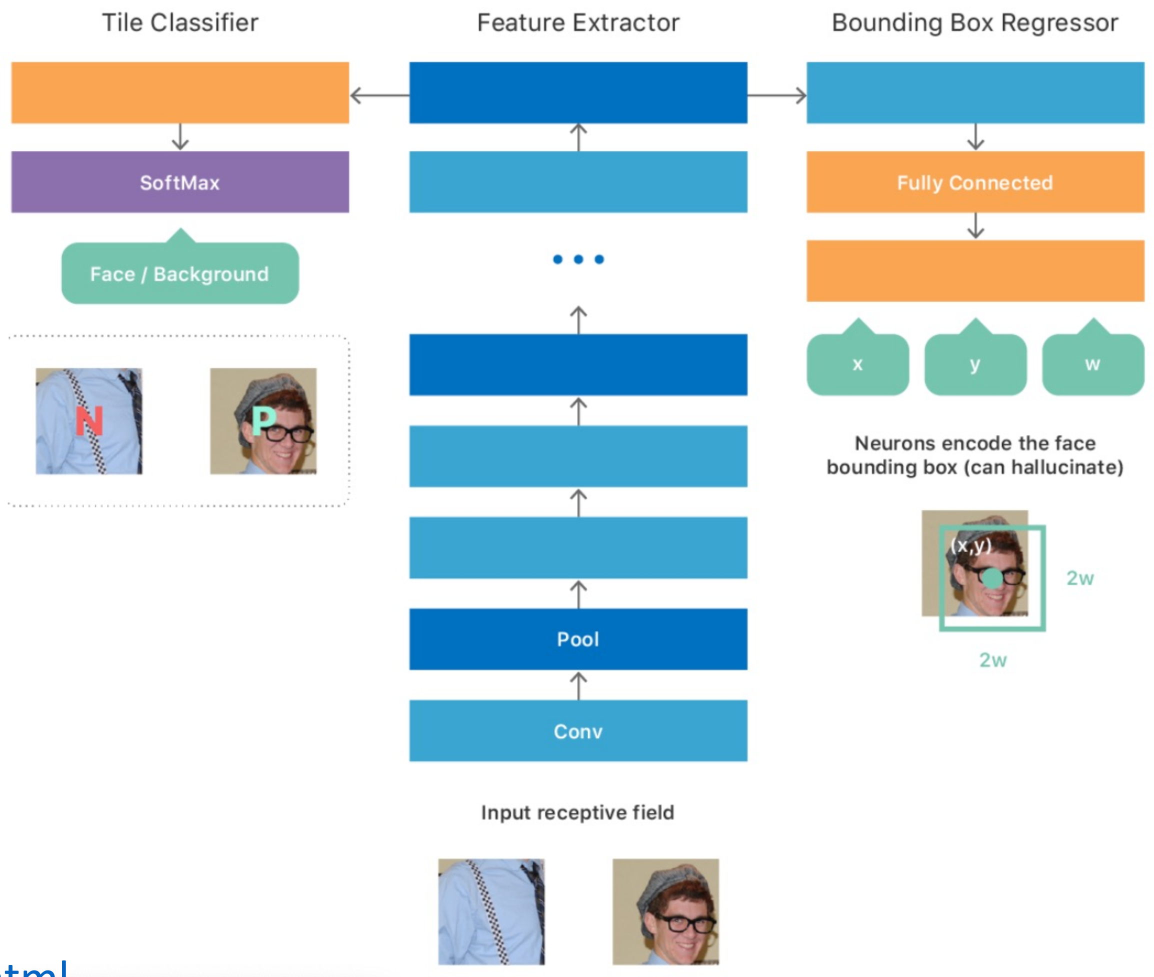

https://machinelearning.apple.com/2017/11/16/face-detection.html

Brad Quinton, Scott Chin

Voice Recognition on Amazon Echo

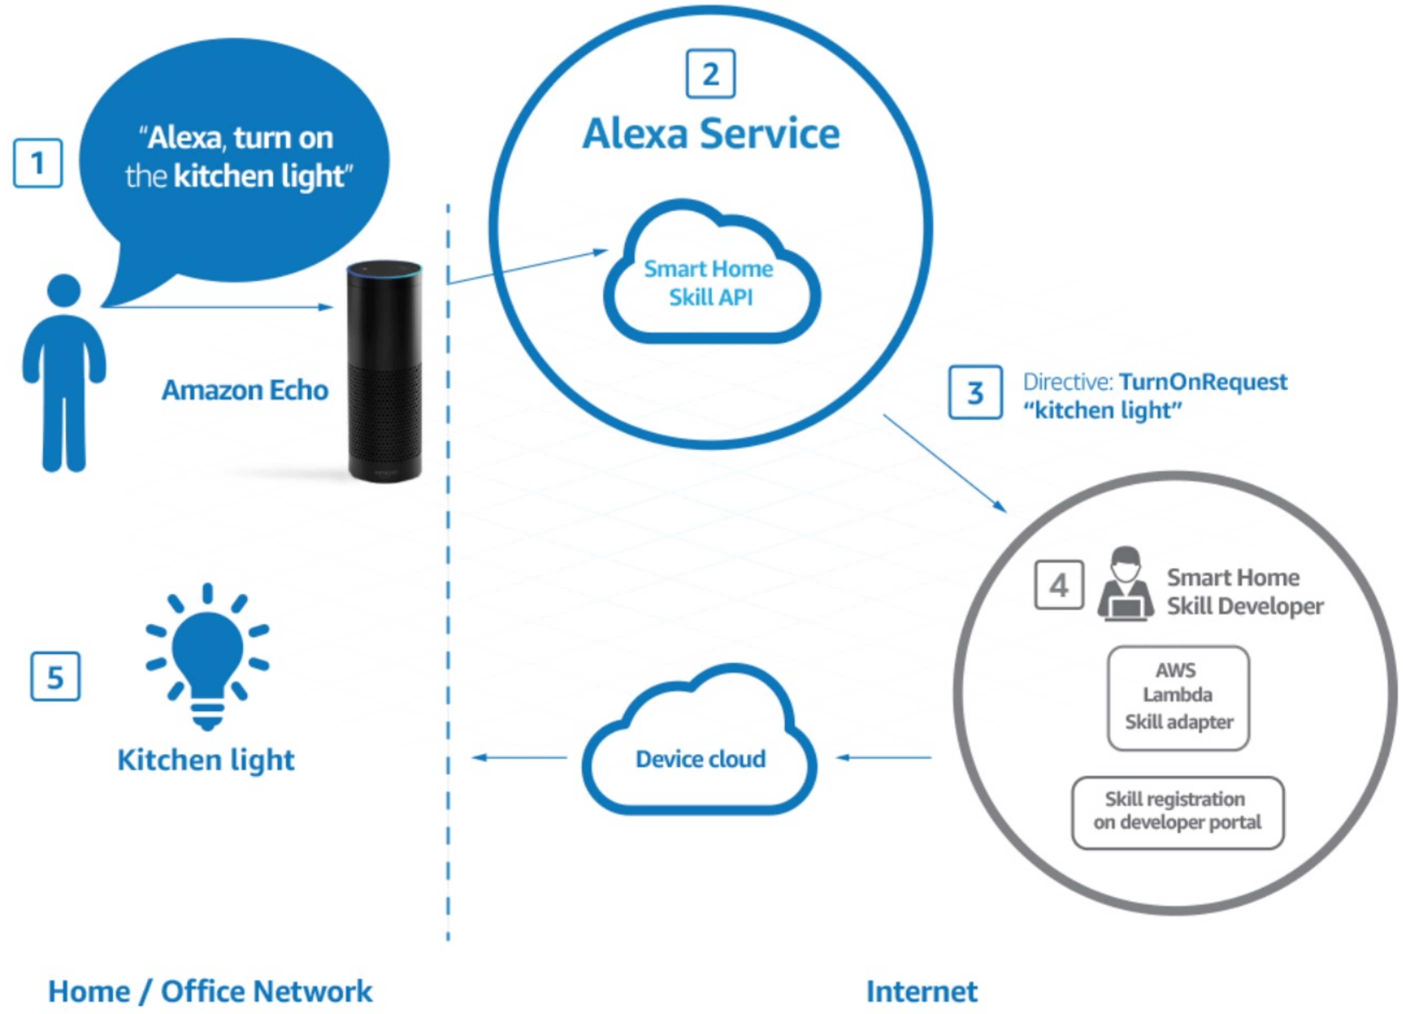

https://towardsdatascience.com/how-amazon-alexa-works-your-guide-to-natural-language-processing-ai-7506004709d3

FaceSwap / DeepFake Detection Challenge

https://deepfakedetectionchallenge.ai

Why is so much $ going into Deep Learning?

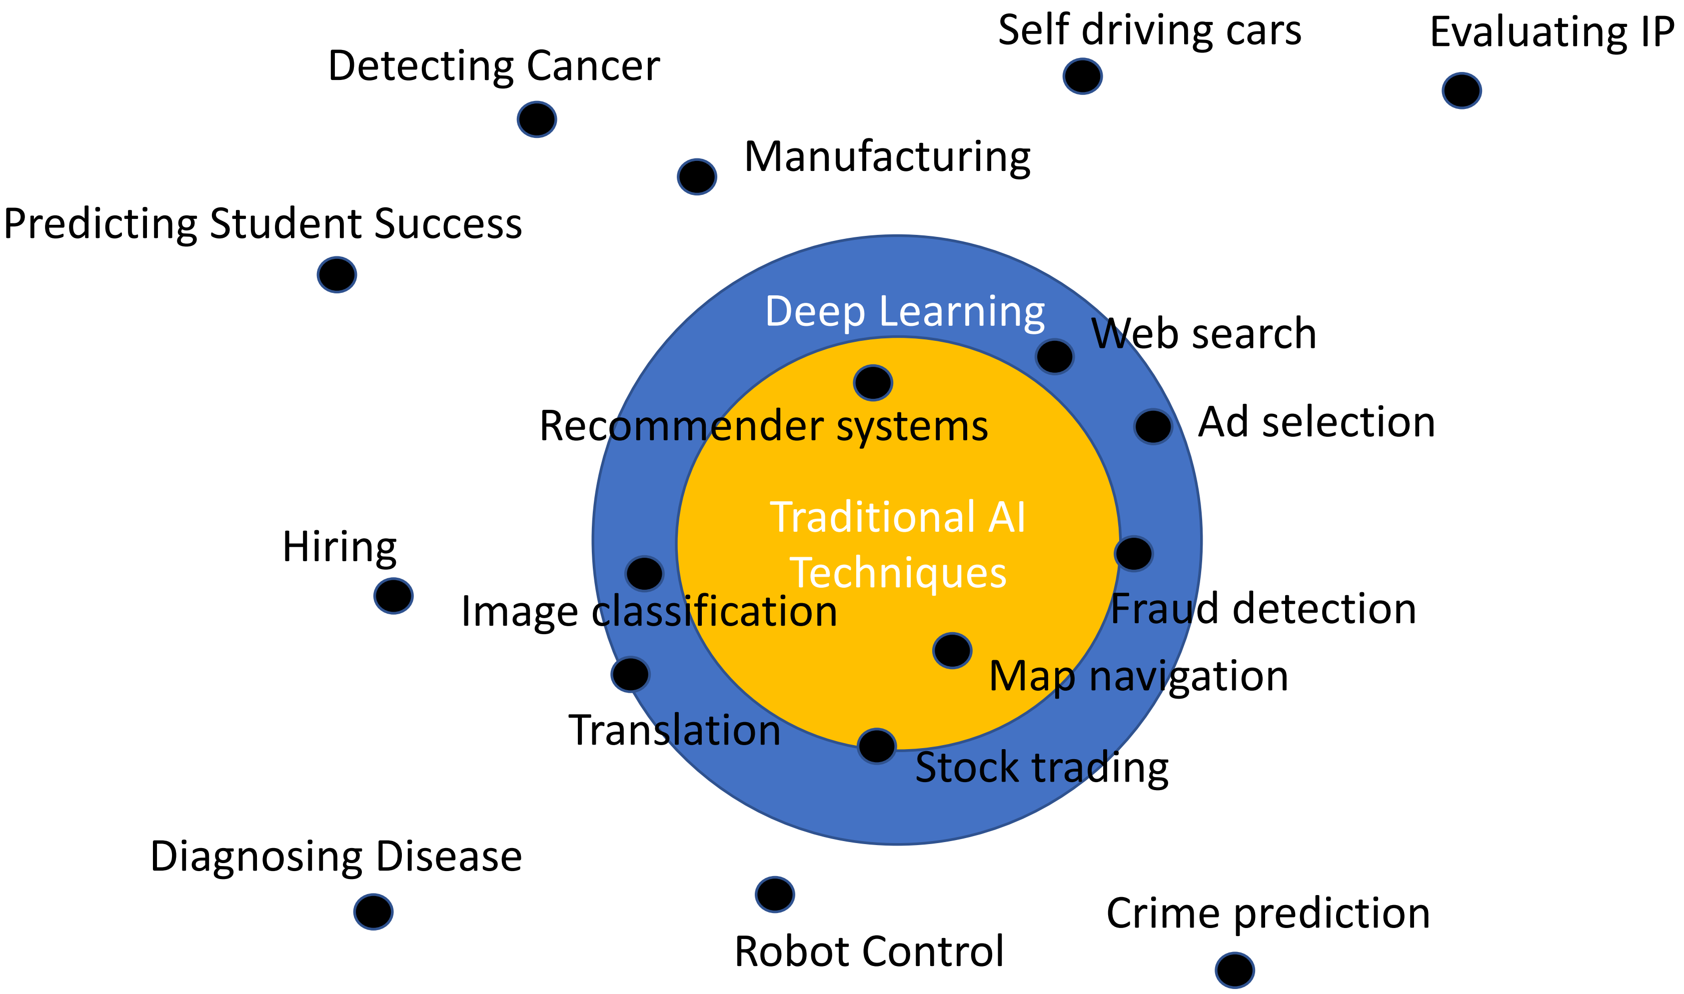

Brad Quinton, Scott Chin

Why is so much $ going into Deep Learning?

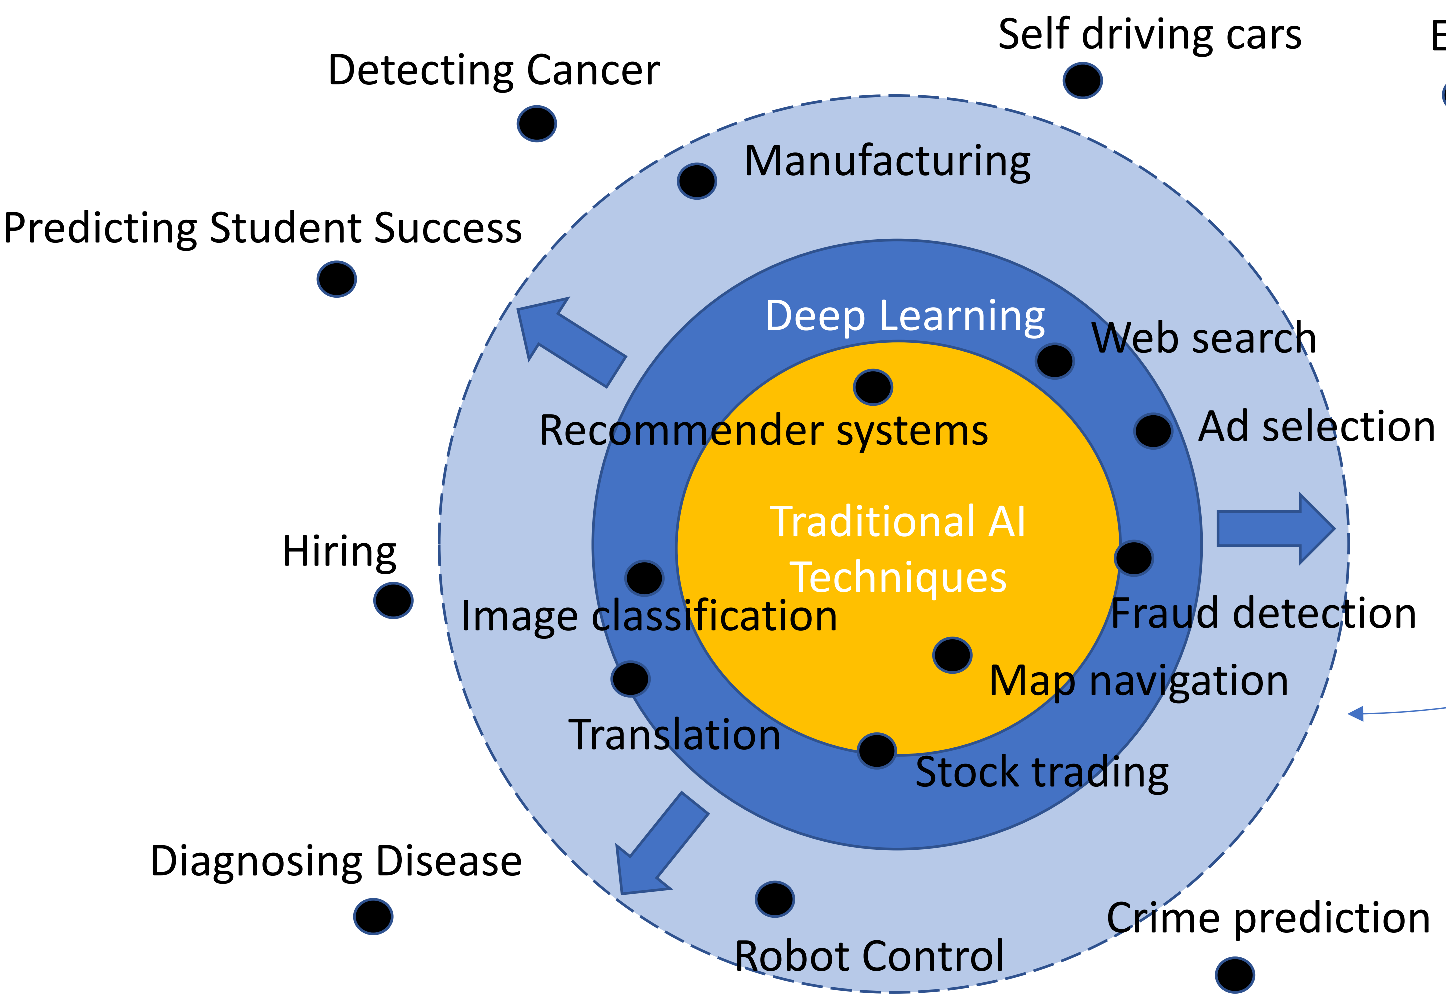

It is currently unclear where this new boundary end up. This drives financial speculation…

Brad Quinton, Scott Chin

On the Maturity of Deep Learning…

- Deep Learning is a very new, very active research area, and there is much that is not yet well understood

- 100 papers/day posted to ArXiv…!

"Deep learning has been somewhat like engineering before physics. Someone writes a paper and says, 'I made this bridge and it stood up!' Another guy has a paper: 'I made this bridge and it fell down—but then I added pillars, and then it stayed up.' Then pillars are a hot new thing. Someone comes up with arches, and it's like, 'Arches are great!'…."

- **David Duvenaud (University of Toronto Professor)**

MIT Technology Review: https://apple.news/ATudzzorPMOKY4-D2FX15Nw

Artificial General Intelligence (AGI) vs AI

- **Artificial Intelligence (AI)** is really just any technique that makes computers act intelligently (or more intelligently)
  - Sorting your e-mail by date
  - WYSWYG editing
  - Auto-correct
  - Automatically classifying cat pictures
  - Almost everything your computer/phone does for you…

- **Artificial General Intelligence (AGI)** is really more of a concept. It is the process of making computers "smart like us"; it is important to note progress in current AI **does not** necessarily imply progress in AGI

Is An AI Takeover Coming?

- Maybe…. but not soon, and Deep Learning alone will not be enough!
- It is important to realize that although there have been some exciting ”breakthroughs” in AI recently, there is still a very long way to go!
- Remember: you could sort cats from dogs when you were 3 years old!

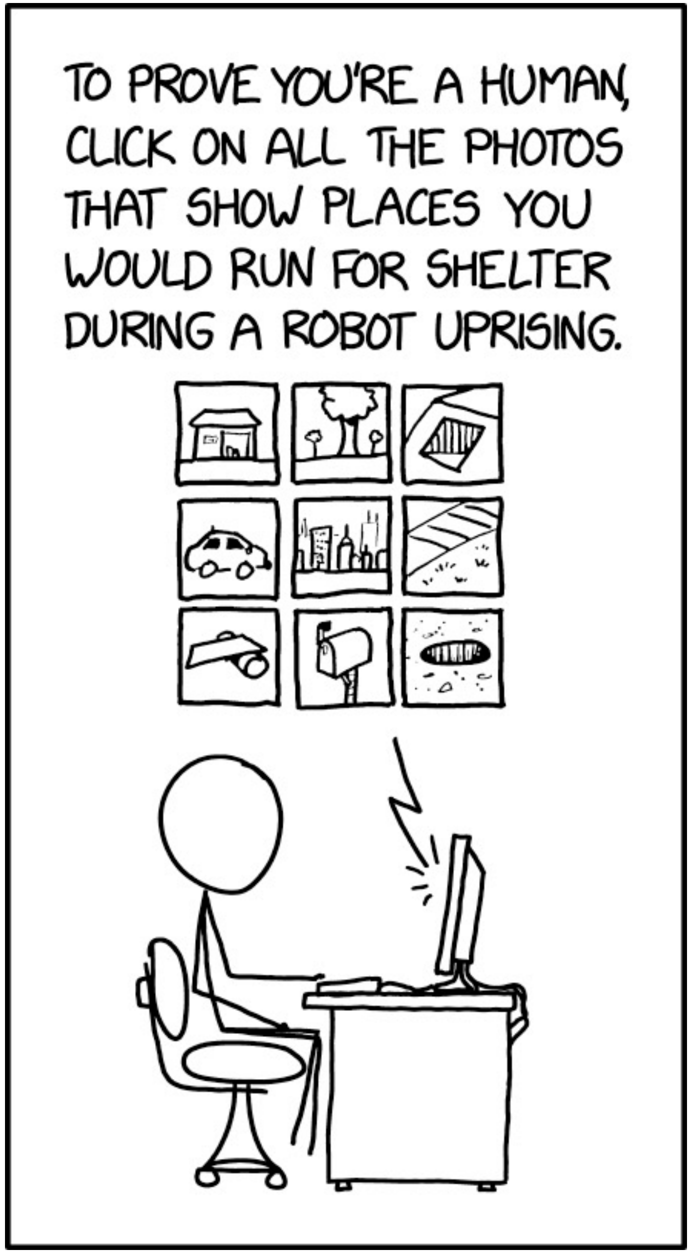

https://xkcd.com/2228/#

Terminology

- **Data Science** – Data Science is the process of using data analysis to build *understanding*. Normally the output of a data science project is a PowerPoint presentation or a report.

- **Machine Learning (ML)** – Process of using example data to create approximate functions, that can then be applied to *new data*. "Understanding" is rarely provided by ML.

- **Neural Networks (NNs)** – ML using an interconnected network of "trainable" artificial neurons (perceptrons) that maps some input X to an output Y.

Terminology

- **Deep Learning** – ML using multi-layered neural networks, which are normally trained with large data sets.
- **Supervised learning** – ML when the example data provides both the expected input and output. You can ”supervise” the training process by identifying and correcting mistakes.
- **Labelled Data** – Example data that includes the expected output. Used in supervised learning.

Terminology

- **Unsupervised learning** – ML when only expected input is provided. In this case the ML system learns relationships between the inputs themselves. Clustering customers into target groups, for example.

- **Unlabeled Data -** Example data that does not include the expected output. Used in unsupervised learning.

- **Reinforcement learning –** ML which uses only high-level goals and repeated trail and error during training. For example training an autonomous helicopter with "don't crash" as the only specific goal.

What are we going to do in this Course?

1. We will learn techniques to **create**, **evaluate** and **improve** Deep Learning systems.

2. Understand both the **capabilities, limitations** and **applications** of Deep Learning.

3. Implement real Deep Learning algorithms with cutting edge industry-standard **tools** and **frameworks.**

4. Have fun with a lot of **interesting problems**!

Key Take-Aways

- Machine Learning is the process of using **examples** to create approximate functions

- Deep Learning is a type of ML using multi-layered “deep” **neural networks**

- Deep Learning can solve many previously intractable problems, but **NOT** all problems