

Class 6

Deep Learning: a glimpse of what
is really happening

Overview

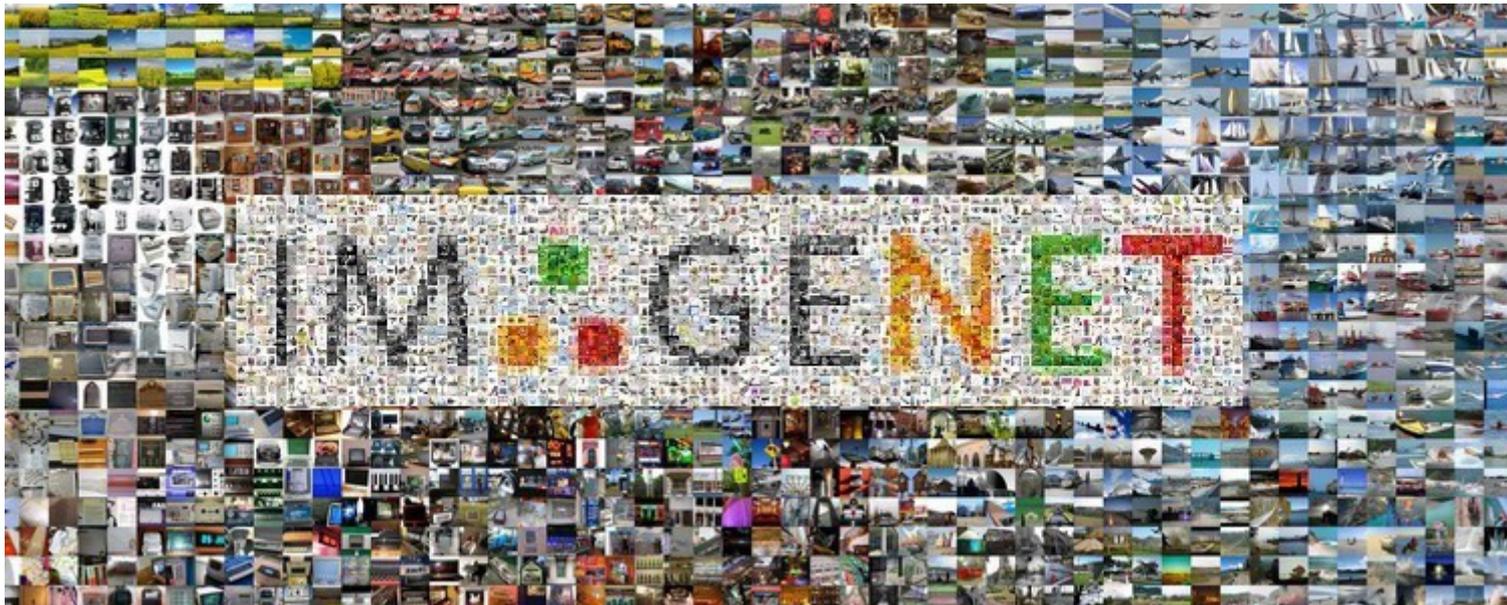
- **Deep** Convolutional Neural Networks
- The Deep Learning Zoo
- Generative Models
- Concluding Remarks

Overview

- **Deep** Convolutional Neural Networks
- The Deep Learning Zoo
- Generative Models
- Concluding Remarks

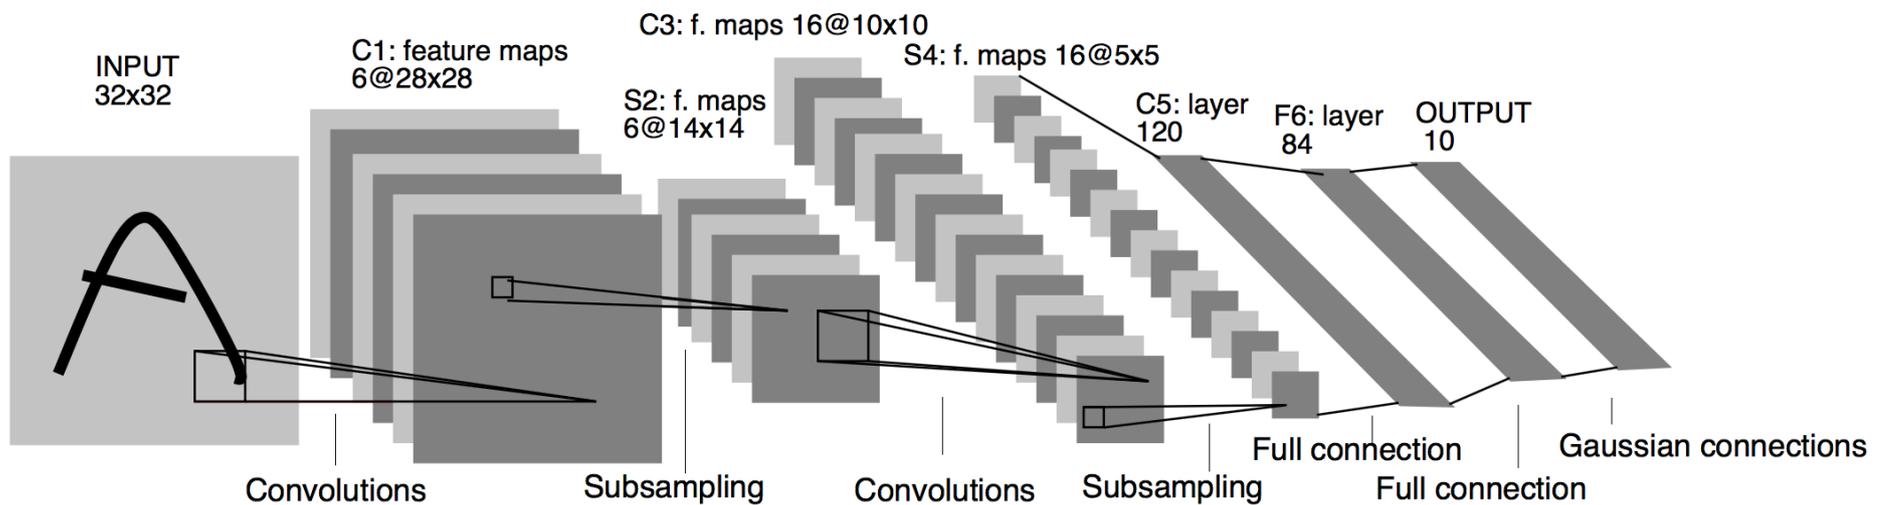
Deep Convolutional NN

- MNIST: 10 classes, 28x28 pixels
- ImageNet: 21841 classes, 256x256 pixels



Toilet paper, yacht, cup, cameleon...

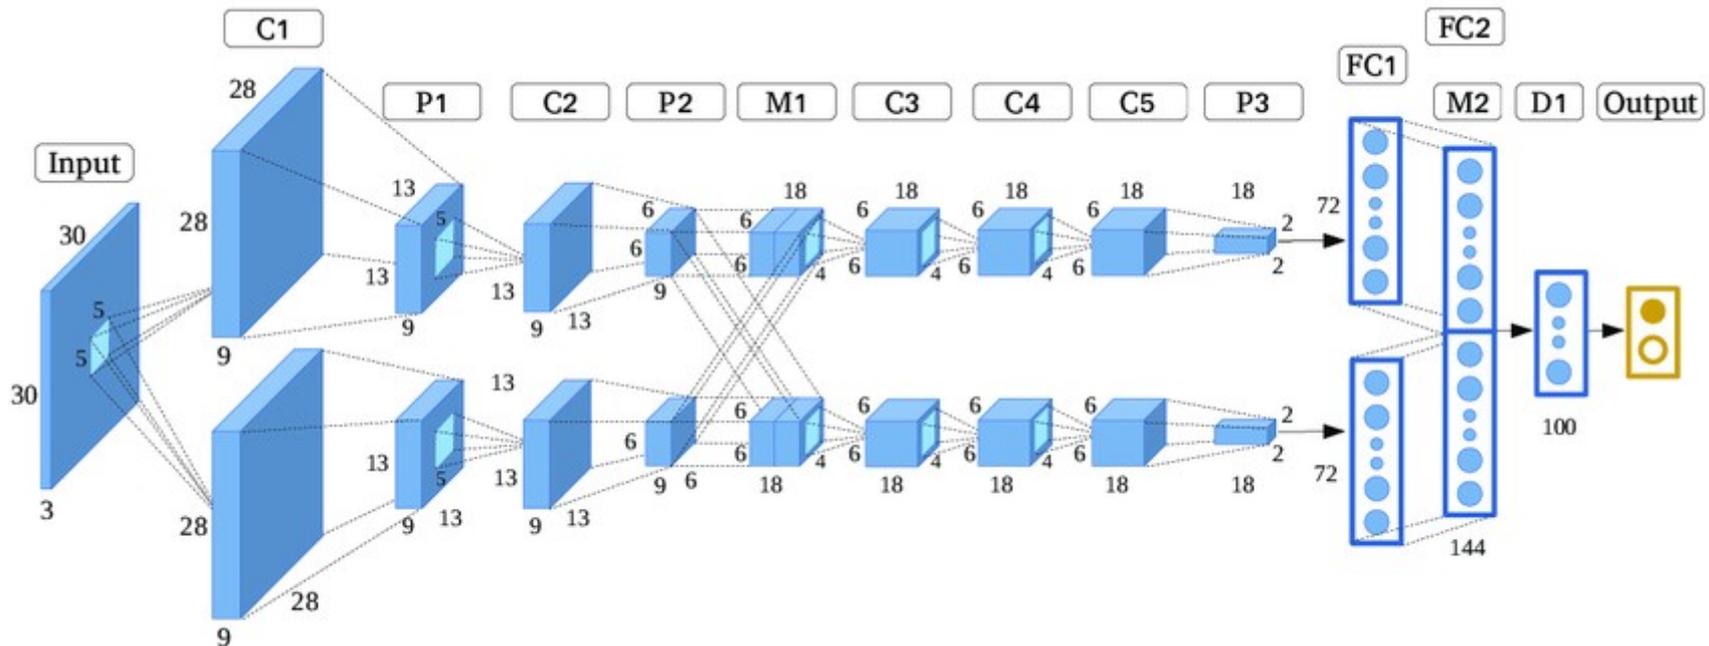
LeNet Architecture



~60k parameters, 5 layers

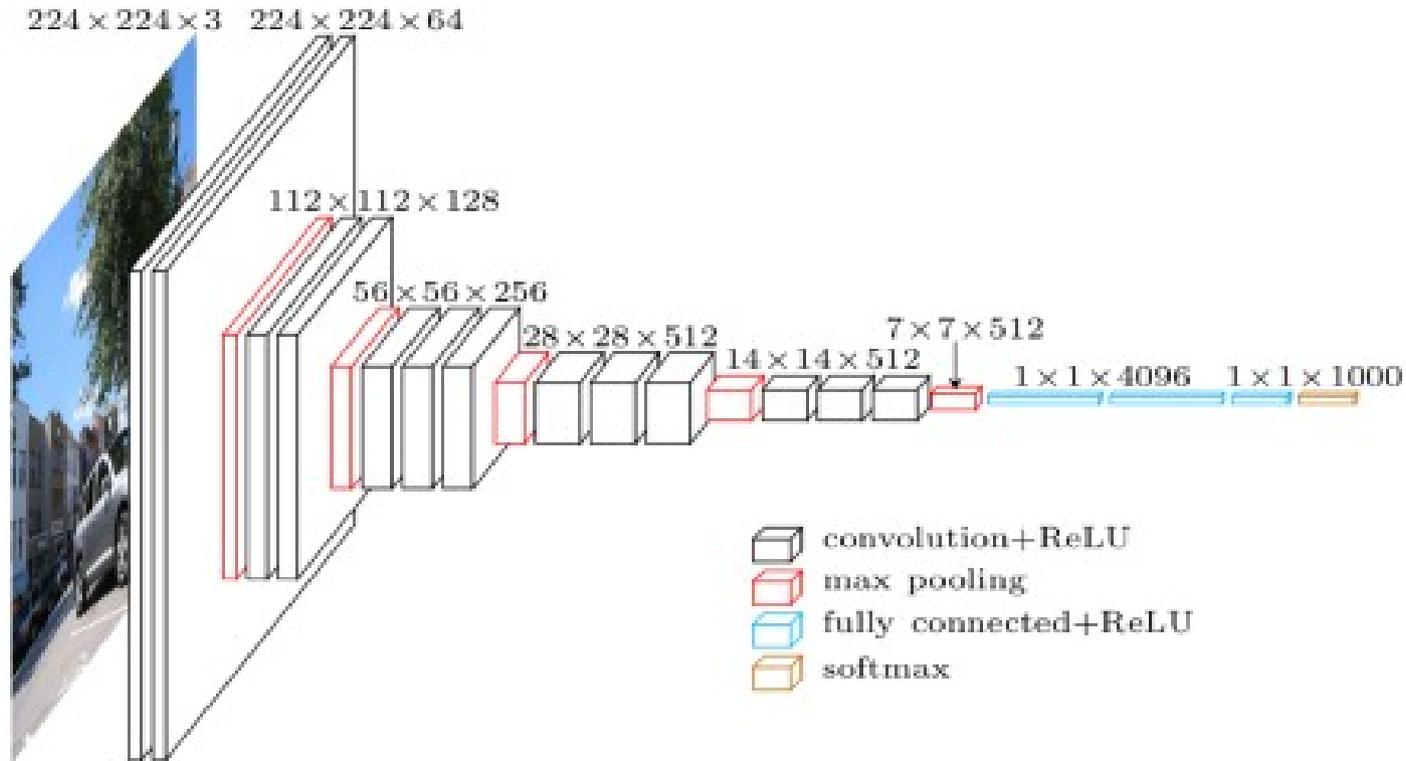
DCNN Architectures

- AlexNet (2012)
 - ~60M parameters, 8 layers



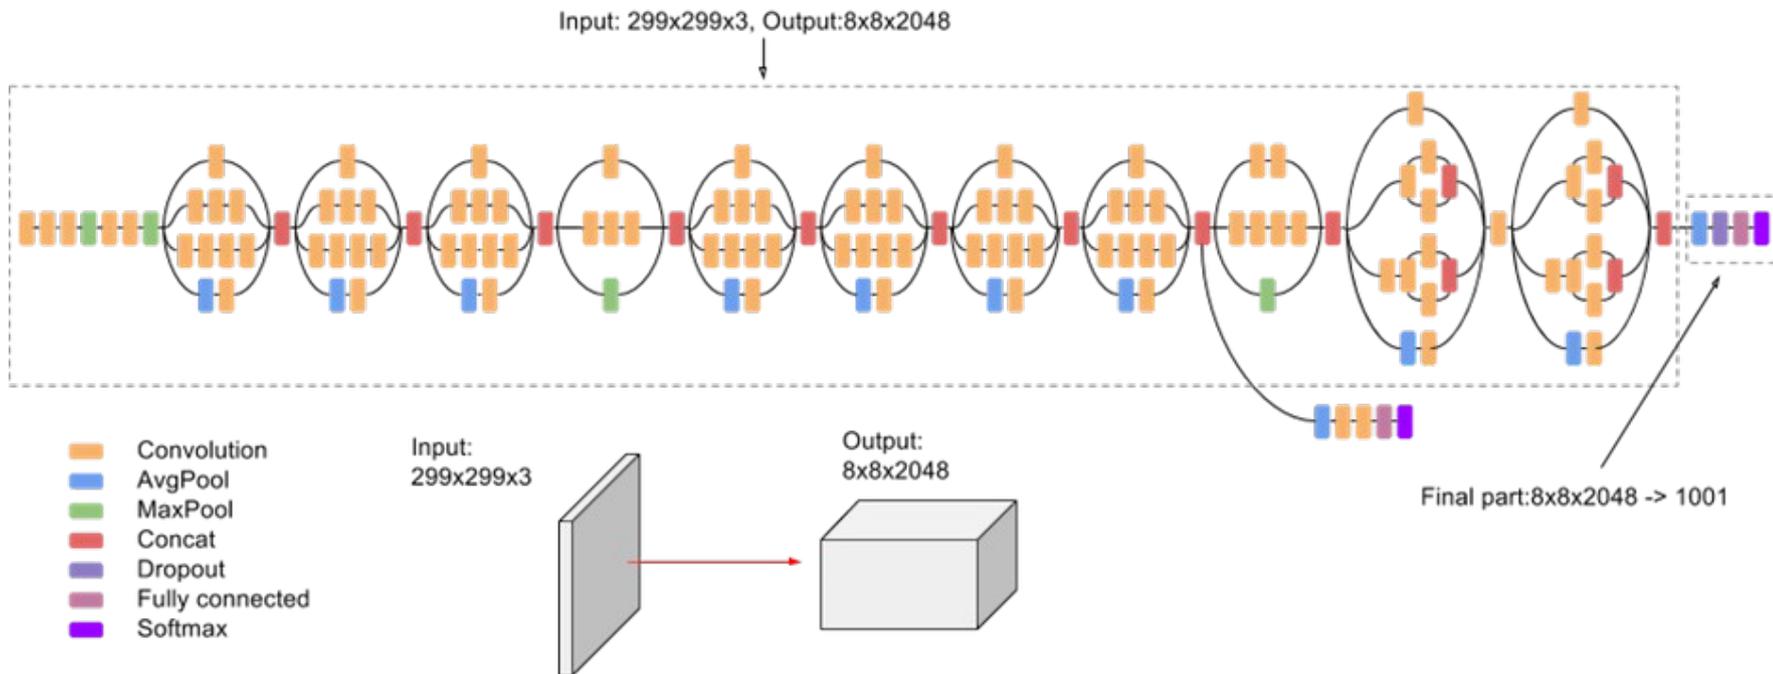
DCNN Architectures

- VGG (2014)
 - ~138M parameters, up to 19 layers



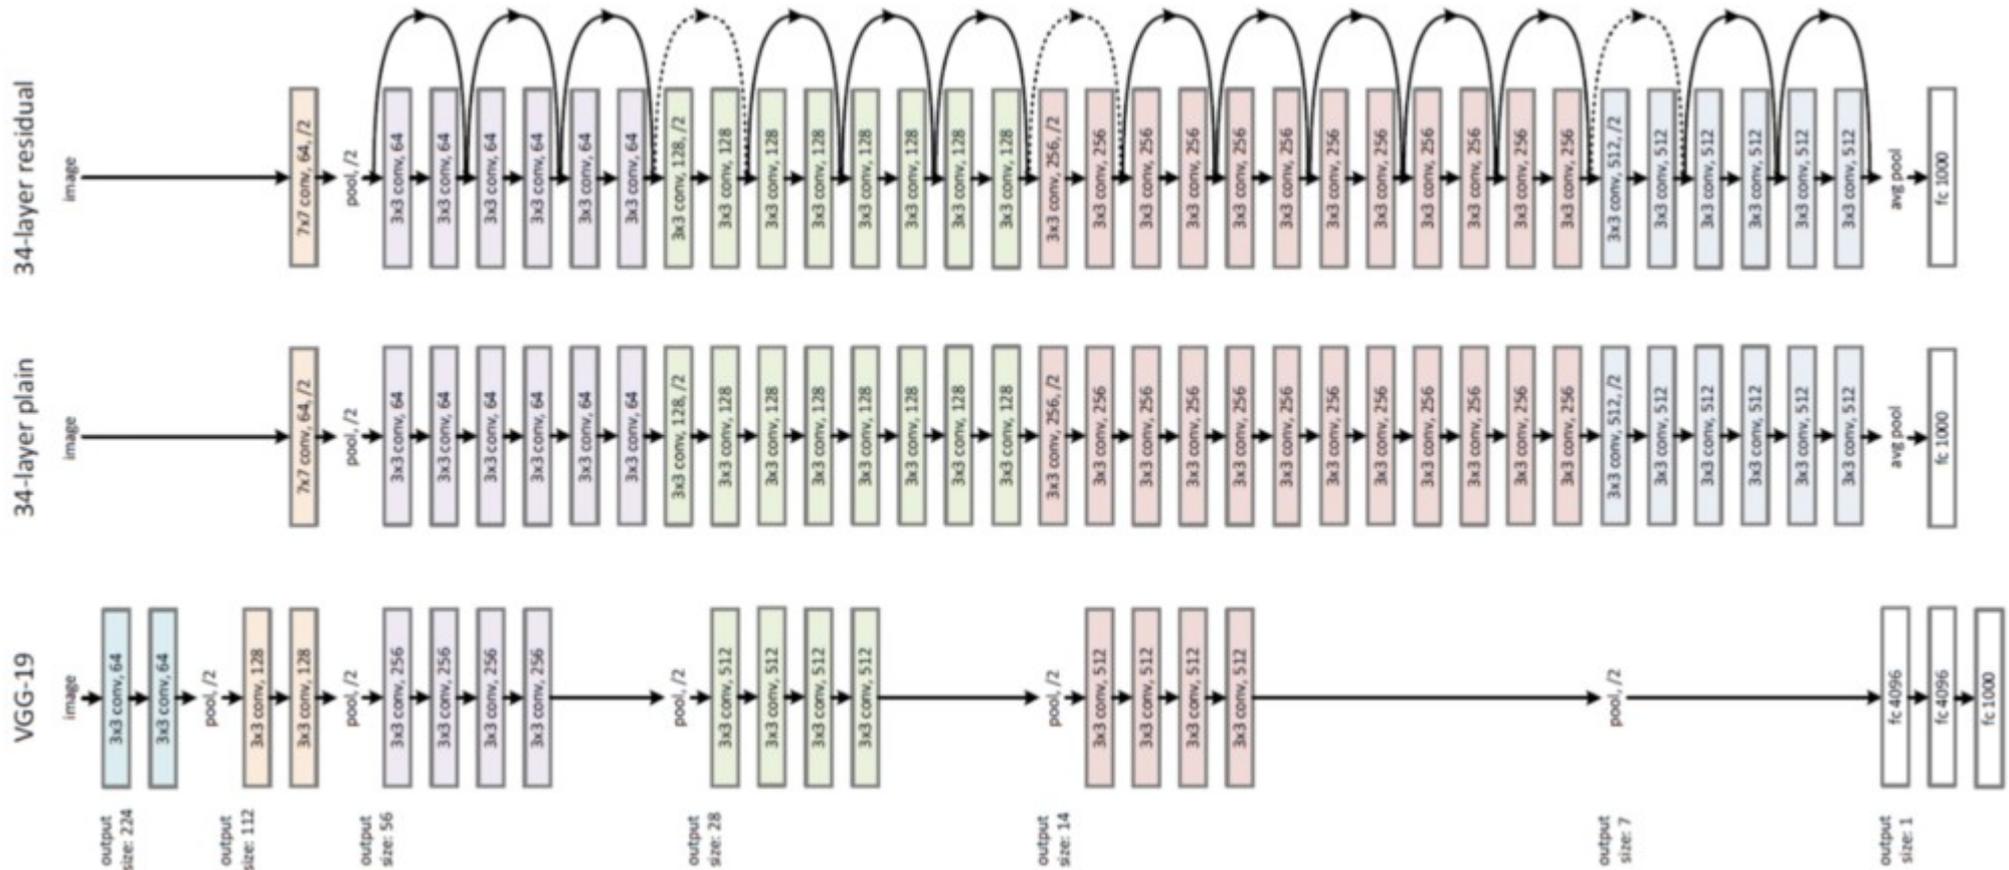
DCNN Architectures

- Inception (2014)
~7M parameters, 22 layers

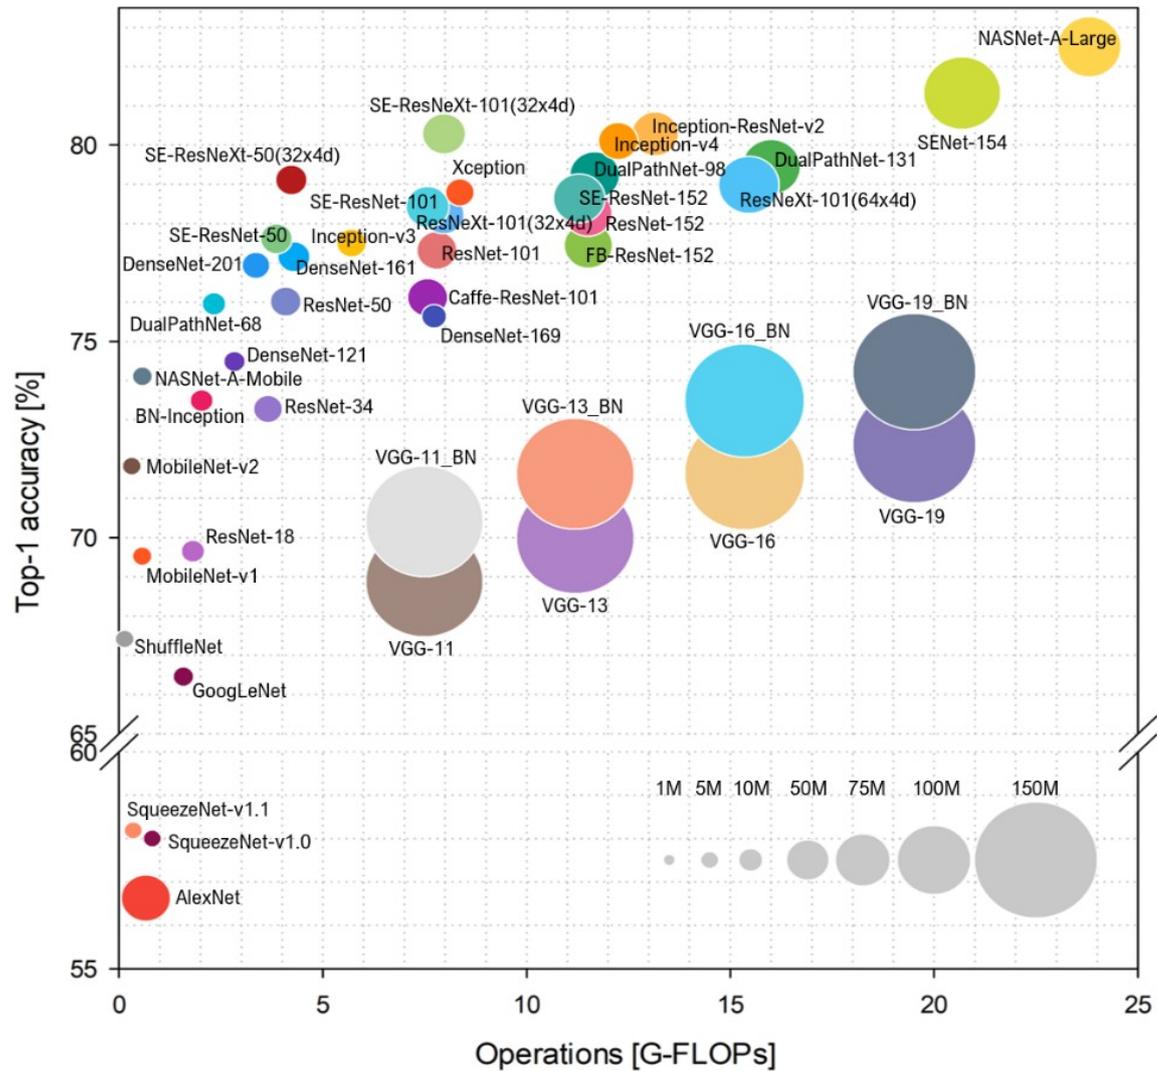


DCNN Architectures

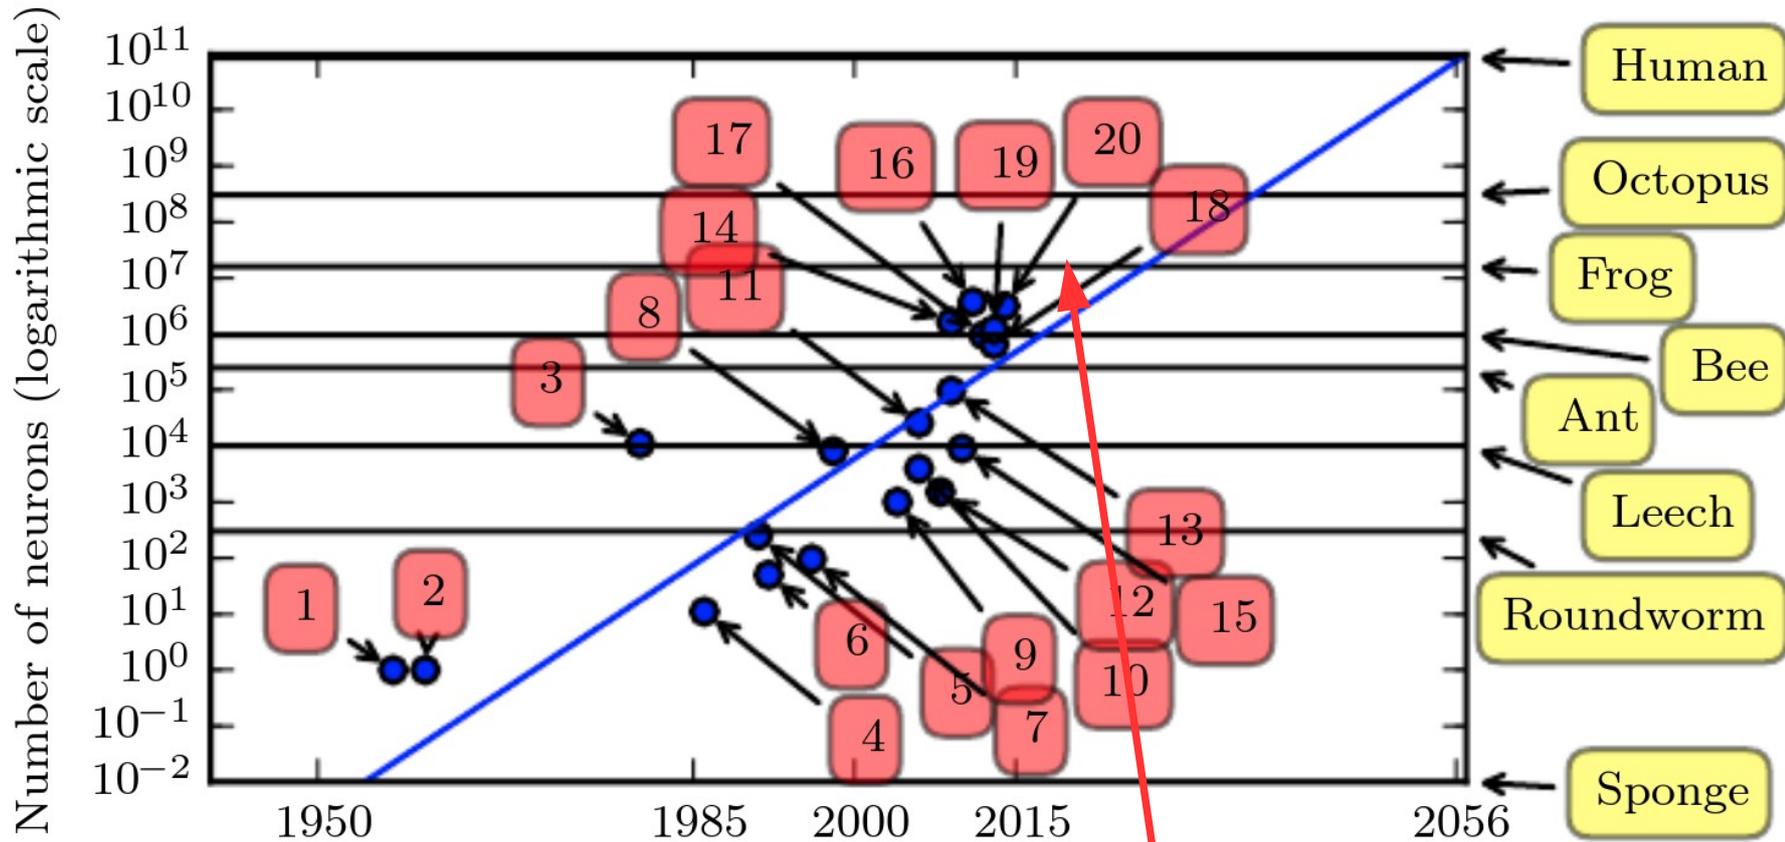
- ResNet (2015)
up to 1000 layers !



A summary

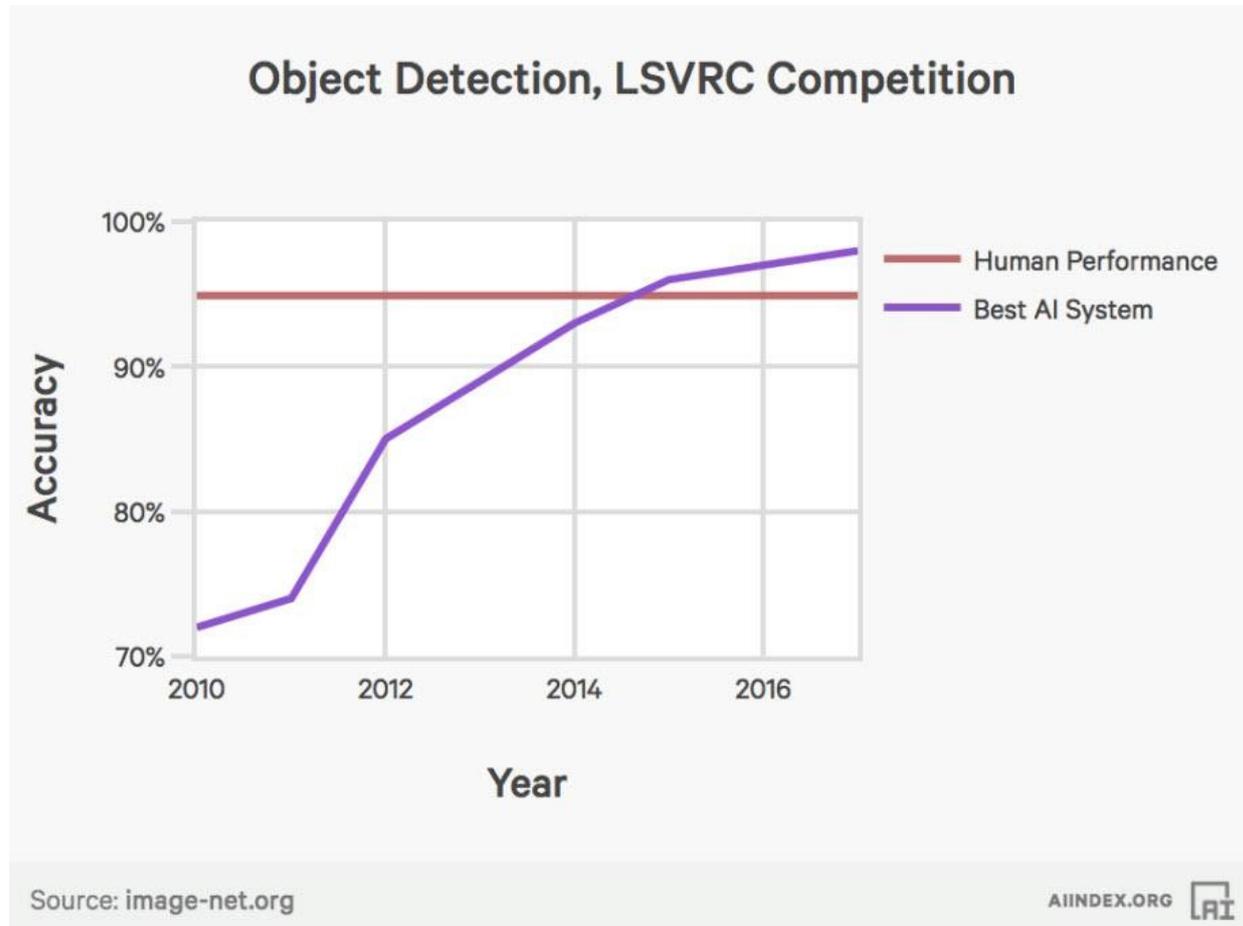


DCNN - Model complexity



**Inception
v3 (25M
neurons)**

Object detection challenge



Object detection - results

GeForce GTX 980M TF model ssd_mobilenet_v1_coco_11_06_2017
Prediction time: 40ms (25.3 fps) AVG: 44ms (22.6 fps)

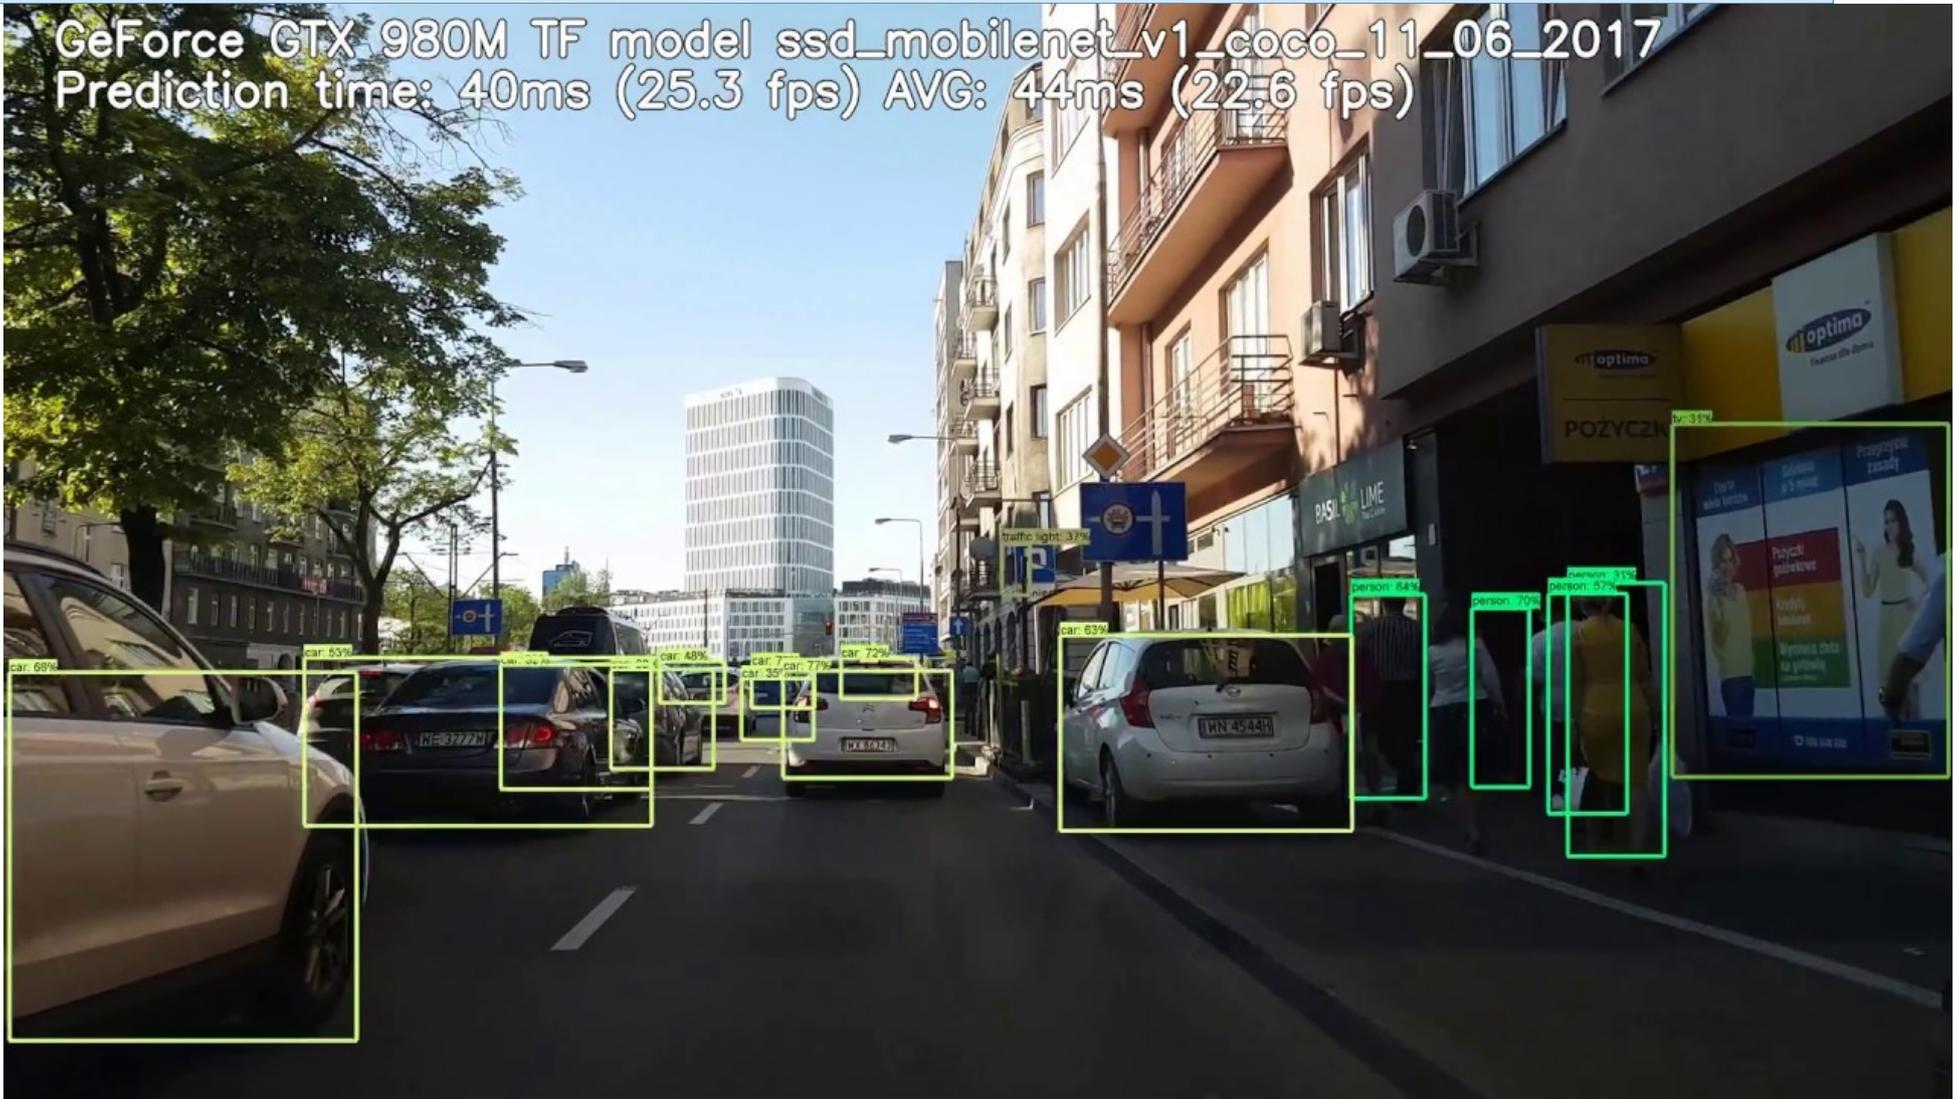
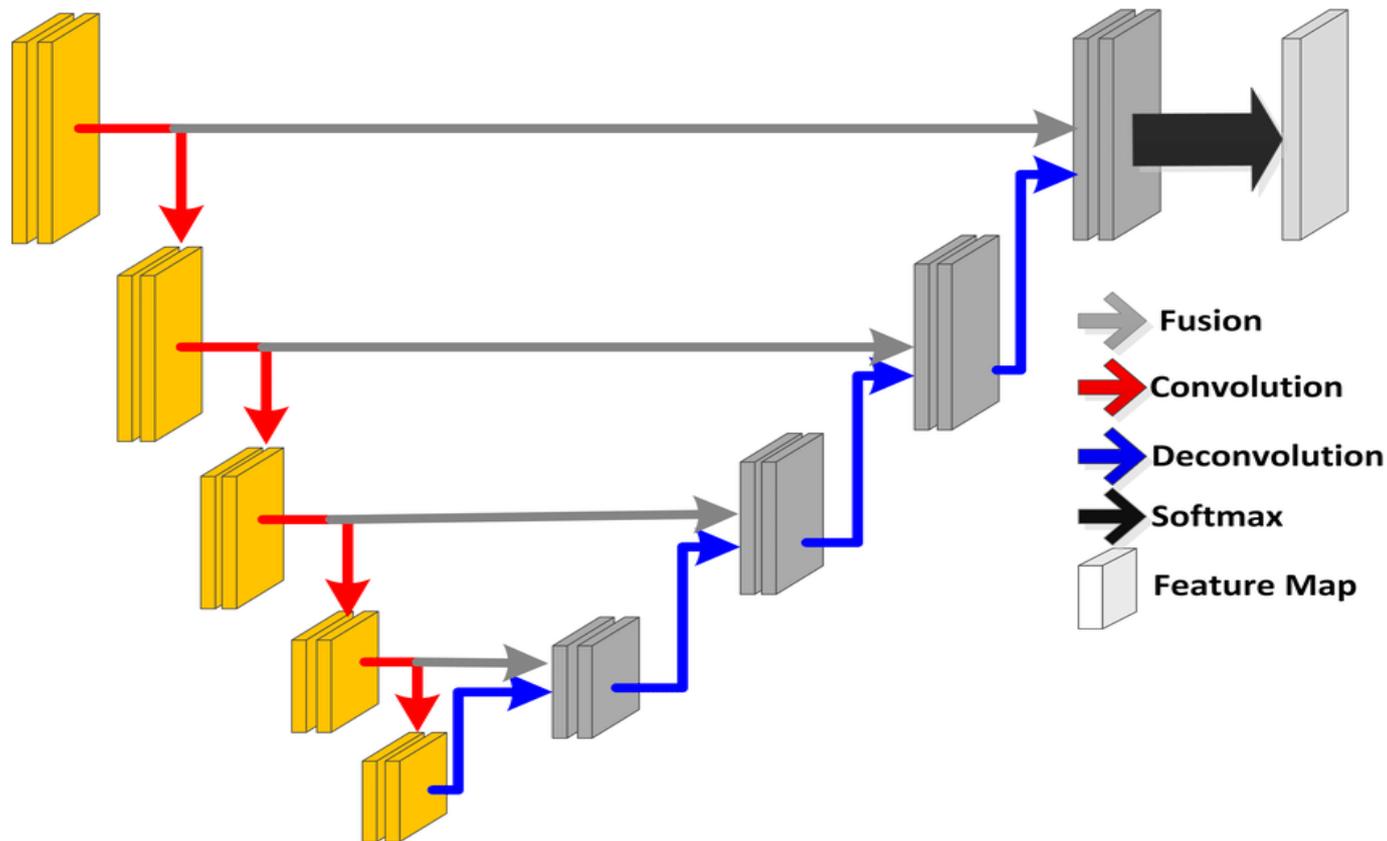


Image Segmentation

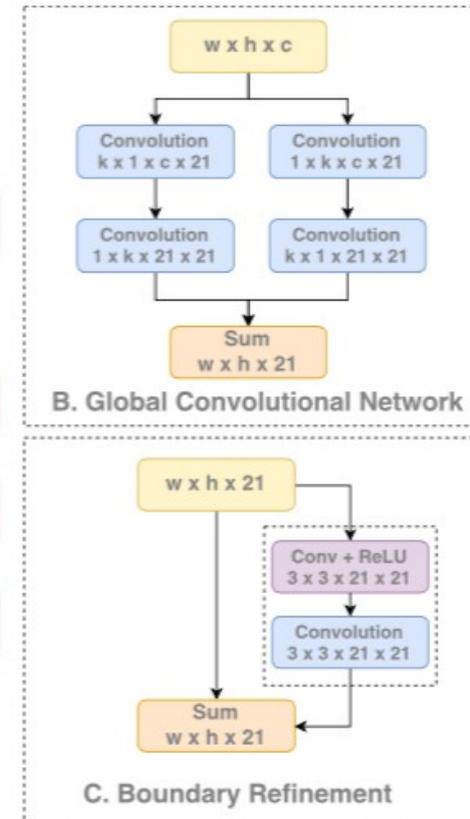
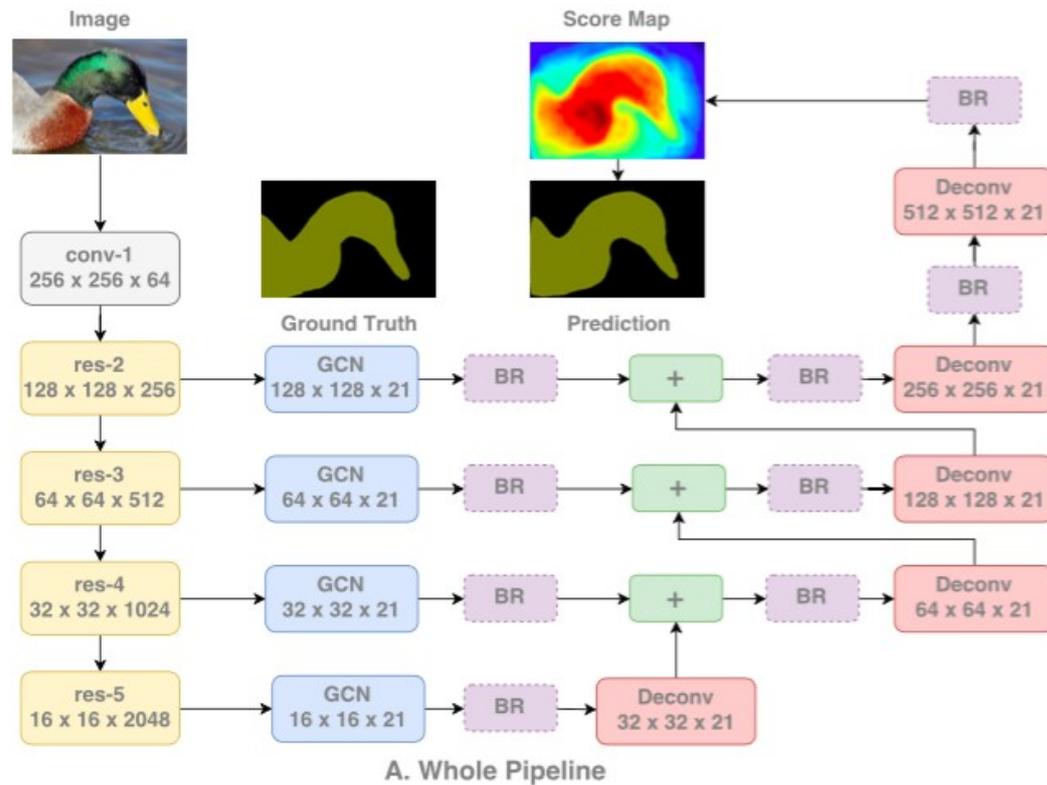


DCNN Architecture for Segmentation

- U-Net



DCNN Architecture for Segmentation

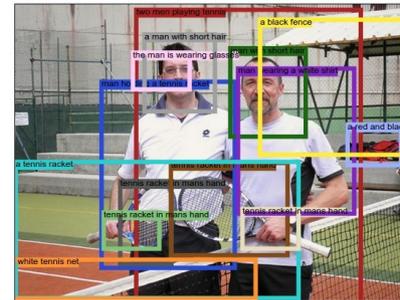


New problems in Computer Vision

- Human Parsing



- Image Captioning



two men playing tennis. man holding a tennis racket. tennis racket in mans hand. man with short hair. tennis racket in mans hand. man wearing a white shirt. a man with short hair. tennis racket in mans hand. a red and black bag. a tennis racket. a white tennis net. a black fence. tennis racket in mans hand. the man is wearing glasses.

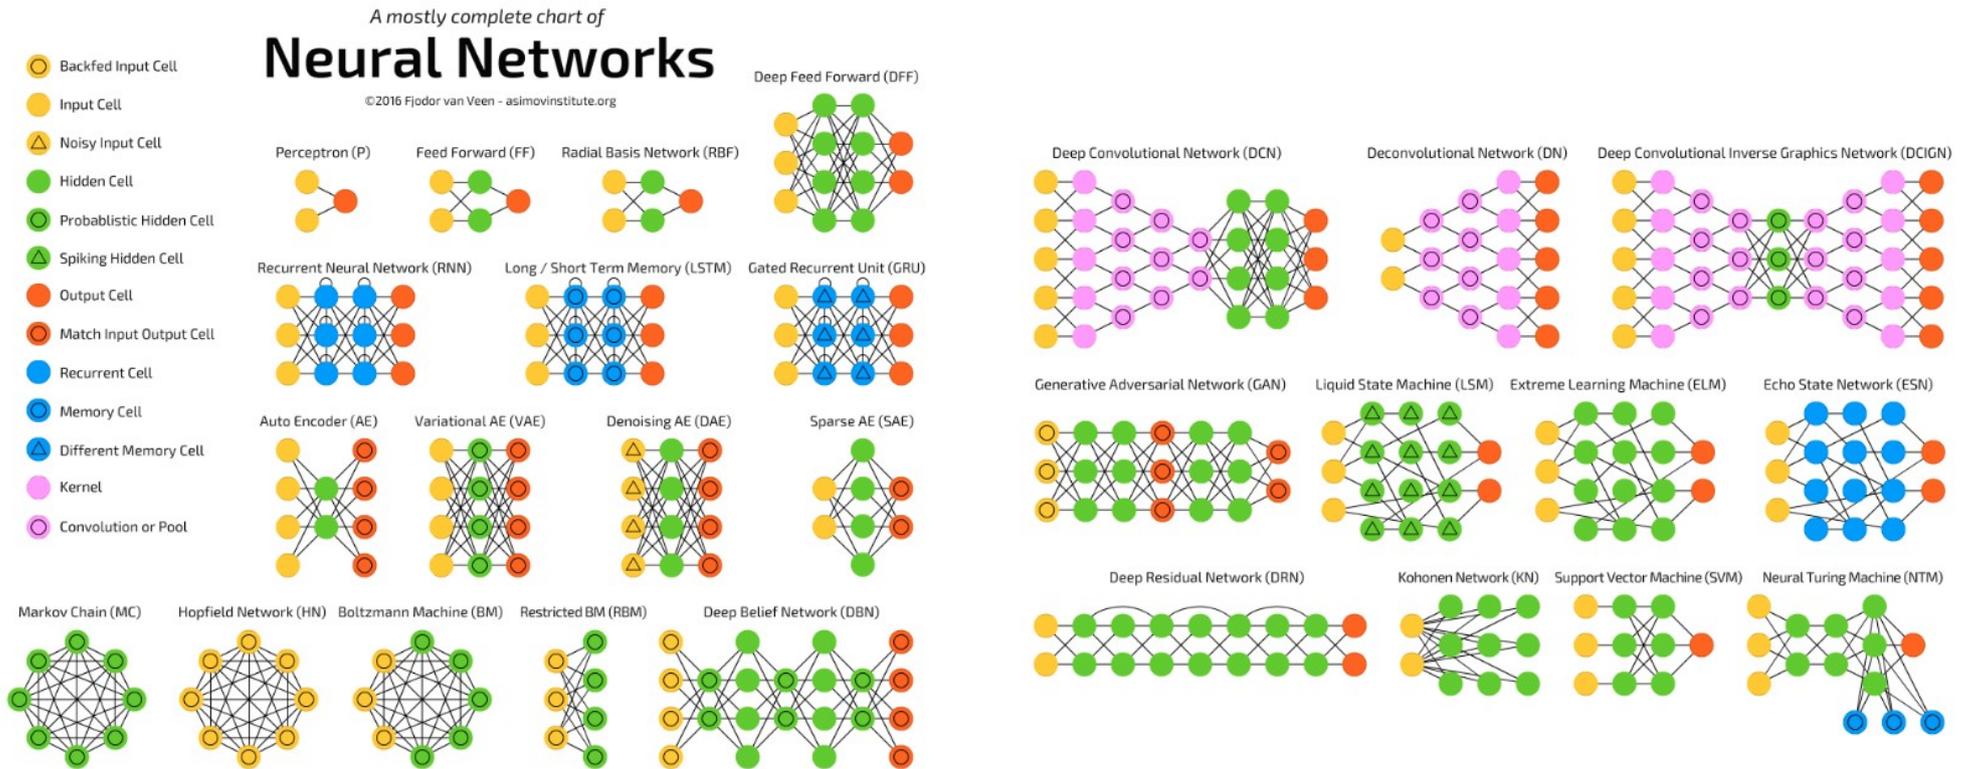
- 3D reconstruction from a single image



Overview

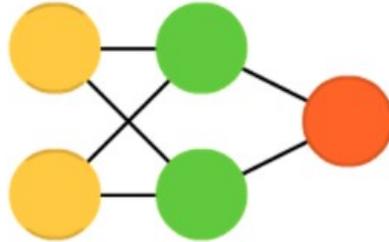
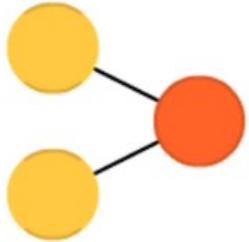
- **Deep** Convolutional Neural Networks
- The Deep Learning Zoo
- Generative Models
- Concluding Remarks

The Deep Learning Zoo



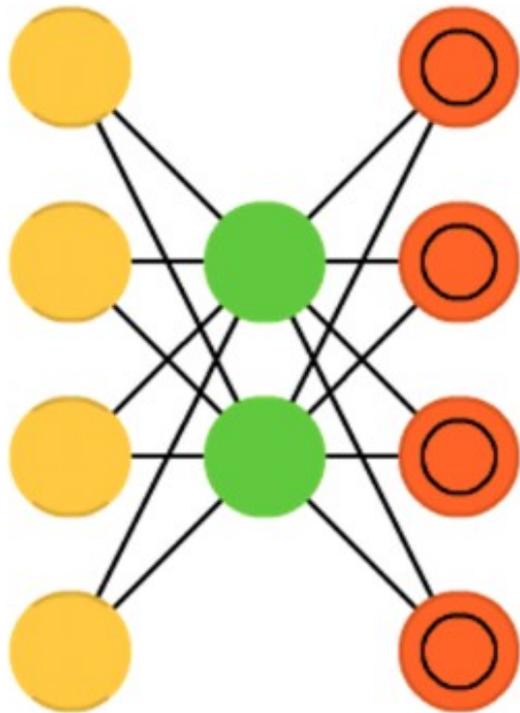
<http://www.asimovinstitute.org>

Perceptron, FFNN



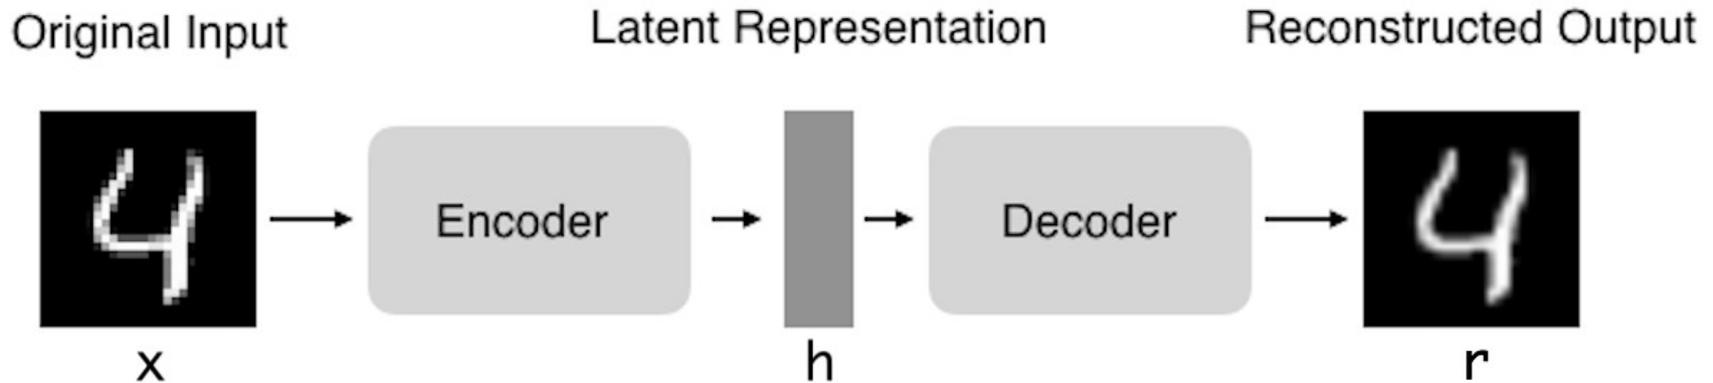
- Backfed Input Cell
- Input Cell
- △ Noisy Input Cell
- Hidden Cell
- Probablistic Hidden Cell
- △ Spiking Hidden Cell
- Output Cell
- Match Input Output Cell
- Recurrent Cell
- Memory Cell
- △ Different Memory Cell
- Kernel
- Convolution or Pool

Auto-encoders



- Backfed Input Cell
- Input Cell
- △ Noisy Input Cell
- Hidden Cell
- Probabilistic Hidden Cell
- △ Spiking Hidden Cell
- Output Cell
- Match Input Output Cell
- Recurrent Cell
- Memory Cell
- △ Different Memory Cell
- Kernel
- Convolution or Pool

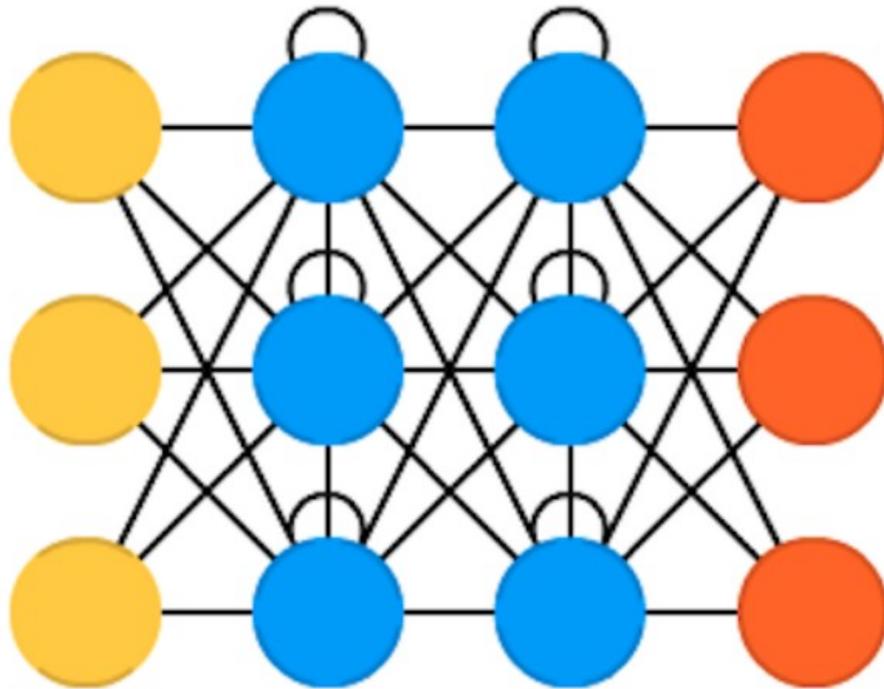
Auto-encoders



- Applications

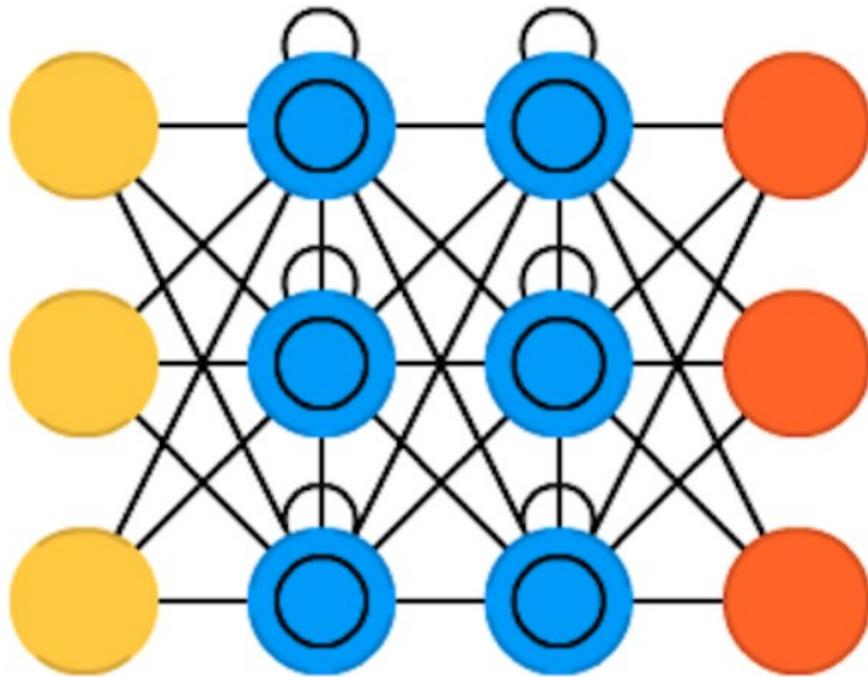
- Denoising
- Dimension Reduction
- Indexation

Recurrent Neural Networks



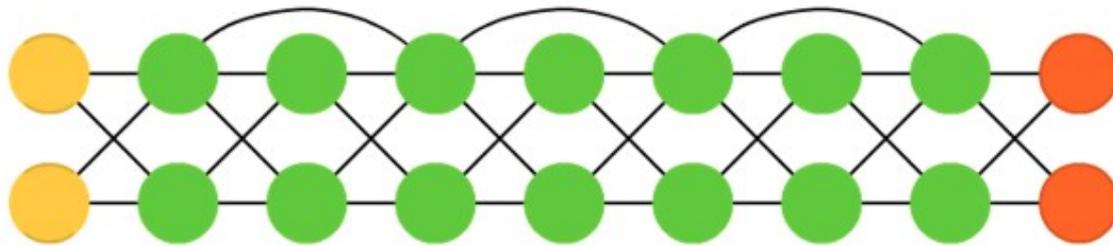
- Backfed Input Cell
- Input Cell
- △ Noisy Input Cell
- Hidden Cell
- Probablistic Hidden Cell
- △ Spiking Hidden Cell
- Output Cell
- Match Input Output Cell
- Recurrent Cell
- Memory Cell
- △ Different Memory Cell
- Kernel
- Convolution or Pool

Long Short-Term Memory (LSTM)



- Backfed Input Cell
- Input Cell
- △ Noisy Input Cell
- Hidden Cell
- Probablistic Hidden Cell
- △ Spiking Hidden Cell
- Output Cell
- Match Input Output Cell
- Recurrent Cell
- Memory Cell
- △ Different Memory Cell
- Kernel
- Convolution or Pool

Deep Residual Networks

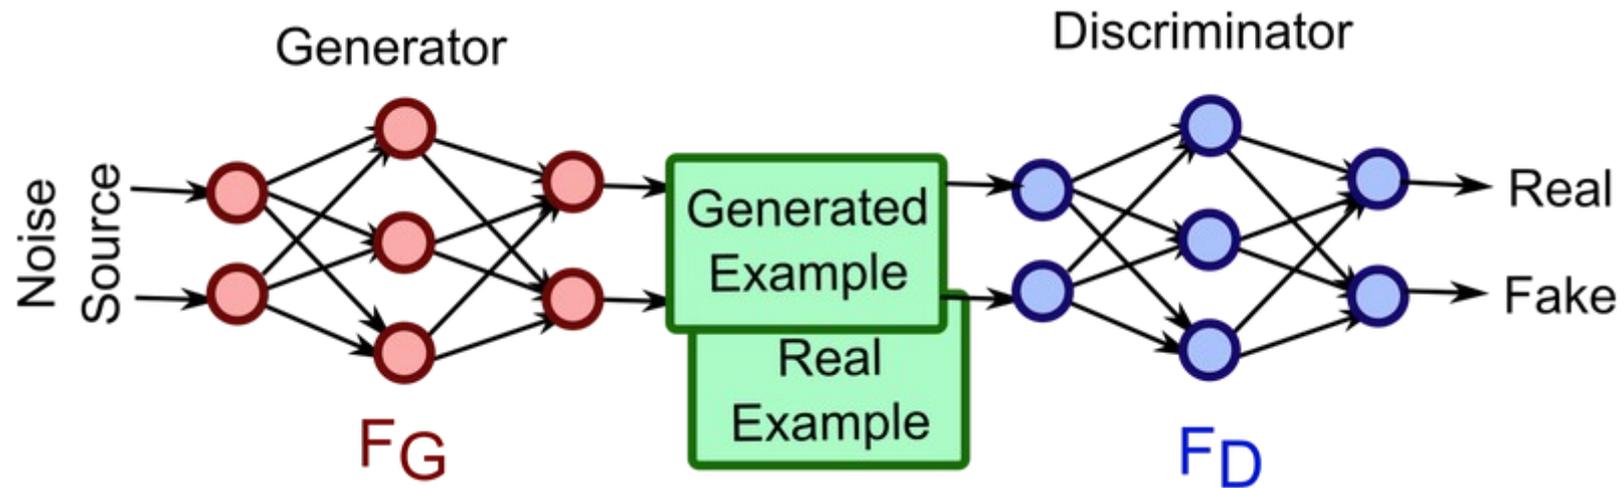


- Backfed Input Cell
- Input Cell
- △ Noisy Input Cell
- Hidden Cell
- Probablistic Hidden Cell
- △ Spiking Hidden Cell
- Output Cell
- Match Input Output Cell
- Recurrent Cell
- Memory Cell
- △ Different Memory Cell
- Kernel
- Convolution or Pool

Overview

- **Deep** Convolutional Neural Networks
- The Deep Learning Zoo
- **Generative Models**
- Concluding Remarks

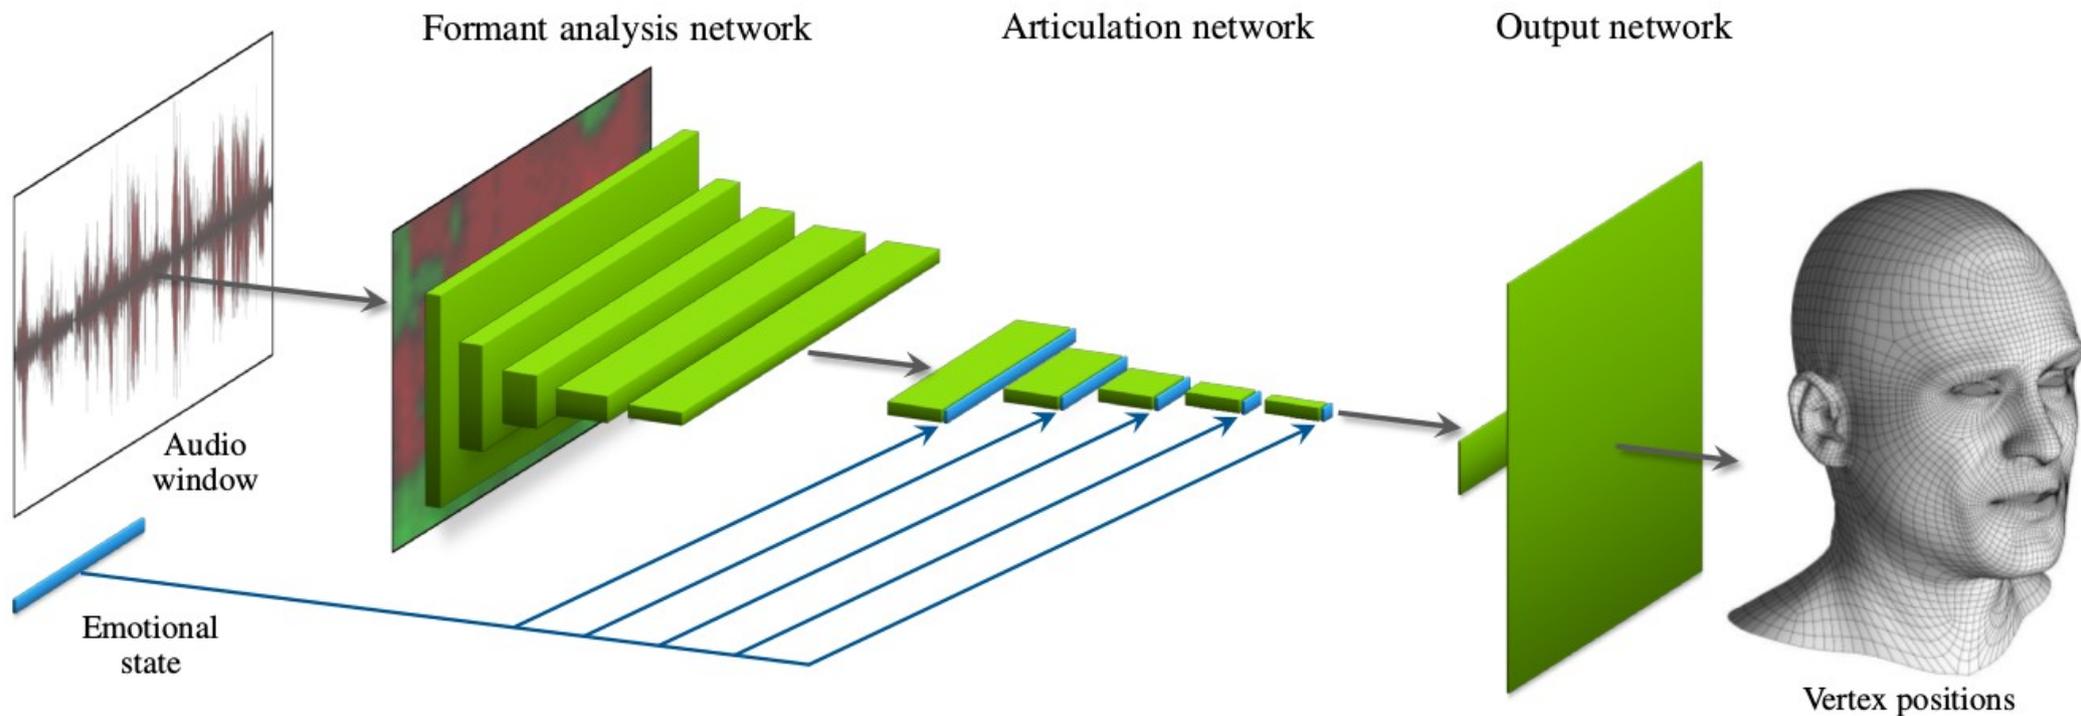
Generative Adversarial Networks



Game theory
Nash equilibrium

Generative Adversarial Networks

- Difficult to train...
- But useful in many applications



Generative Adversarial Networks

- Example Video

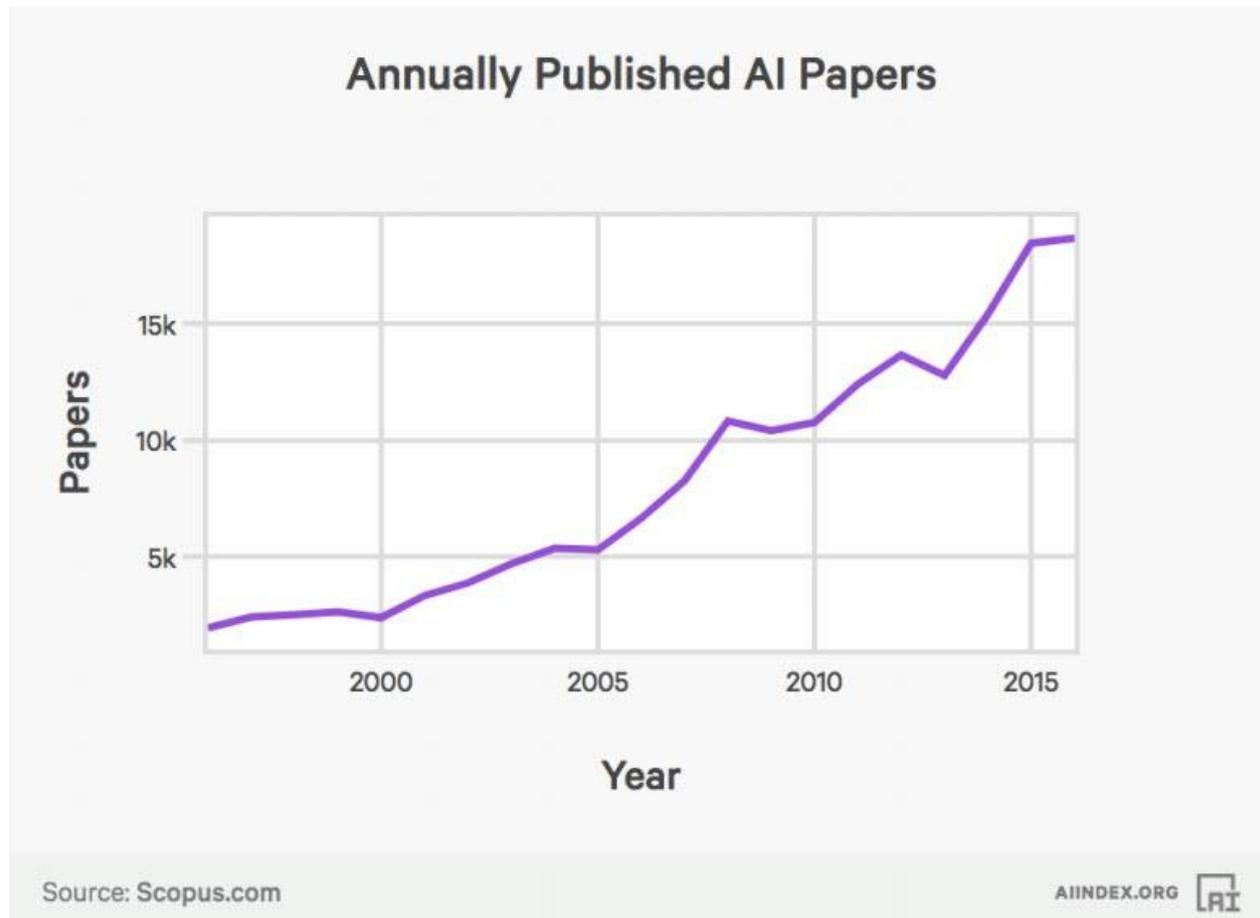
Generative Adversarial Networks



Overview

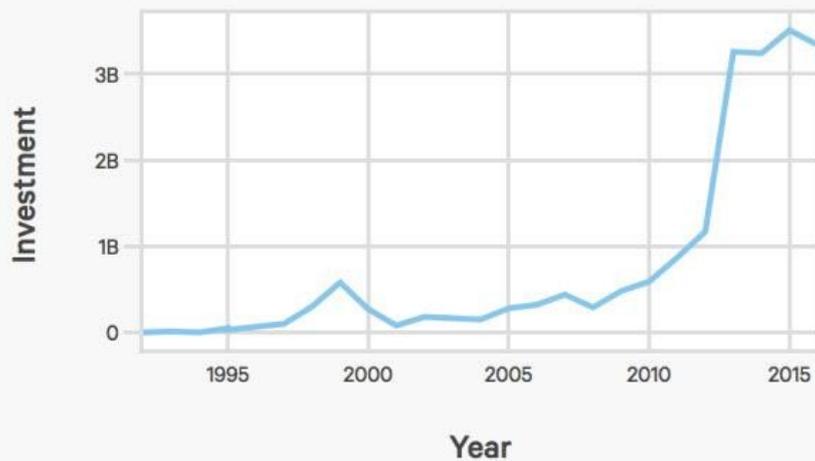
- **Deep** Convolutional Neural Networks
- The Deep Learning Zoo
- Generative Models
- Concluding Remarks

About AI, Deep Learning...

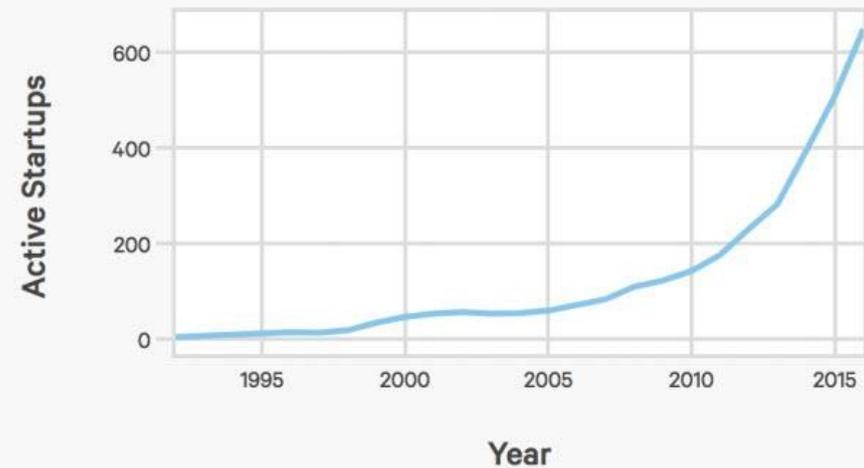


About AI, Deep Learning...

Annual VC Investment in AI Startups



Startups Developing AI Systems



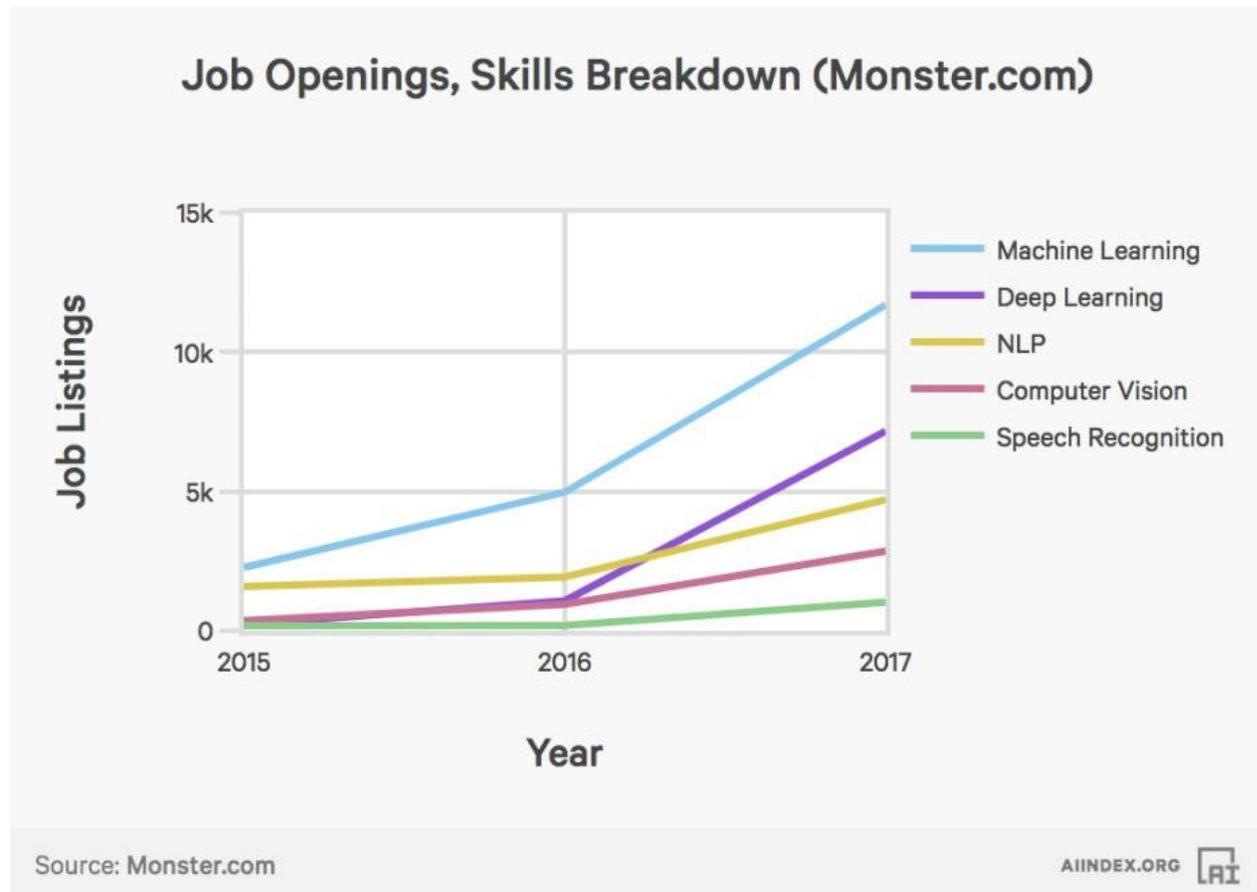
Sources: Crunchbase, VentureSource, Sand Hill Econometrics

AIINDEX.ORG 

Sources: Crunchbase, VentureSource, Sand Hill Econometrics

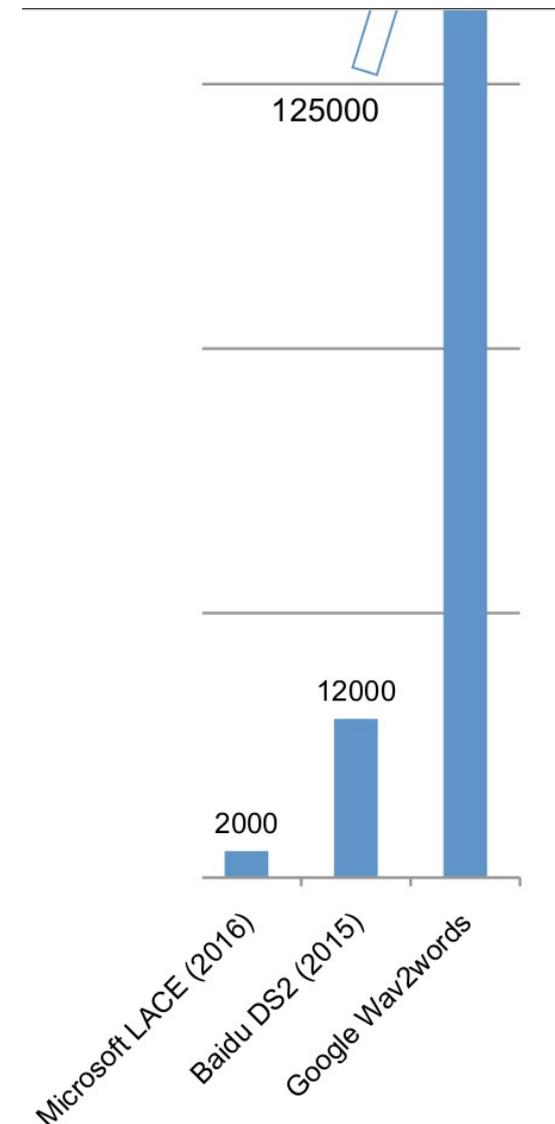
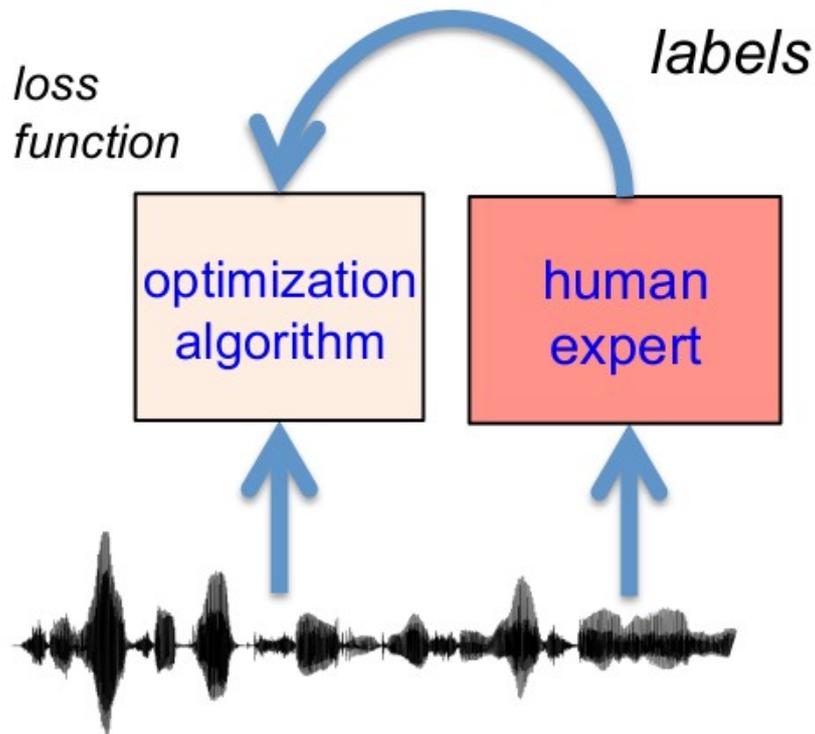
AIINDEX.ORG 

About AI, Deep Learning...



Challenges - Supervision

« She had your dark suit in greasy wash water all year »



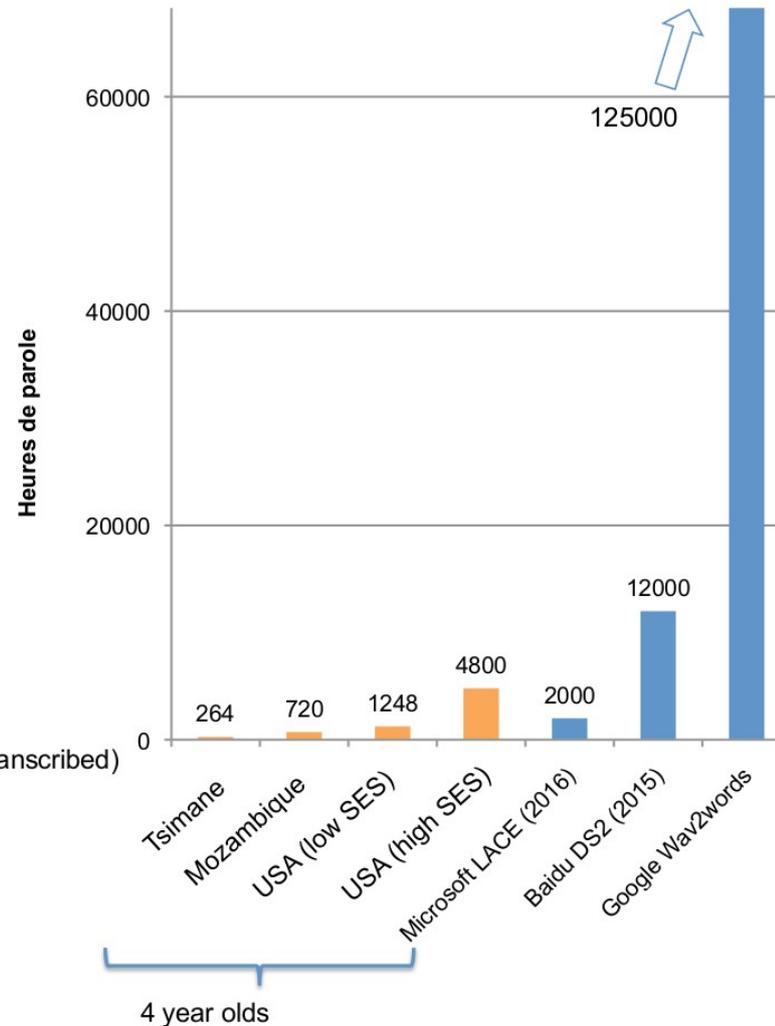
Challenges - Supervision

How do infants do?



Supervised (orthographically transcribed)

Cristia et al, 2017



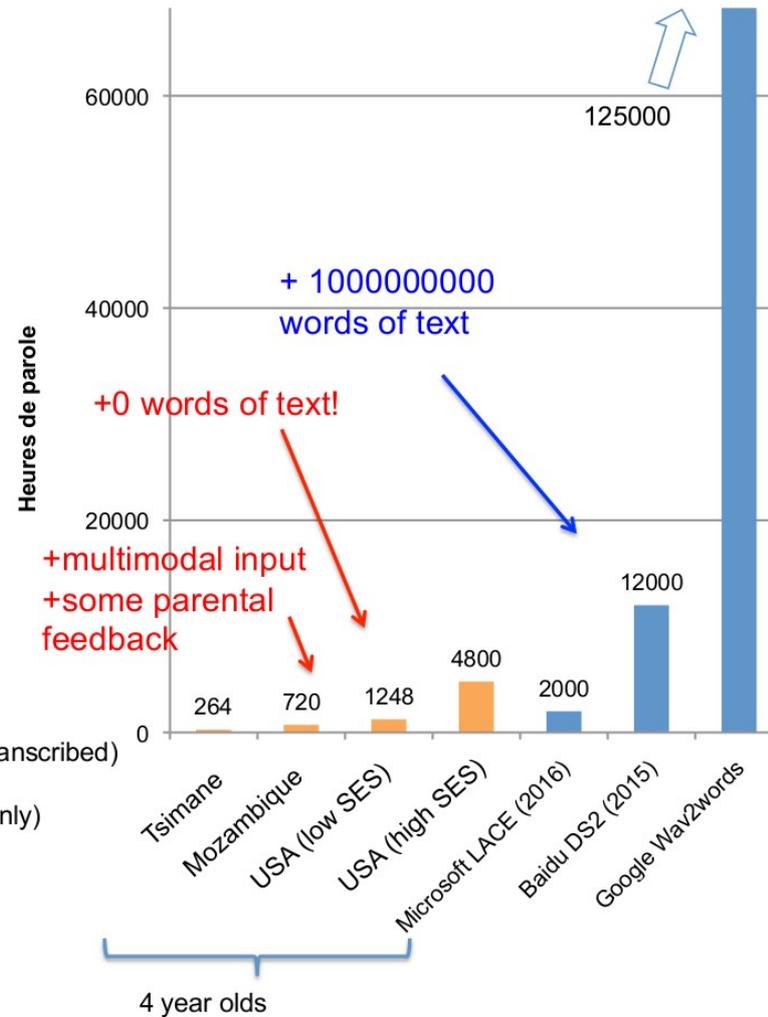
Challenges - Supervision

How do infants learn language?

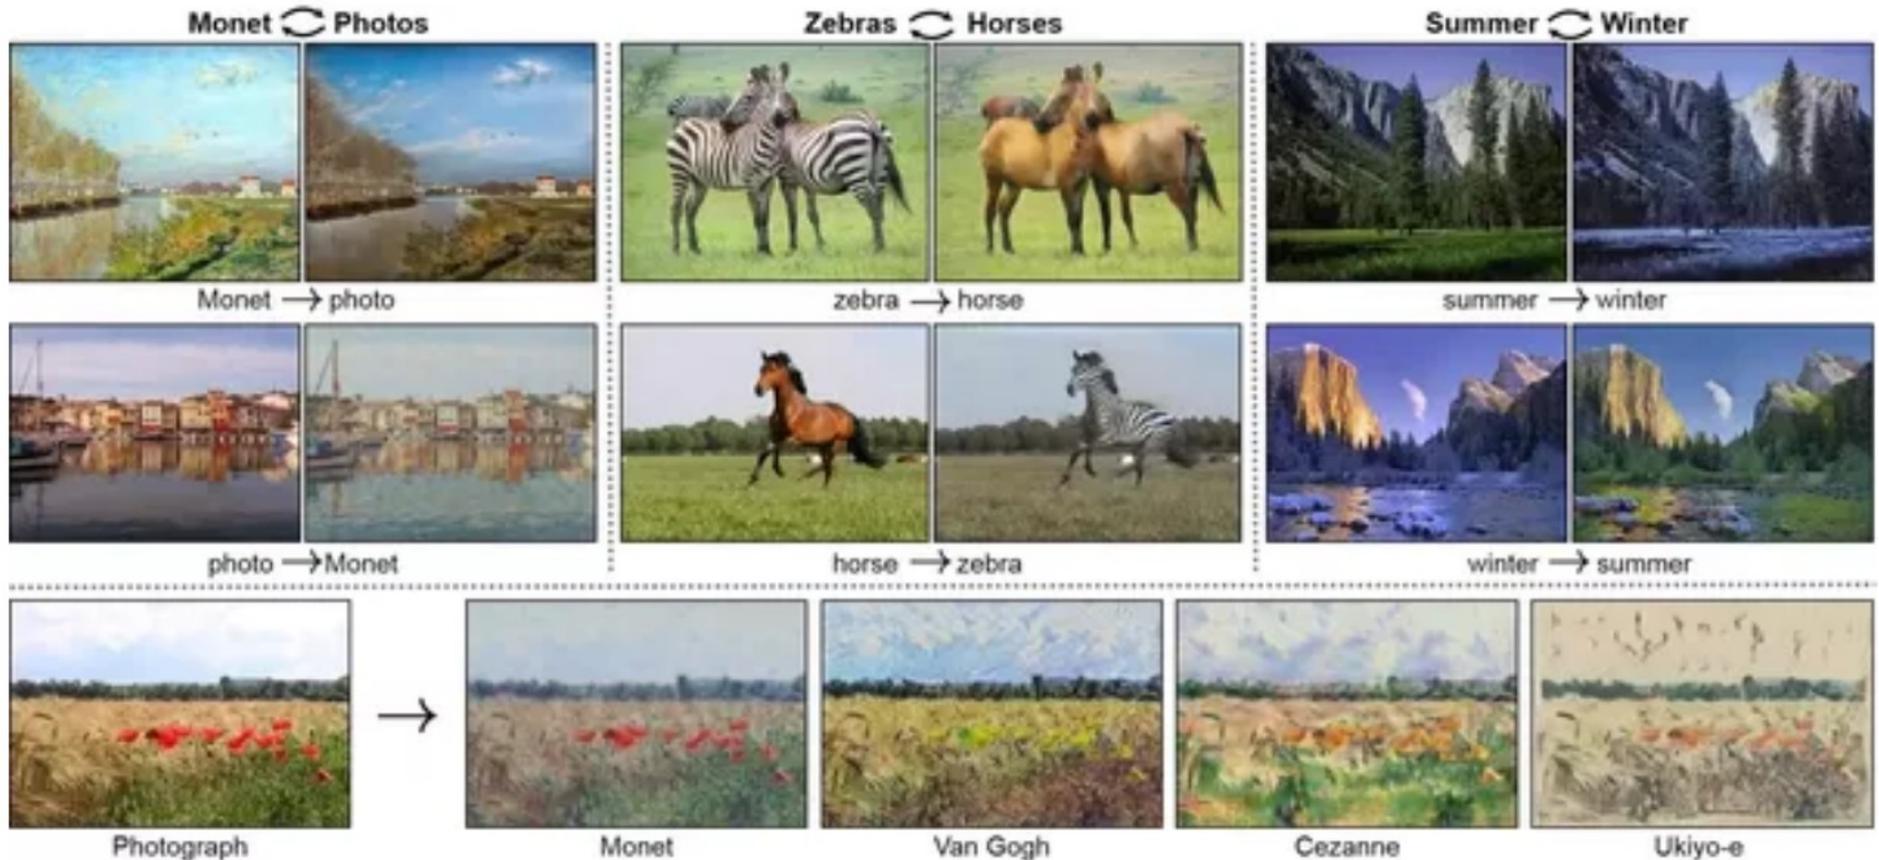


- Supervised (orthographically transcribed)
- Unsupervised (sensory data only)

Cristia et al, 2017



Challenges - Domain Adaptation



Challenges - Domain Adaptation

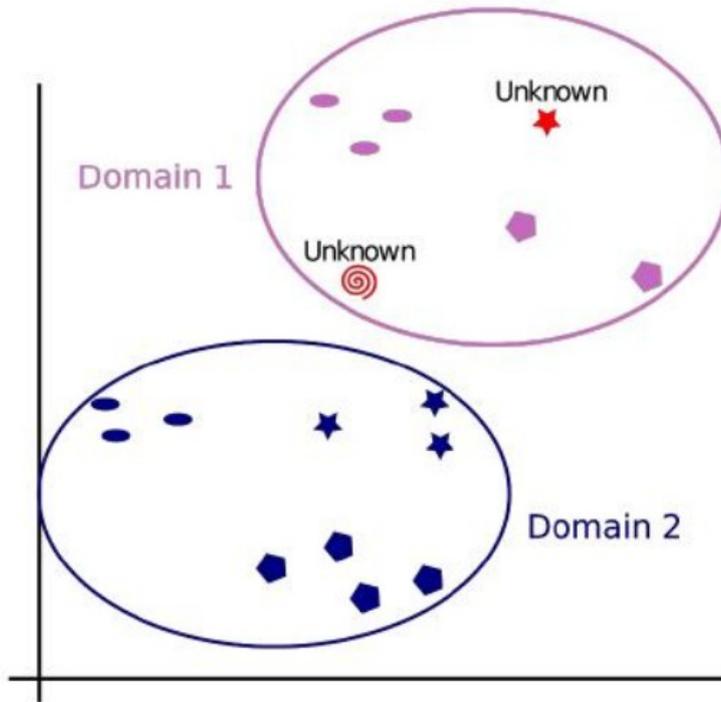


Challenges - Domain Adaptation

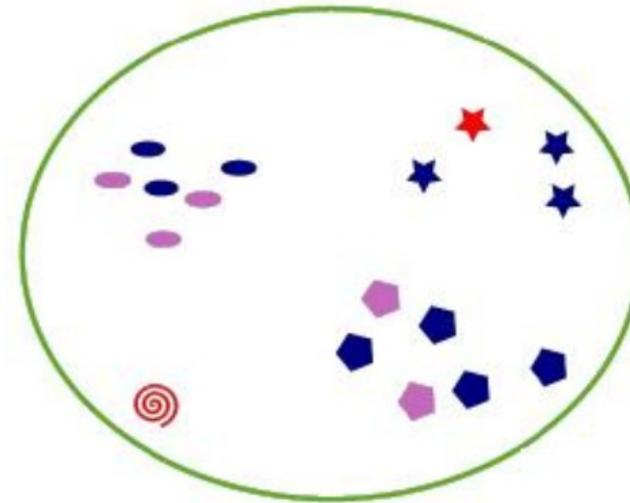


Challenges - Domain Adaptation

What we have

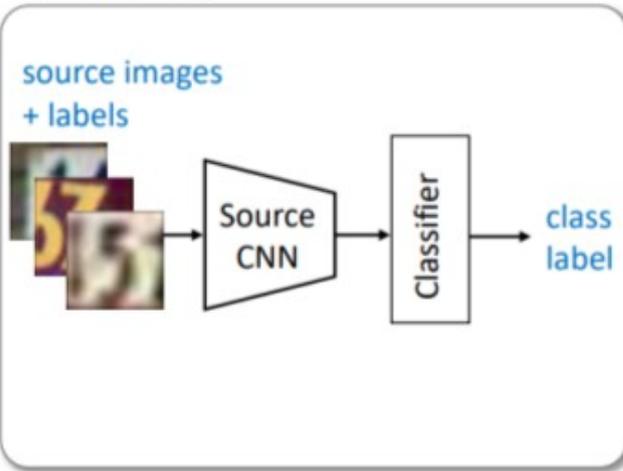


What we want

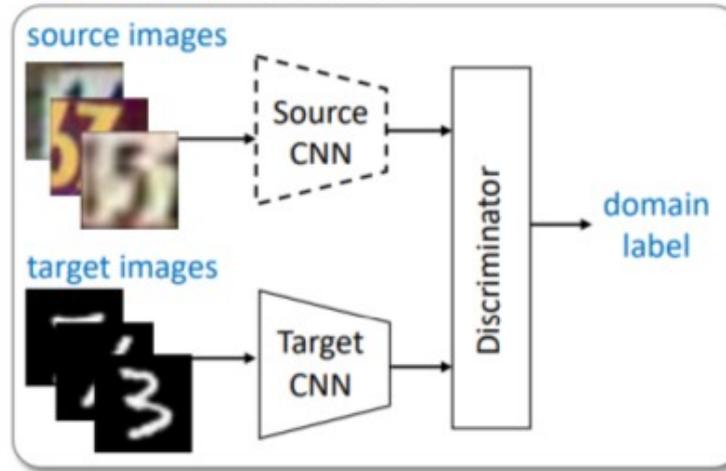


Challenges - Domain Adaptation

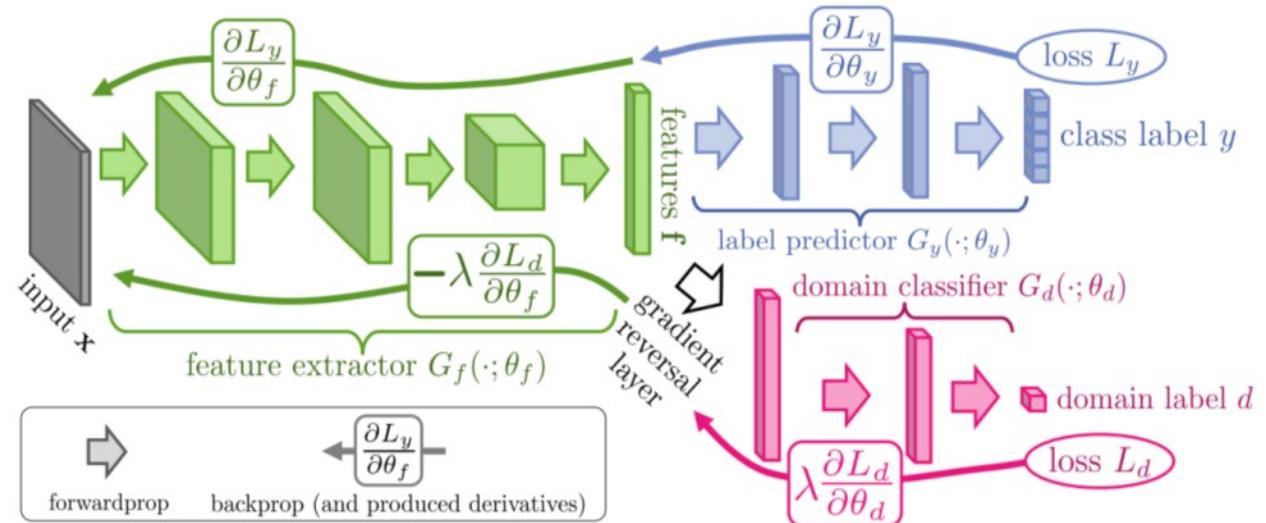
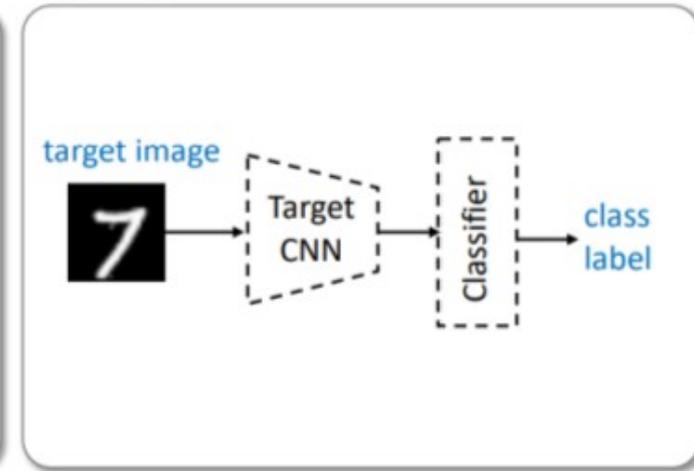
Pre-training



Adversarial Adaptation



Testing



Multiple Task Learning

MENTAL DEVELOPMENT	Average age skills begin	3 months	6 months	9 months	1 year	2 years	3 years	5 years
Communication and language	cries when wet or hungry	coos when comfortable	makes simple sounds	uses certain sounds for different things	begins to use simple single words	begins to use words together	uses simple sentences	
Social Behavior		smiles when smiled at		begins to understand and respond to "NO!"	begins to do simple things when asked	likes to be praised after completing simple tasks	interacts with both adults and children	
Self-care	sucks breast	takes everything to mouth		chews solid food	drinks alone from glass	takes off simple clothes	toilet trained	helps with simple work
Attention and interest	smiles when smiled at	brief interest in toys and sounds		develops strong attachments to caretakers	takes longer interest in toys and activities	sorts different objects	builds playthings with several pieces	
Play	grasps things placed in hand	plays with own body	plays with simple objects	begins to enjoy first social games (peek-a-boo)	imitates and copies people	begins to play with other children	plays independently with children and toys	
Intelligence and learning	cries when hungry or uncomfortable	recognizes mother	recognizes several people	looks for toys that fall out of sight	copies simple actions	points to things when asked	follows simple instructions	follows multiple instructions

What is Intelligence?

- Curiosity
- Adaptation to new situations
- Knowledge transfert
- ...