Deep learning: supplementary materials

Machine Learning, spring 2021

Yoonsuck Choe

The Rise of Deep Learning (early/mid 2010-present)

- CNN, RNN, Attention, Deep Reinforcement Learning.

CNN model

Learned Features

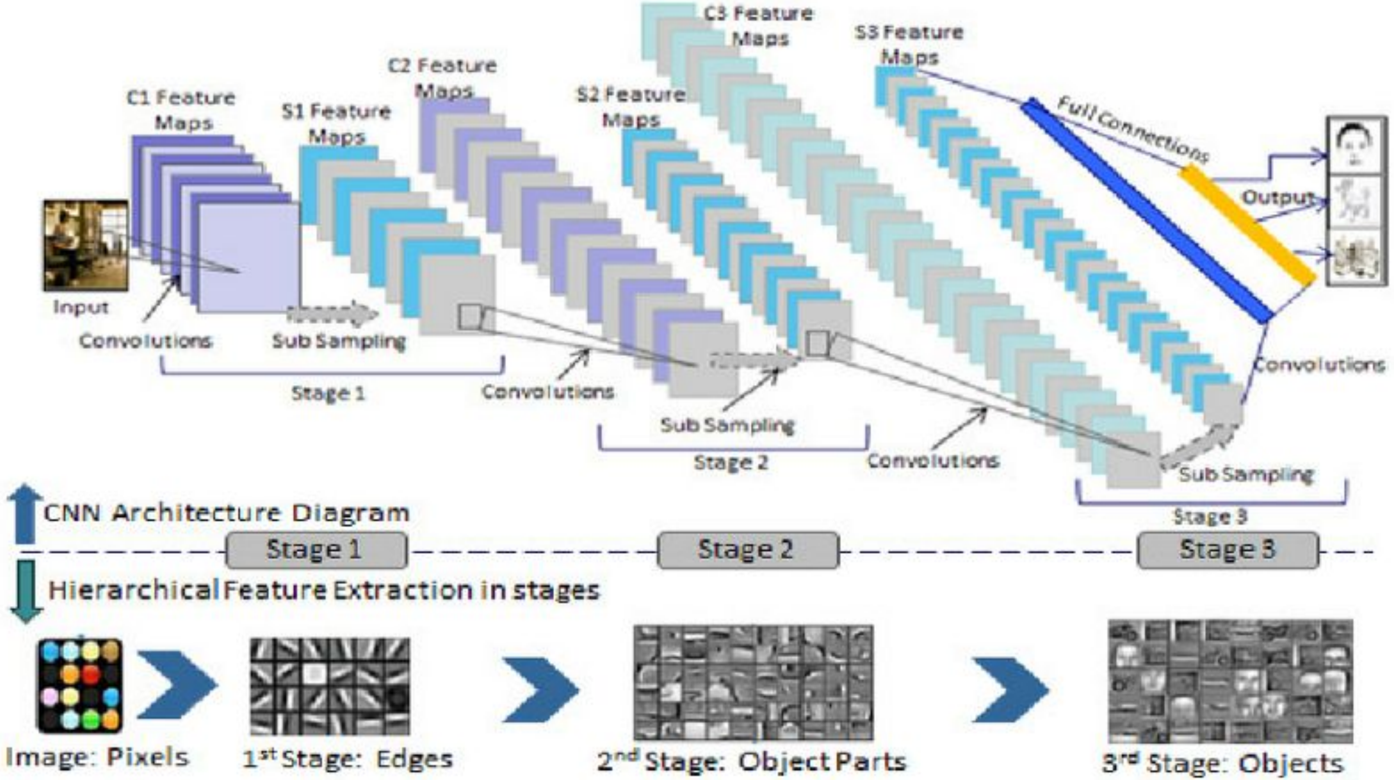

The deeper the better!

- Major factors in deep learning's success:
  - Very deep neural networks
  - Big data
  - Massive computing (GPU)

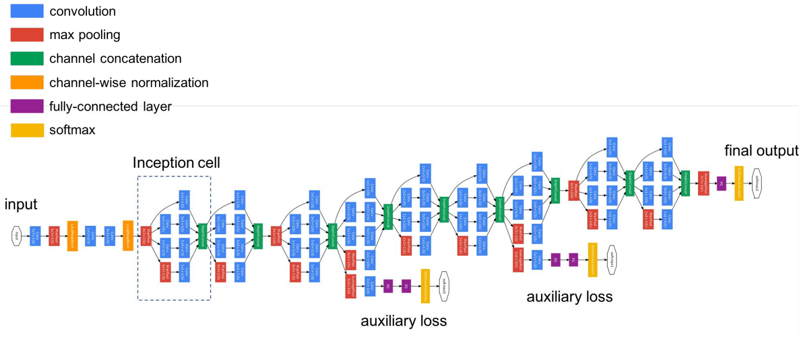

CNN model (Google's Inception)

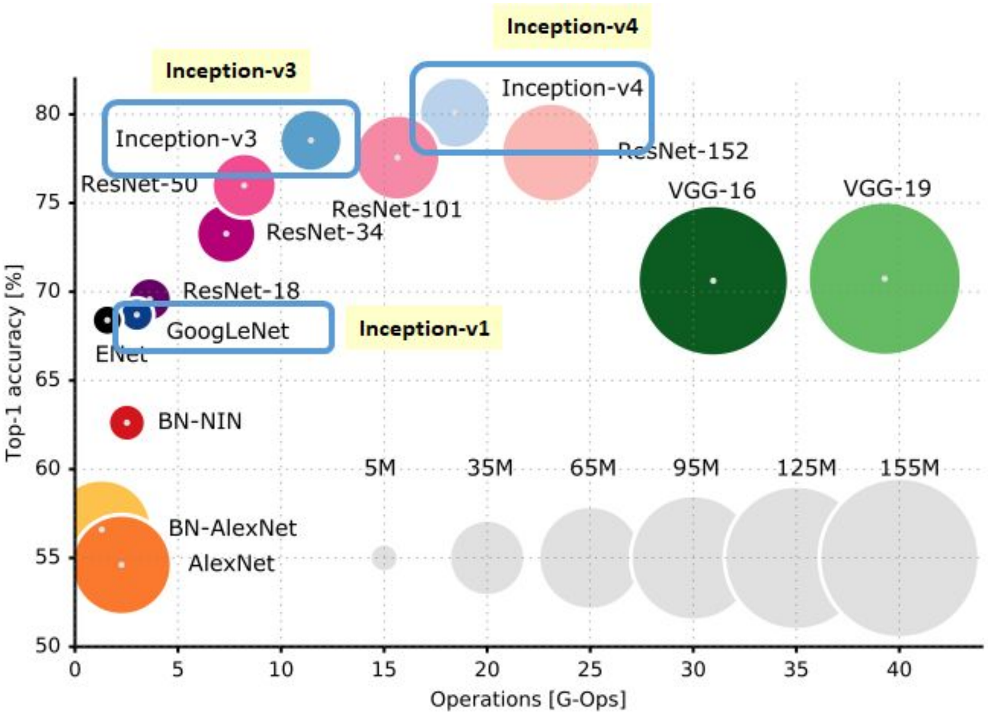

Advances in Deep Learning

- Vision/Speech, NLP, NMT - Superhuman performance in select tasks
  - Beyond Classification: Detection, segmentation
  - Multimodal: Visual Question Answering

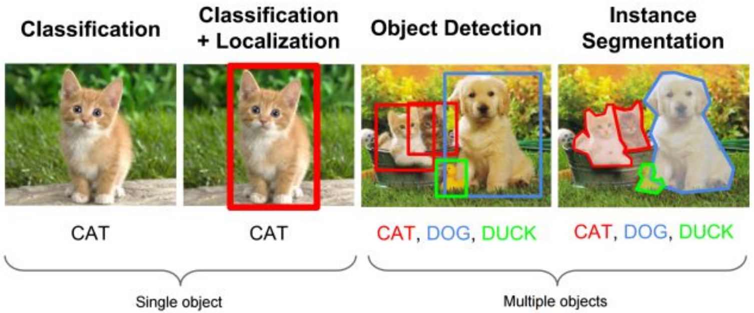

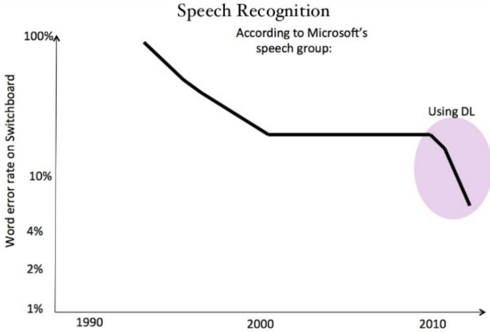

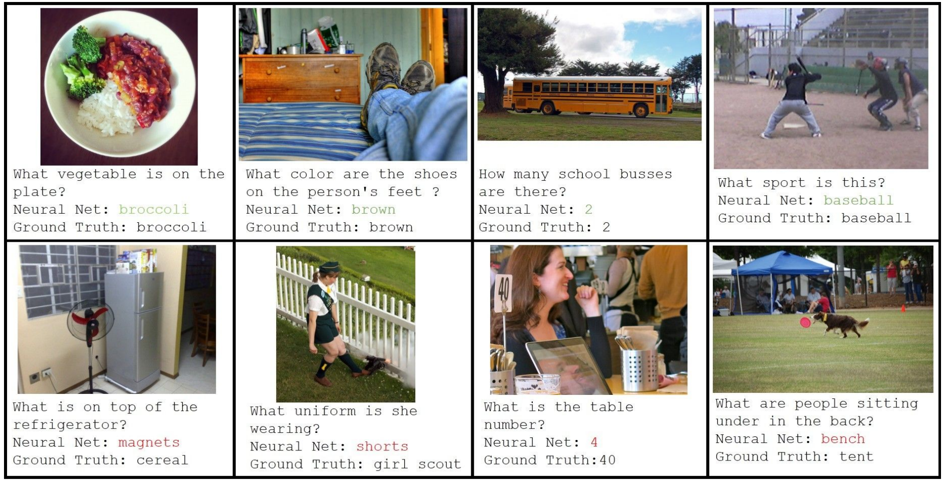

What vegetable is on the plate?
Neural Net: broccoli
Ground Truth: broccoli

What color are the shoes on the person's feet ?
Neural Net: brown
Ground Truth: brown

How many school busses are there?
Neural Net: 2
Ground Truth: 2

What sport is this?
Neural Net: baseball
Ground Truth: baseball

What is on top of the refrigerator?
Neural Net: magnets
Ground Truth: cereal

What uniform is she wearing?
Neural Net: shorts
Ground Truth: girl scout

What is the table number?
Neural Net: 4
Ground Truth:40

What are people sitting under in the back?
Neural Net: bench
Ground Truth: tent

Deep Reinforcement Learning

- Video games, robot control, Deep RL (AlphaGo, AlphaStar) :
  - Analyze visual input and generation action and learn based on reward.

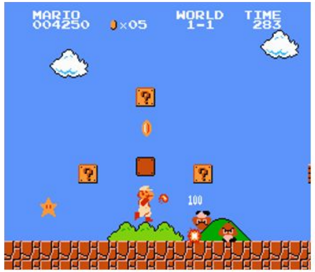

Super Mario Bros

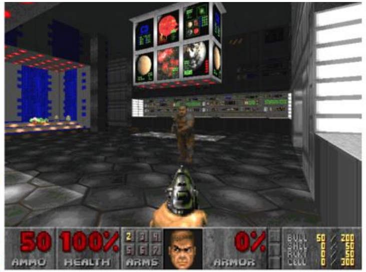

VizDoom

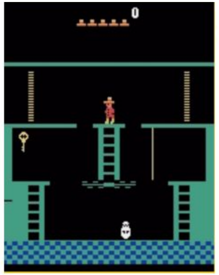

Montezumas Revenge

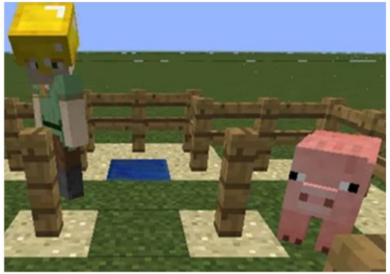

Minecraft

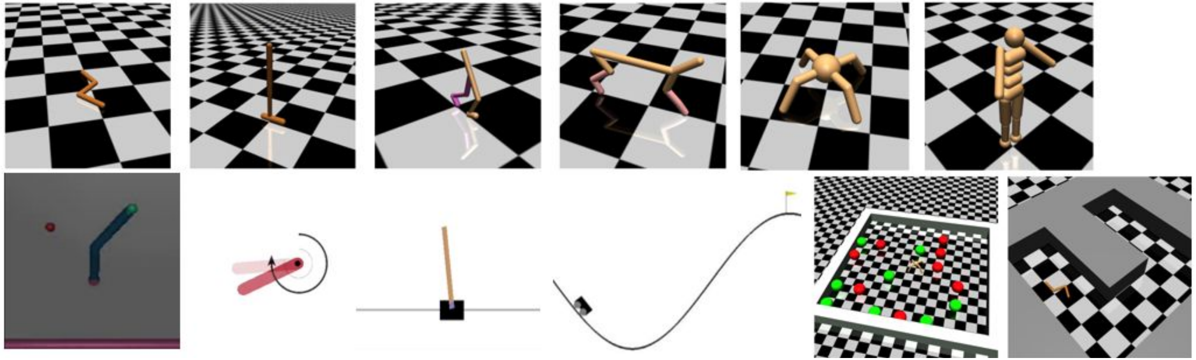

Limitations of Deep Learning

- Data hungry, Can't do complex reasoning,
- Lack of common sense
- Explainability, Sensitive to noise/Adversarial input

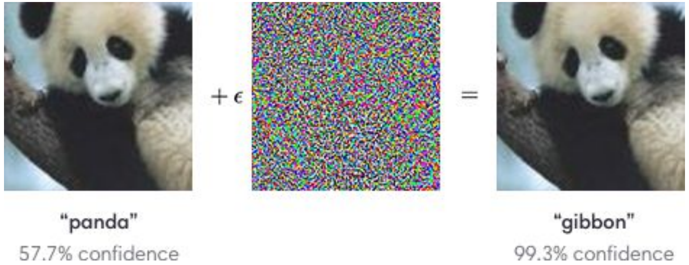

Sensitivity to noise / Adversarial input

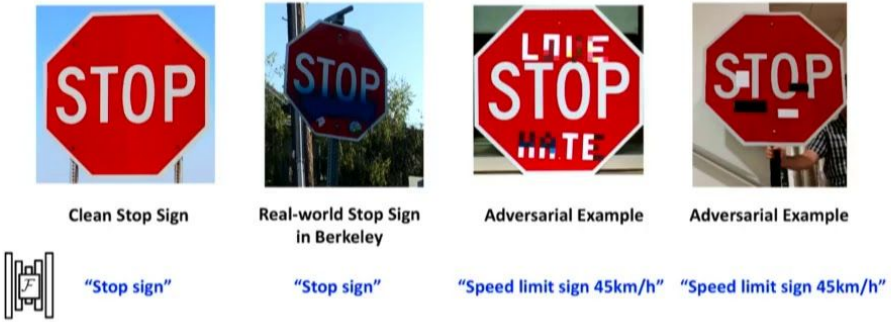
Clean Stop Sign
Real-world Stop Sign in Berkeley
Adversarial Example
Adversarial Example
"Stop sign"
"Stop sign"
"Speed limit sign 45km/h"
"Speed limit sign 45km/h"

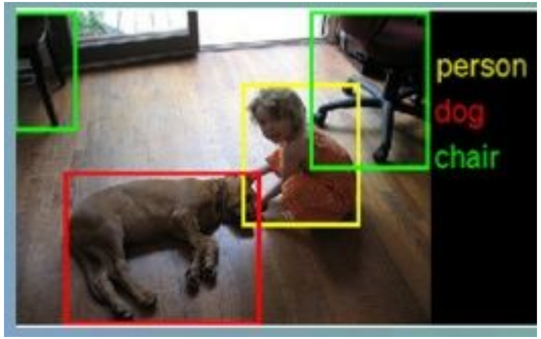

What to rescue first when there's a fire?

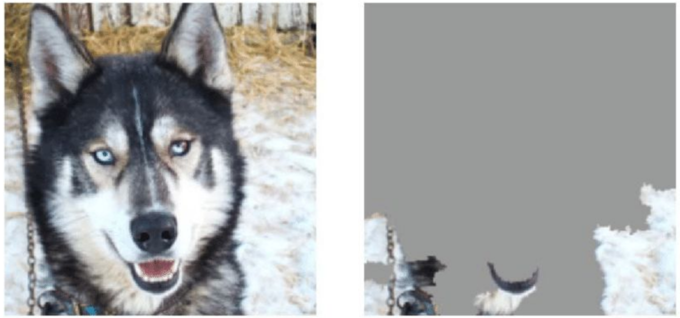

Learned snow field feature, not husky feature