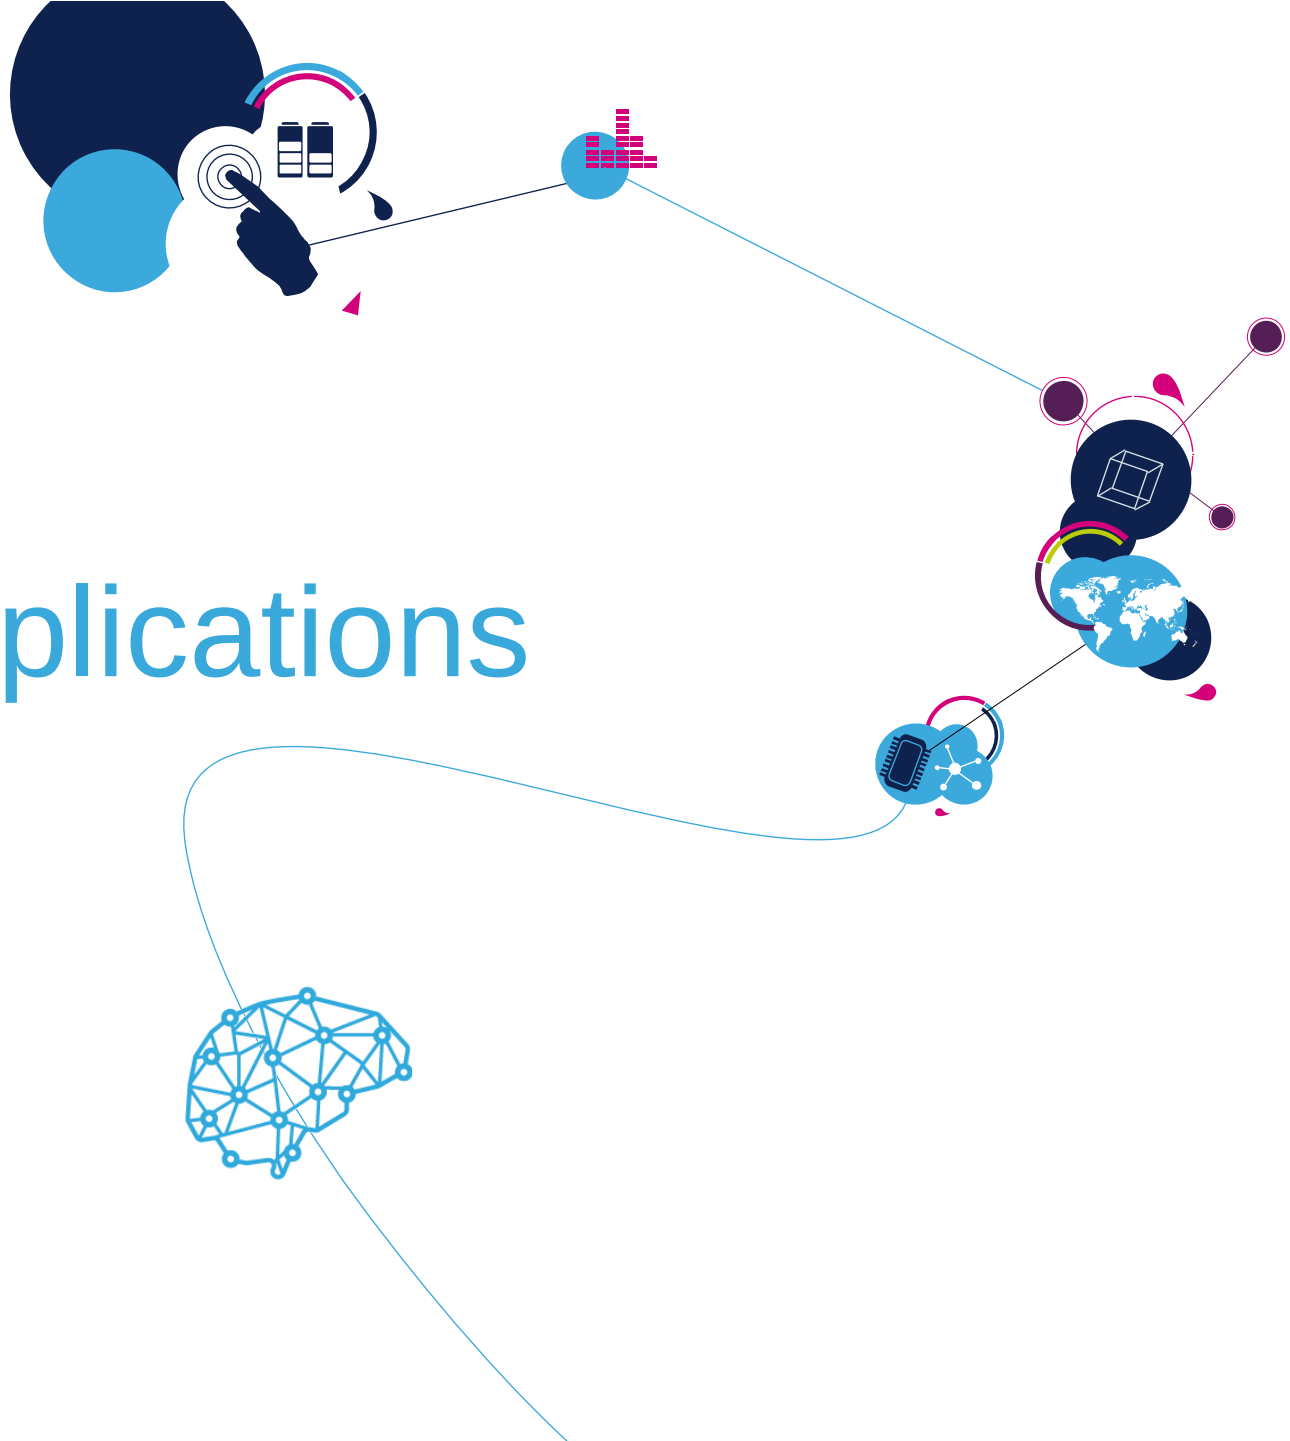

Deep Learning Applications

Danilo Pau

Advanced System Technology

Agrate Brianza

life.augmented

Applications of Neural networks

- **Financial**
  - Stock Market Prediction
  - Credit Worthiness
  - Credit Rating
  - Bankruptcy Prediction
  - Property Appraisal
  - Fraud Detection
  - Price Forecasts
  - Economic Indicator Forecasts

life.augmented

Applications of Neural networks

- **Financial**
  - Stock Market Prediction
  - Credit Worthiness
  - Credit Rating
  - Bankruptcy Prediction
  - Property Appraisal
  - Fraud Detection
  - Price Forecasts
  - Economic Indicator Forecasts
- **Medical**
  - Medical Diagnosis
  - Detection and Evaluation of Medical Phenomena
  - Patient's Length of Stay Forecasts
  - Treatment Cost Estimation

Applications of Neural networks

- **Financial**
  - Stock Market Prediction
  - Credit Worthiness
  - Credit Rating
  - Bankruptcy Prediction
  - Property Appraisal
  - Fraud Detection
  - Price Forecasts
  - Economic Indicator Forecasts
- **Medical**
  - Medical Diagnosis
  - Detection and Evaluation of Medical Phenomena
  - Patient's Length of Stay Forecasts
  - Treatment Cost Estimation
- **Industrial**
  - Process Control
  - Quality Control
  - Temperature and Force Prediction
  - Maintenance prediction
  - Faults detection

Applications of Neural networks

- **Financial**
  - Stock Market Prediction
  - Credit Worthiness
  - Credit Rating
  - Bankruptcy Prediction
  - Property Appraisal
  - Fraud Detection
  - Price Forecasts
  - Economic Indicator Forecasts
- **Medical**
  - Medical Diagnosis
  - Detection and Evaluation of Medical Phenomena
  - Patient's Length of Stay Forecasts
  - Treatment Cost Estimation
- **Industrial**
  - Process Control
  - Quality Control
  - Temperature and Force Prediction
  - Maintenance prediction
  - Faults detection
- **Science**
  - Pattern Recognition
  - Recipes and Chemical Formulation Optimization
  - Chemical Compound Identification
  - Physical System Modeling
  - Ecosystem Evaluation
  - Polymer Identification
  - Recognizing Genes
  - Botanical Classification
  - Signal Processing: Neural Filtering
  - Biological Systems Analysis
  - Ground Level Ozone Prognosis
  - Odor Analysis and Identification

life.augmented

Applications of Neural networks

- **Financial**
  - Stock Market Prediction
  - Credit Worthiness
  - Credit Rating
  - Bankruptcy Prediction
  - Property Appraisal
  - Fraud Detection
  - Price Forecasts
  - Economic Indicator Forecasts
- **Medical**
  - Medical Diagnosis
  - Detection and Evaluation of Medical Phenomena
  - Patient's Length of Stay Forecasts
  - Treatment Cost Estimation
- **Industrial**
  - Process Control
  - Quality Control
  - Temperature and Force Prediction
  - Maintenance prediction
  - Faults detection
- **Science**
  - Pattern Recognition
  - Recipes and Chemical Formulation Optimization
  - Chemical Compound Identification
  - Physical System Modeling
  - Ecosystem Evaluation
  - Polymer Identification
  - Recognizing Genes
  - Botanical Classification
  - Signal Processing: Neural Filtering
  - Biological Systems Analysis
  - Ground Level Ozone Prognosis
  - Odor Analysis and Identification
- **Educational**
  - Teaching Neural Networks
  - Neural Network Research
  - College Application Screening
  - Predict Student Performance
- **Data Mining**
  - Prediction
  - Classification
  - Change and Deviation Detection
  - Knowledge Discovery
  - Response Modeling
  - Time Series Analysis
- **Sales and Marketing**
  - Sales Forecasting
  - Targeted Marketing
  - Service Usage Forecasting
  - Retail Margins Forecasting
- **Operational Analysis**
  - Retail Inventories Optimization
  - Scheduling Optimization
  - Managerial Decision Making
  - Cash Flow Forecasting

life.augmented

Applications of Neural networks

- **Financial**
  - Stock Market Prediction
  - Credit Worthiness
  - Credit Rating
  - Bankruptcy Prediction
  - Property Appraisal
  - Fraud Detection
  - Price Forecasts
  - Economic Indicator Forecasts
- **Medical**
  - Medical Diagnosis
  - Detection and Evaluation of Medical Phenomena
  - Patient's Length of Stay Forecasts
  - Treatment Cost Estimation
- **Industrial**
  - Process Control
  - Quality Control
  - Temperature and Force Prediction
- **Science**
  - Pattern Recognition
  - Recipes and Chemical Formulation Optimization
  - Chemical Compound Identification
  - Physical System Modeling
  - Ecosystem Evaluation
  - Polymer Identification
  - Recognizing Genes
  - Botanical Classification
  - Signal Processing: Neural Filtering
  - Biological Systems Analysis
  - Ground Level Ozone Prognosis
  - Odor Analysis and Identification
- **Educational**
  - Teaching Neural Networks
  - Neural Network Research
  - College Application Screening
  - Predict Student Performance
- **Data Mining**
  - Prediction
  - Classification
  - Change and Deviation Detection
  - Knowledge Discovery
  - Response Modeling
  - Time Series Analysis
- **Sales and Marketing**
  - Sales Forecasting
  - Targeted Marketing
  - Service Usage Forecasting
  - Retail Margins Forecasting
- **Operational Analysis**
  - Retail Inventories Optimization
  - Scheduling Optimization
  - Managerial Decision Making
  - Cash Flow Forecasting
- **HR Management**
  - Employee Selection and Hiring
  - Employee Retention
  - Staff Scheduling
  - Personnel Profiling
- **Energy**
  - Electrical Load Forecasting
  - Energy Demand Forecasting
  - Short and Long-Term Load Estimation
  - Predicting Gas/Coal Index Prices
  - Power Control Systems
  - Hydro Dam Monitoring
- **Other**
  - Sports Betting
  - Making Horse and Dog Racing Picks
  - Quantitative Weather Forecasting
  - Games Development
  - Optimization Problems, Routing
  - Agricultural Production Estimates

Applications of Neural networks

Google ASSISTANT

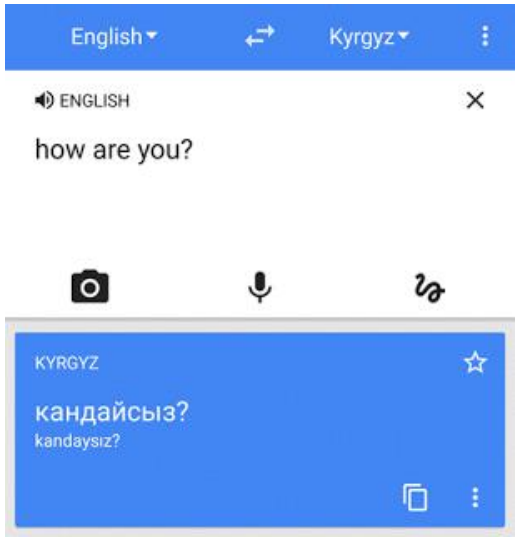

Tap to speak

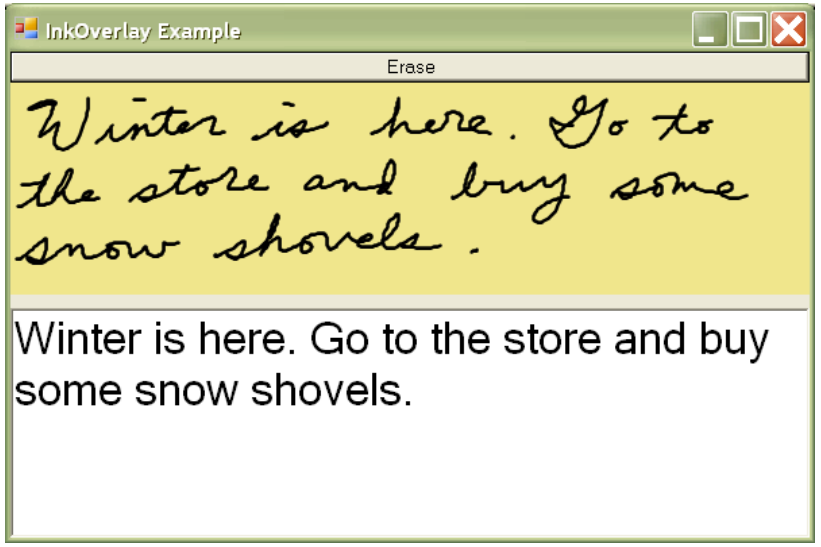

life.augmented

Some others

Artificial Intelligence Offers New Ways of Working!!!

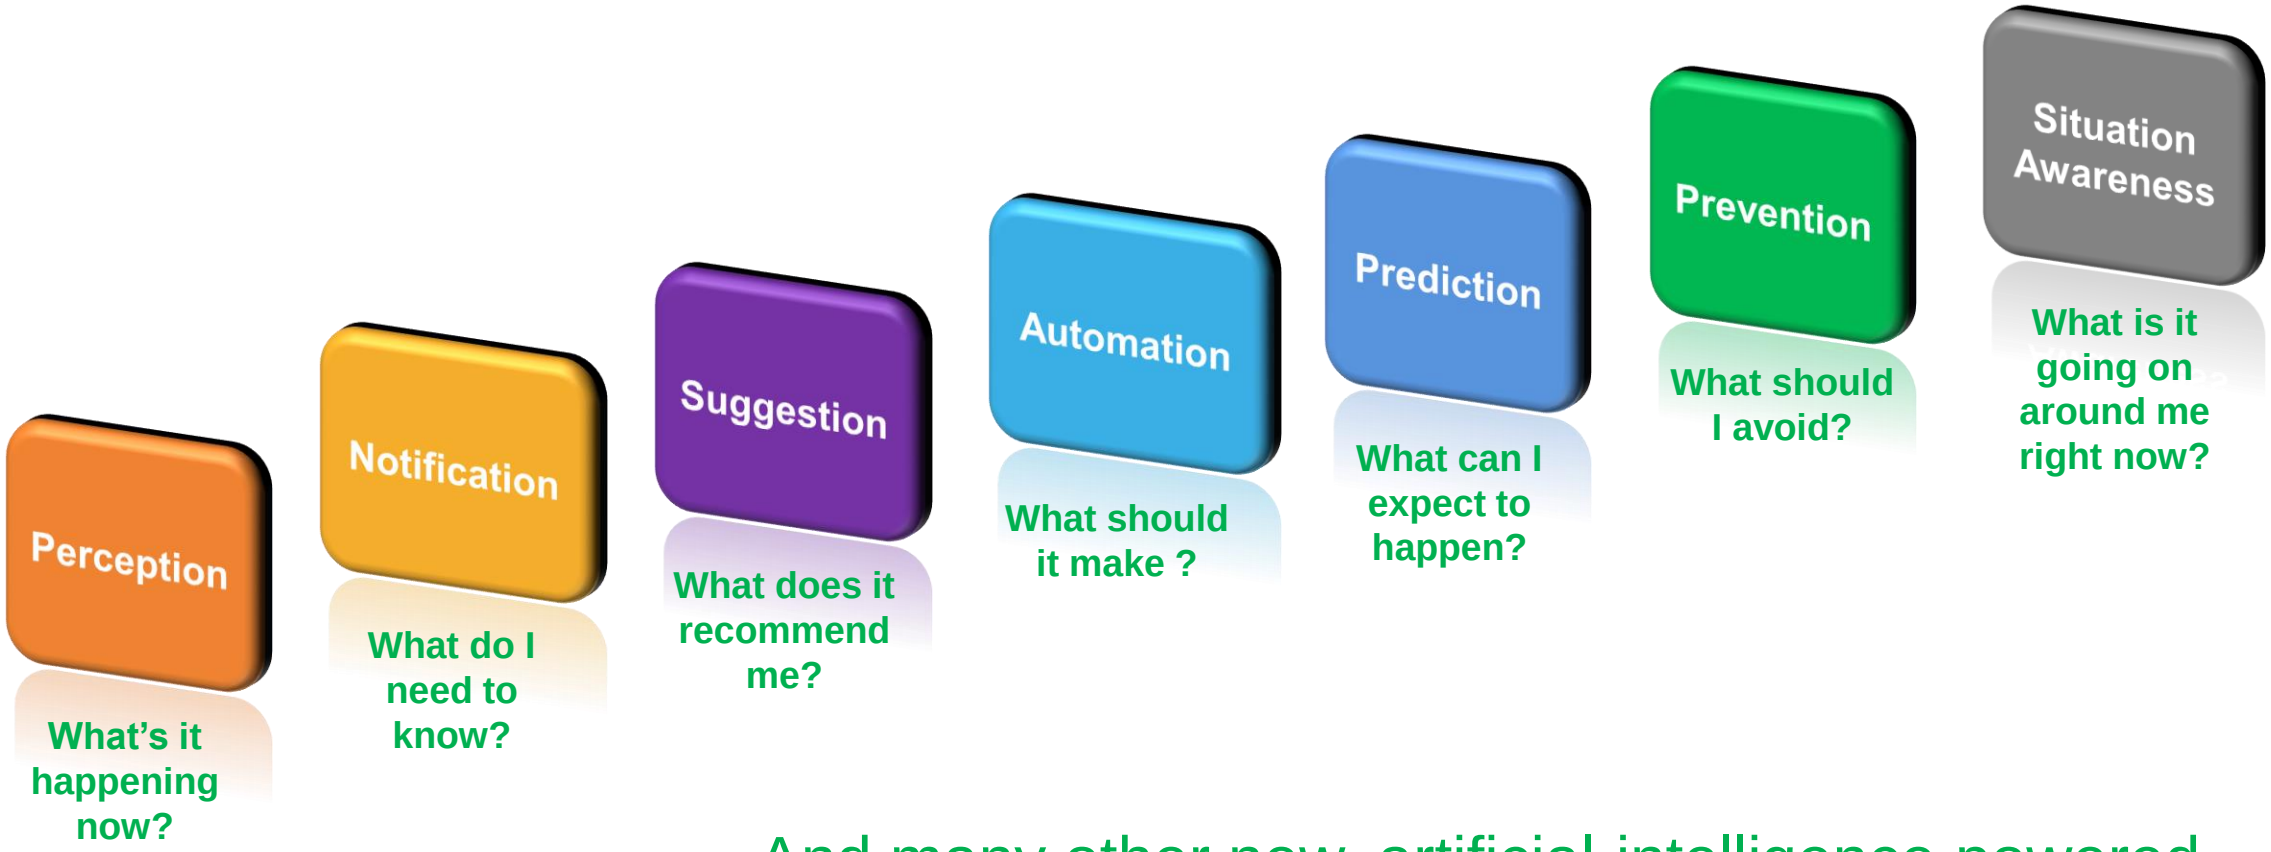

And many other new, artificial-intelligence powered applications…appearing without pace ….

life.augmented

ImageNet

- The ImageNet Large Scale Visual Recognition Challenge (ILSVRC)

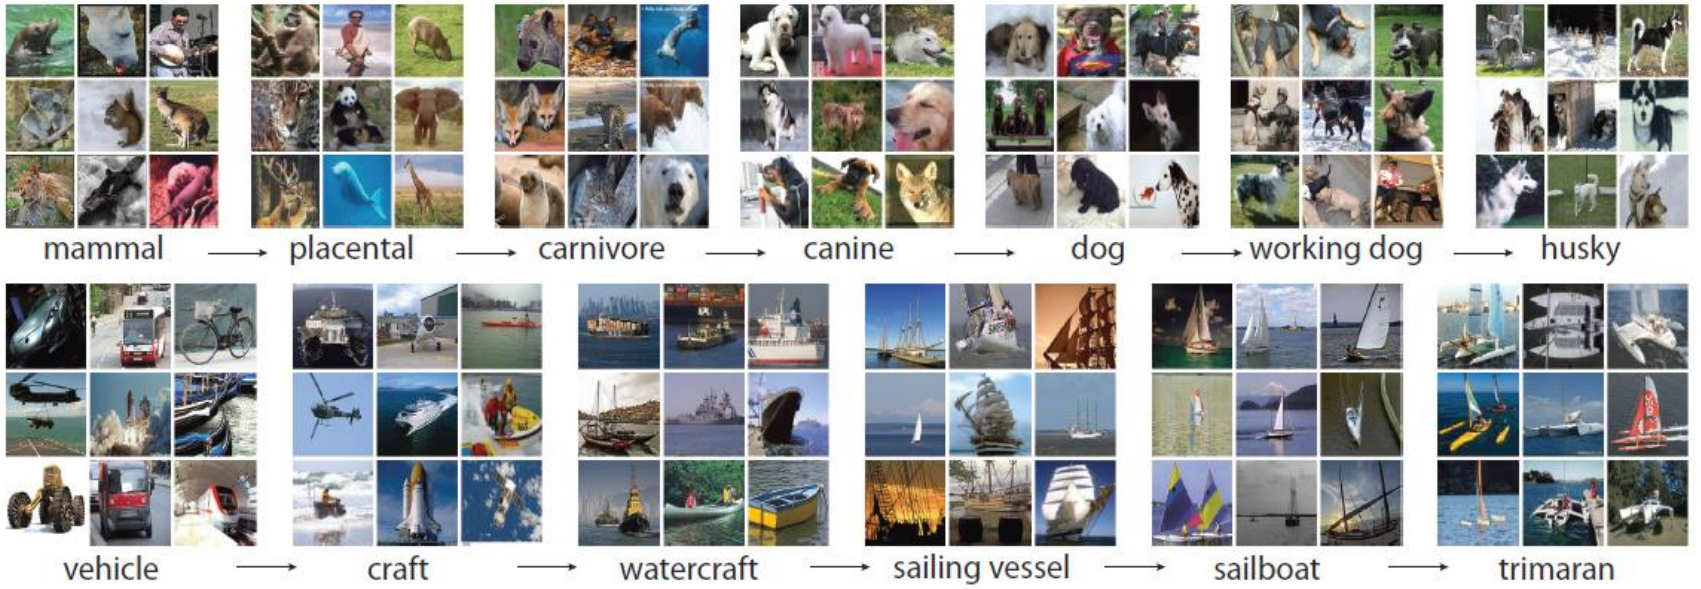

- 1,461,406 full resolution images
- Complex and multiple textual annotation
- Hierarchy of 1000 object classes along several dimensions
- The image classification challenge is run annually since 2010
- Sponsors were: Google, Stanford, Facebook and the University of North Carolina
- Main participants were: Google, Adobe, Microsoft, Samsung, Lenovo, Orange, Toyota, and several major universities.

Some Examples from ImageNet

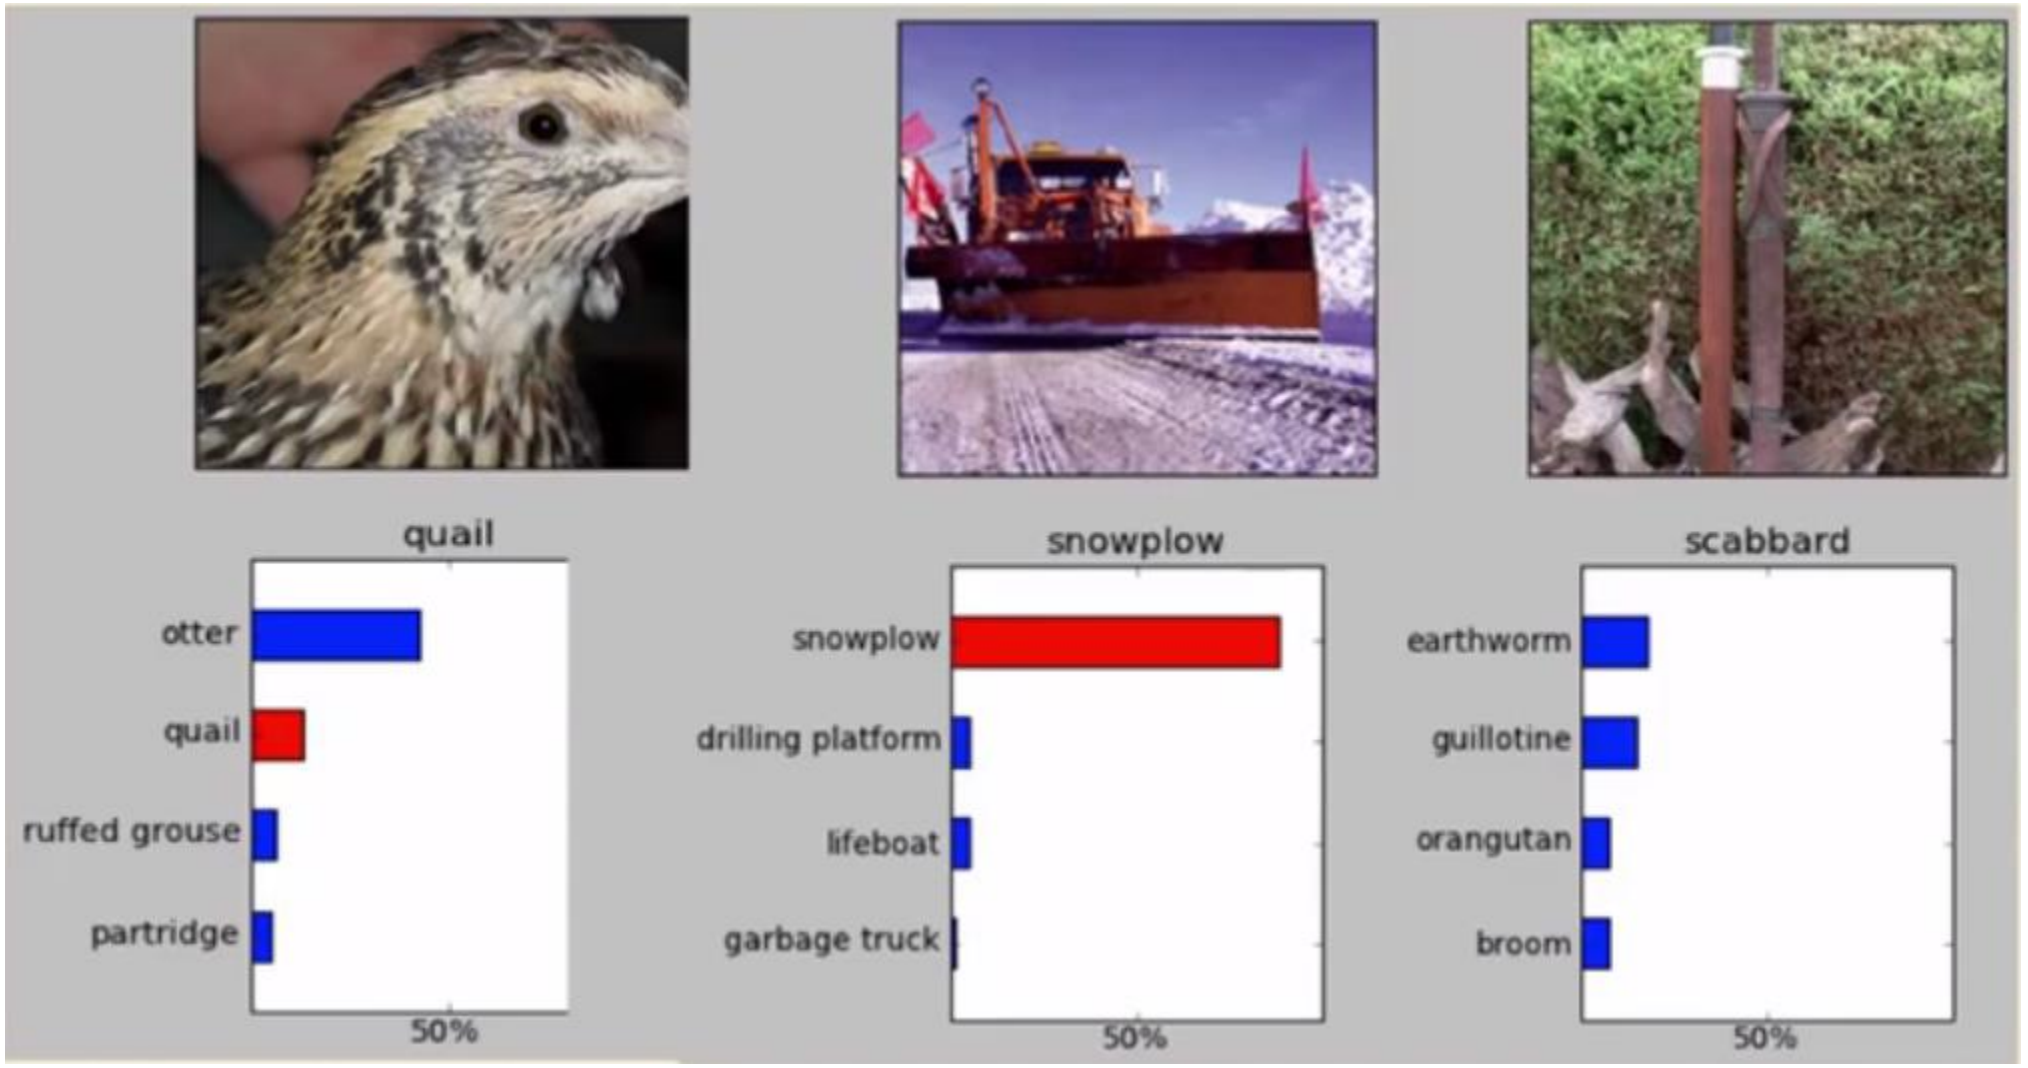

Some Examples from ImageNet (2)

Some Examples from ImageNet (3)

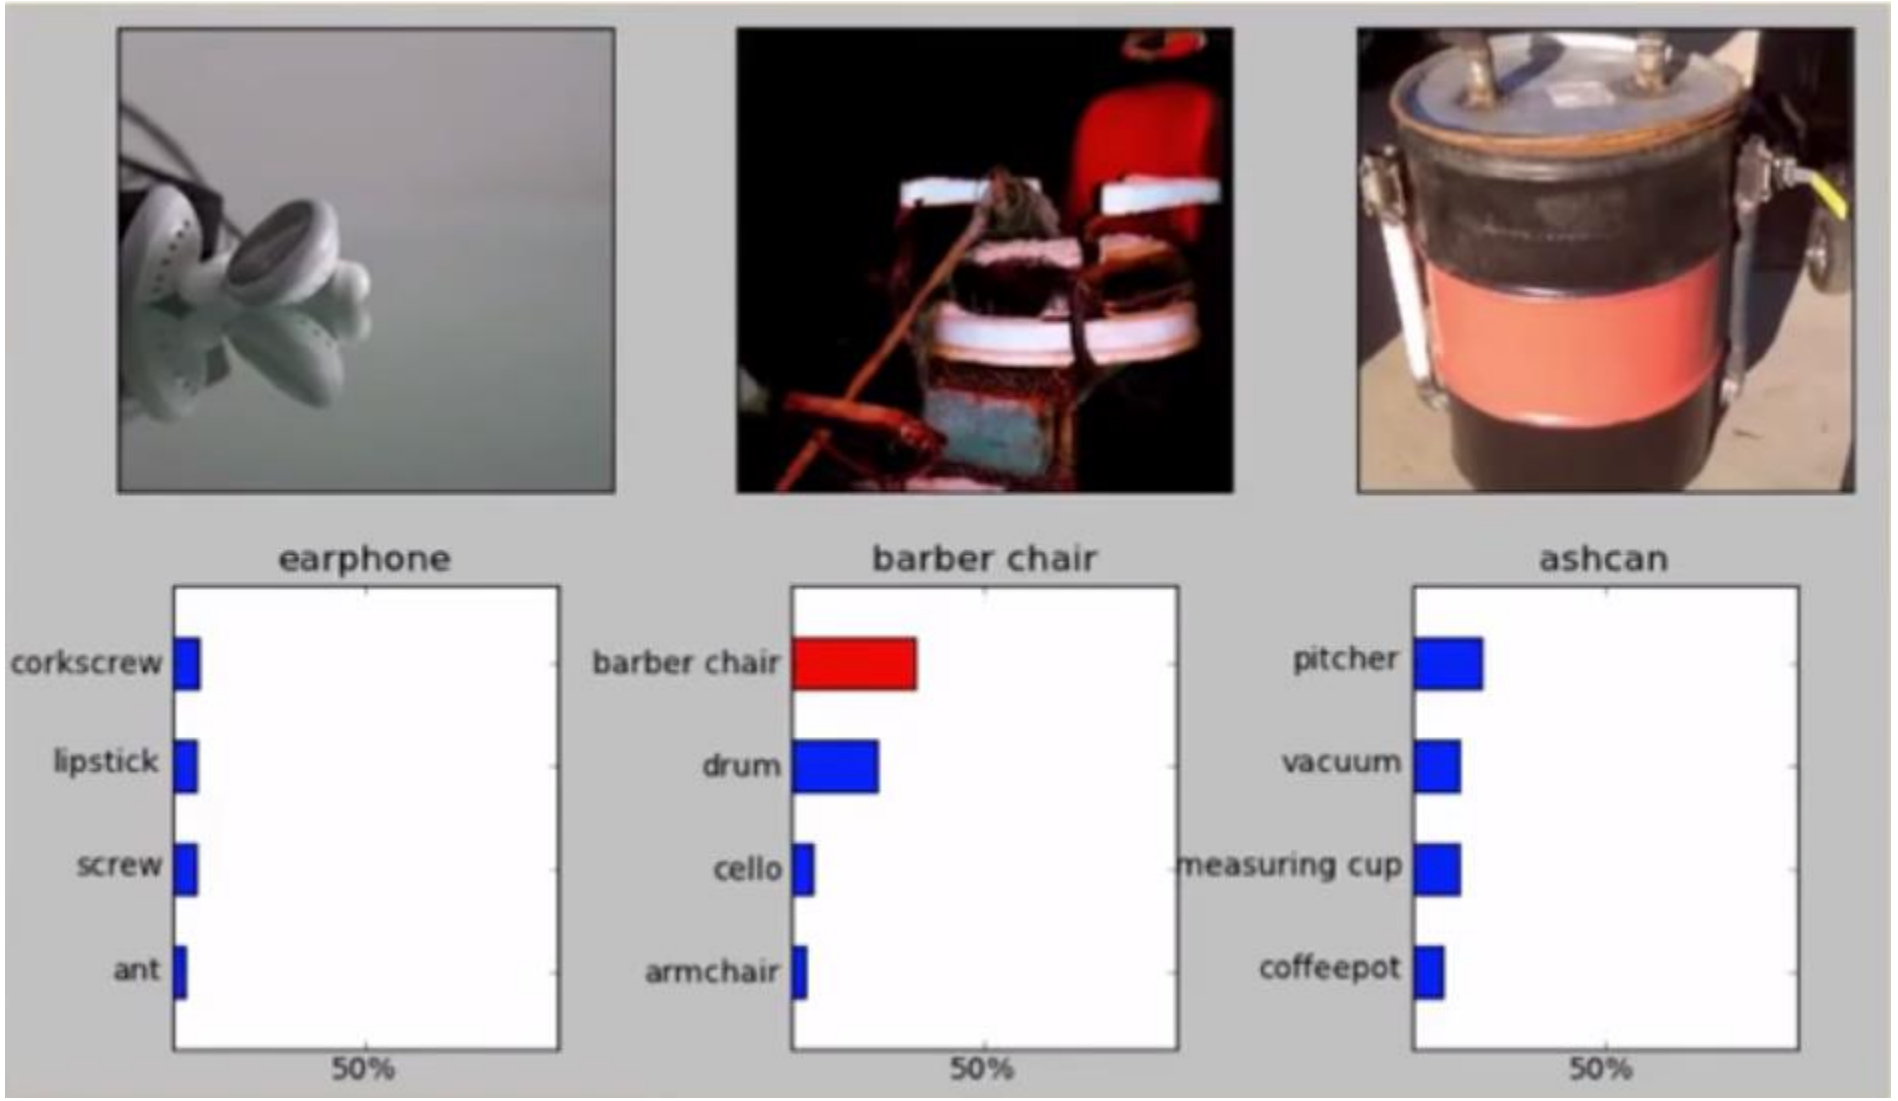

Image Classification using AlexNet

Top 1 = 1.0000, revolver, six-gun
Top 2 = 0.0000, rifle
Top 3 = 0.0000, assault rifle, assault gun

Top 1 = 1.0000, brambling, Fringilla montifringilla
Top 2 = 0.0000, ruddy turnstone
Top 3 = 0.0000, junco, snowbird

Top 1 = 1.0000, balloon
Top 2 = 0.0000, airship, dirigible
Top 3 = 0.0000, water tower

Top 1 = 0.9393, bonnet, poke bonnet
Top 2 = 0.0394, wool, woolen, woollen
Top 3 = 0.0075, mitten

Top 1 = 0.9850, horse cart, horse-cart
Top 2 = 0.0075, jinrikisha, ricksha, rickshaw
Top 3 = 0.0046, oxcart

Top 1 = 0.8428, French loaf
Top 2 = 0.0596, hotdog, hot dog, red hot
Top 3 = 0.0132, dough

Top 1 = 0.9512, pineapple, ananas
Top 2 = 0.0129, brain coral
Top 3 = 0.0099, sea cucumber, holothurian

Top 1 = 0.9857, panpipe, pandean pipe, syrinx
Top 2 = 0.0060, perfume, essence
Top 3 = 0.0019, cuirass

Transfer Learning

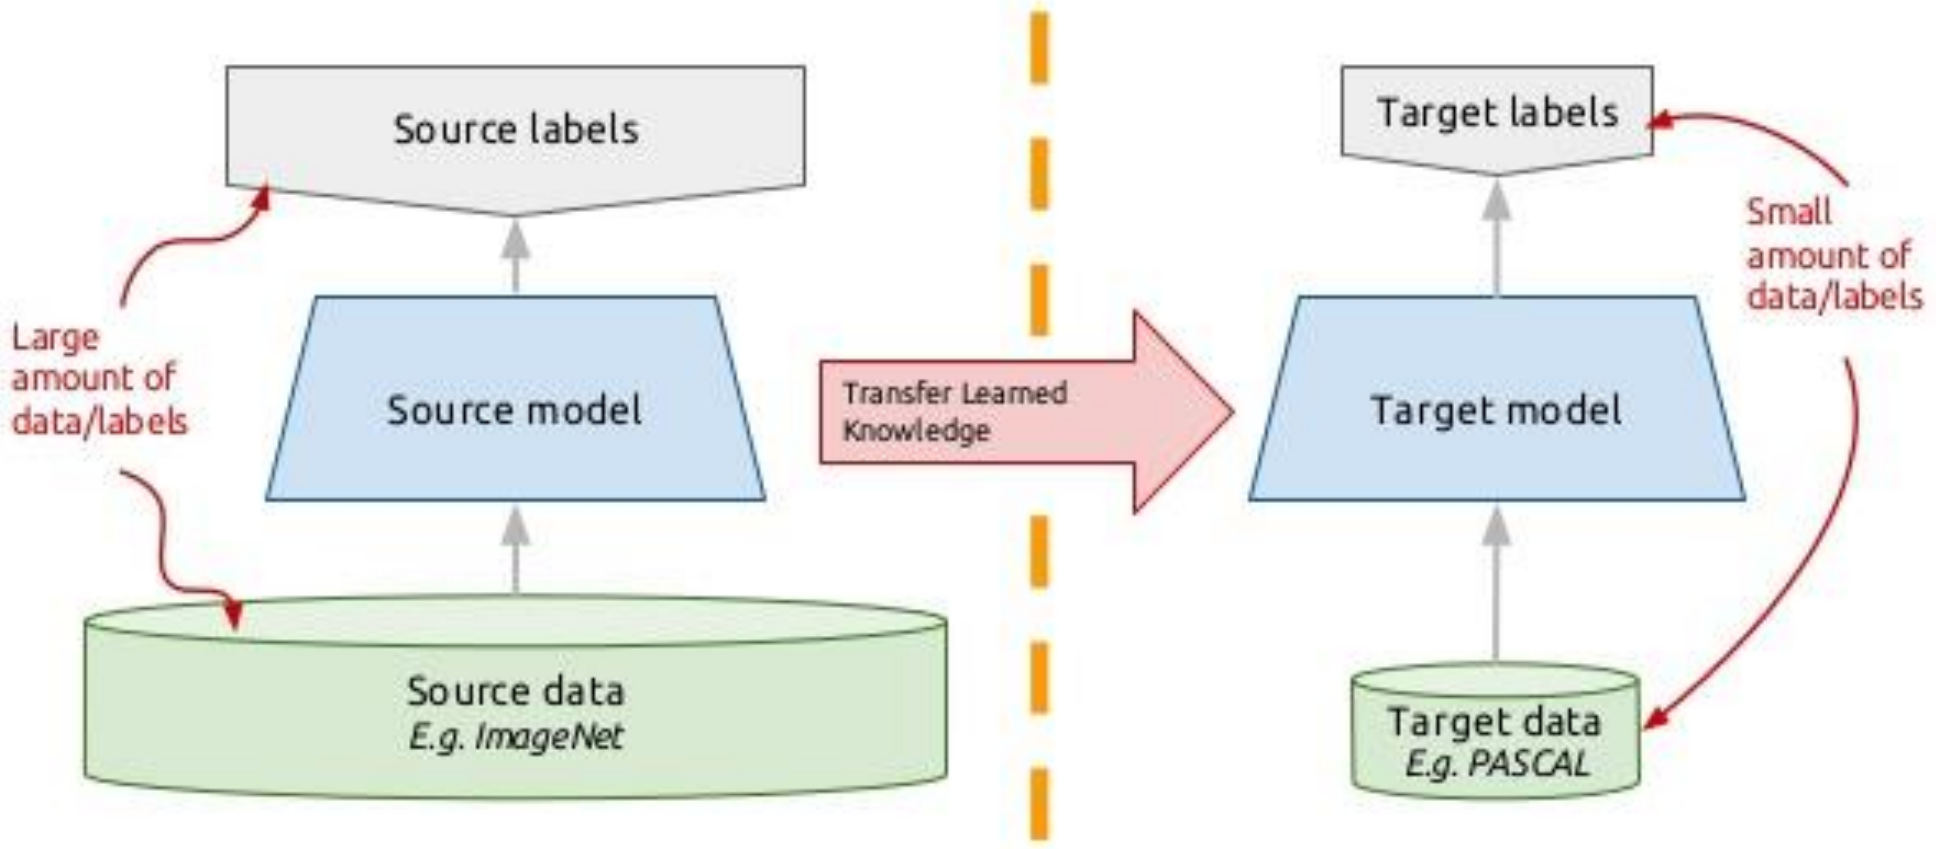

Transfer Learning

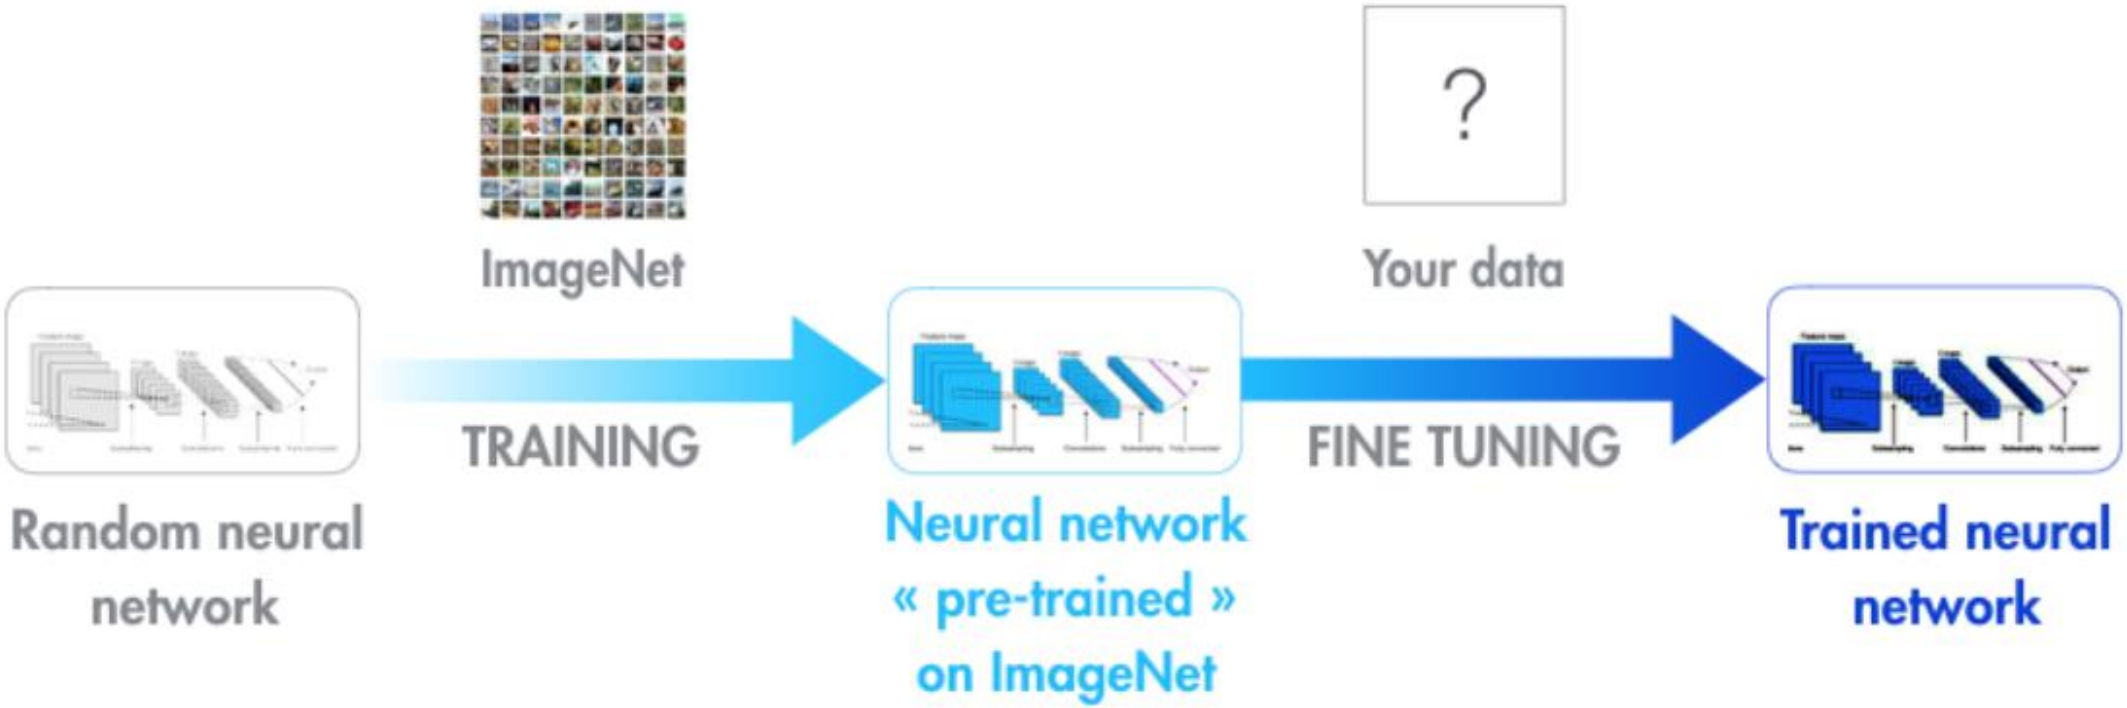

Solving Object Recognition

IMAGENET **Large Scale Visual Recognition Challenge (ILSVRC)**

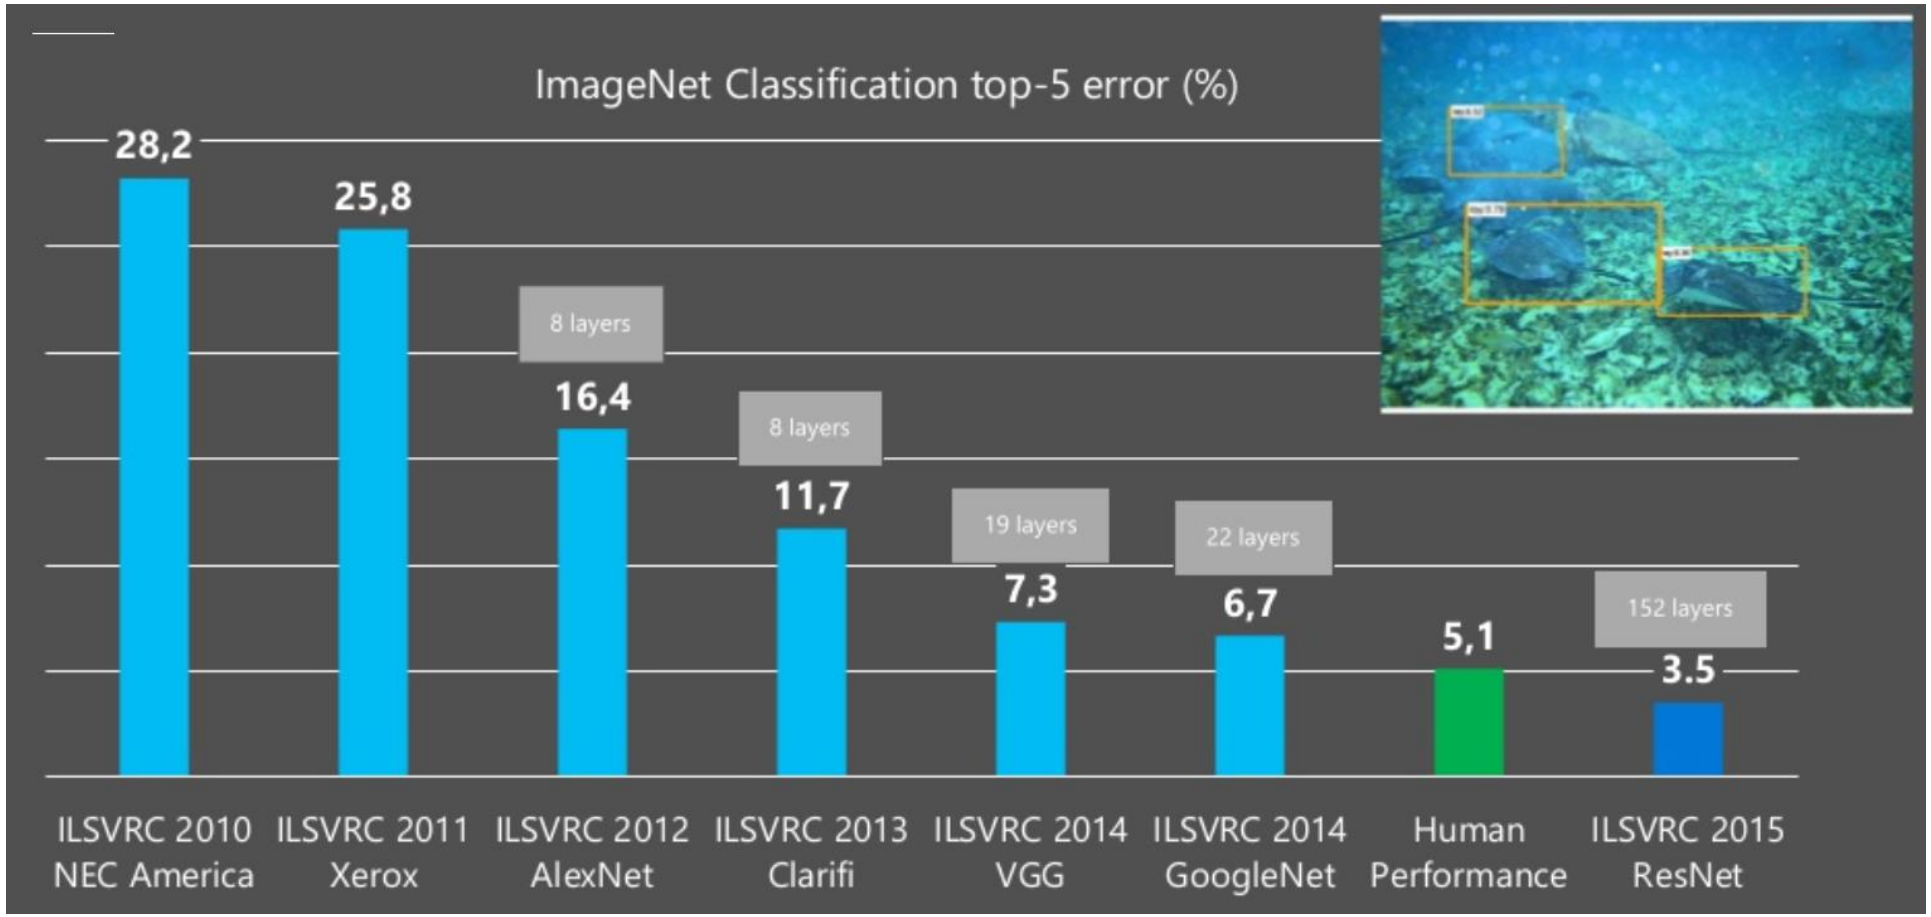

life.augmented

Comparing Different DCNNs

- Comparative charts at Top-1 accuracy
  - *i.e. how often the DCNN is right with ImageNet with its top prediction*

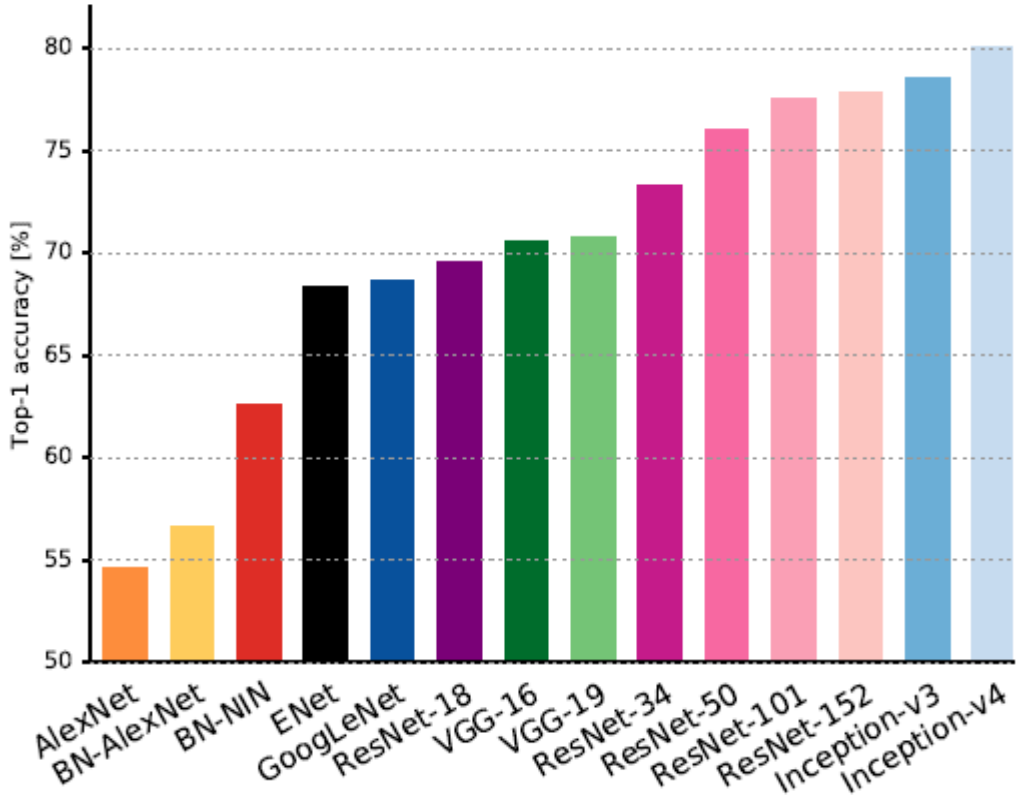

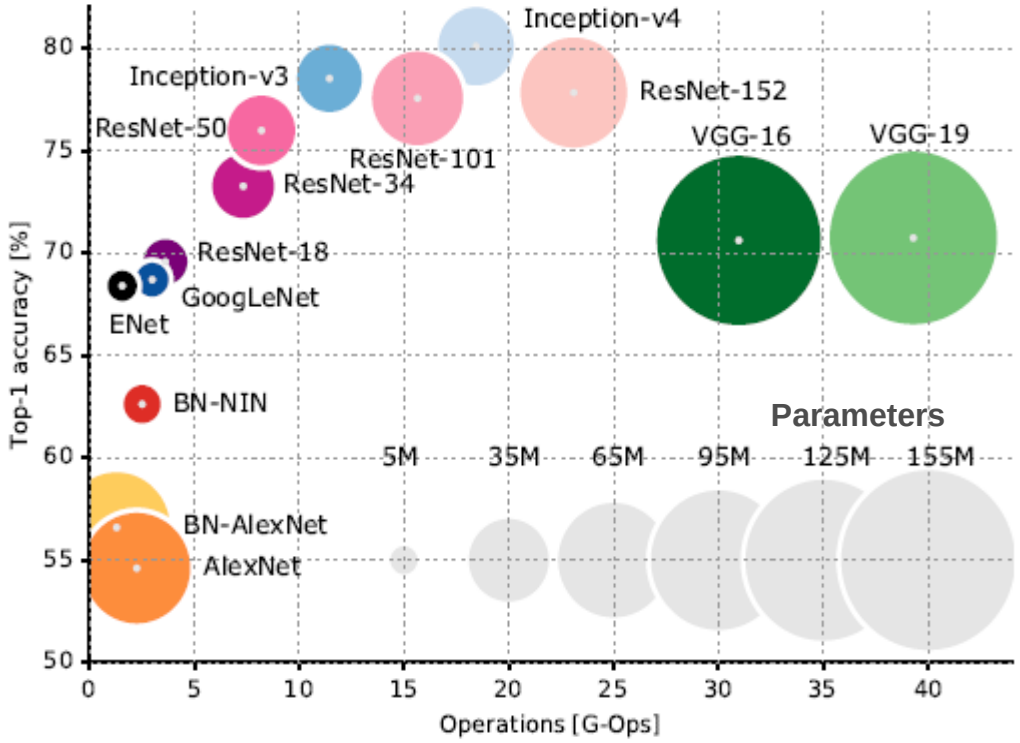

Image from [https://arxiv.org/abs/1605.07678, 2017]

Beyond Image Classification

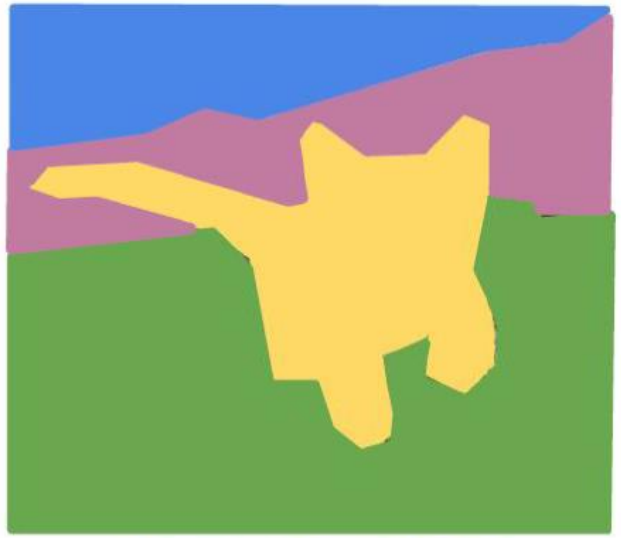

GRASS, CAT, TREE, SKY

No objects, just pixels

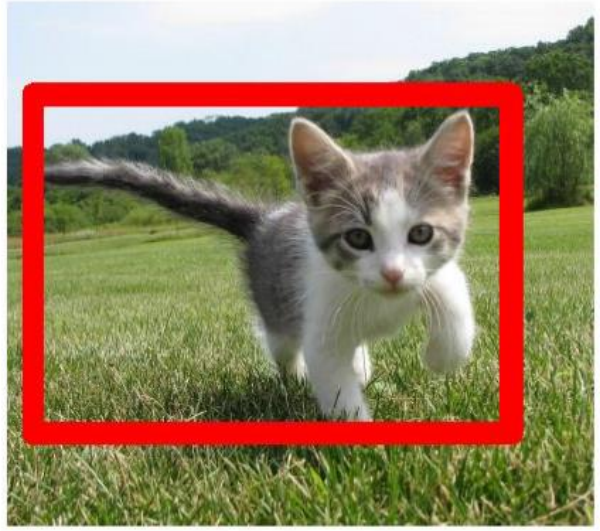

CAT

Single Object

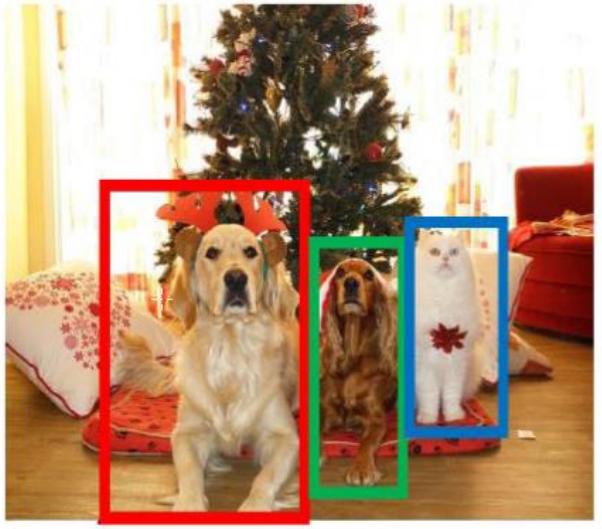

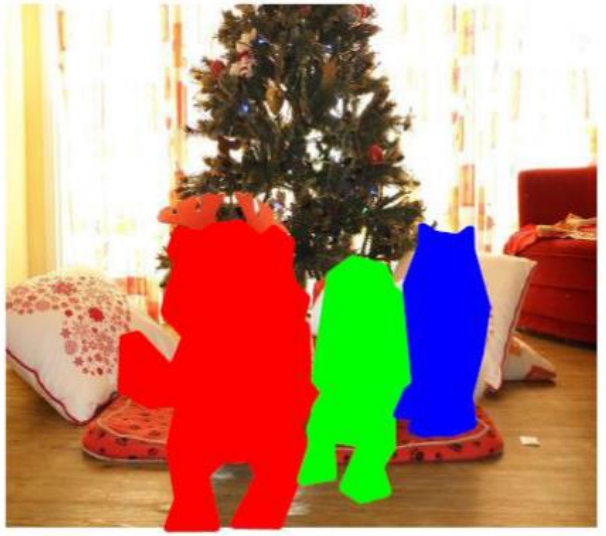

DOG, DOG, CAT

DOG, DOG, CAT

Multiple Object

Object detection and positioning

- Generate boxes and classifications

- **Two-stage Process**
  - Generate bounding box candidates
  - Pass each candidate through a DCNN
  - Select those candidates that are classified with higher certainty

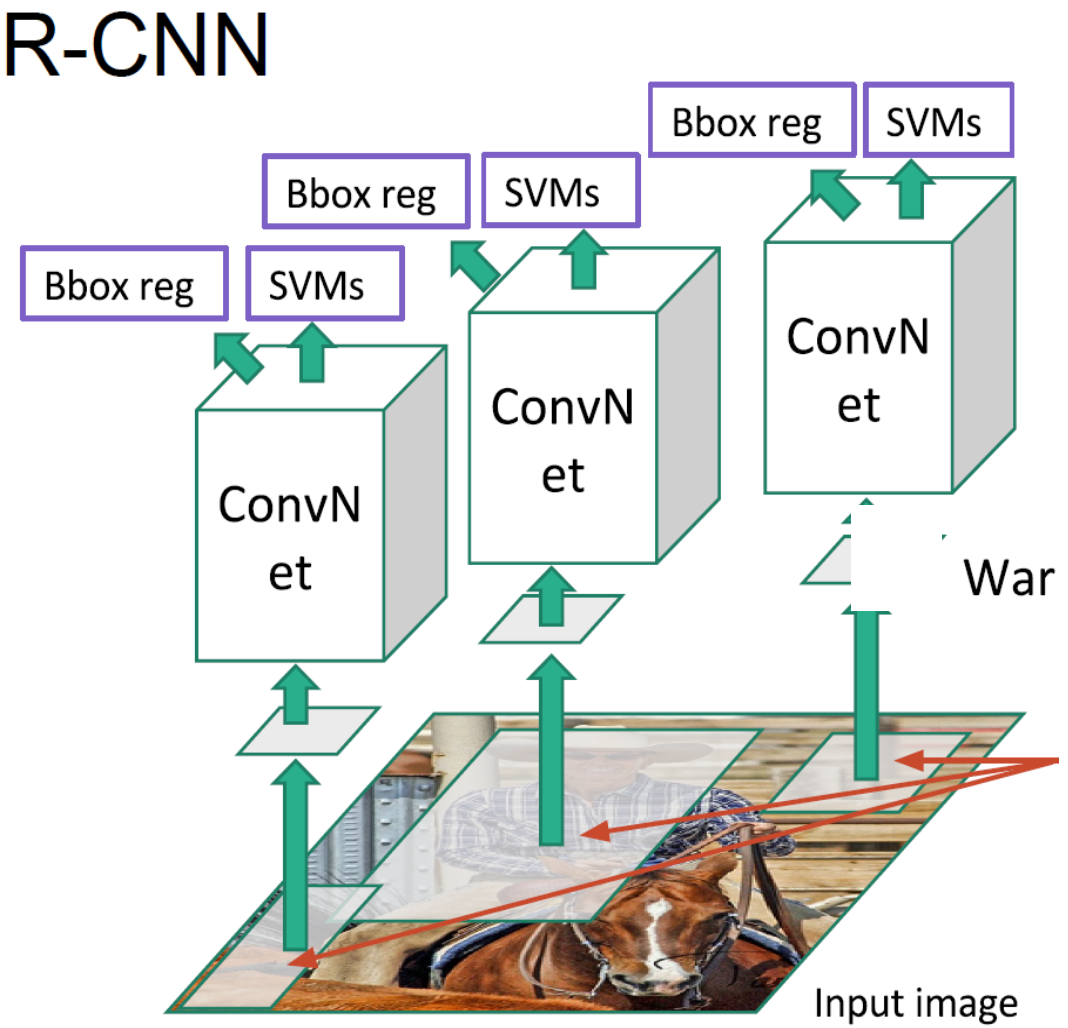

[images from http://cs231n.stanford.edu/slides/2017/cs231n_2017_lecture11.pdf]

Image Detection using Yolo

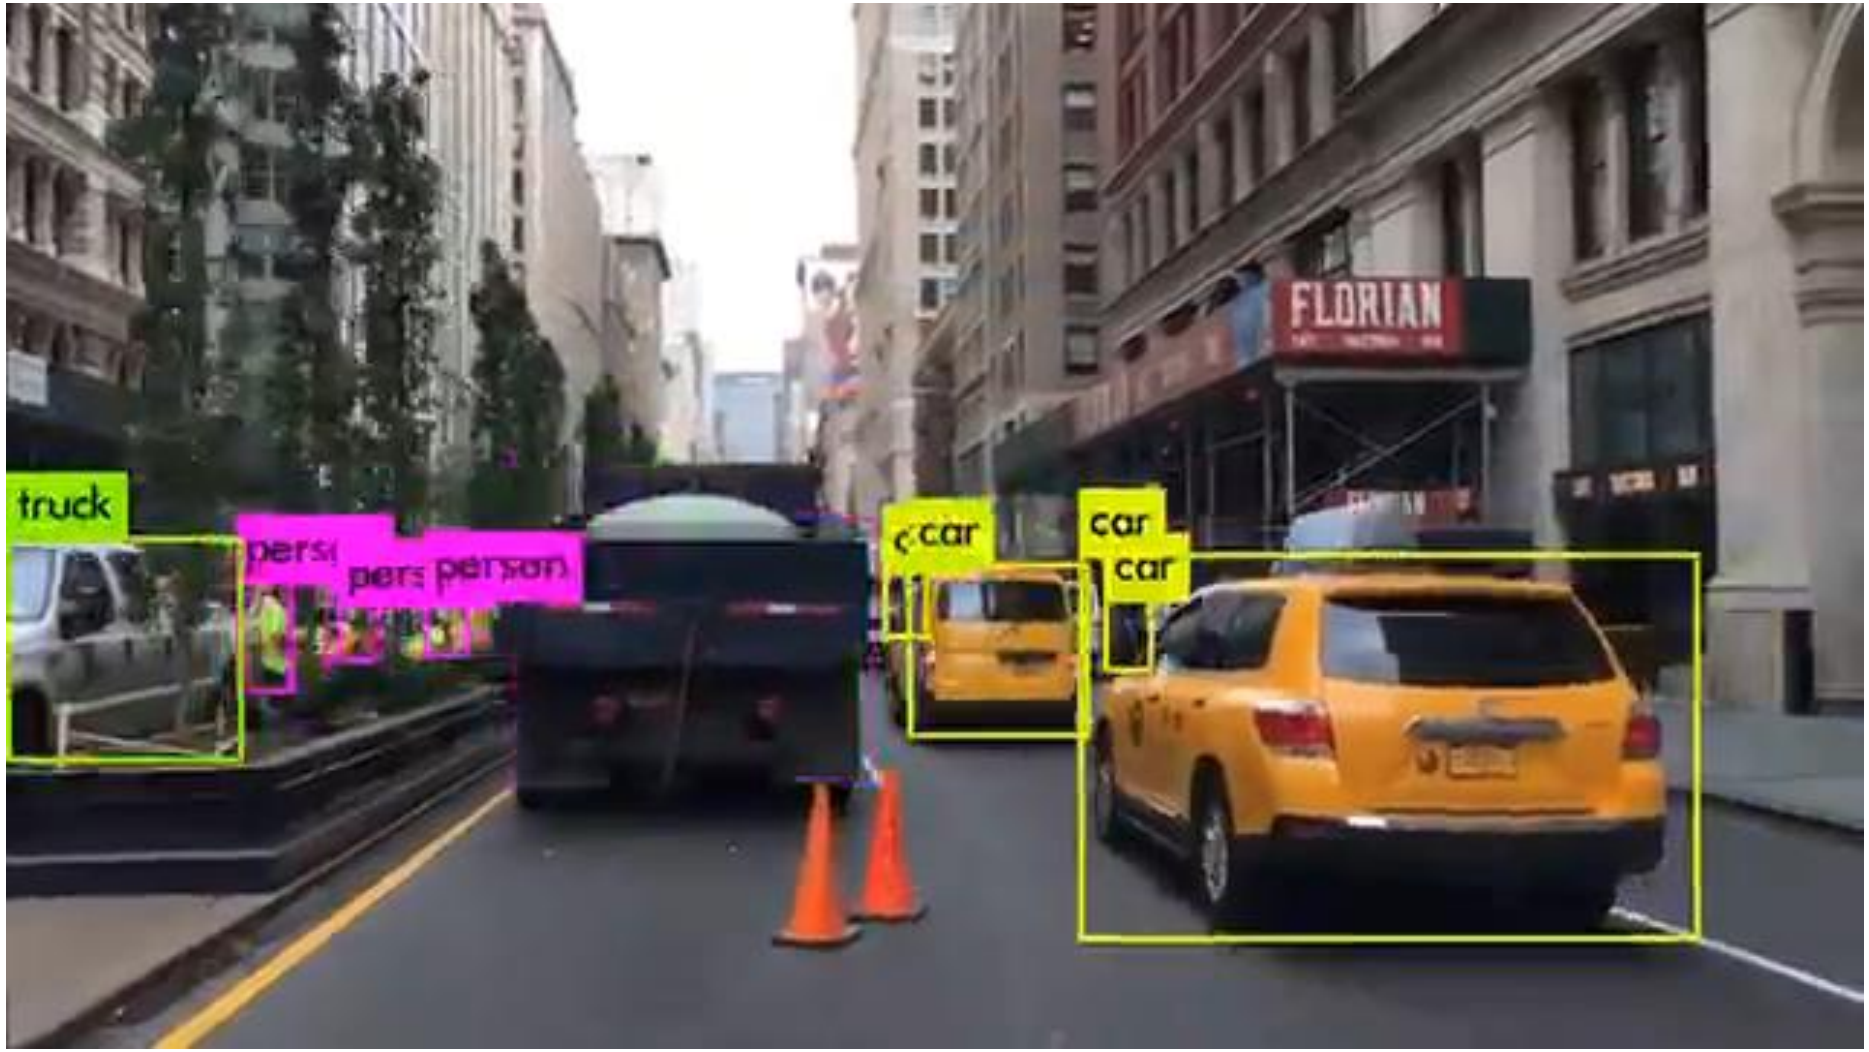

https://www.youtube.com/watch?v=YmbhRxQkLMg

Semantic Segmentation

- Deep Convolutional Neural Networks
- First downsample, then upsample

Input:
3 x H x W

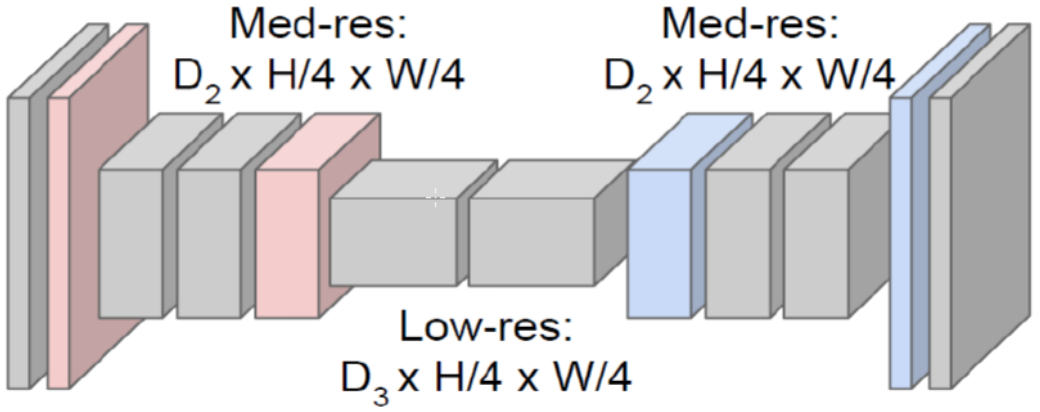

High-res:
D_1 x H/2 x W/2

High-res:
D_1 x H/2 x W/2

Predictions:
H x W

Object per Pixel Labelling

https://www.youtube.com/watch?v=qWl9idsCuLQ

The Power of Abstraction

- Different Layers of a Deep Convolutional Neural Network store many kind of information

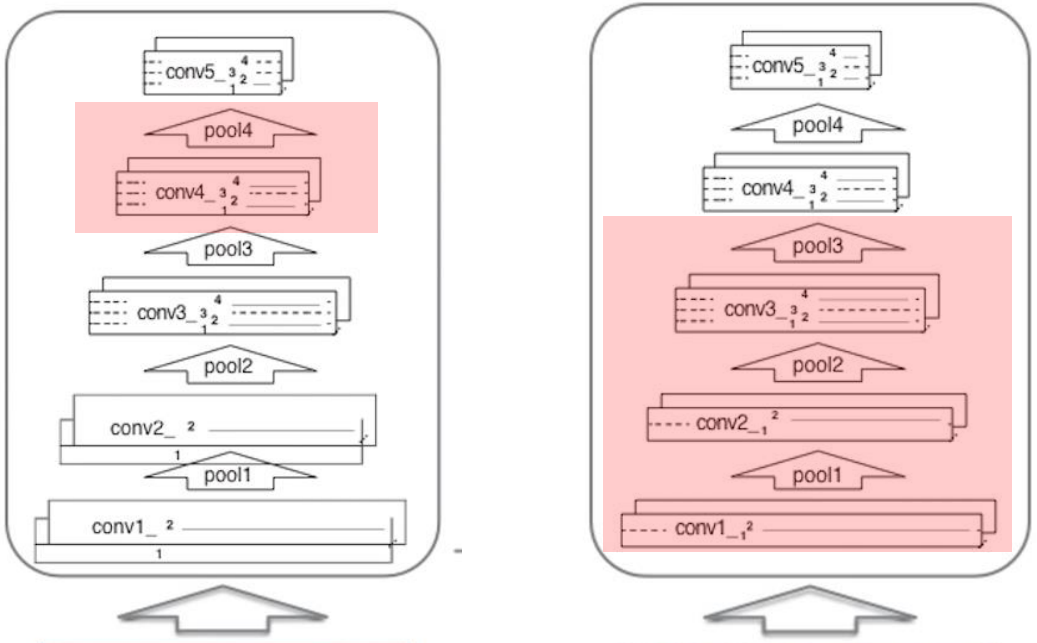

'Content' image

'Style' image

Create a new image by combining more of the 'Content' top layer and more of 'Style' low layers

Result

life.augmented

The Power of Abstraction

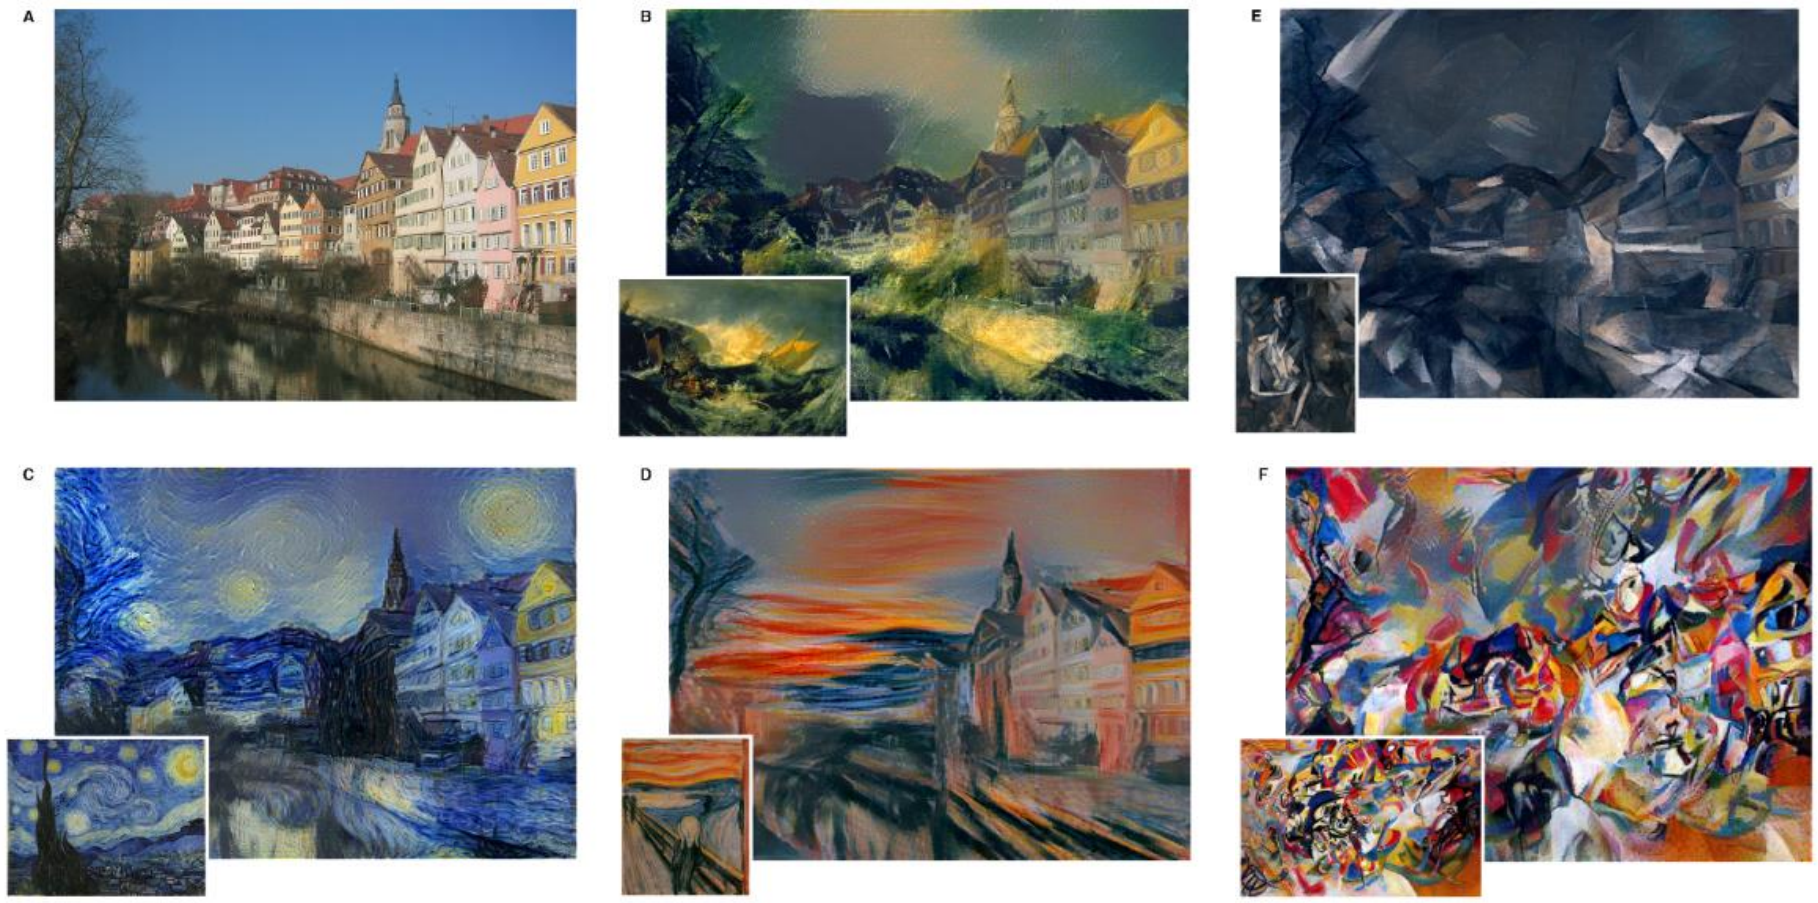

A tool for Artists

- Enhancing lower layers whatever they detected.
- For example, lower layers tend to produce strokes or simple ornament-like patterns, because sensitive to basic features such as edges and their orientations

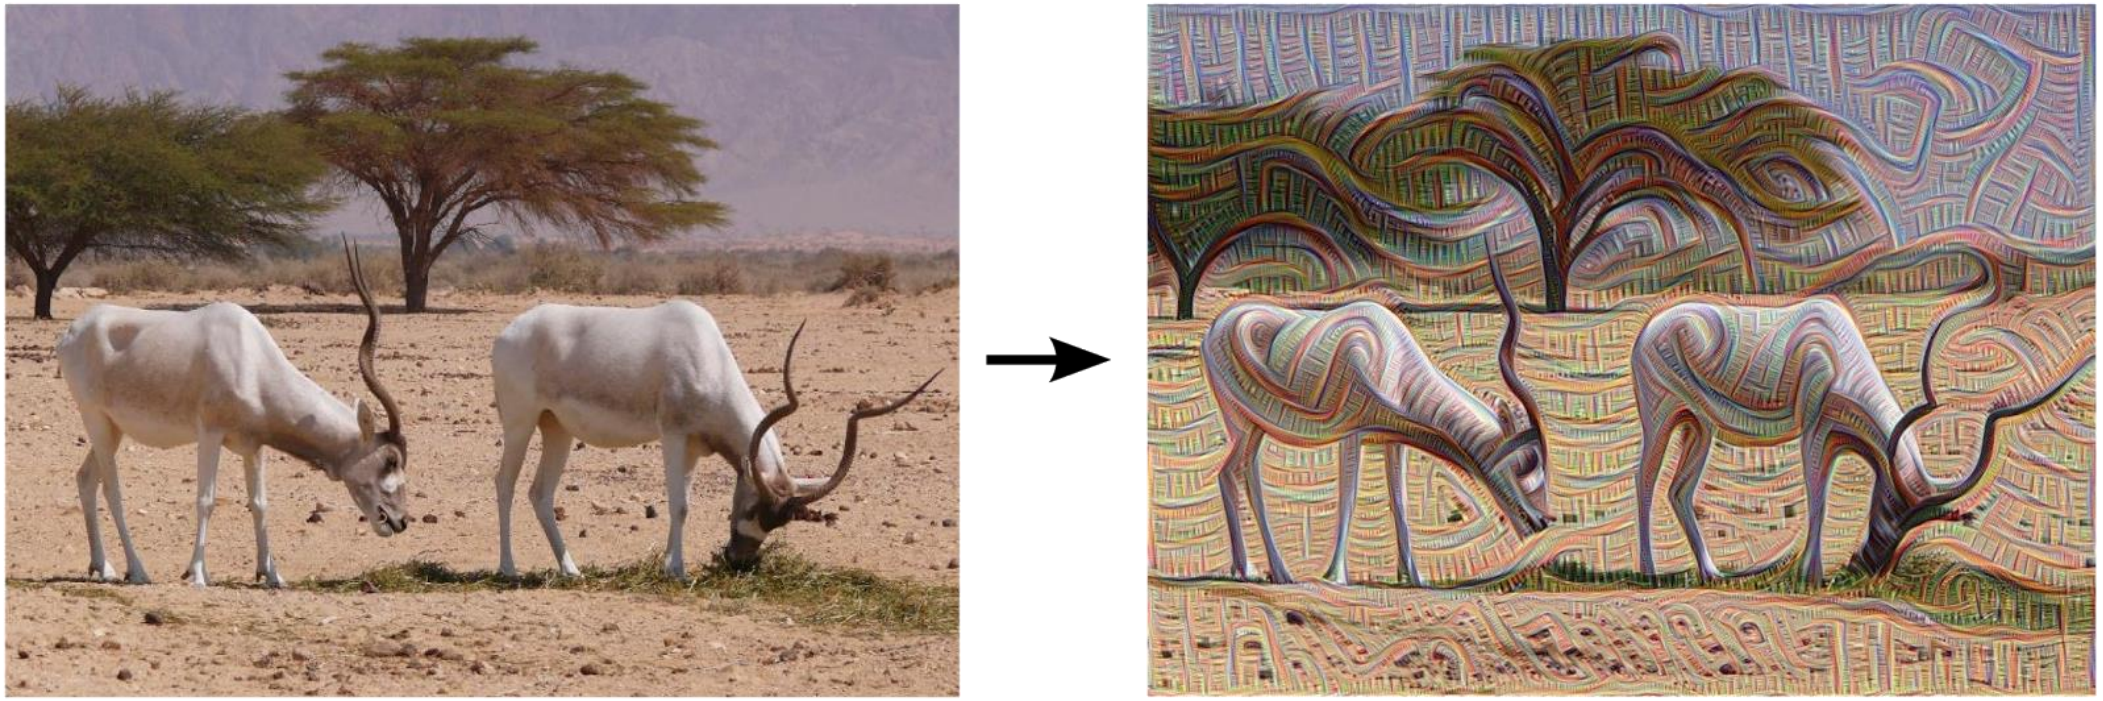

[images from https://research.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html]

A tool for Artists

- Enhancing lower layers whatever they detected.

[images from https://research.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html]

A tool for Artists

- Enhancing upper layers which identify more sophisticated features
- If a cloud looks a little bit like a bird, the network will make it look more like a bird. In turn will make the network recognize the bird even more strongly on the next pass and so forth, until a highly detailed bird appears, seemingly out of nowhere.

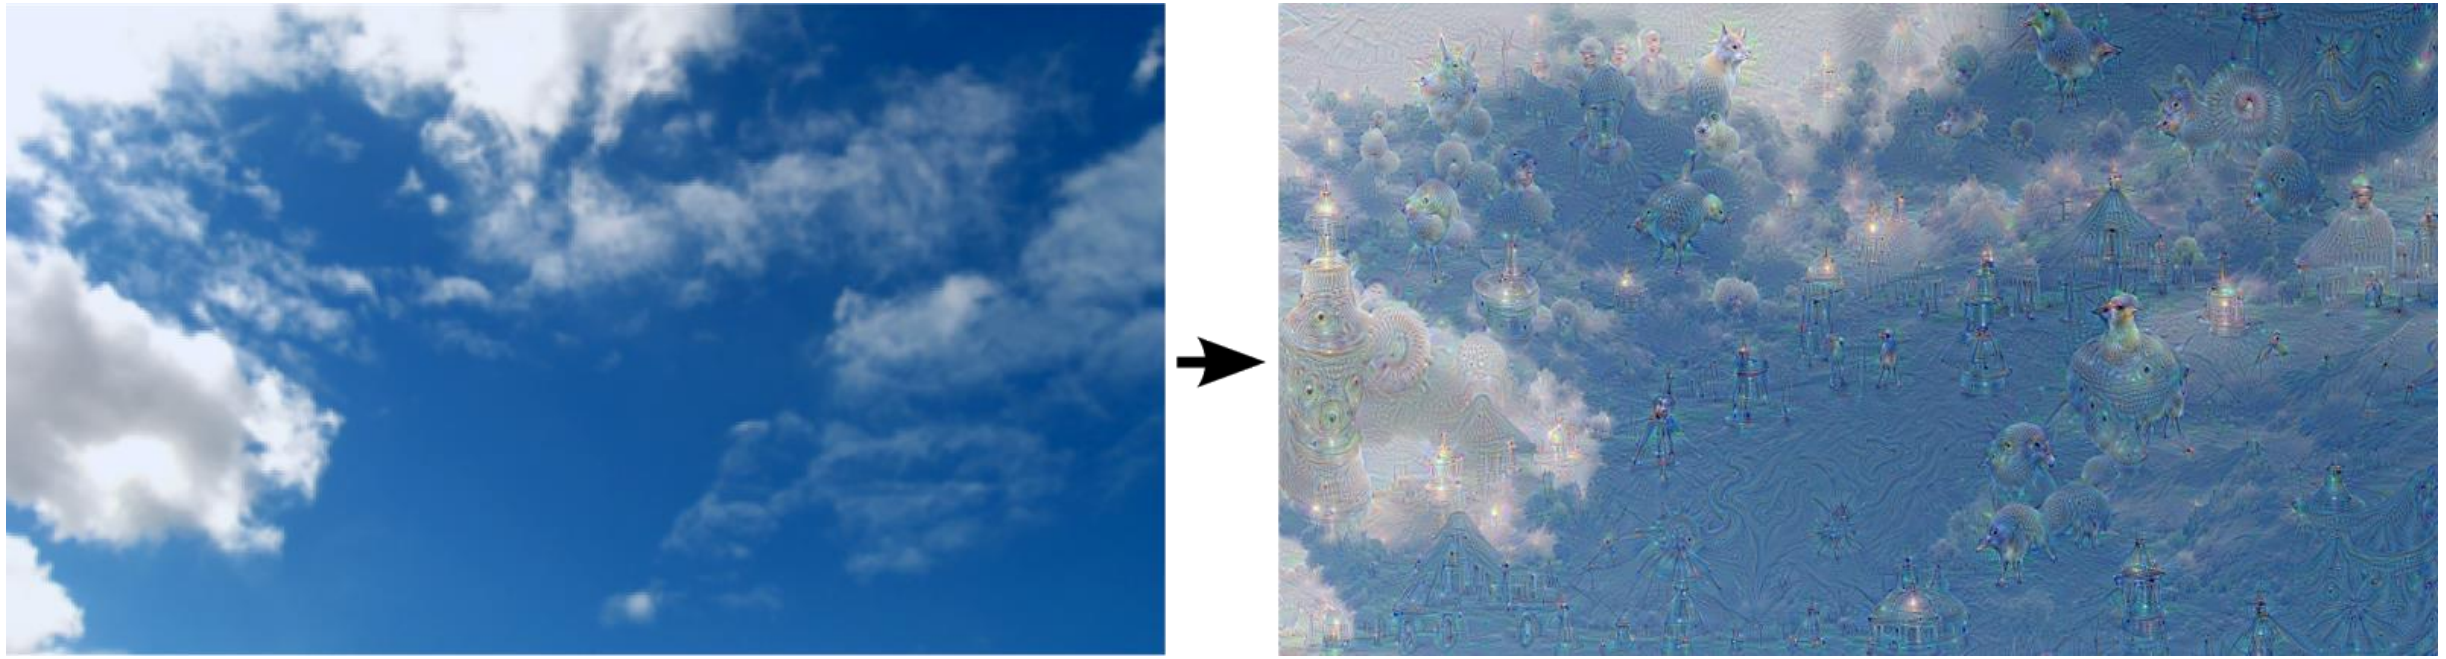

[images from https://research.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html]

A tool for Artists

- Let the DCNN to go on its own

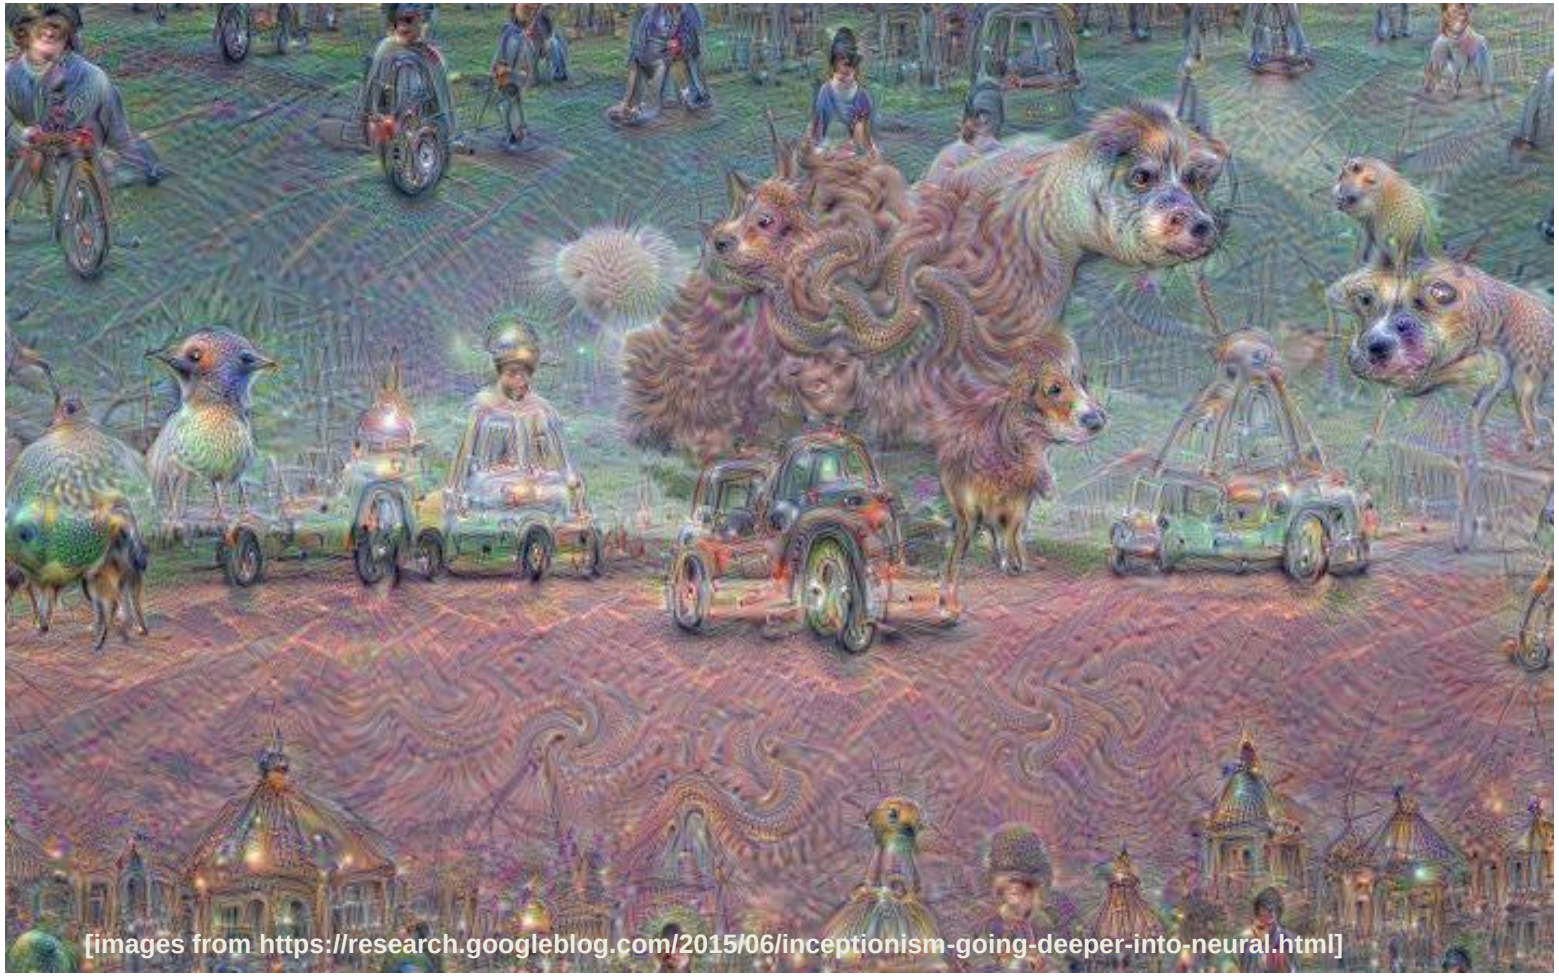

[images from https://research.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html]

A tool for Artists

- E.g. apply the algorithm iteratively on its own outputs and apply some zooming after each iteration, we get an endless stream of new impressions, exploring the set of things the network knows about.

[images from https://research.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html]

Reproducing Non-Existing Masterpieces

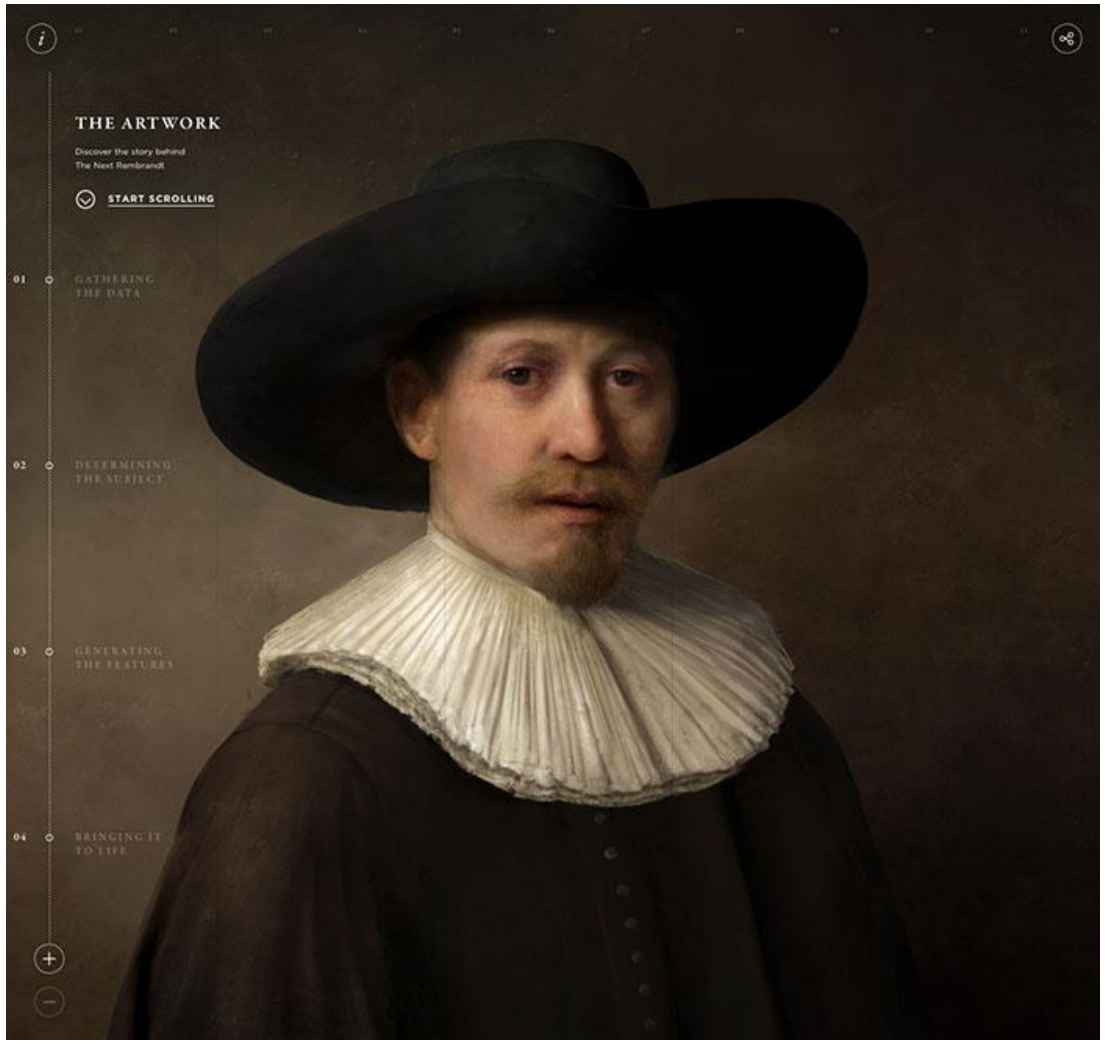

[figures from the website]

- *Next Rembrandt project* [https://www.nextrembrandt.com/, 2016].

Translating music across musical instruments

https://export.arxiv.org/pdf/1805.07848
https://youtu.be/vdxCqNWTpUs

Image Captioning

NeuralTalk Sentence Generation Results

Showing results for coco on 1000 images

https://cs.stanford.edu/people/karpathy/deepimagesent/generationdemo/

a group of people sitting at a table with wine glasses
logprob: -6.71

a baseball player is swinging a bat at a ball
logprob: -6.79

Image Captioning

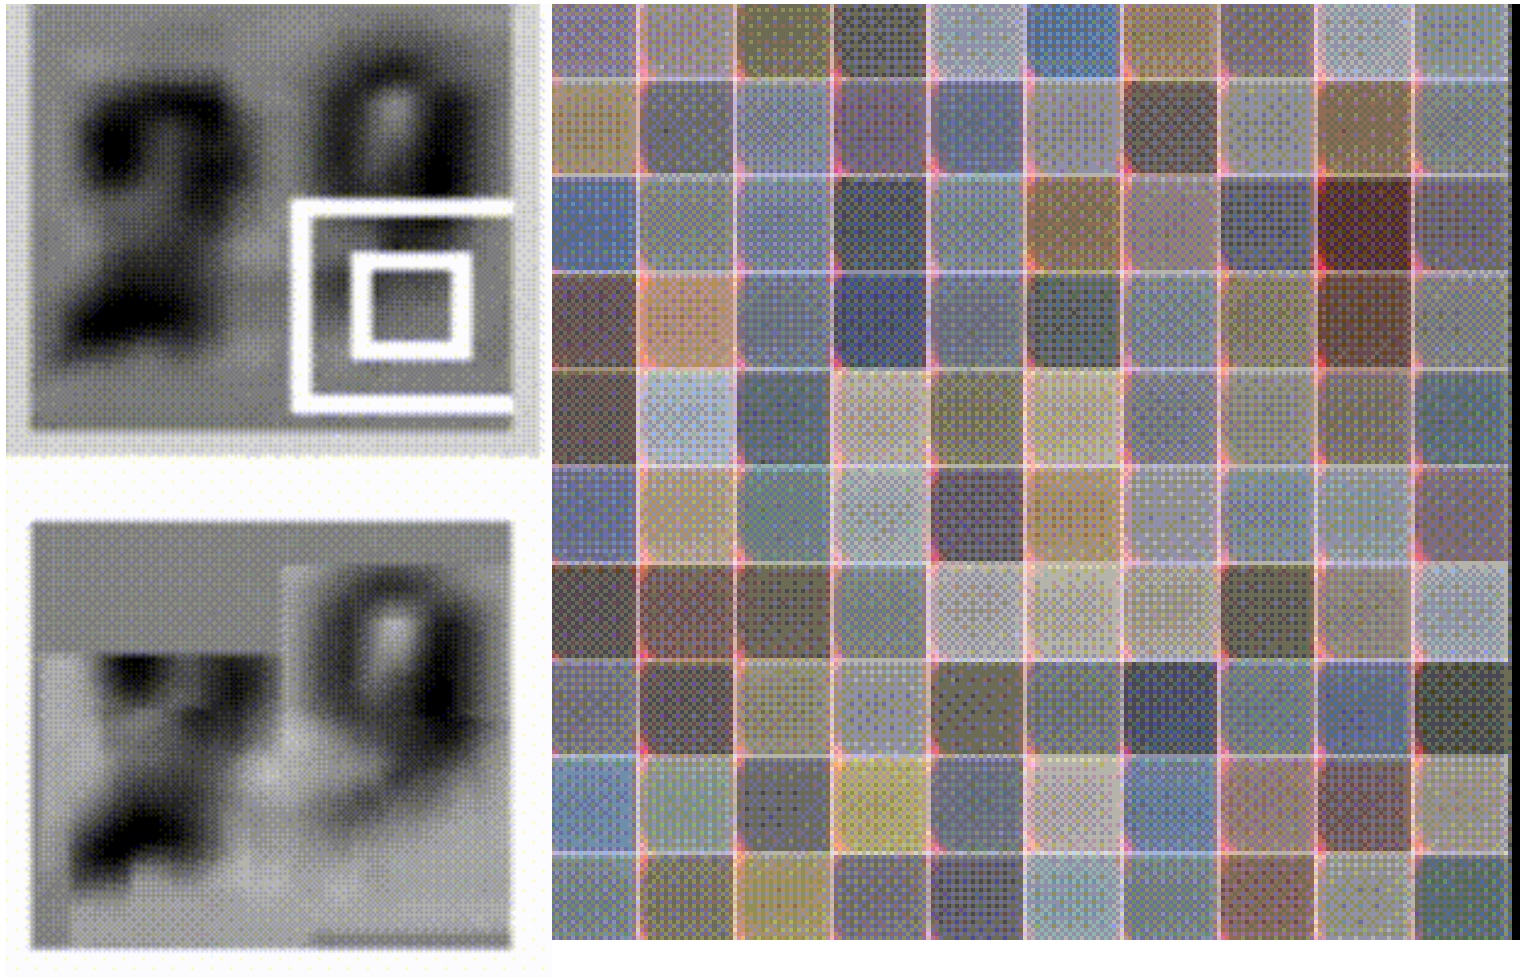

RNN learns to read house numbers.

RNN learns to paint house numbers.

http://karpathy.github.io/2015/05/21/rnn-effectiveness/

Visual Chatbot

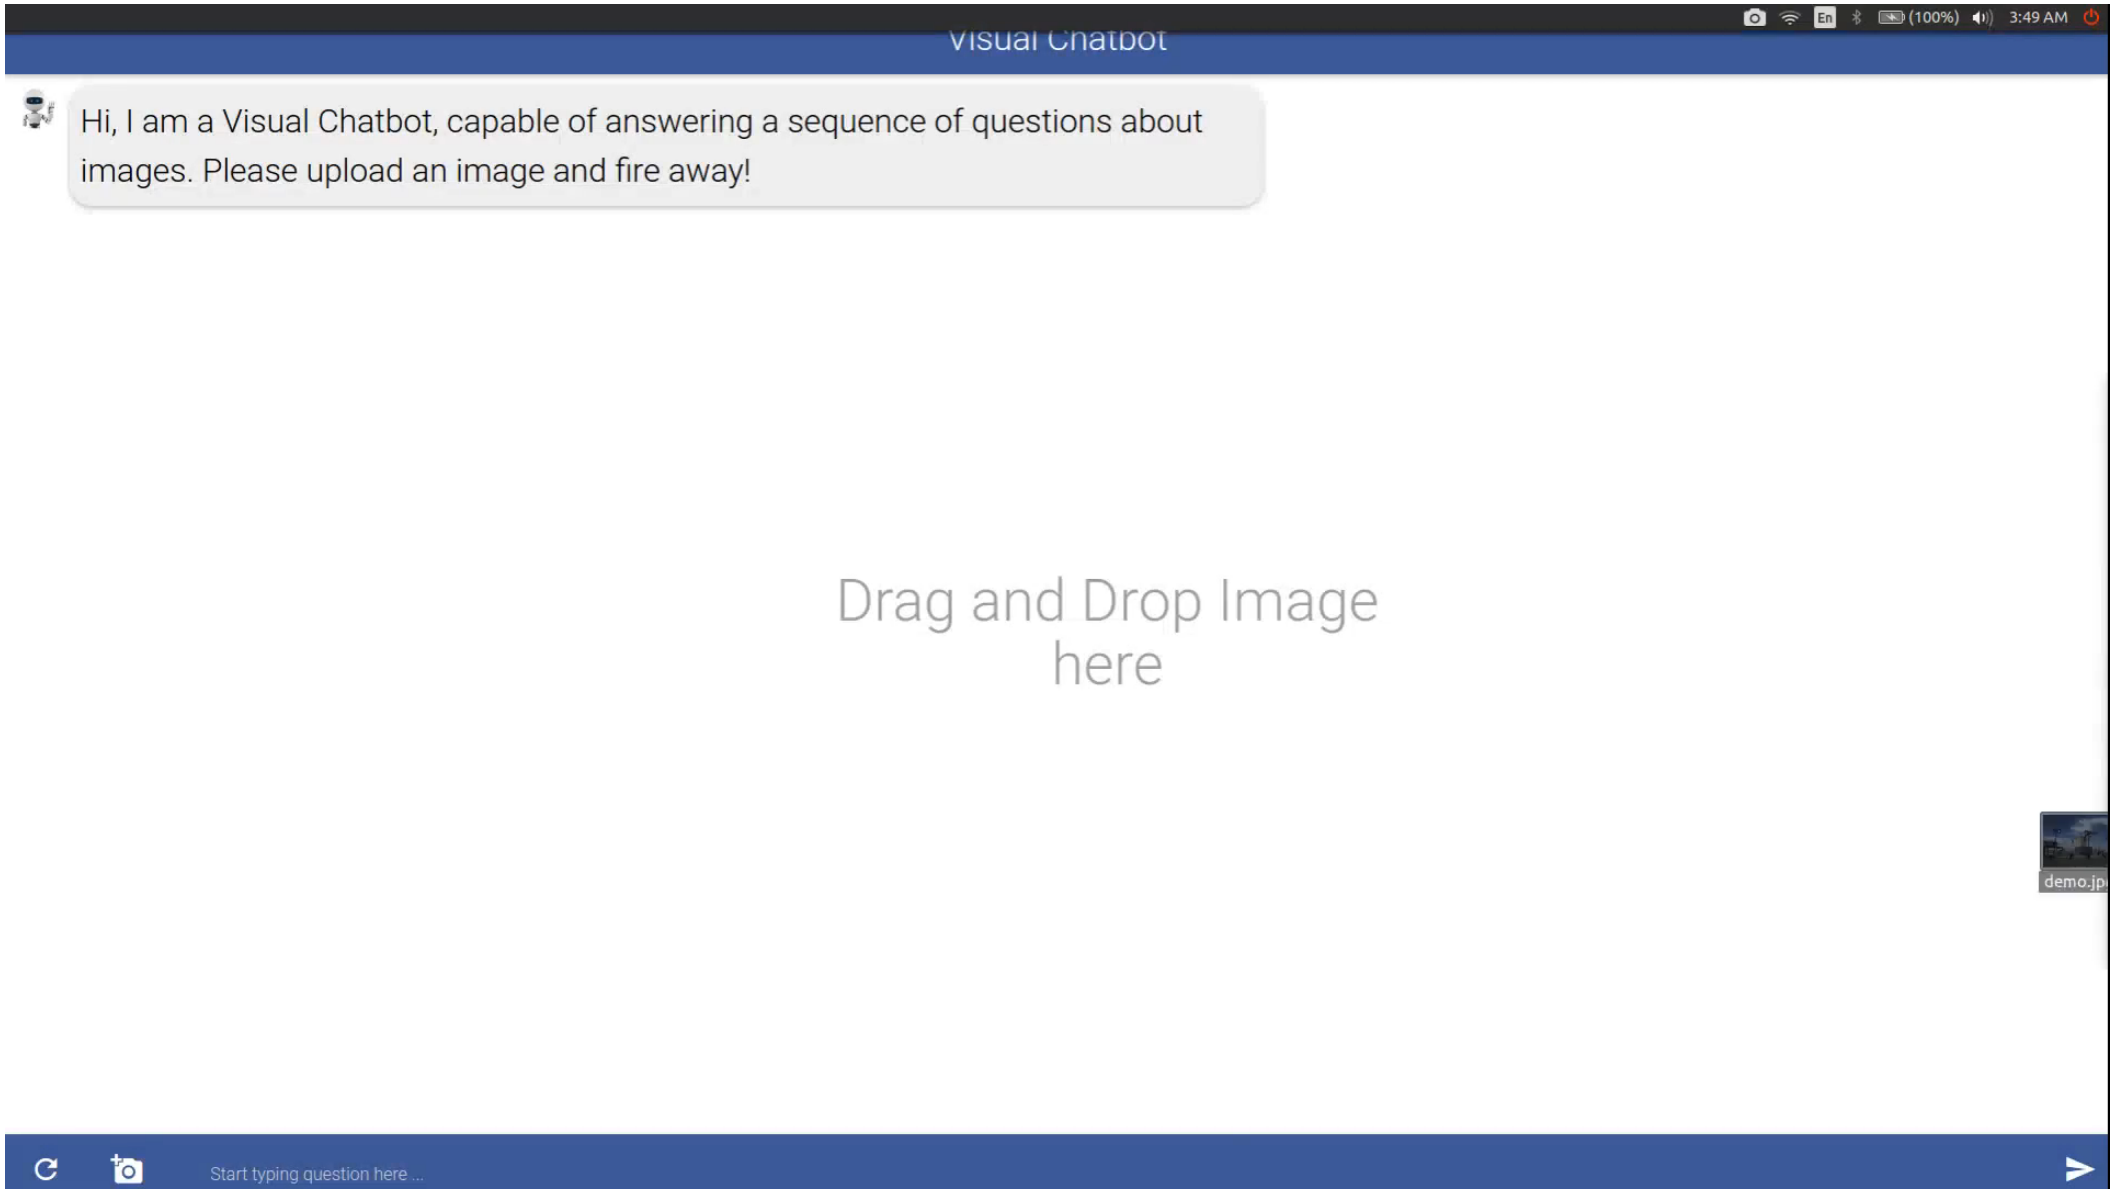

Image Compression, storage

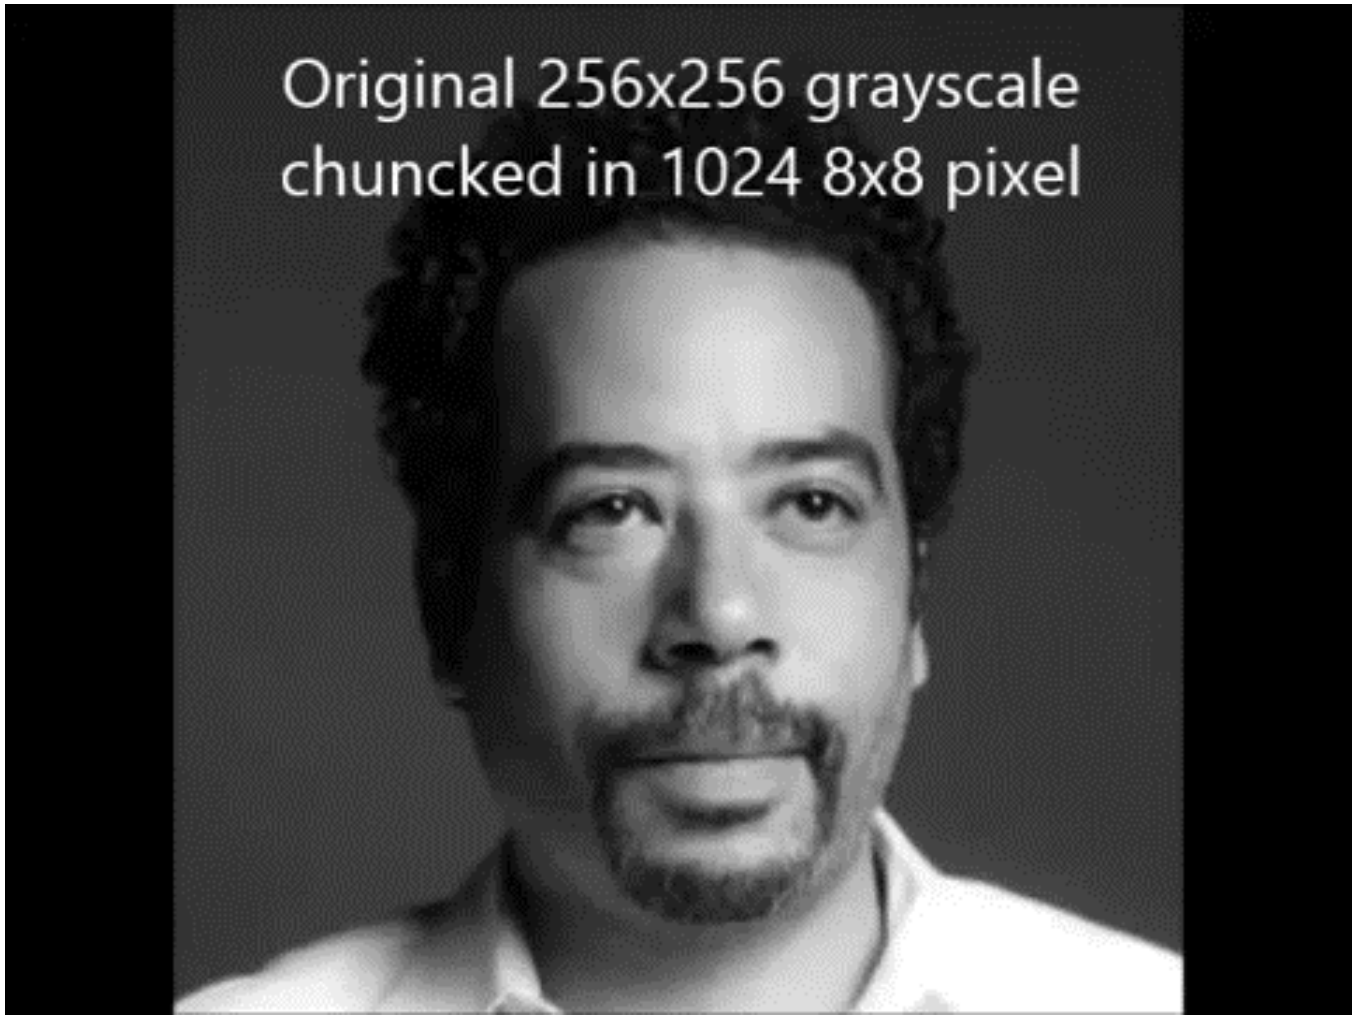

Artificial Neural Network that learns how to compress grayscale images - 100 milions training epocs

https://www.youtube.com/watch?v=Z-7SBgpturk

Extreme GAN Learned Image Compression

Our Algorithm (2379 Bytes) vs BPG (2565 Bytes)

https://data.vision.ee.ethz.ch/aeirikur/extremecompression/

Learning to see in the dark

Left, a photo brightened with traditional photo editing software. Right, the same image brightened with deep learning. (Intel/UIUC)

https://qz.com/1279913/artificial-intelligence-is-learning-to-see-in-the-dark/
https://github.com/cchen156/Learning-to-See-in-the-Dark

Applications – Voice generation

- Speech synthesis directly from text
- Google Assistant

"She earned a doctorate in sociology at Columbia University."

- Tacotron 2 or Human?

https://google.github.io/tacotron/publications/tacotron2/index.html

Google Duplex Assistant

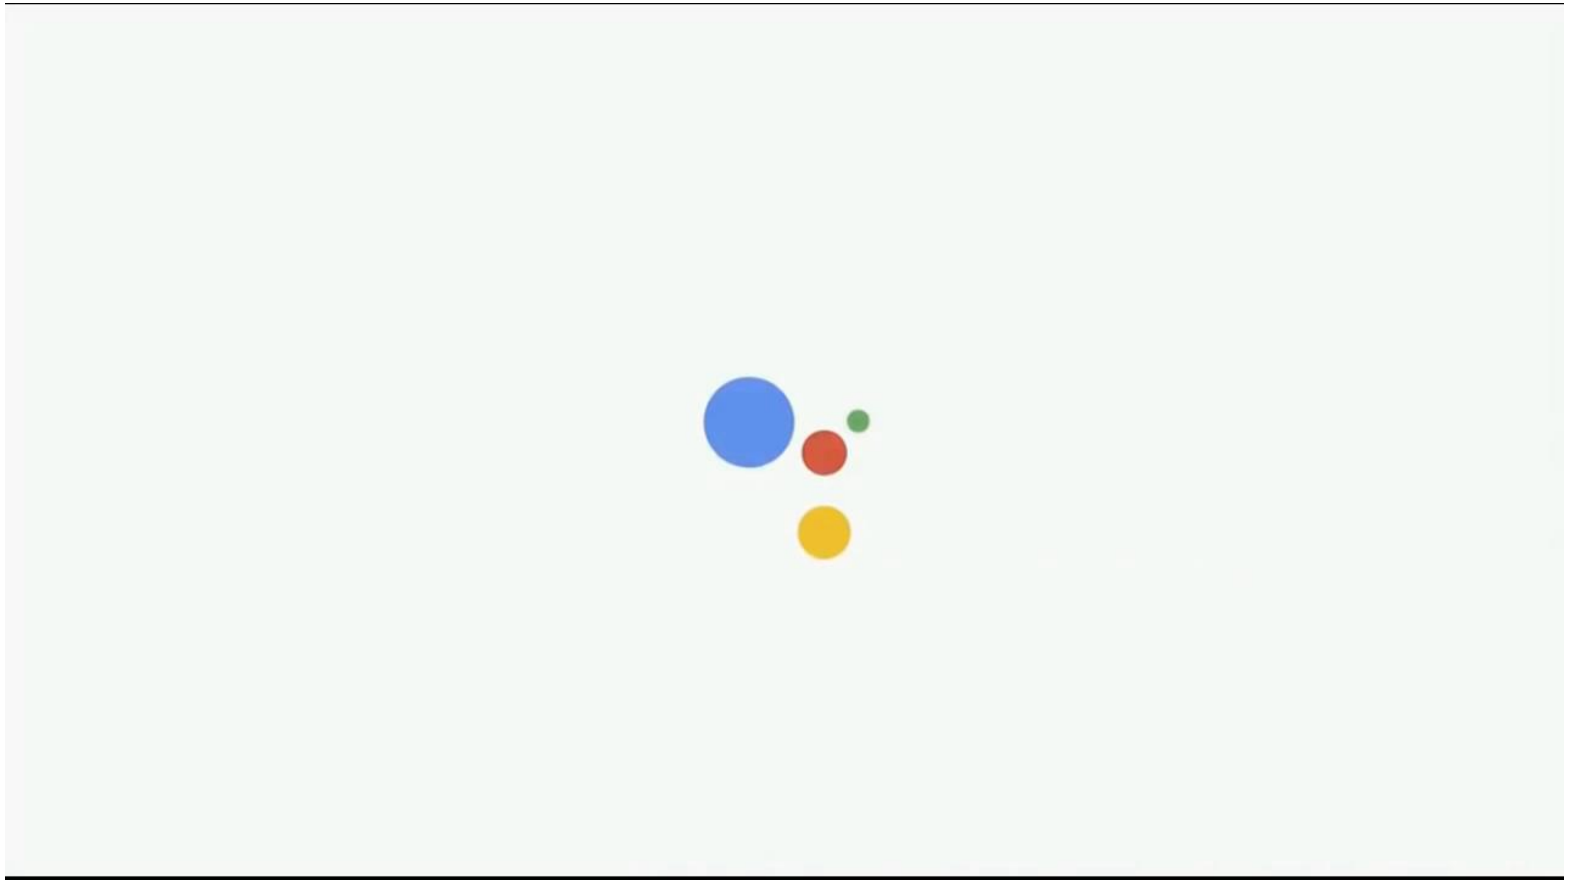

https://www.youtube.com/watch?v=D5VN56jQMWM&feature=youtu.be

Speech Recognition

1784

Wolfgang von Kempelen creates the **Acoustic-Mechanical Speech Machine** in Vienna

Thomas Edison invents the first dictation machine

1879

1952

Bell Labs releases **Audrey**, capable of recognizing spoken digits with **90% accuracy** - but only when spoken by its inventor

IBM Shoebox can understand 16 English words

1962

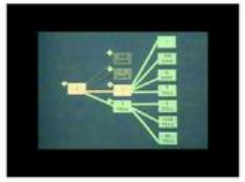

1971

Harpy, created at Carnegie Mellon University, can comprehend **1,011 words** - and some phrases

IBM Tangora, using the Hidden Markov Model, predicts upcoming phonemes in speech

1986

2006

The National Security Agency (NSA) starts using speech recognition to isolate key words in recorded speech

Google launches a voice search app, bringing speech recognition to mobile devices

2008

2011

Apple announces Siri, ushering in the age of the voice-enabled digital assistant

ClickZ

life.augmented

Speech Recognition using DNN

- A speech recognition system has many stages:
  - Pre-processing: Convert the sound wave into a vector of acoustic coefficients.
  - The acoustic model: Use a few adjacent vectors of these coefficients to place bets on which part of which phoneme is being spoken.
  - Decoding: Find the sequence of the bets which does the best job to fit the acoustic data and kind of things that people say.

- DNN has replaced the previous machine learning algorithms for speech recognition.
  - Previous methods: HMM (Hidden Markov Model), DTW (Dynamic Time Warping), GMM (Gaussian Mixture Models) etc.

Speech Recognition using DNN

Word Error rate from MSR, IBM & Google

The Task	Hours of training data	Deep neural network	Gaussian Mixture model	Gaussian Mixture model with more data
Switchboard (Microsoft research)	309	18.5 %	27.4	18.6
English Broadcast news (IBM)	50	17.5 %	18.8	
Google voice search (Android >4.1)	5,870	12.3 % (and falling)		16 % (>> 5,870)

Speech Recognition

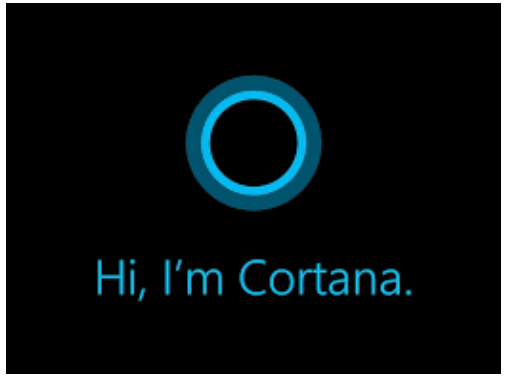

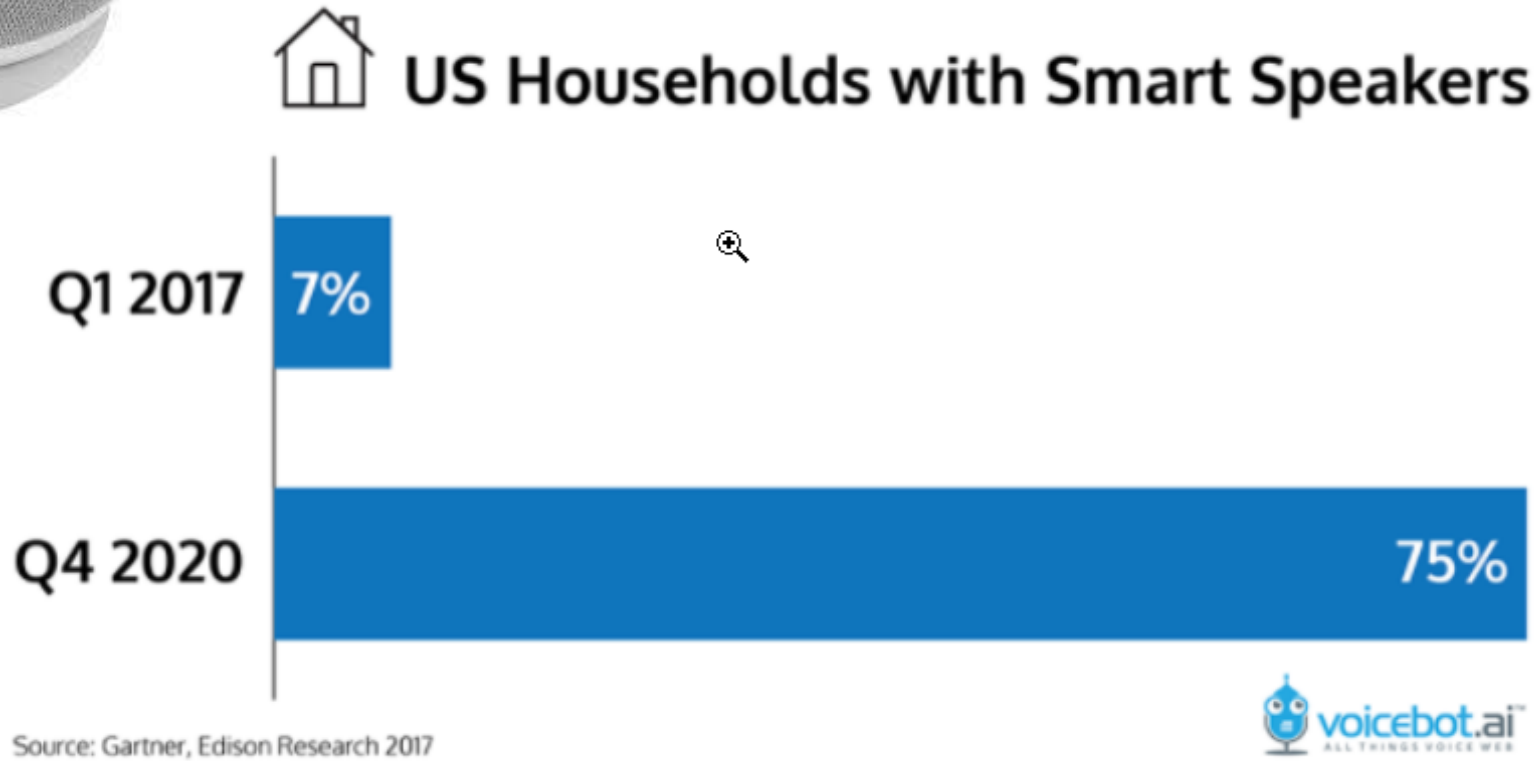

life.augmented

Which Personal Assistant is the Smartest?

Personal Assistant	% Questions Answered	100% Complete & Correct
The Google Assistant on Google Home	68.1%	90.6%
Cortana	56.5%	81.9%
Siri	21.7%	62.2%
Alexa on the Amazon Echo	20.7%	87.0%

Google search (for comparison purposes)	% Questions Answered	100% Complete & Correct
Google Search	74.3%	97.4%

https://www.stonetemple.com/digital-personal-assistants-test

Other Applications

- Speech recognition

- Language modeling

English: $p_1 > p_2 > p_3 > p_4$

P_1 = P("a quick brown dog")
P_2 = P("dog quick a brown")
P_3 = P("un chien quick brown")
P_4 = P("un chien brun rapide")

- Translation

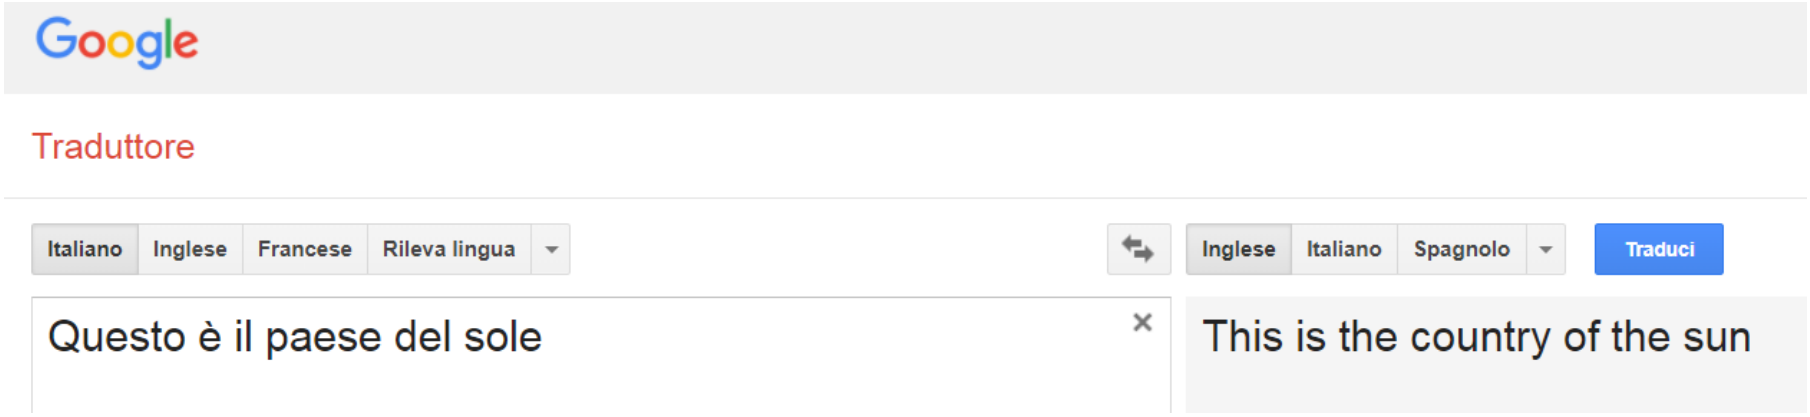

Playing Atari with Deep Reinforcement Learning

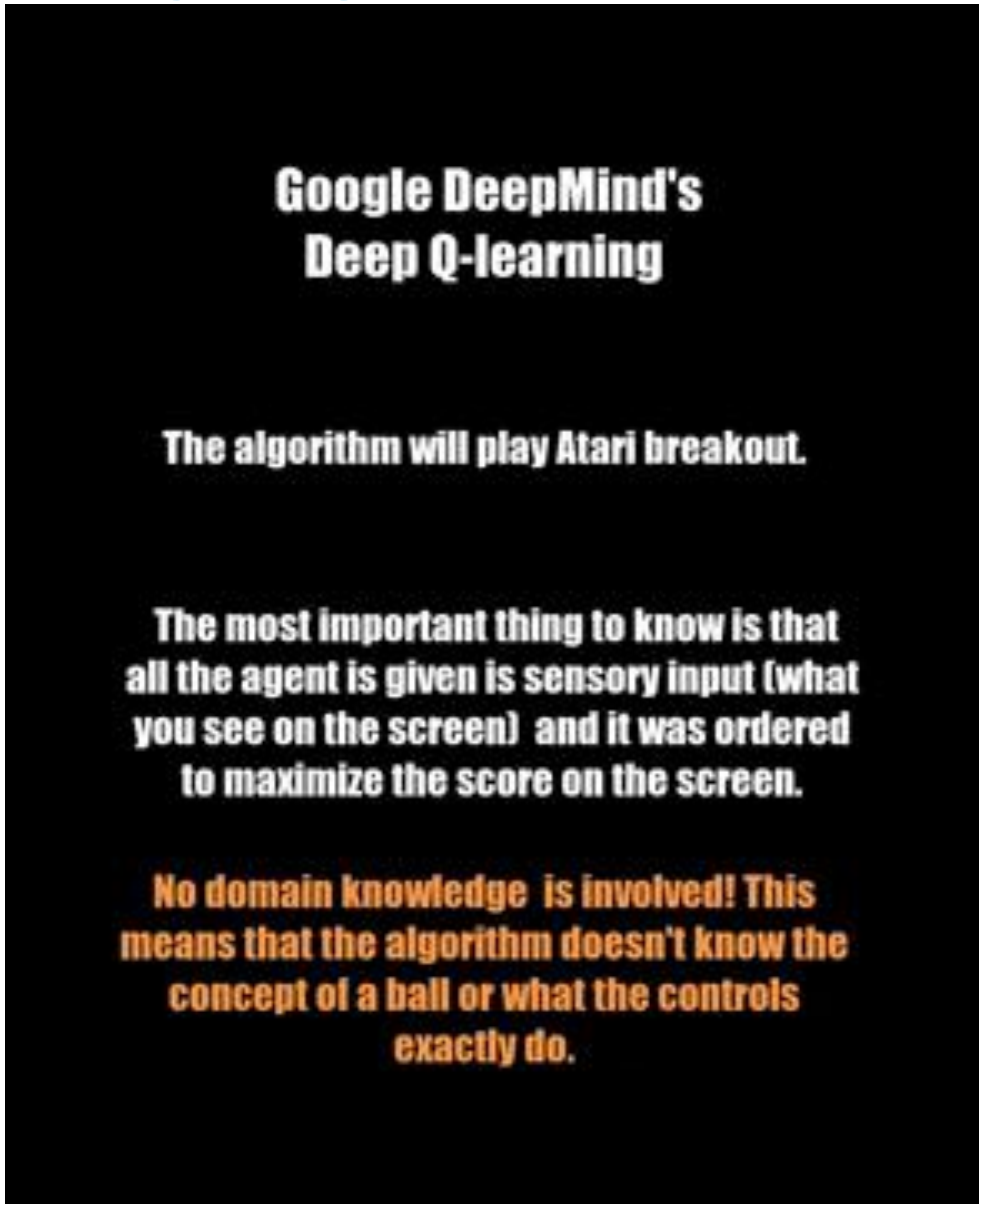

It is *autonomous*

It learns by itself, it receives no human expertise as input

In many cases, it outperforms human players

DeepMind AlphaGo

- “The game of Go originated in China more than 2,500 years ago.
  The rules of the game are simple: Players take turns to place black or white stones on a board, trying to capture the opponent's stones or surround empty space to make points of territory. As simple as the rules are, Go is a game of profound complexity. There are more possible positions in Go than there are atoms in the universe.”
    - On March 2016, AlphaGo won 4-1 against the legendary Lee Sedol, one of the top three players in the world.

life.augmented

Apples Gathering

- Things went smoothly as there were enough apples to gather, but as soon as the apples began to dwindle, the two agents became aggressive, using laser beams to knock each other out of the game to steal all the apples.

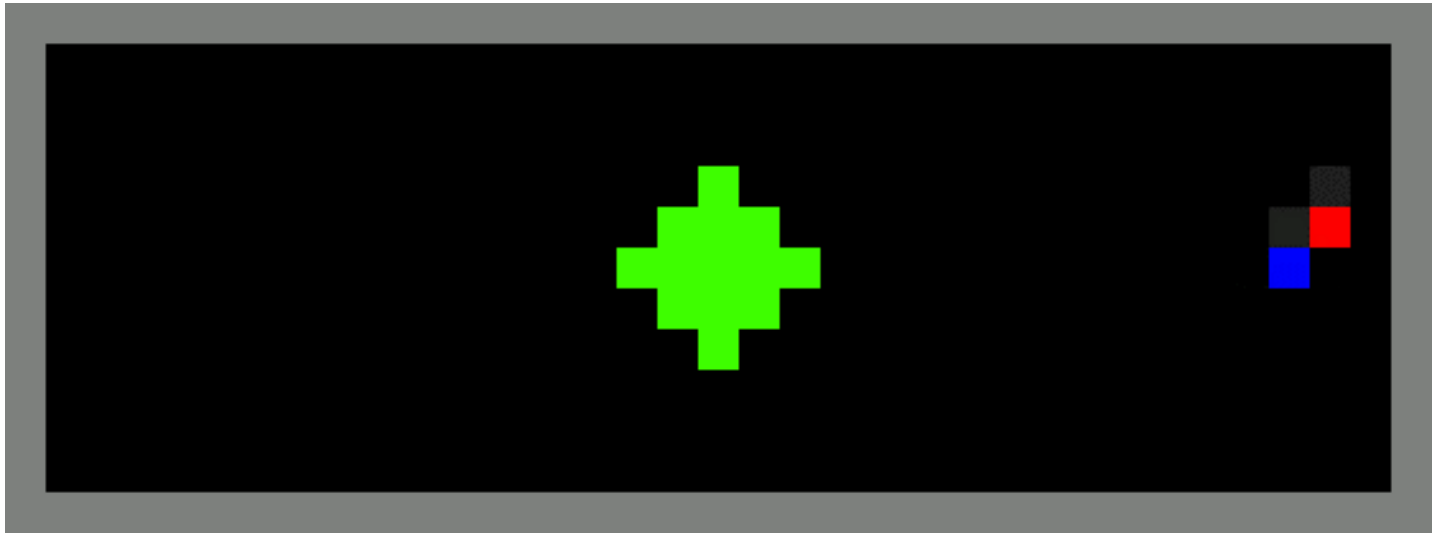

AI agents in blue and red;
Aapples in green;
Laser beams in yellow

https://www.sciencealert.com/google-deep-mind-has-learned-to-become-highly-aggressive-in-stressful-situations

Wolfpack

- Actively encouraged co-operation: if both wolves were near the prey when it was captured, they both received a reward - regardless of which one actually took it down.
- Lone wolf can overcome it, but is at risk of losing the prey to scavengers

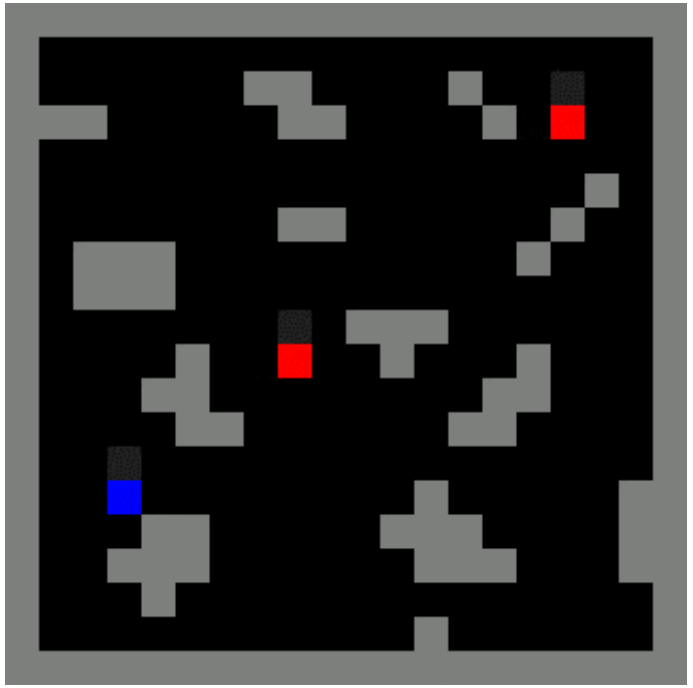

AI agents - two as (red) wolves, and one as the (blue) prey.

https://www.sciencealert.com/google-deep-mind-has-learned-to-become-highly-aggressive-in-stressful-situations

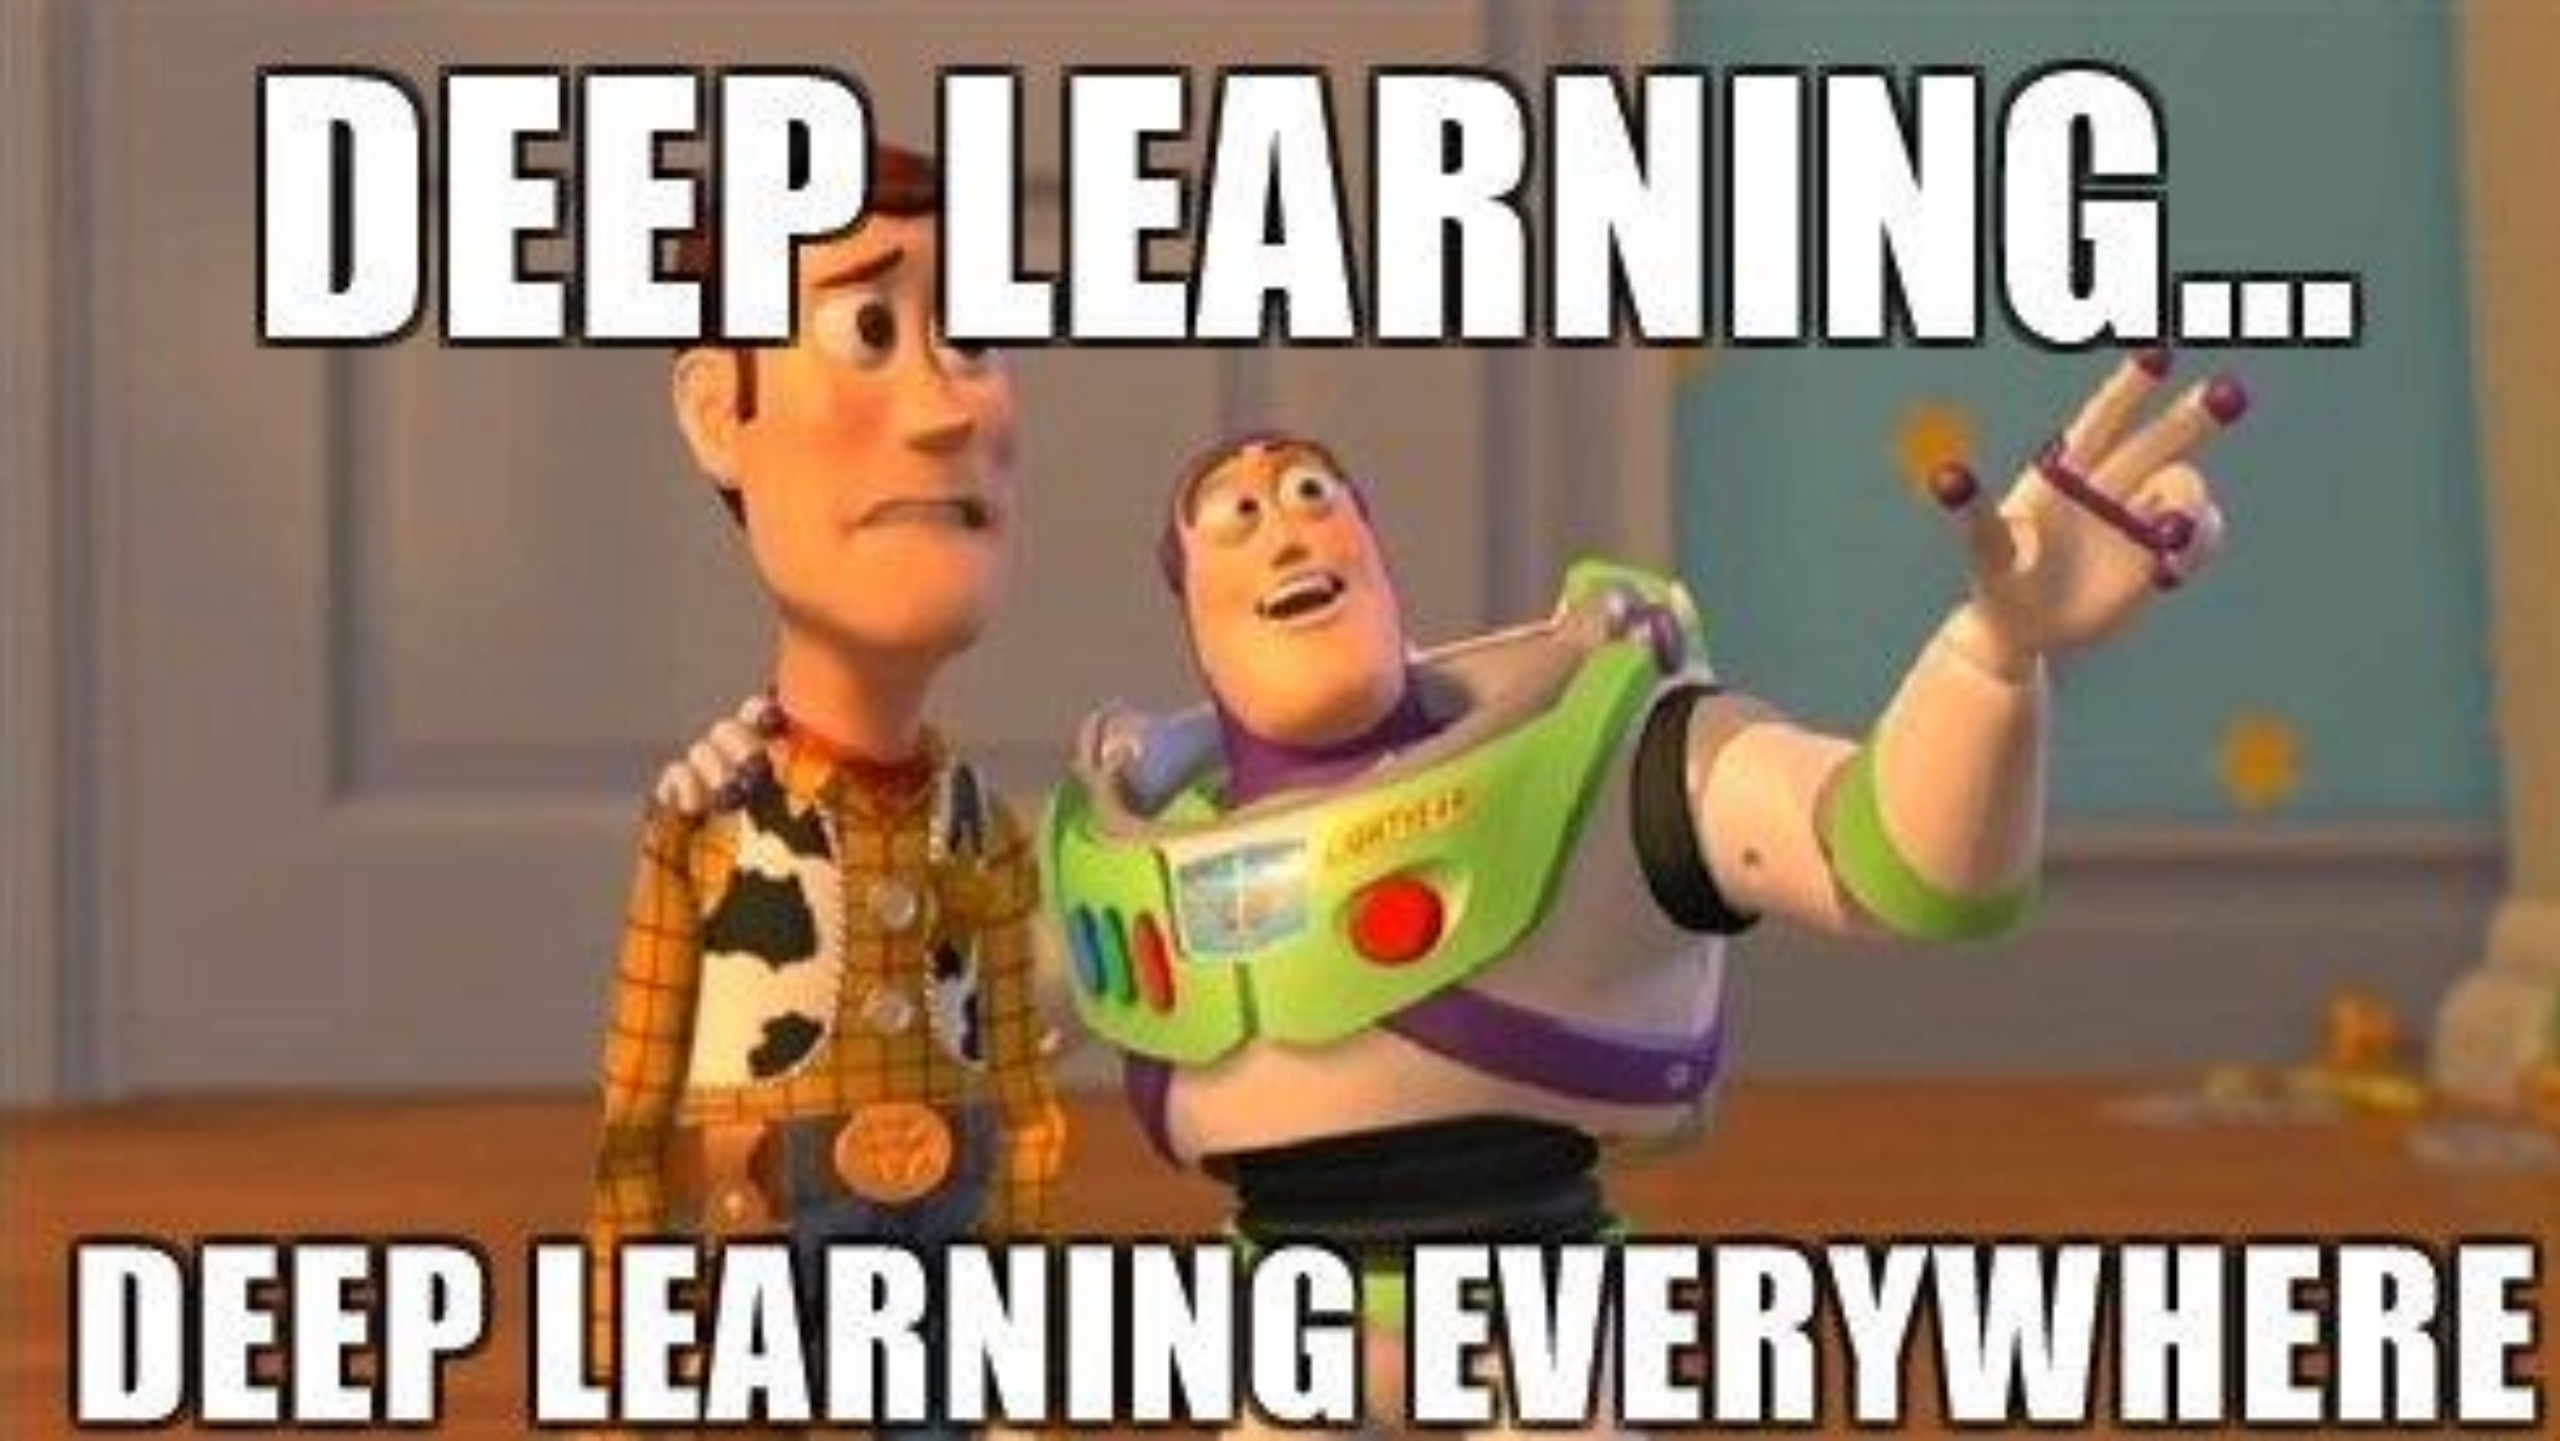
DEEP LEARNING...
DEEP LEARNING EVERYWHERE