Deep Learning: Theory and Practice

Applications of Deep Learning

25-04-2019

Natural Language Processing

Recent Trends in Deep Learning Based Natural Language Processing

Tom Young, Devamanyu Hazarika, Soujanya Poria, Erik Cambria

(Submitted on 9 Aug 2017 (v1), last revised 25 Nov 2018 (this version, v8))

Deep learning methods employ multiple processing layers to learn hierarchical representations of data and have produced state-of-the-art results in many domains. Recently, a variety of model designs and methods have blossomed in the context of natural language processing (NLP). In this paper, we review significant deep learning related models and methods that have been employed for numerous NLP tasks and provide a walk-through of their evolution. We also summarize, compare and contrast the various models and put forward a detailed understanding of the past, present and future of deep learning in NLP.

https://arxiv.org/abs/1708.02709#

Natural Language Processing

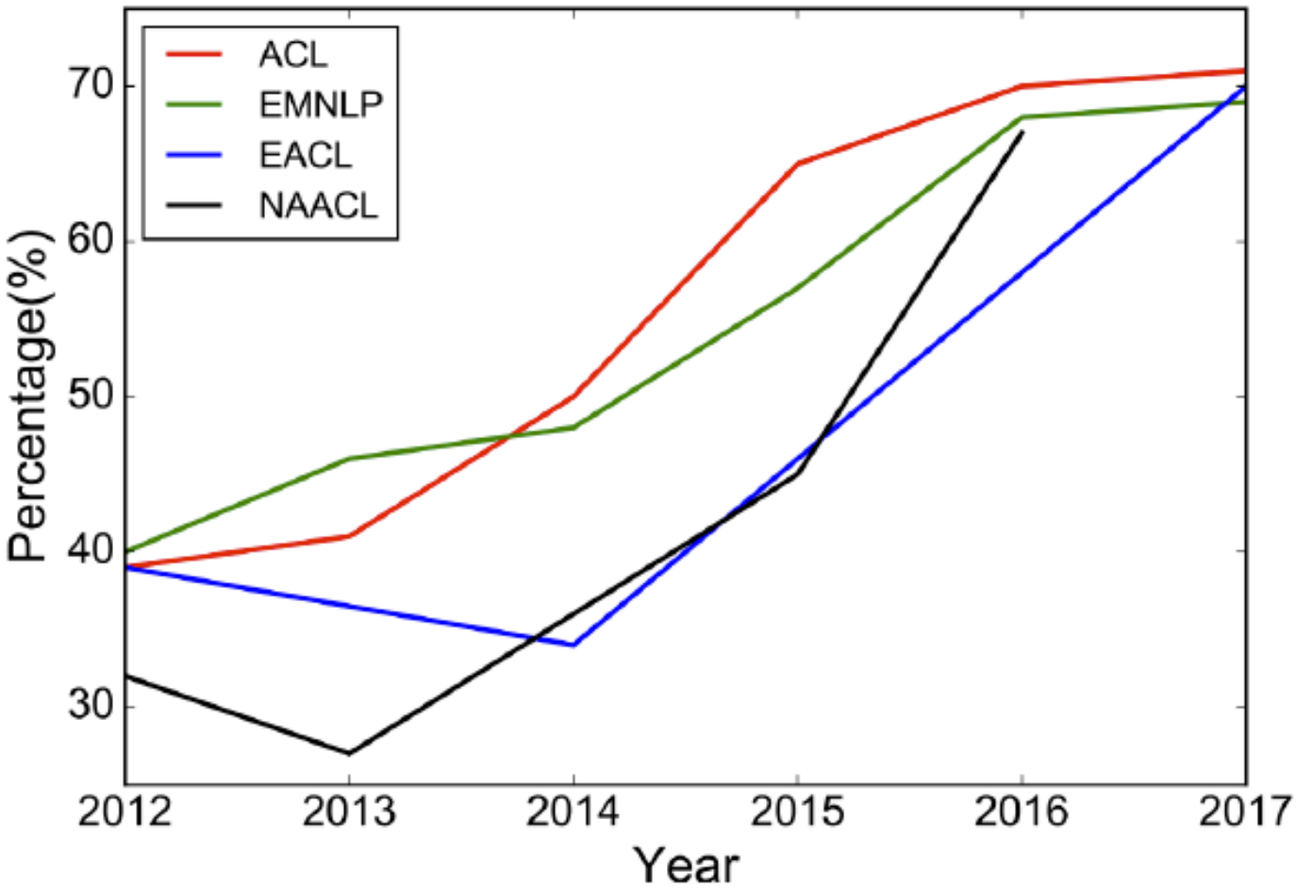

Fig. 1: Percentage of deep learning papers in ACL, EMNLP, EACL, NAACL over the last 6 years

Word Embeddings

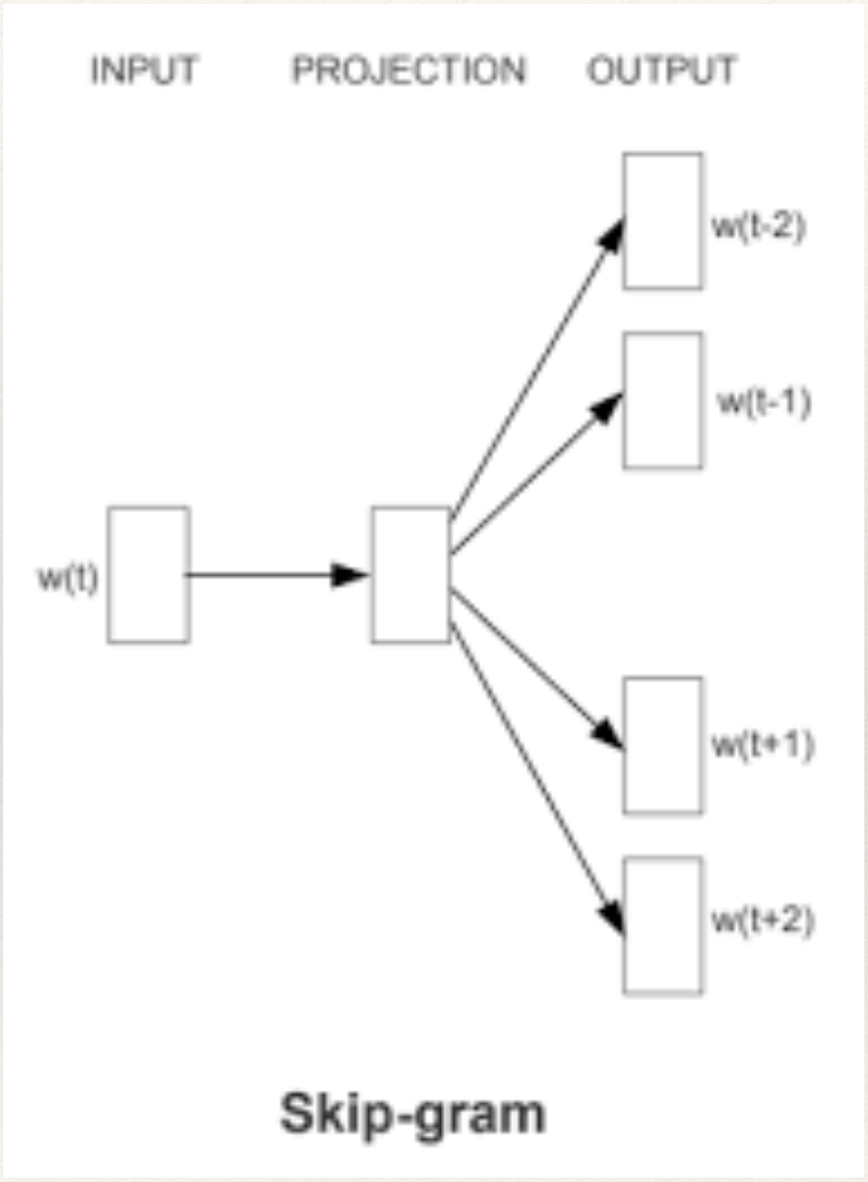

Skip-gram

Word	Cosine distance
norway	0.760124
denmark	0.715460
finland	0.620022
switzerland	0.588132
belgium	0.585835
netherlands	0.574631
iceland	0.562368
estonia	0.547621
slovenia	0.531408

Learn a vector representation for each word

Language Modeling

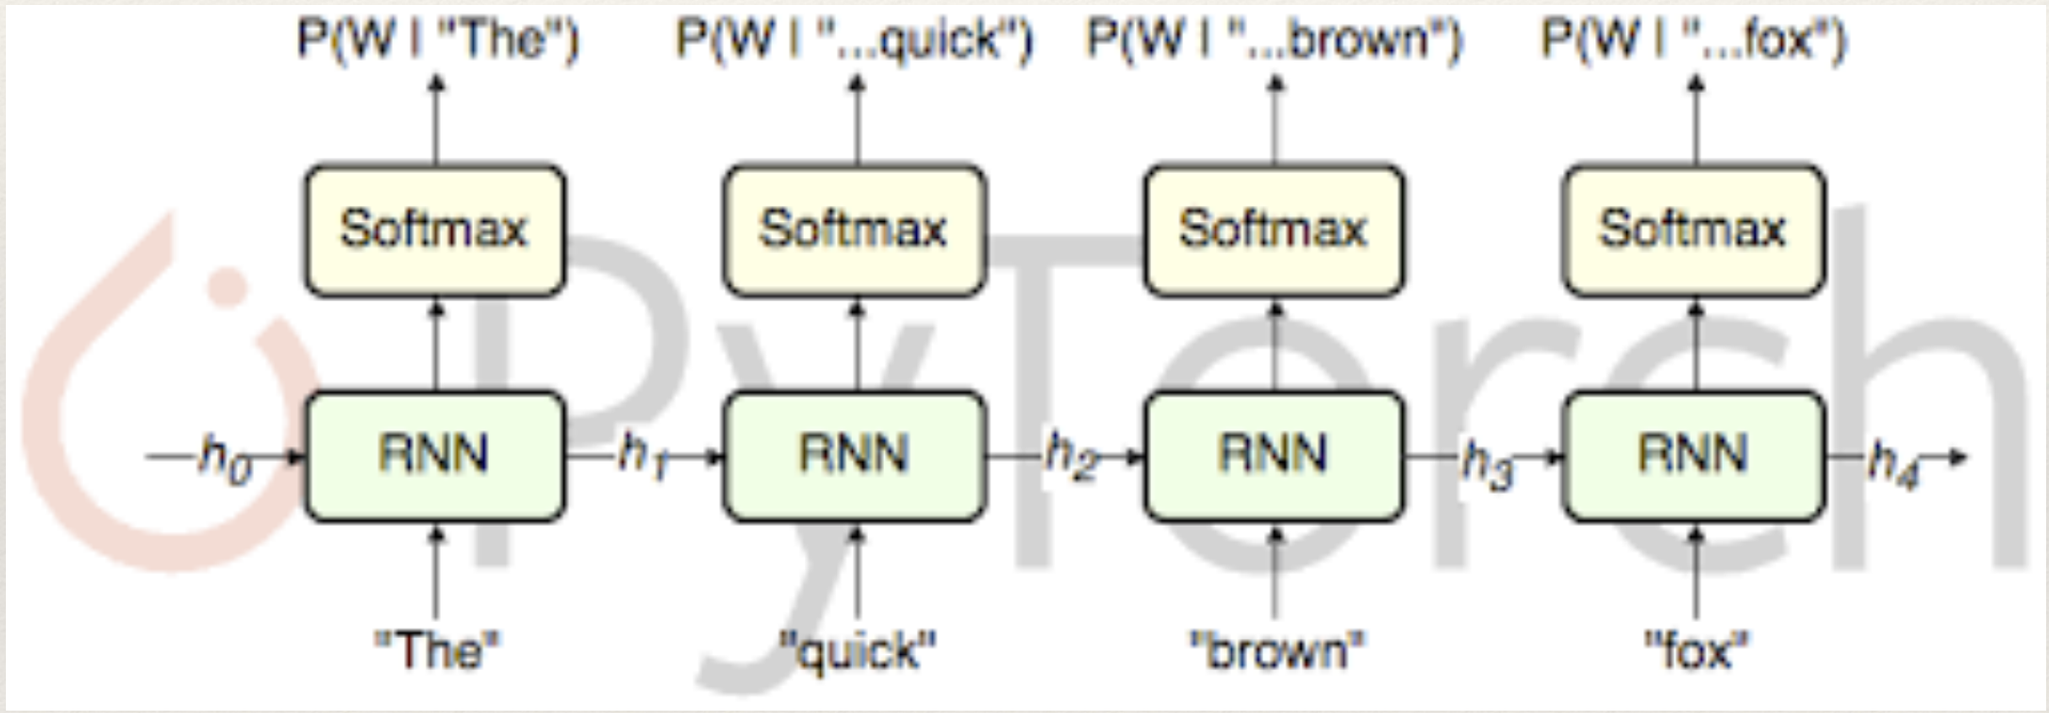

Sentiment Analysis/Summarization

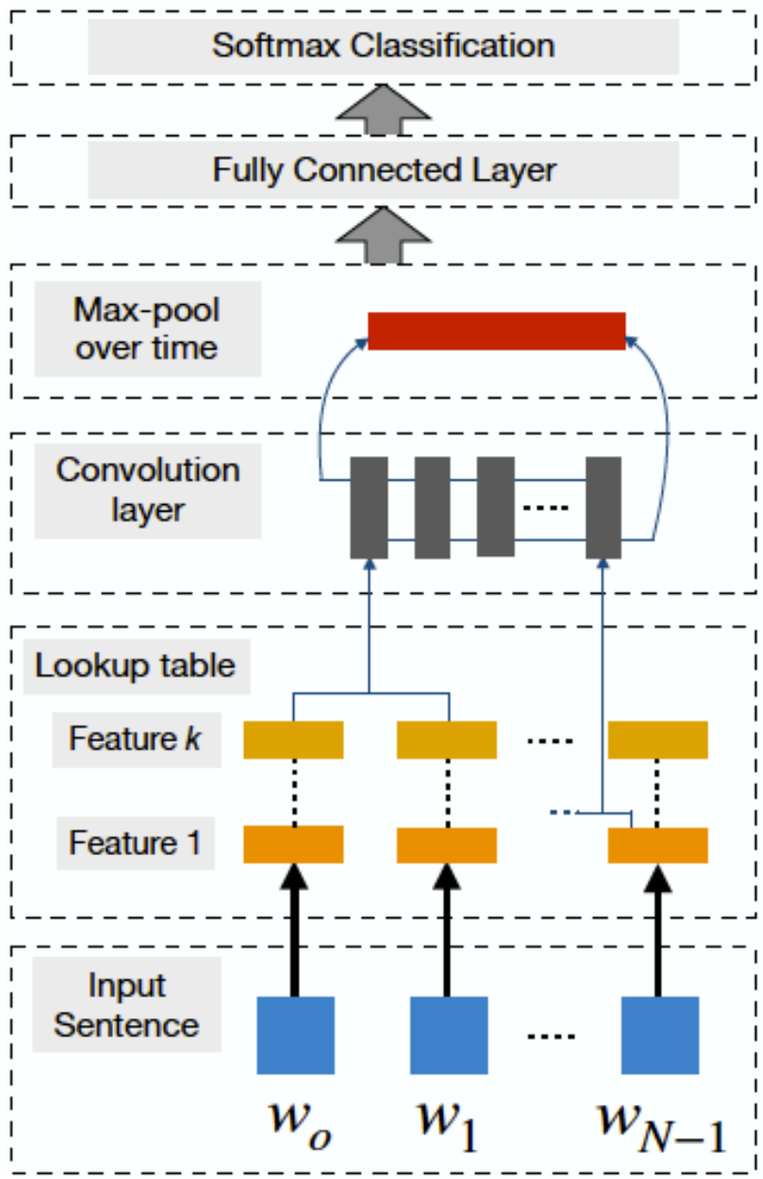

Fig. 5: CNN framework used to perform word wise class prediction (Figure source: Collobert and Weston [19])

Machine Translation

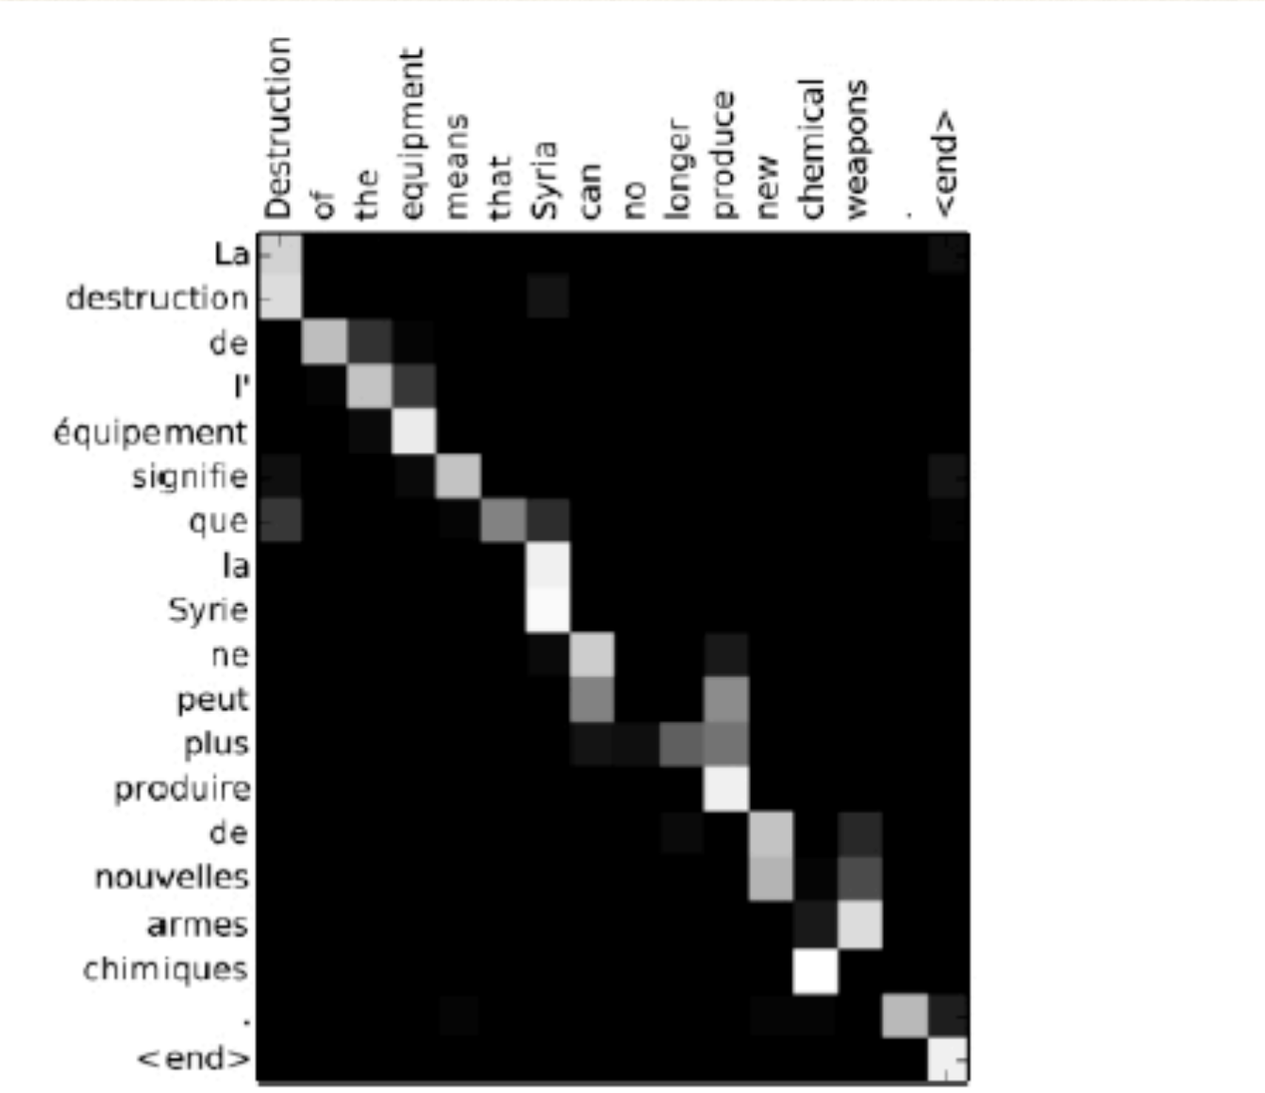

Fig. 14: Word alignment matrix (Figure source: Bahdanau et al. [103])

Architectures

Recurrent Neural Networks

Attention Based Models

Speech Processing

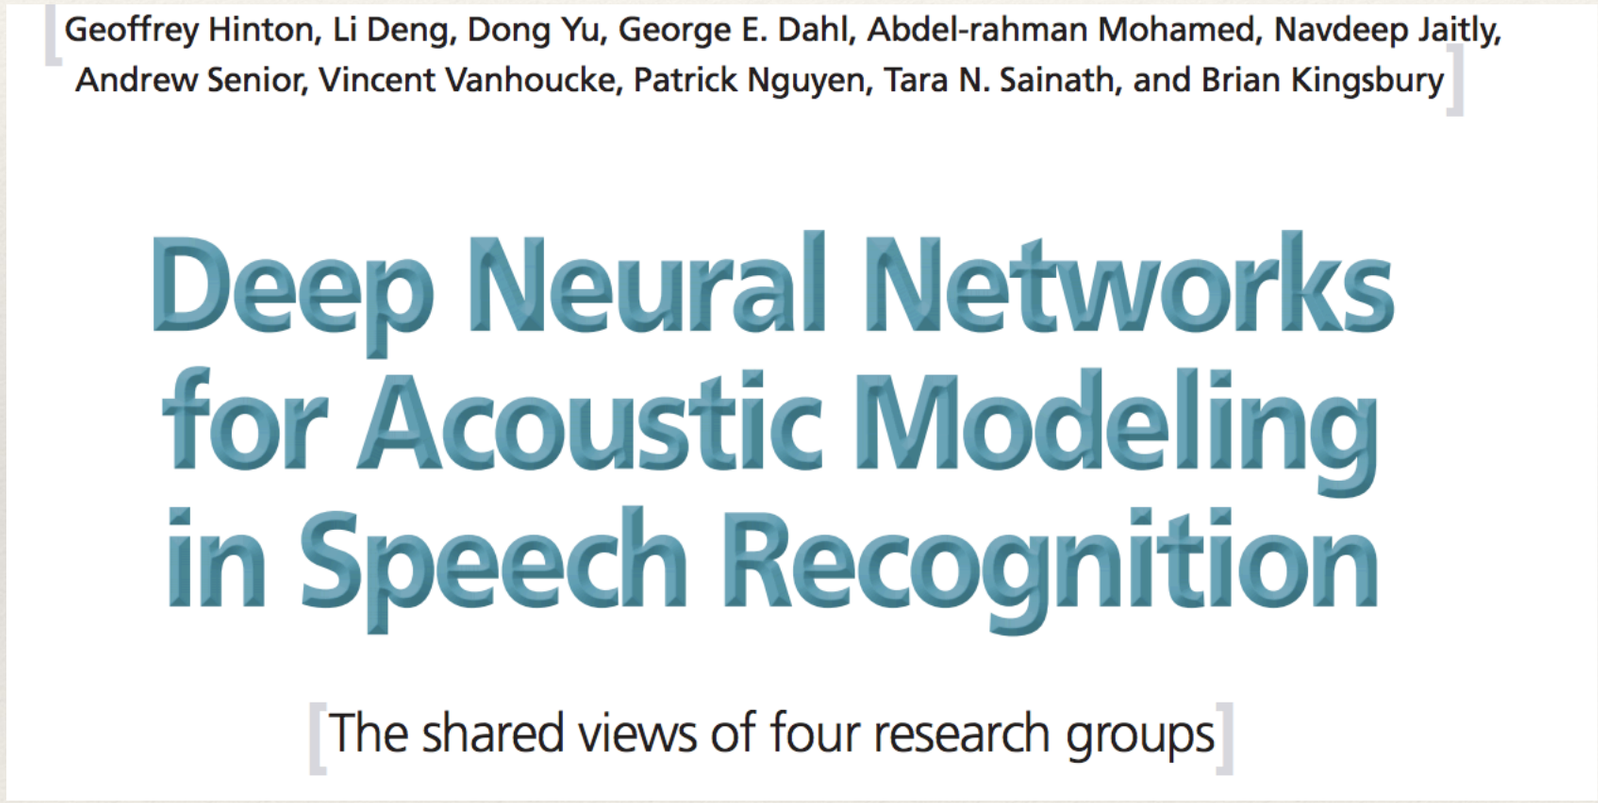

https://www.microsoft.com/en-us/research/wp-content/uploads/2016/02/HintonDengYuEtAl-SPM2012.pdf

Speech Recognition

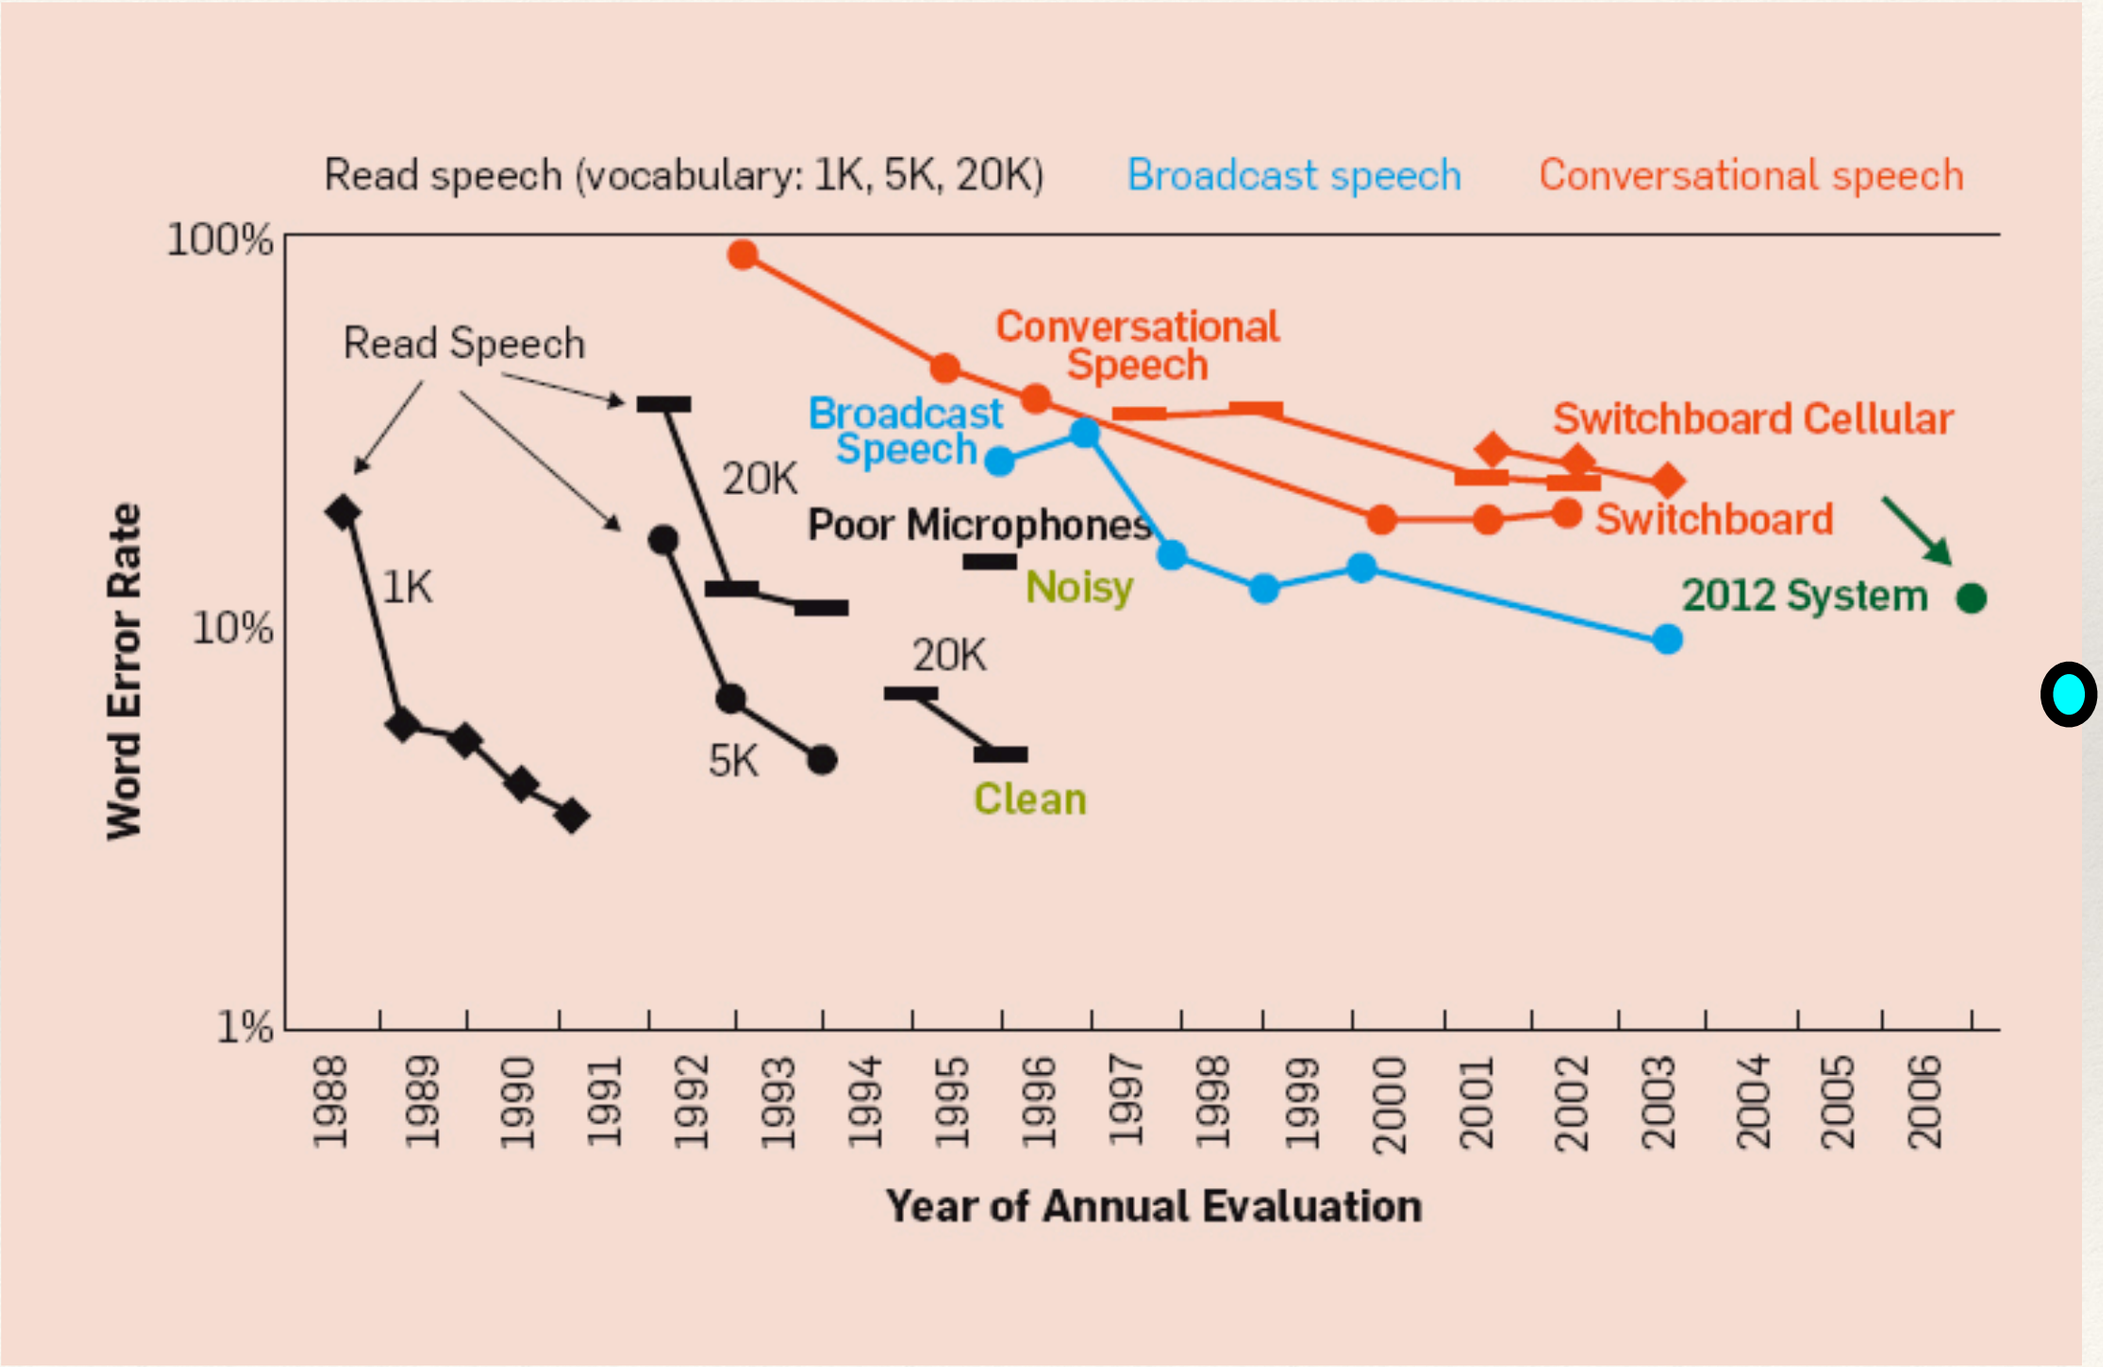

2018

5.3%

Claims of human parity !!!

Speaker Recognition Embedding

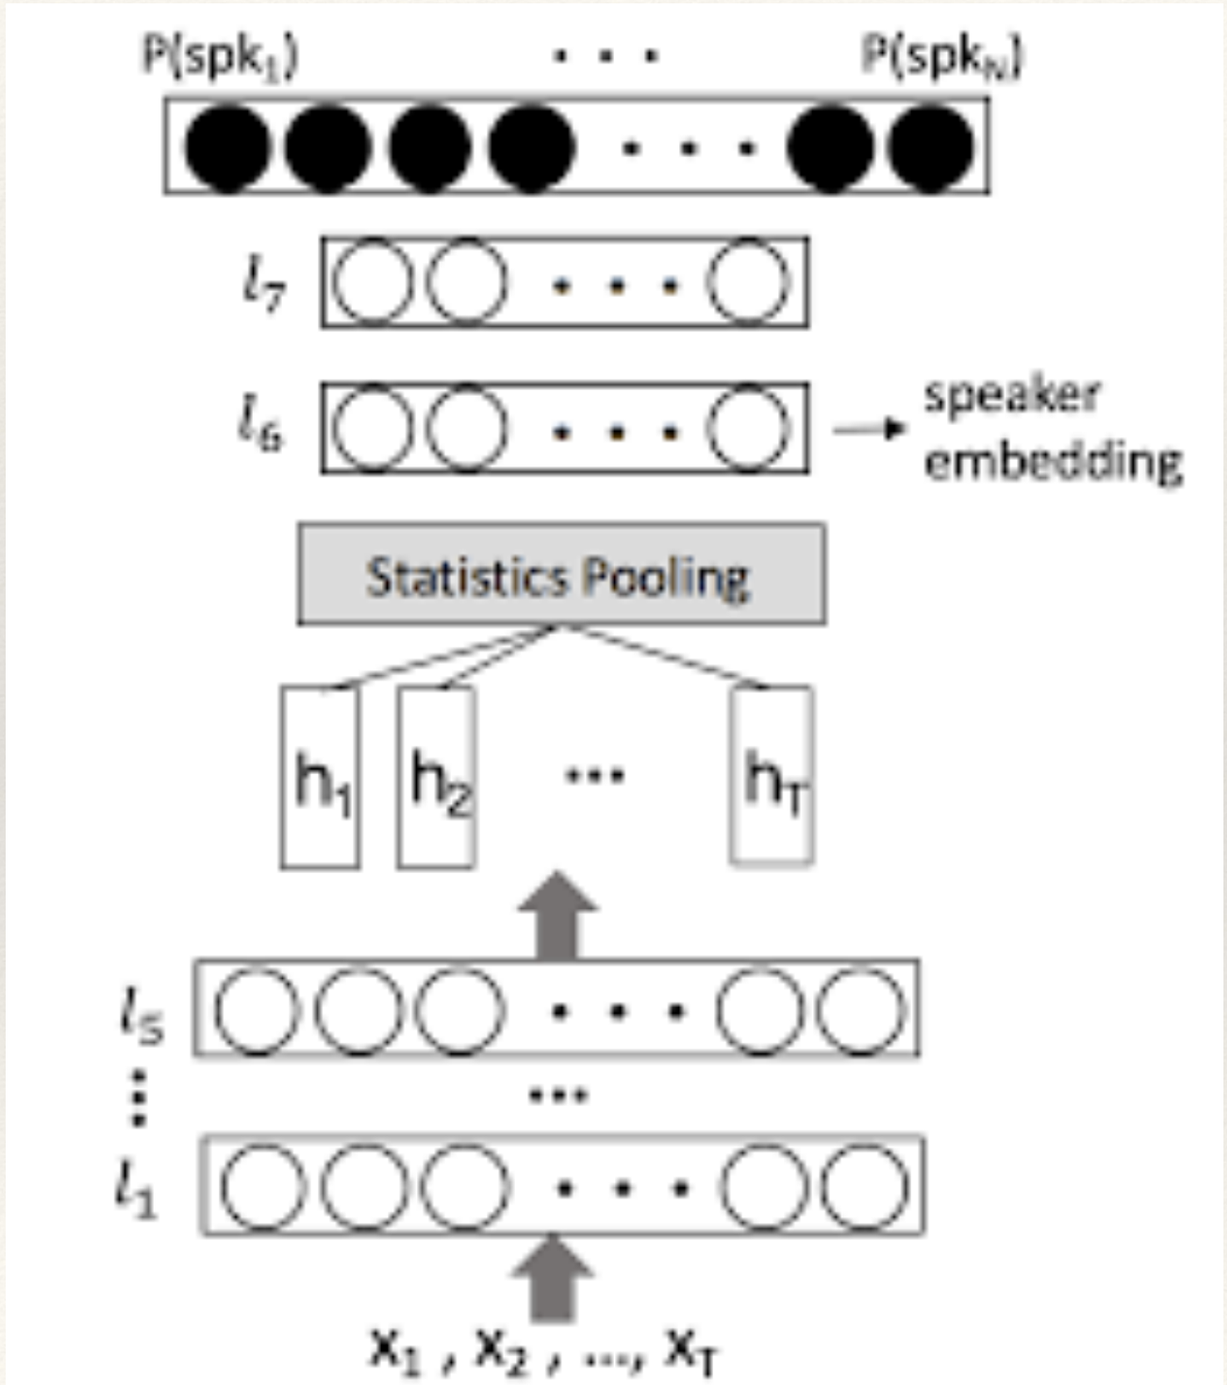

Embeddings used for verification task.

Speech Enhancement

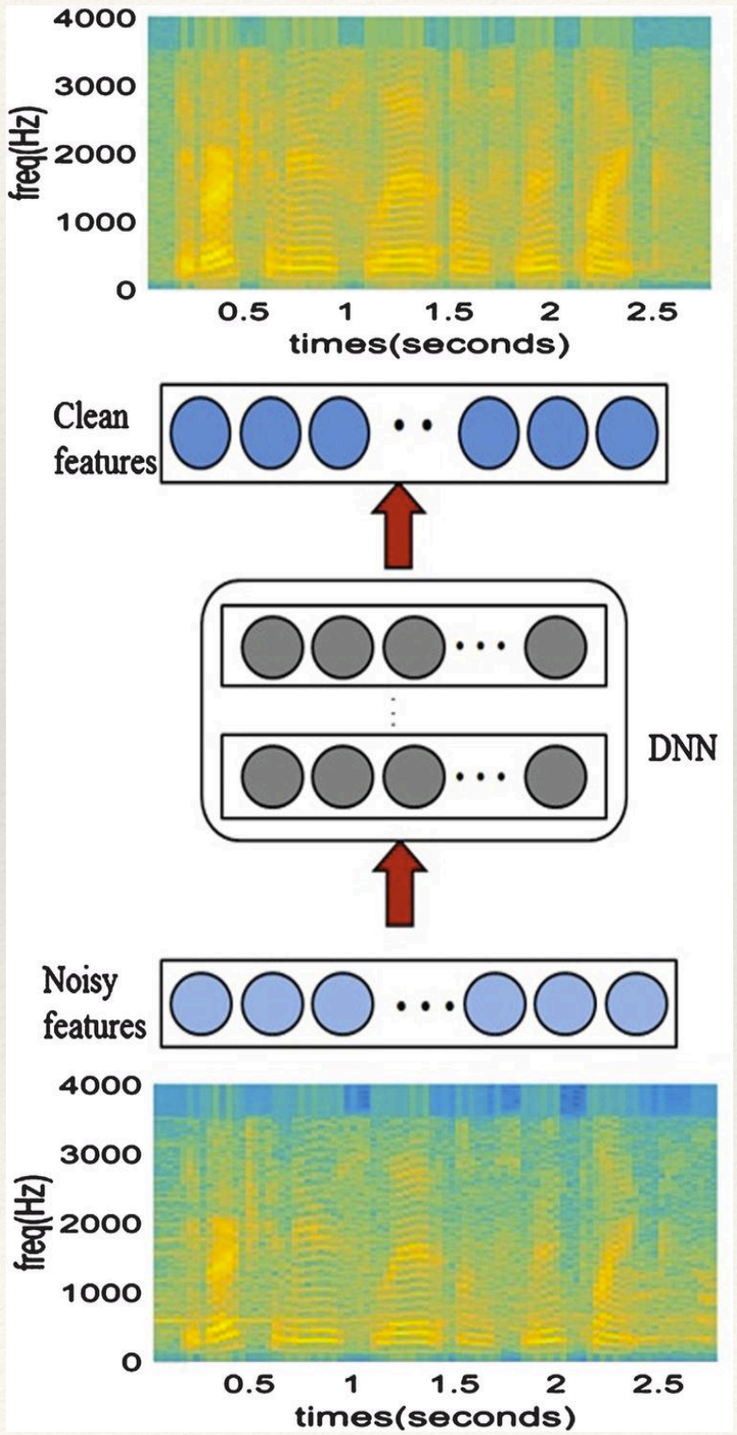

Removing Noise Using a Supervised Training Mechanism

Architectures

Deep neural networks

Bidirectional LSTM Networks

Image Processing

Object Identification

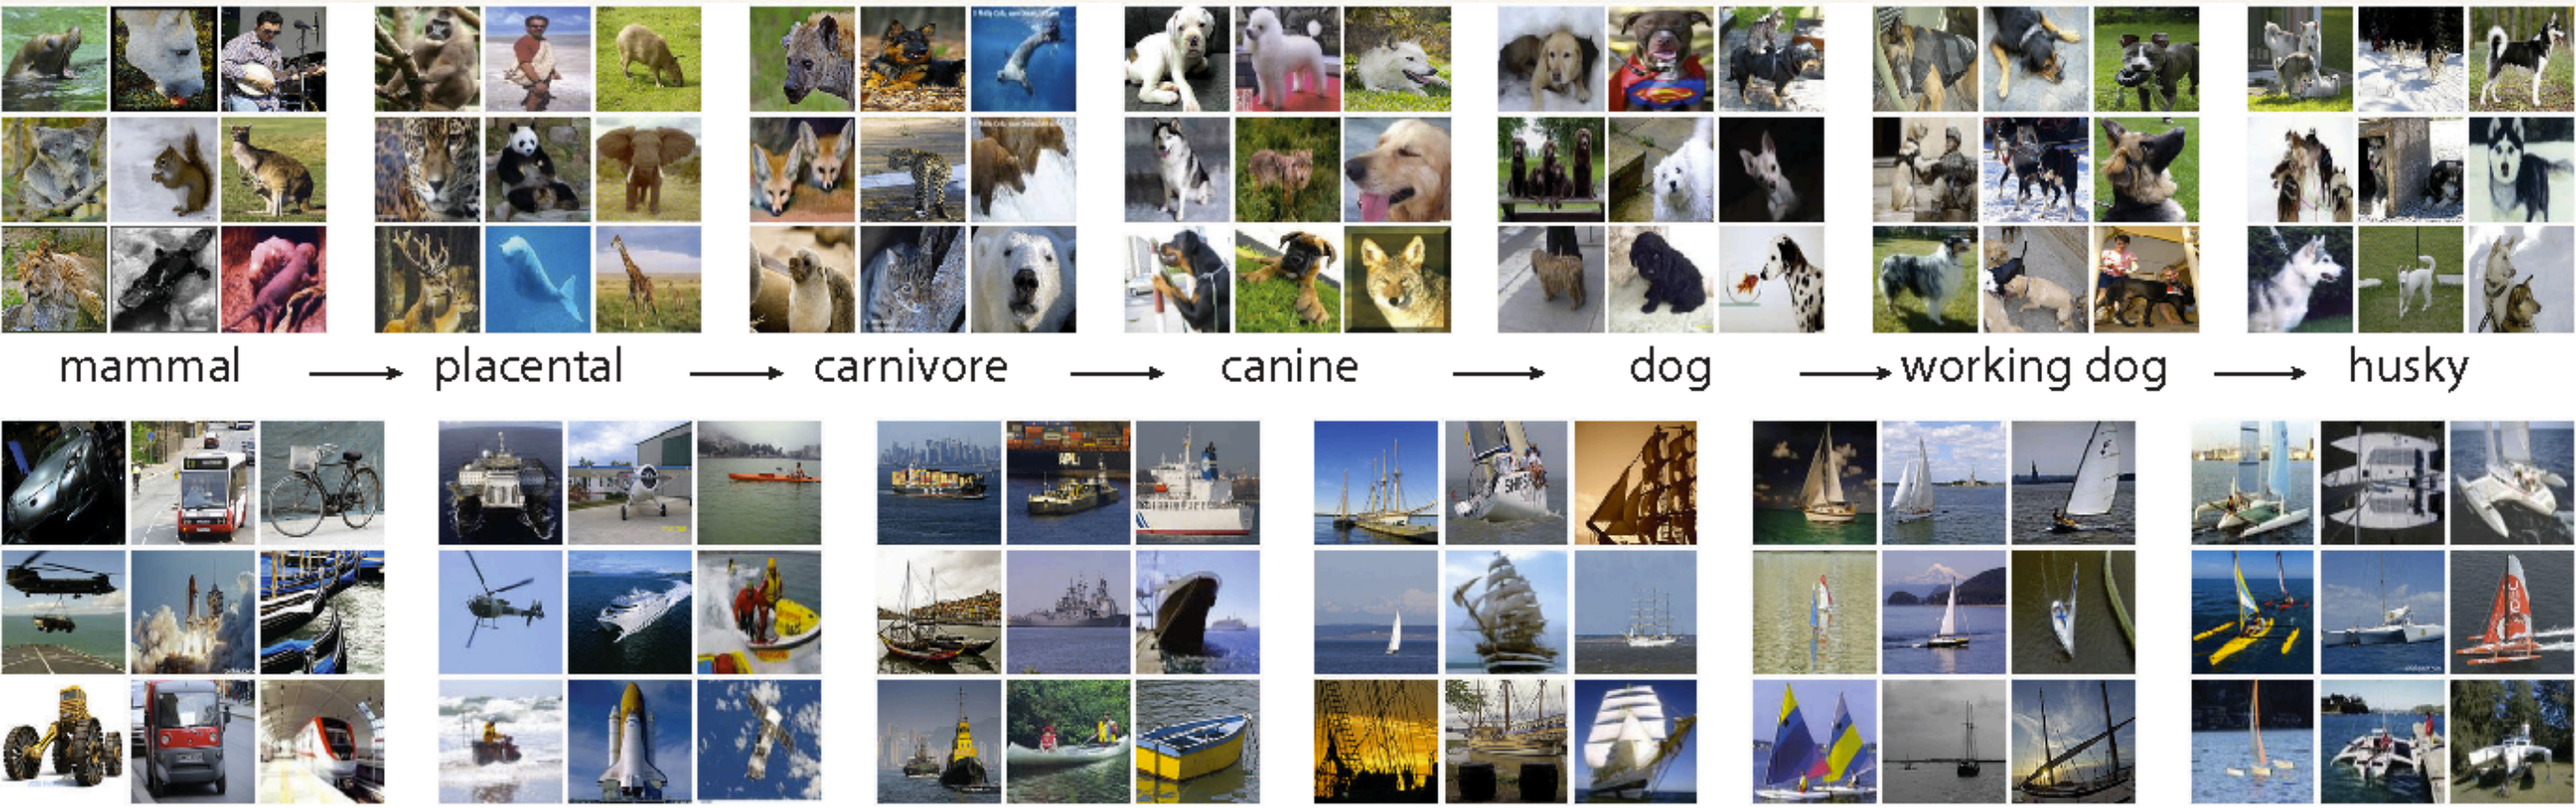

Object Segmentation

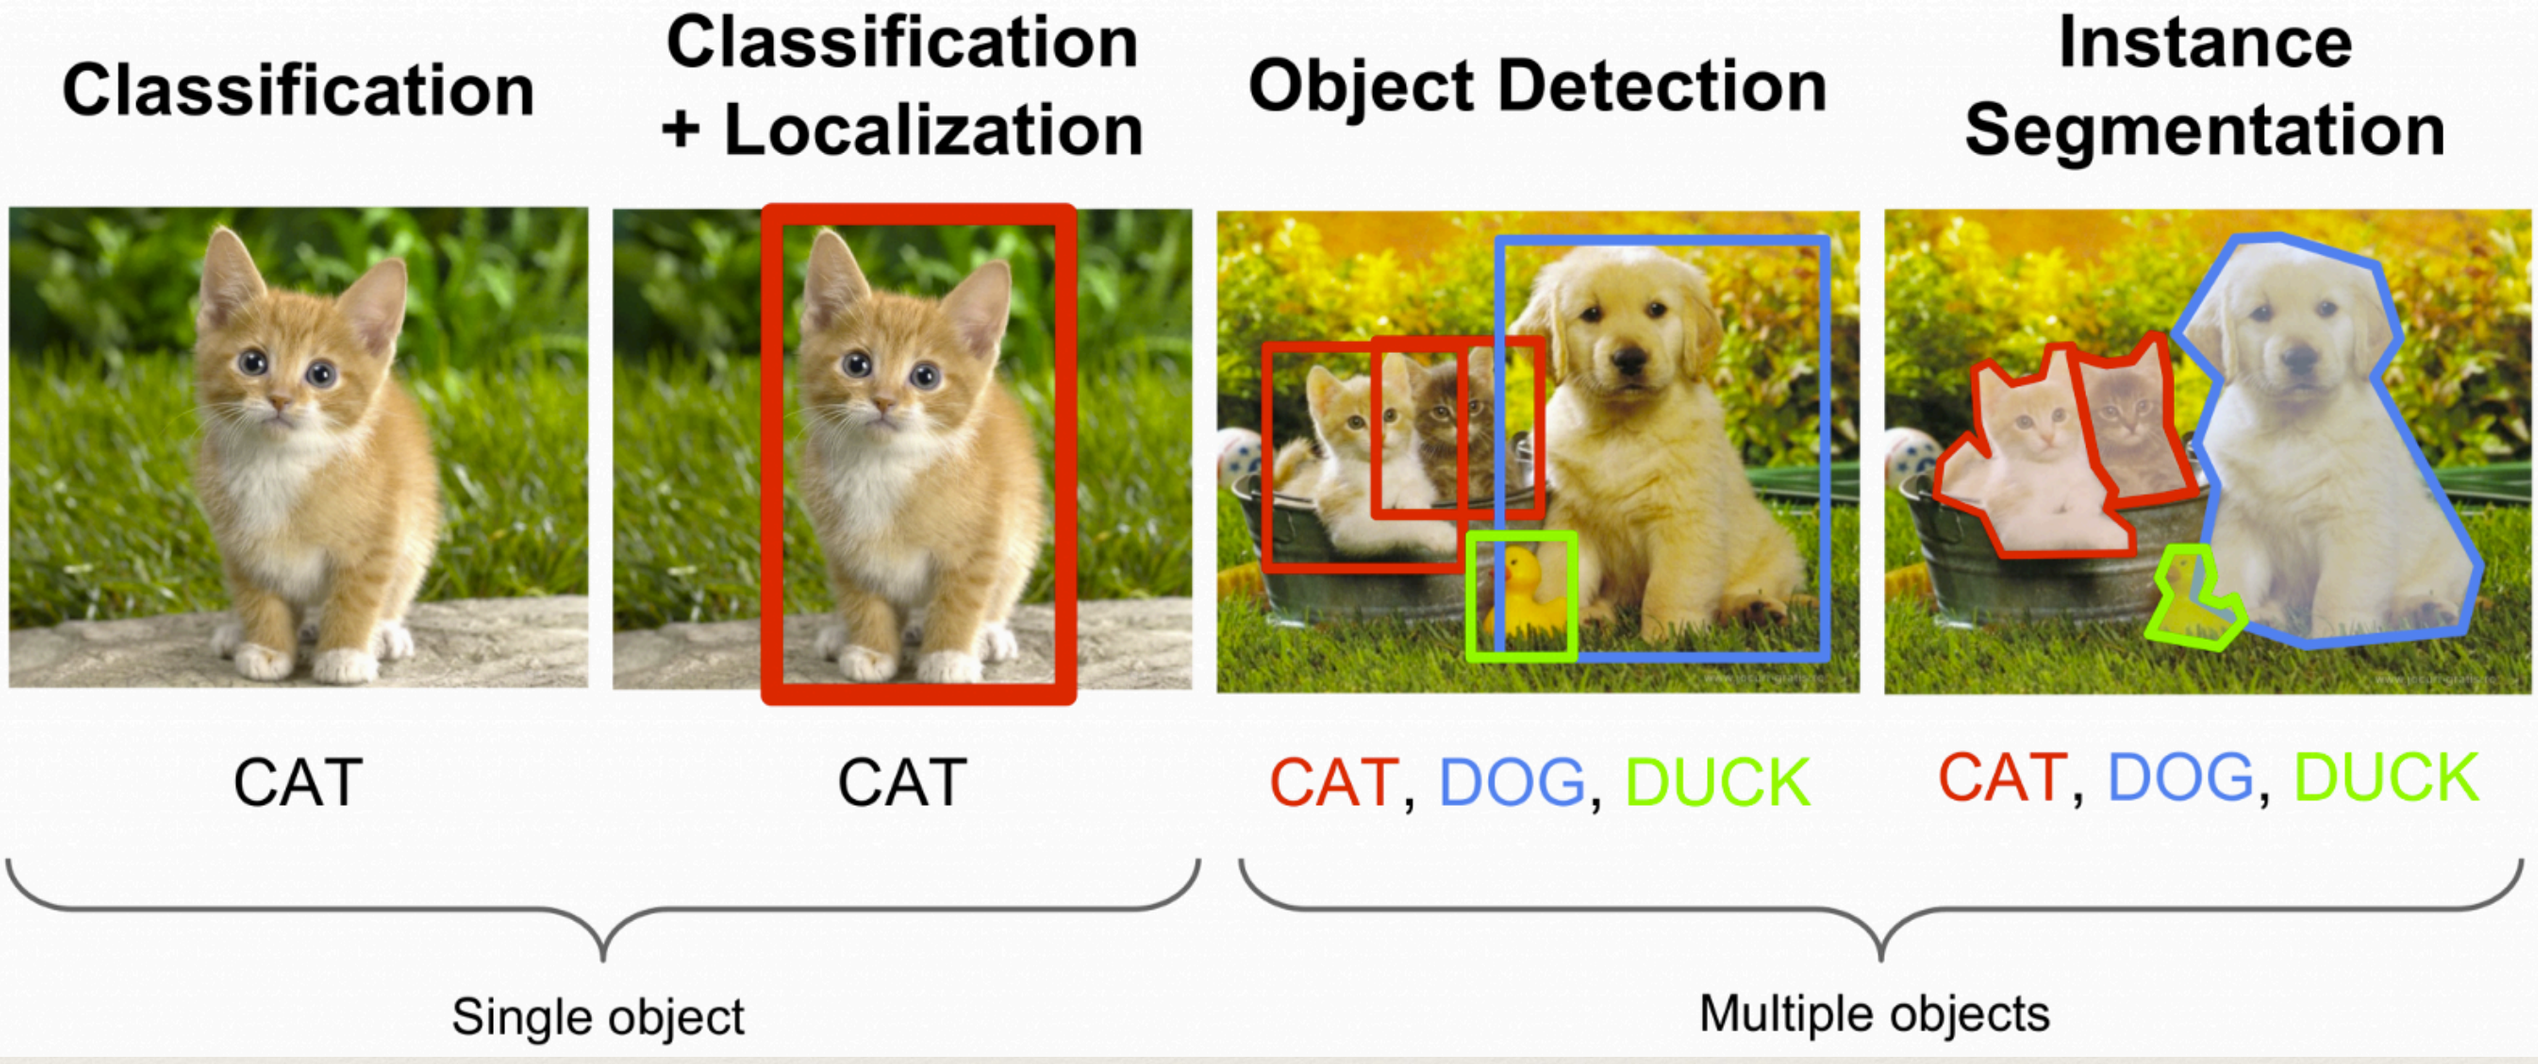

Face Detection

Siamese network

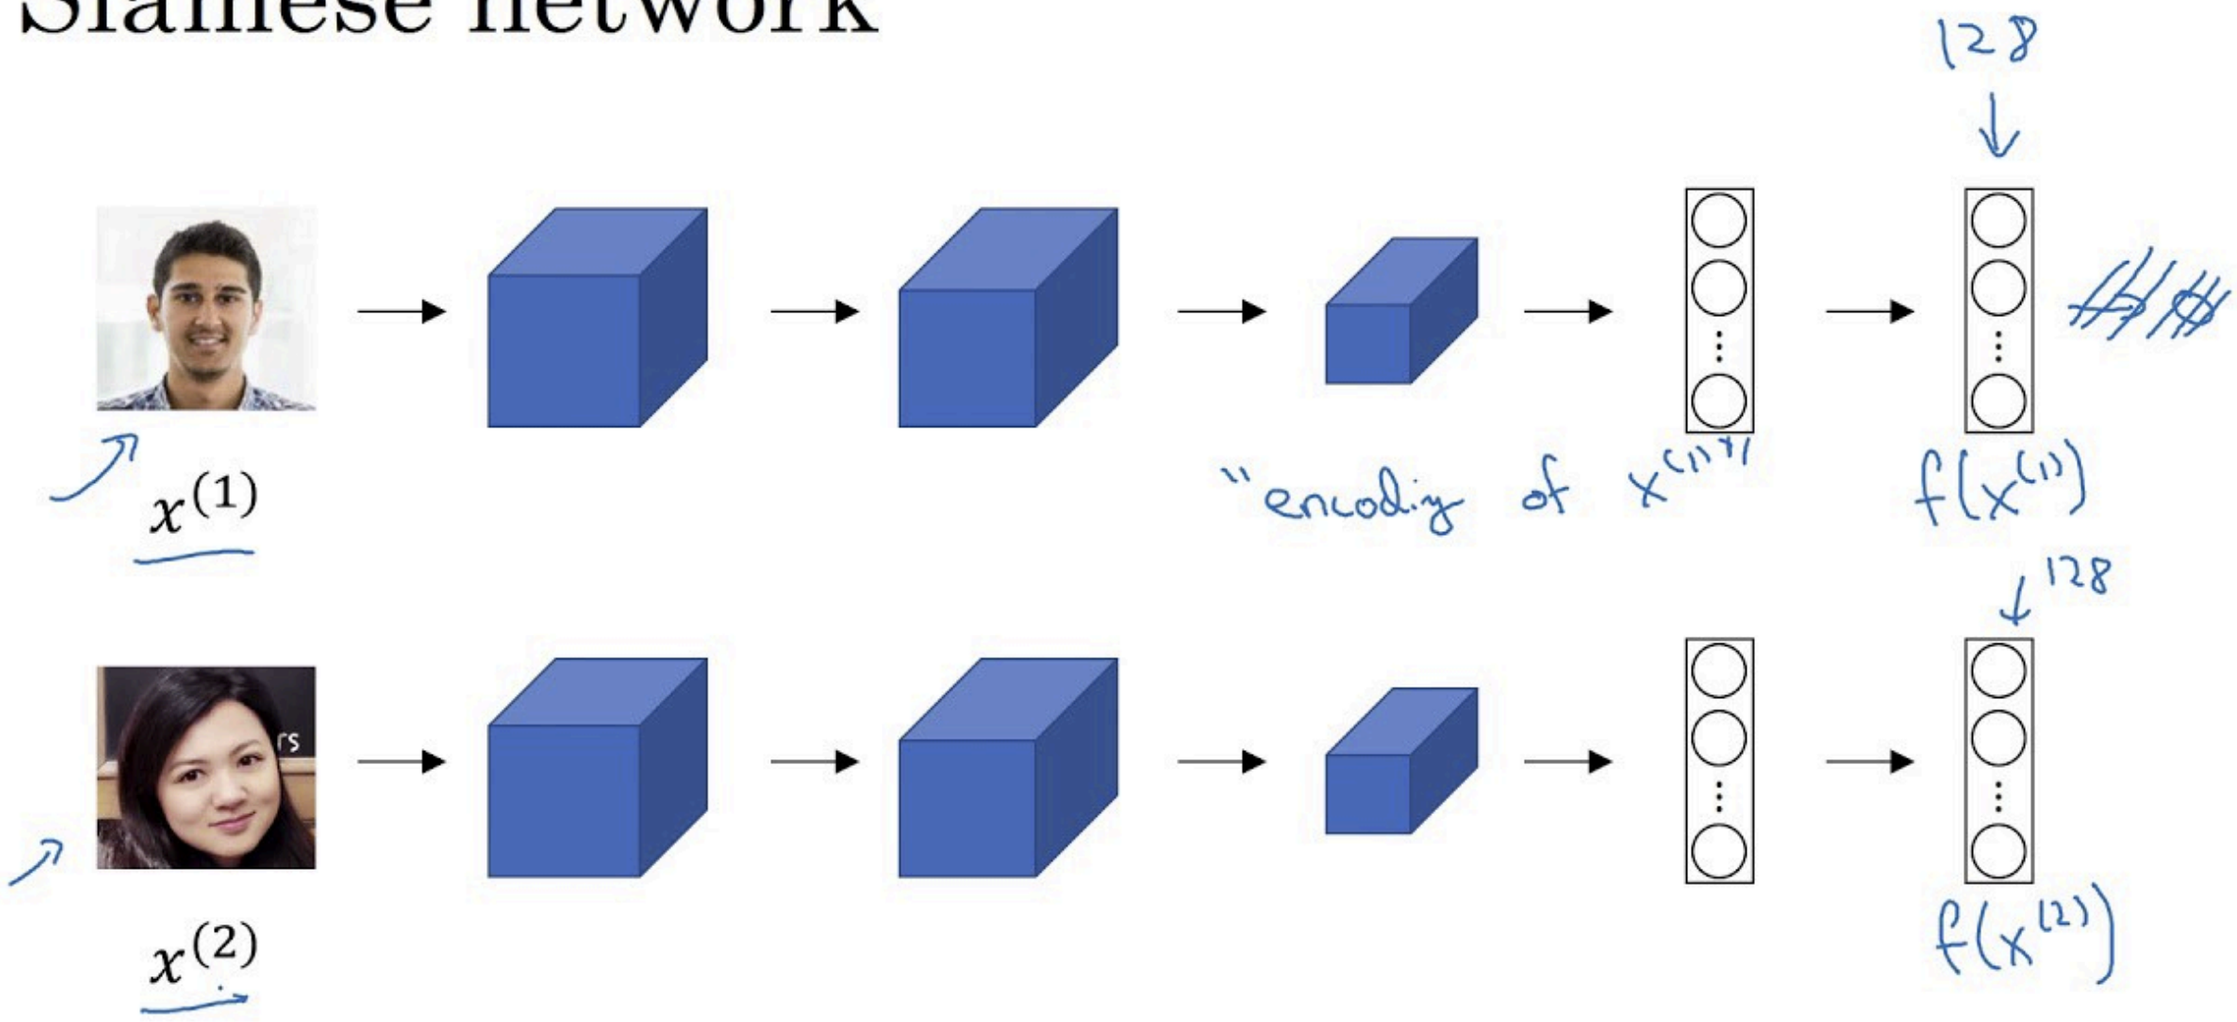

[Taigman et. al., 2014. DeepFace closing the gap to human level performance]

Architectures Commonly Used

Convolutional networks

Resnet Architectures

Future Directions

Unsupervised and Semi-supervised learning

- How little of labelled data is good enough ?
- Transfer Learning Principles

Deep Generative Models

- Using deep learning for generative modeling

Understanding And Insights Into Deep Networks

- Understanding parameters and representations.

Software Tools

Theano

http://deeplearning.net/software/theano/

TensorFlow

https://www.tensorflow.org/

PyTorch

https://pytorch.org/

Kaldi

http://kaldi-asr.org/