# Deep Learning

Samuel Cheng

# Project presentation

- Start as usual class time
- Please prepare 20-25 minutes presentation. Be quick with problem statement. Focus on your approach and result
  - Take a format similar to conference presentation
- ~5 min Q/A
- Grading
  - Presentation: clarity, structure, references, etc. (10/40)
  - Technical: correctness, depth, innovation, etc. (15/40)
  - Evaluation and results: sound evaluation metric, thoroughness in analysis and experimentation, results and performance (15/40)
- Expectation
  - National conference quality (4/4), research day quality (3/4), Tuesday meeting quality (2/4), just show up (1/4)

# We have went through a lot …

- Backprop
- Regularization, weight initialization
- CNN
  - R-CNN, faster R-CNN
- Seq2seq models
  - Neural machine translations
  - Chatbots
- Memory networks, neural Turing machines
- Autoencoders
- GANs
- Deep Q-learning

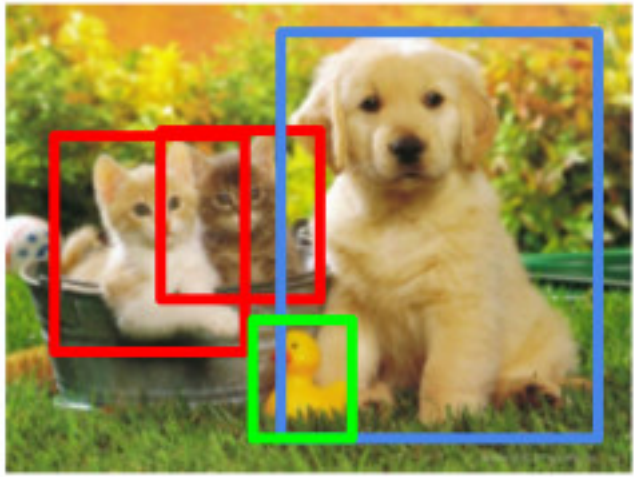

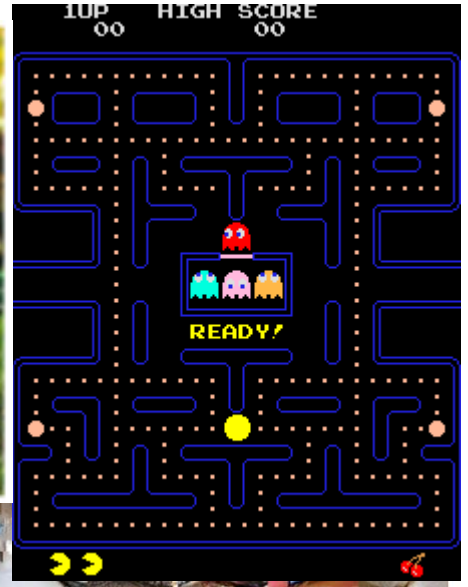


# Many ideas were not new…

- Mainly two things happened
  - Inexpensive computational power
    - GPUs
  - Large dataset available
    - ImageNet
    - MS COCO
    - Kaggle …
- Persistent efforts of many researchers
  - Hinton (Toronto, Google)
  - Yann Lecun (NYU, Facebook)
  - Bengio (Montreal)
  - Andrew Ng (Stanford, Google, Baidu)
  - …


IMAGENET

# Three Types of Learning

- Reinforcement Learning
  - The machine predicts a scalar reward given once in a while
  - A few bits for some samples
- Supervised Learning
  - The machine predicts a category or a few numbers for each input
  - 10→10,000 bits per sample
- Unsupervised Learning
  - The machine predicts any part of its input for any observed part
  - Predicts future frame in videos
  - Millions of bits per sample


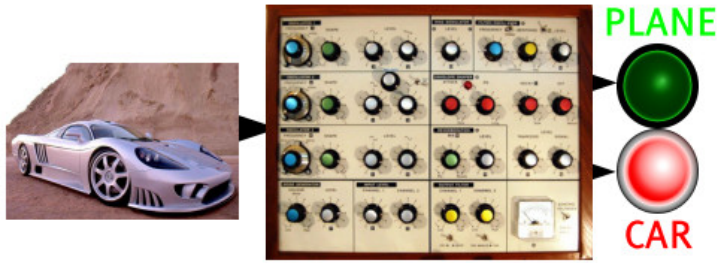


For Yann Lecun: father of CNN, director of Facebook AI Research, NYU professor

# How much information does machine needs to predict

**Reinforcement Learning** (cherry)

- The machine predicts a scalar reward given once in a while.
- **A few bits for some samples**

**Supervised Learning** (icing)

- The machine predicts a category or a few numbers for each input
- **10→10,000 bits per sample**

**Unsupervised Learning** (cake)

- The machine predicts any part of its input for any observed part.
- Predicts future frames in videos
- **Millions of bits per sample**

For Yann Lecun: father of CNN, director of Facebook AI Research, NYU professor

# Four missing pieces of AI (by Lecun)

- Theoretical Understanding for Deep Learning
  - What is the geometry of the objective function in deep networks?
  - Why the ConvNet architecture works so well? [Mallat, Bruna, Tygert…[
- Integrating Representation/Deep Learning with Reasoning, Attention, Planning and Memory
  - A lot of recent work on reasoning/planning, attention, memory, learning "algorithms"
  - Memory-augmented neural nets
  - "Differentiable" algorithm
- Integrating supervised, unsupervised and reinforcement learning into a single "algorithm"
  - Boltzmann machines would be nice if they worked
  - Stacked What-Where Auto-Encoders, Ladder Networks…
- Effective ways to do unsupervised learning
  - Discovering the structure and regularities of the world by observing it and living in it like animals and human do

# Information Theory and Statistical Learning

## A shameless advertisement of my fall course

- Will look into (shallow) machine learning models not discussed in this class
  - SVM
  - Decision trees
  - Sparse coding
  - GMM, MRF, CRF …
- Why relevant?
  - They are still very useful when you just need a quick and dirty way that does not need to have state-of-the-art accuracy
  - New ideas almost never came from scratch. They all are just some modification of old ideas
    - Standing on the shoulders of giants

# Epilogue

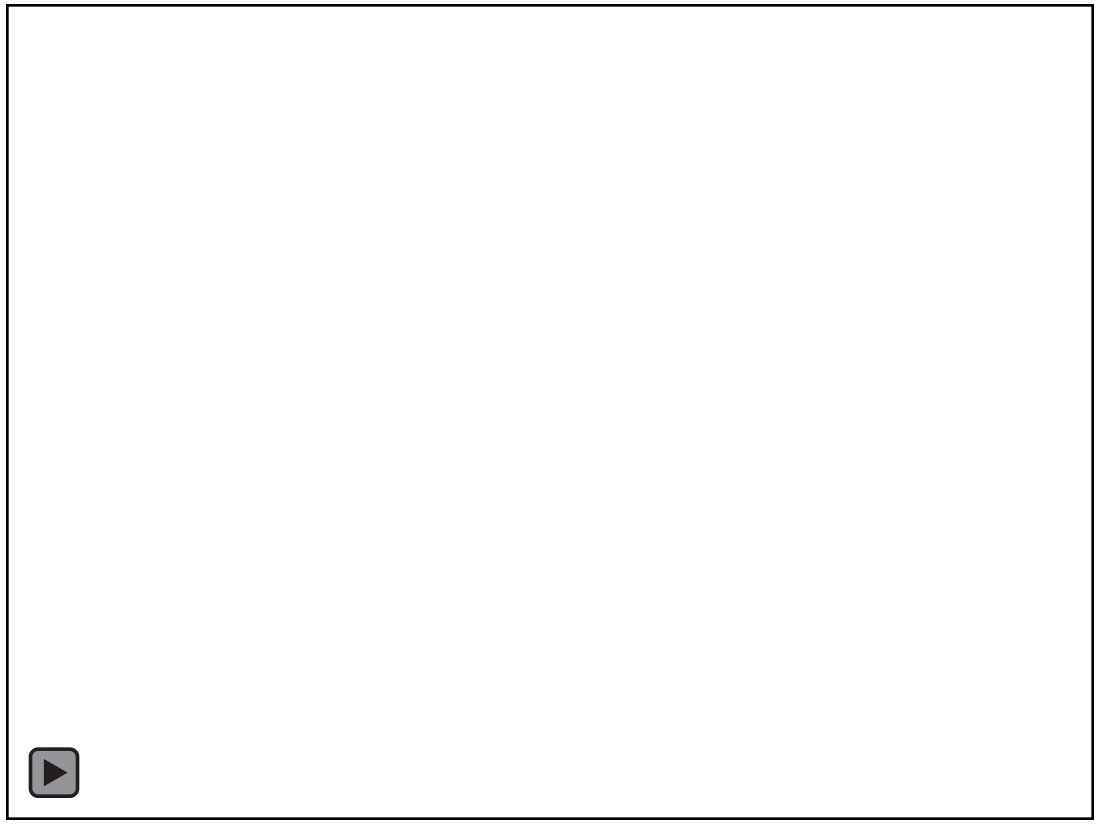

- ‘When I was doing my Ph.D., my advisor would tell me that (I was wasting my time) every week. And I would say, “give me six months and I will prove you that it works.” And every six months, I’d say that again’ - Geoffrey Hinton
- Don’t easily believe something wouldn’t work just because someone told you so
  - Try it yourself!
- If you really believe in it, be persistent and enjoy your last laugh

Wish you all good luck with your finals and presentations!
And have a fruitful sem-break!
Don’t forget to fill in evaluation!