Deep Learning

Jana Kosecka
George Mason University

4444 Research II
kosecka@gmu.edu , 3-1876
Office hours Tue 2-3pm

Slides courtesy L. Lazebnik (Univ. of Illinois CS498) and others

Logistics

- **Grading:** Homeworks, projects, class participation
- **Prerequisites:** basic statistical concepts, probability, machine learning, strong Python programming skills
- **Course web page** cs.gmu.edu/~kosecka/cs747/
- Homeworks 30%
- Class presentations 10%
- Paper summaries 10%
- Class participation 10%
- Project Proposal/Presentation: 10%
- Project Report: 30%

Recommended Text

- Deep Learning I. Goodfellow et al
- deeplearningbook.org

- University of Illinois Introduction to Deep Learning
- Stanford CS231n: Convolutional Neural Networks for Visual Recognition
- Stanford CS230: Deep Learning
- Princeton COS 495: Introduction to Deep Learning
- IDIAP EE559: Deep Learning

Introduction to Deep Learning

- What is “learning”?
  - Improving performance through experience
  - Getting a computer to do well on a task without manually building in competence
- What is “deep”?
  - Learning using multi-layer neural networks

- What is the relationship between deep learning, ML, and AI?

An incomplete timeline of deep learning

- 1943: McCulloch and Pitts neurons
- 1958: Rosenblatt's perceptron

1958 New York Times…

NEW NAVY DEVICE LEARNS BY DOING

Psychologist Shows Embryo of Computer Designed to Read and Grow Wiser

WASHINGTON, July 7 (UPI) —The Navy revealed the embryo of an electronic computer today that it expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence.

The embryo—the Weather Bureau's $2,000,000 "704" computer—learned to differentiate between right and left after fifty attempts in the Navy's demonstration for newsmen.

The service said it would use this principle to build the first of its Perceptron thinking machines that will be able to read and write. It is expected to be finished in about a year at a cost of $100,000.

Dr. Frank Rosenblatt, designer of the Perceptron, conducted the demonstration. He said the machine would be the first device to think as the human brain. As do human beings, Perceptron will make mistakes at first, but will grow wiser as it gains experience, he said.

Dr. Rosenblatt, a research psychologist at the Cornell Aeronautical Laboratory, Buffalo, said Perceptrons might be fired to the planets as mechanical space explorers.

Without Human Controls

The Navy said the perceptron would be the first non-living mechanism "capable of receiving, recognizing and identifying its surroundings without any human training or control."

The "brain" is designed to remember images and information it has perceived itself. Ordinary computers remember only what is fed into them on punch cards or magnetic tape.

Later Perceptrons will be able to recognize people and call out their names and instantly translate speech in one language to speech or writing in another language, it was predicted.

Mr. Rosenblatt said in principle it would be possible to build brains that could reproduce themselves on an assembly line and which would be conscious of their existence.

In today's demonstration, the "704" was fed two cards, one with squares marked on the left side and the other with squares on the right side.

Learns by Doing

In the first fifty trials, the machine made no distinction between them. It then started registering a "Q" for the left squares and "O" for the right squares.

Dr. Rosenblatt said he could explain why the machine learned only in highly technical terms. But he said the computer had undergone a "self-induced change in the wiring diagram."

The first Perceptron will have about 1,000 electronic "association cells" receiving electrical impulses from an eye-like scanning device with 400 photo-cells. The human brain has 10,000,000,000 responsive cells, including 100,000,000 connections with the eyes.

An incomplete timeline of deep learning

- 1943: McCulloch and Pitts neurons
- 1958: Rosenblatt's perceptron
- 1969: Minsky and Papert Perceptrons book
- 1960-1986: Back-propagation
- 1980: Fukushima's Neocognitron
- 1989: Convolutional neural networks (LeNet)
- 2012: AlexNet

- For much, much, *much* more detail, see Schmidhuber's historical overview

Successes of Deep Learning

- Computer Vision
- Speech and Language
- Game Playing
- Robotics

Deep Learning

Deep learning in computer vision - notable progress in image classification

Continue to mark its progress in various robot perception or control tasks

- ~14 million images, 20k classes
- Images gathered from Internet
- Human labels via Amazon Turk

IMAGENET

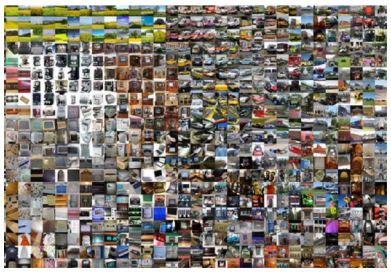

[NIPS 2012]

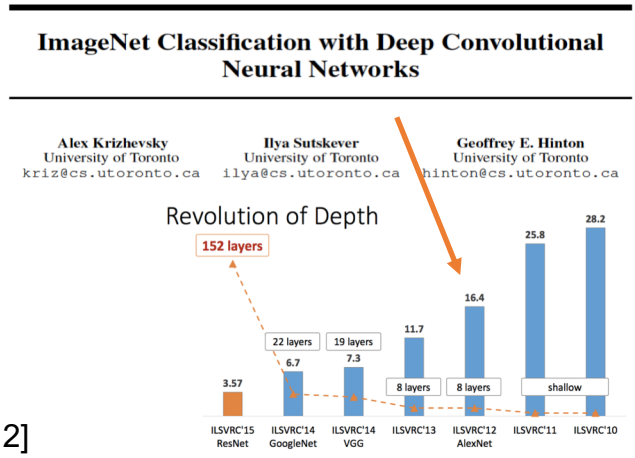

Convolutional neural networks

Beginnings: [LeCun *et al.* 1998]

Stack multiple stages of simple cells / complex cells layers

Higher stages compute more global, more invariant features

Classification layer on top

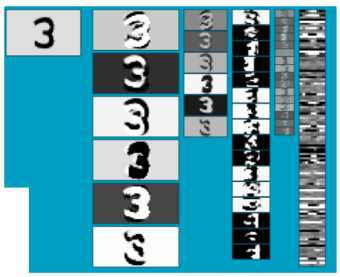

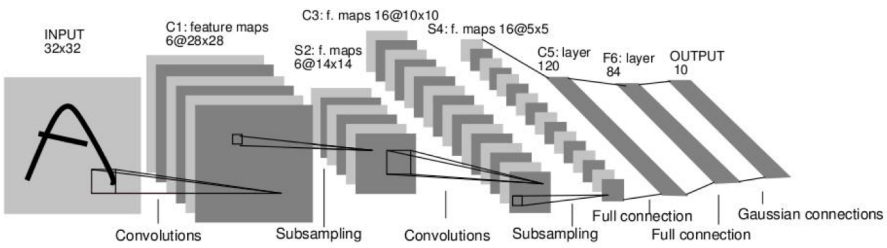

Training traditional supervised approach - Back-propagation, Lots of labeled data. Initial approach was run on CPUs

Perception Problems

- Need to process multiple-sensing modalities, video

What ? Where ? How ?

- 3D reconstruction, visual odometry
- Object instance and category classification
- Pose estimation (humans, objects)
- Semantic segmentation
- Matching, recognition and localization

- Robotics - active perception and control

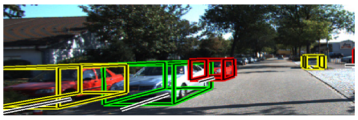


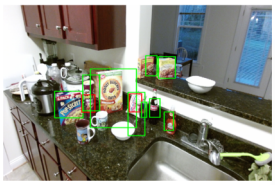
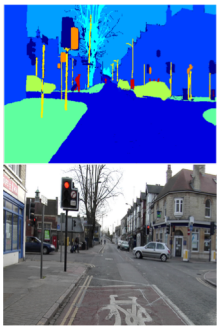

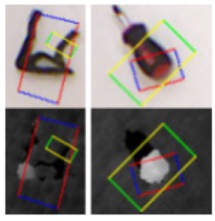
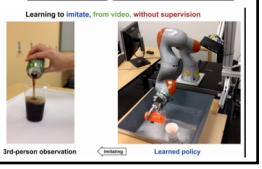

Computer Vision

- Object detection, instance segmentation

K. He, G. Gkioxari, P. Dollar, and R. Girshick, Mask R-CNN,
ICCV 2017 (Best Paper Award)

Games

- 2013: DeepMind uses deep reinforcement learning to beat humans at some Atari games

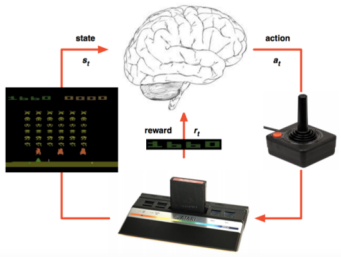

- 2016: DeepMind's AlphaGo system beats Go grandmaster Lee Sedol 4-1
- 2017: AlphaZero learns to play Go and chess from scratch

Robotics

- End-to-end training of deep visuomotor policies

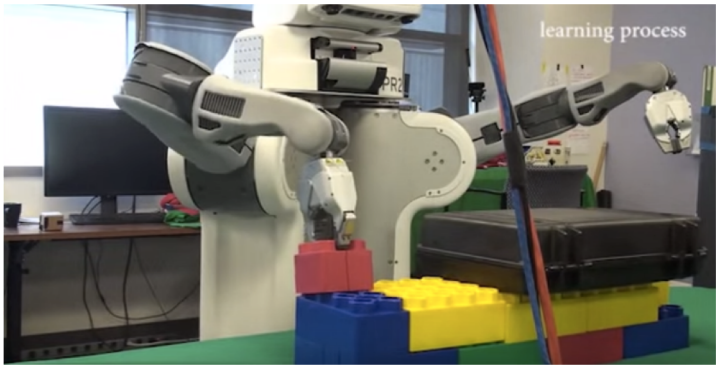

Overview video, training video

Self-driving cars

- Deep learning crucial for the global success of automotive autonomy – Automotive World, 6/26/2018

Image source

This class

ML basics, linear classifiers

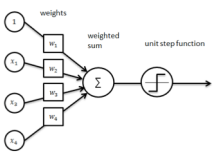

Multilayer neural networks, backpropagation

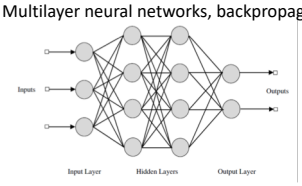

Convolutional networks and applications

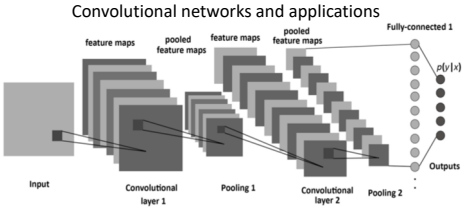

Recurrent networks and applications

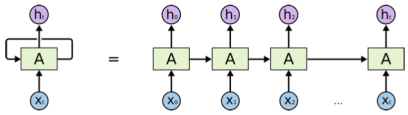

Deep reinforcement learning

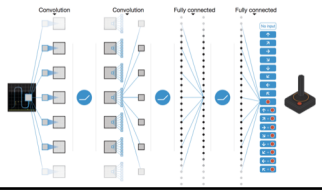

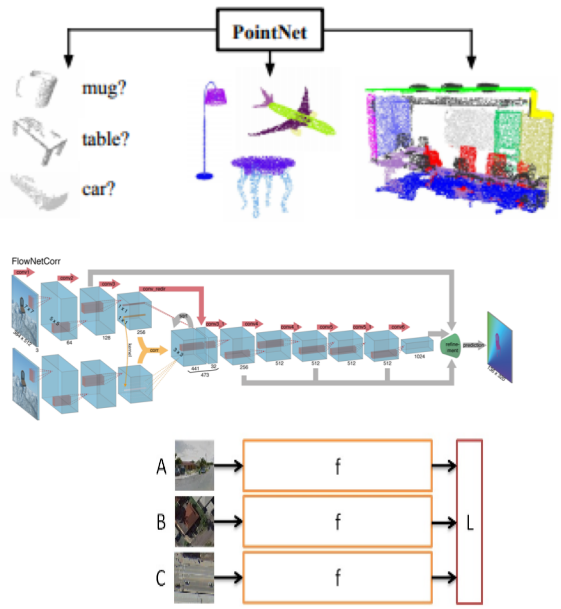

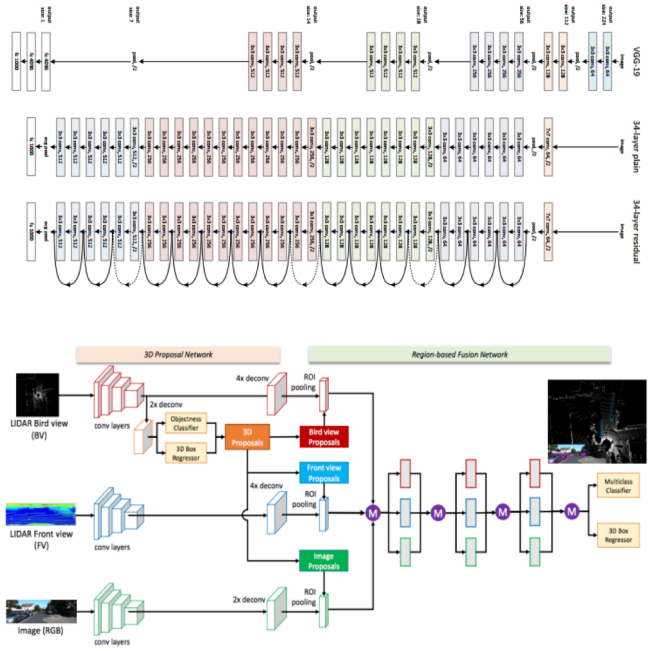

Two key components that made DNN perception viable for robotics:

Accuracy

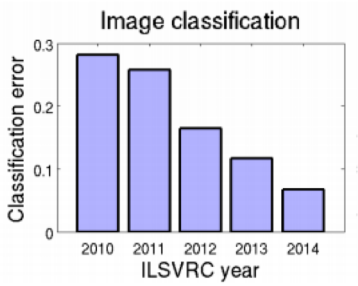

Significant Imagenet error rate decrease
Russakovsky et al'14

Speed/Real-time

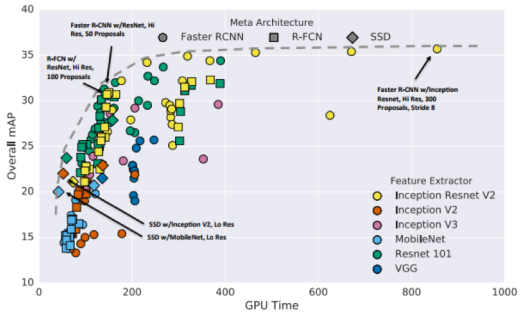

Detection algorithms in 50-150 ms
Huang et al'17