UNIVERSITY OF TORONTO

VECTOR INSTITUTE

Video 1: What is a ML task?

Introduction to Machine Learning

Prof. Nicolas Papernot

Material used in this course is adapted from several prior iterations of similar courses taught by others. This includes CSC321 by Prof. Grosse and Coursera's ML course by Prof. Ng.

What is machine learning?

Program algorithms that solve specific problems

Program a single algorithm that learns from data

```
def algorithm (x)
   return y
```

```
def learning_algorithm (X,Y)
   return algorithm

algorithm(x) = y
```

Why use machine learning?

- For many problems, it's difficult to program the correct behavior by hand (i.e., find the algorithm that solves a task)

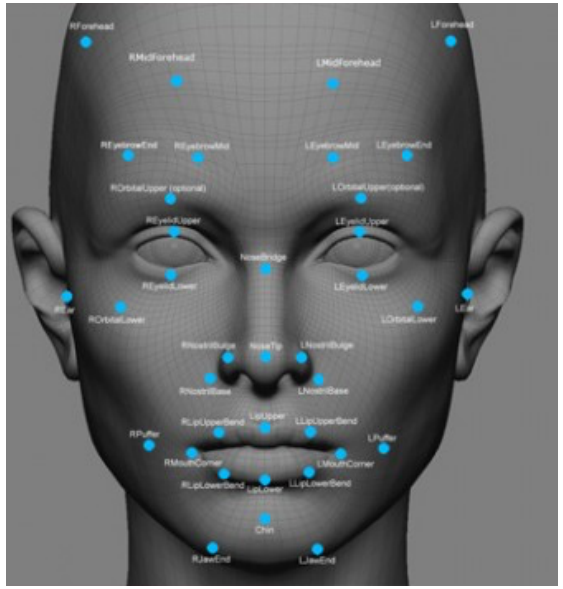

Source extremetech.com

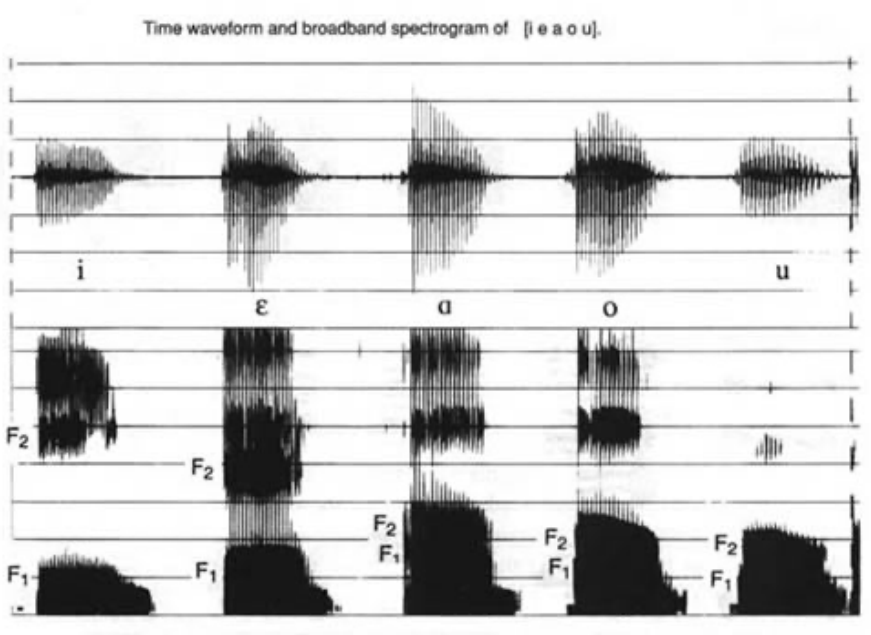

Source ncvs.org

Why use machine learning?

- For many problems, the algorithm may need to change throughout time (i.e., the algorithm needs to adapt to a changing environment)

On 2020-01-06, 9:40 AM, "Garth Gibson" <oyinkanadesanya@gmail.com> wrote:

Hello,

I'm in a conference right now, can't make
any phone conversation right now
but let me know if you get my
message and if you do, kindly
reply me with your
Personal Number to text you.
Thanks

Garth Gibson
President and CEO at Vector Institute

sent from my iphone

Source spam.com

Why use machine learning?

- We might be interested in an algorithm that performs better than human (programmers)

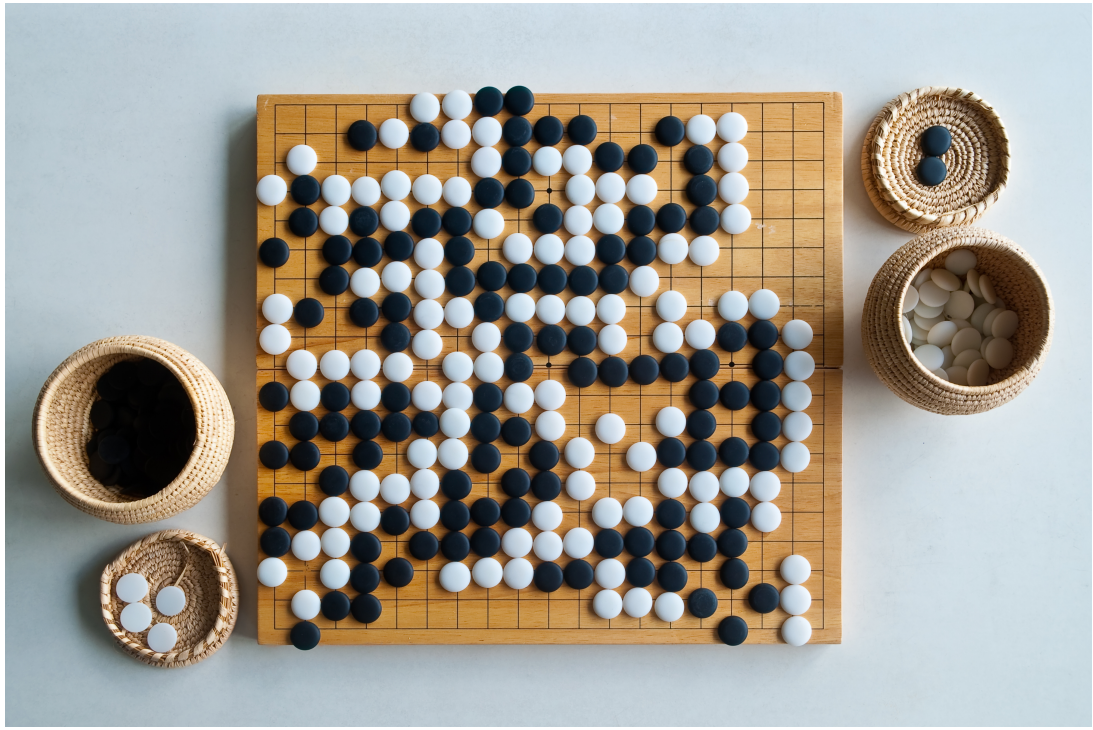

Source fortune.com

Types of machine learning

Supervised learning	Reinforcement learning	Unsupervised learning
Labeled data	Reward signal	Unlabeled data
Goal: predict correct label	Maximize reward signal	Varies (typically looking for interesting patterns in data)

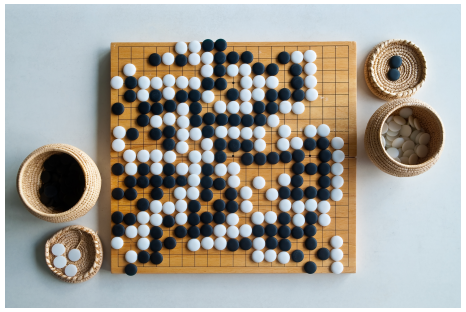

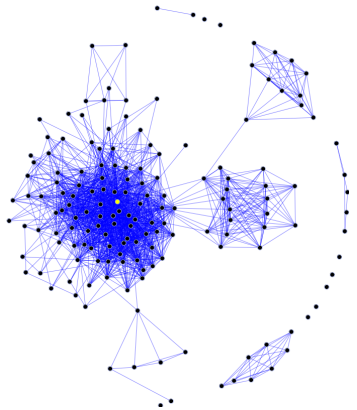

Source: Wikipedia.org

“fruit fly” of supervised learning research

Le Cun et al.

MNIST

- Task: given an image of a handwritten digit, predict the digit class
  - Input: the image
    - Pixels?
    - Output of feature extraction (e.g., edge/shape detection?)
  - Target: the digit class
- Data 70K images labeled by humans
  - Training set: first 60K
  - Test set: last 10K
- Can achieve 99%+ accuracy since 90s

Still learning from MNIST…

Why is this classified as a 3?

Circa 2013-2015

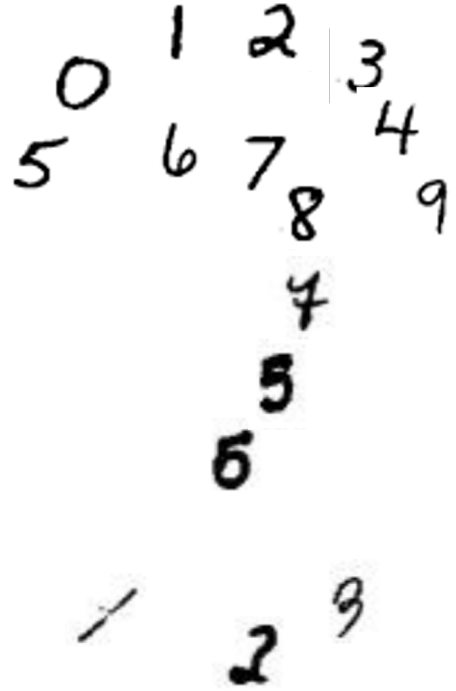

Why are some examples easier to learn, in particular with privacy?

Circa today

Object recognition

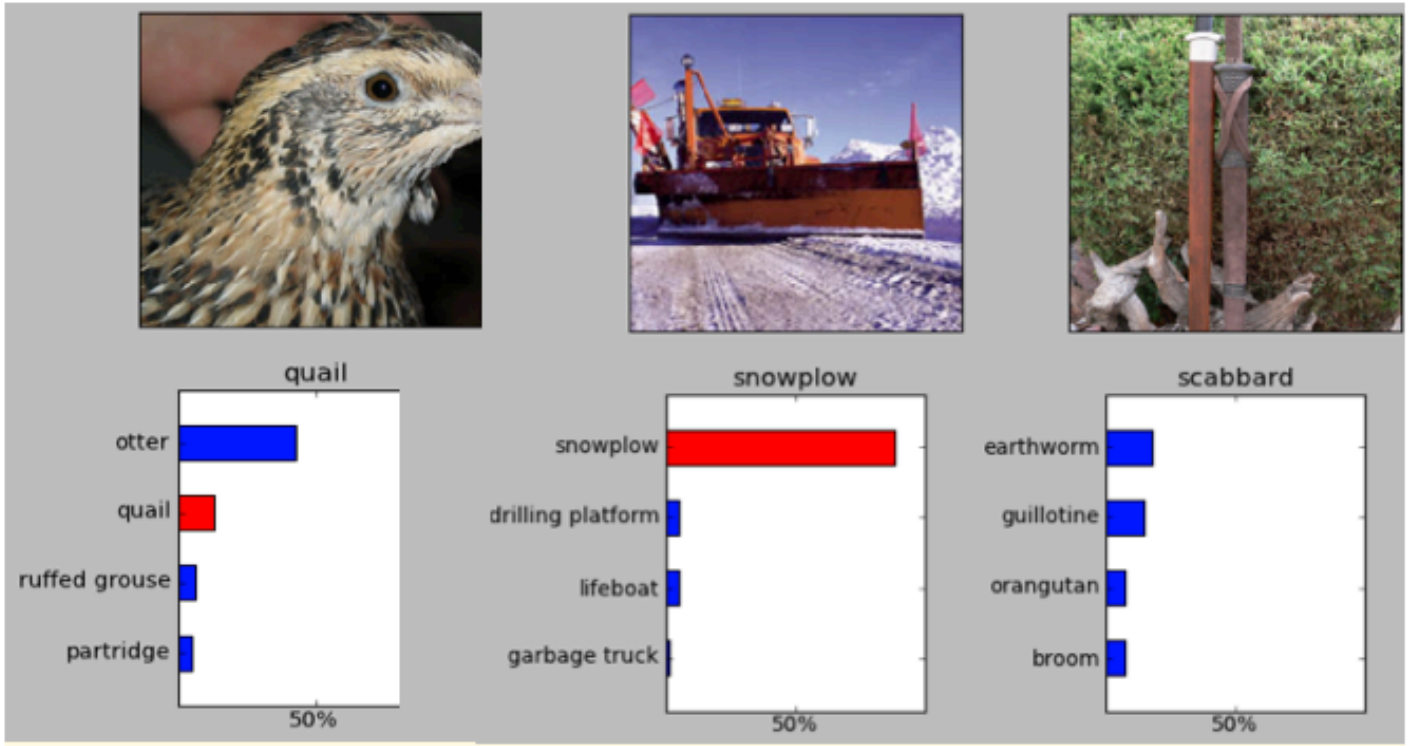

(Krizhevsky and Hinton, 2012)

ImageNet dataset: thousands of categories, millions of labeled images
Lots of variability in viewpoint, lighting, etc.
Performance measured through top5 and top1

Imagenet top5

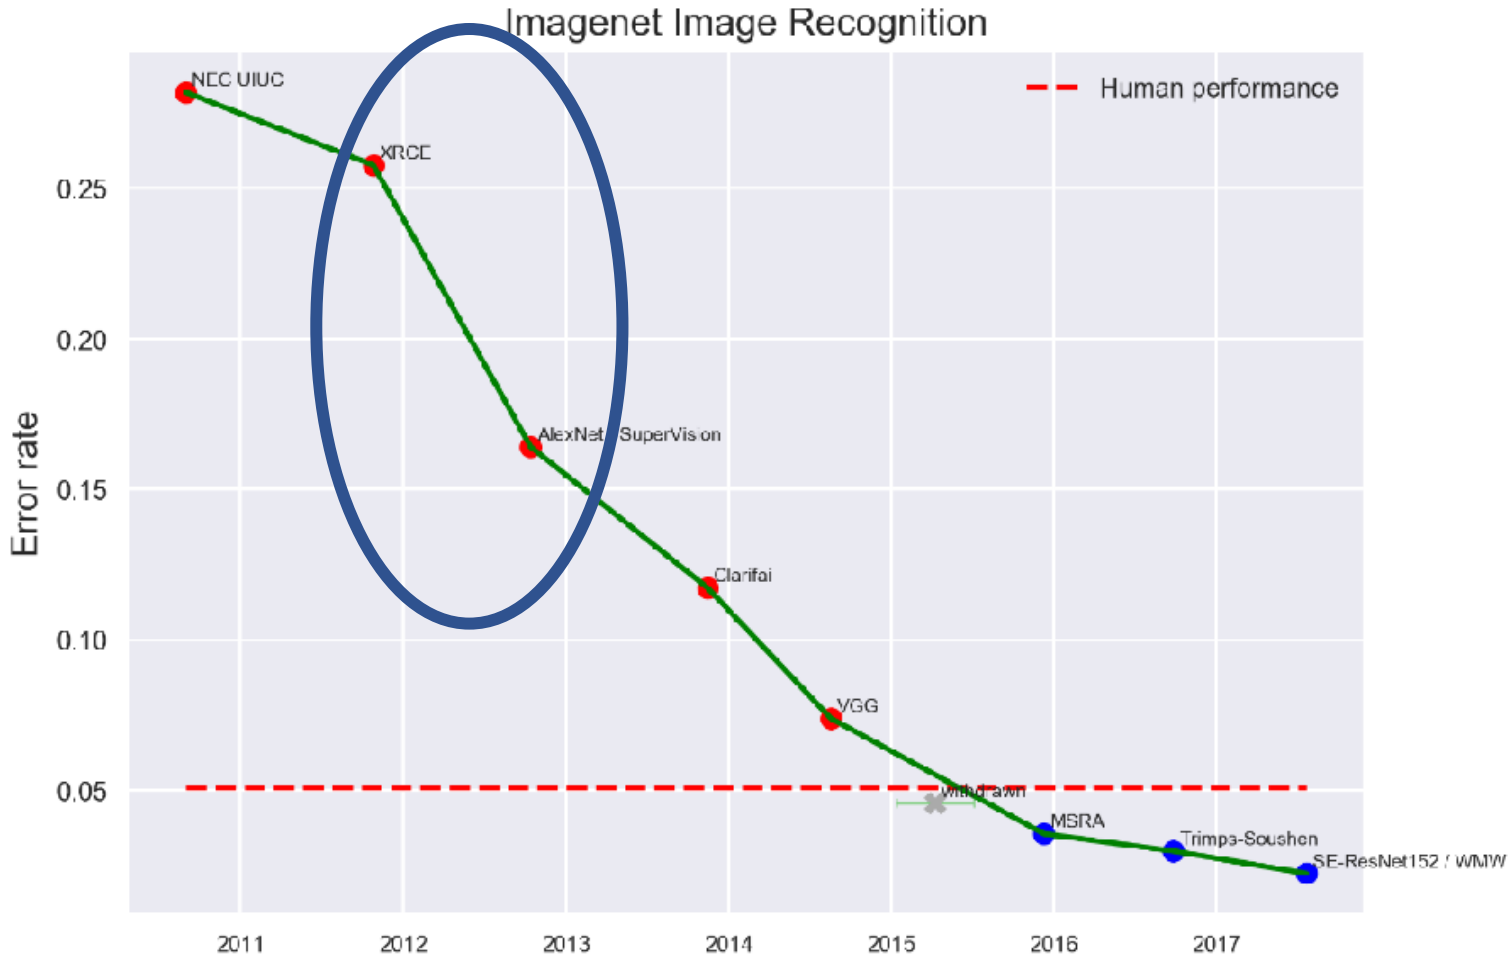

Source semanticscholar.org

Caption generation

TAGS:

frisbees frisbee pushups golfers kickball

Nearest Neighbor Sentence:

- several people that are playing in a frisbee game .

Top-5 Generated:

- a group of girls are playing a game of frisbee .
- a group of girls are playing a soccer game .
- a group of girls playing on a soccer game .
- a group of people playing a game of frisbee .
- the young people are playing a game of frisbee .

Given: dataset of Flickr images with captions

Reinforcement learning

- An agent interacts with an environment (e.g. game of Breakout)
- In each time step,
  - the agent receives observations (e.g. pixels) which give it information about the state (e.g. positions of the ball and paddle)
  - the agent picks an action (e.g. keystrokes) which affects the state
- The agent periodically receives a reward (e.g. points)
- The agent wants to learn a policy, or mapping from observations to actions, which maximizes its average reward over time

Unsupervised learning: generative modeling

- Learn distribution of dataset (e.g., natural images)
- Evaluate human perception of data sampled from model
- These results were considered impressive in 2014:

Denton et al.

Unsupervised learning: generative modeling

https://thispersondoesnotexist.com/

A typical ML pipeline

1. Input representation: what each dimension of x contains
2. Model hypothesis class: y=g(wx+b)
3. Training algorithm to find w and b
4. Test model

Why take this class?

- Debugging learning algorithms requires sophisticated detective work, which requires understanding what goes on beneath the hood.
- That's why we derive things by hand in this class!