



Mathematical
Institute

Deep learning application highlights.

THEORIES OF DEEP LEARNING: C6.5, VIDEO 2

Prof. Jared Tanner

Mathematical Institute

University of Oxford

Oxford
Mathematics



Recent breakthroughs in deep learning applications:

Much of why we are interested in deep learning is its remarkable efficacy

Deep learning is now widespread in applications, showing remarkable abilities to perform complex tasks.

- ▶ Computer vision; image classification, Imagenet challenges.
- ▶ Complex strategy games such as Go.
- ▶ AI for art with style transfer, sound generation from videos, and text generation.
- ▶ Deep learning is now increasingly used in scientific applications.
- ▶ And many many more applications: e.g. medical diagnostics....

ImageNet Large Scale Visual Recognition Challenge (ILSVRC)

Image classification, localization, and detection:



<http://image-net.org>

ImageNet was first presented in 2009 to help benchmark image classification algorithms in the ILSVRC.

2010-14: Image classification; 1.2 million training labeled images

2011-14: Single object localisation; 524,000 training labeled bbox

2013-14: All object classification per scene; 456,000 training set

<https://link.springer.com/article/10.1007/s11263-015-0816-y>

ImageNet Large Scale Visual Recognition Challenge

Image classification, localization, and detection: complex set of similar data



“The ILSVRC dataset contains many more fine-grained classes compared to the standard PASCAL VOC benchmark; for example, instead of the PASCAL “dog” category there are 120 different breeds of dogs in ILSVRC2012-2014 classification and single-object localization tasks.”

<https://link.springer.com/article/10.1007/s11263-015-0816-y>

ImageNet Large Scale Visual Recognition Challenge

Image classification, localization, and detection

Image classification



Ground truth

Steel drum
Folding chair
Loudspeaker

Accuracy: 1

Scale
T-shirt
Steel drum
Drumstick
Mud turtle

Accuracy: 1

Scale
T-shirt
Giant panda
Drumstick
Mud turtle

Accuracy: 0

Single-object localization



Ground truth



Accuracy: 1



Accuracy: 0




Accuracy: 0


Object detection



Ground truth



AP: 1.0 1.0 1.0 1.0



AP: 0.0 0.5 1.0 0.3




AP: 1.0 0.7 0.5 0.9

2013-14: All object classification per scene; 456,000 training set
<https://arxiv.org/pdf/1409.0575.pdf>

ImageNet Large Scale Visual Recognition Challenge

Image classification, localization, and detection: "super-human performance."



<https://arxiv.org/pdf/1409.0575.pdf>

2012 ILSVRC classification won using 7 layer CNN by Krizhevsky, Sutskever, and Hinton; users in widespread use of ConvNets.

This success marked the start of DNNs widespread use.

Mastering the game of Go with deep neural networks and tree search

David Silver^{1*}, Aja Huang^{1*}, Chris J. Maddison¹, Arthur Guez¹, Laurent Sifre¹, George van den Driessche¹, Julian Schrittwieser¹, Ioannis Antonoglou¹, Veda Panneershelvam¹, Marc Lanctot¹, Sander Dieleman¹, Dominik Grewe¹, John Nham², Nal Kalchbrenner¹, Ilya Sutskever², Timothy Lillicrap¹, Madeleine Leach¹, Koray Kavukcuoglu¹, Thore Graepel¹ & Demis Hassabis¹

The game of Go has long been viewed as the most challenging of classic games for artificial intelligence owing to its enormous search space and the difficulty of evaluating board positions and moves. Here we introduce a new approach to computer Go that uses ‘value networks’ to evaluate board positions and ‘policy networks’ to select moves. These deep neural networks are trained by a novel combination of supervised learning from human expert games, and reinforcement learning from games of self-play. Without any lookahead search, the neural networks play Go at the level of state-of-the-art Monte Carlo tree search programs that simulate thousands of random games of self-play. We also introduce a new search algorithm that combines Monte Carlo simulation with value and policy networks. Using this search algorithm, our program AlphaGo achieved a 99.8% winning rate against other Go programs, and defeated the human European Go champion by 5 games to 0. This is the first time that a computer program has defeated a human professional player in the full-sized game of Go, a feat previously thought to be at least a decade away.

<https://storage.googleapis.com/deepmind-media/alphago/AlphaGoNaturePaper.pdf>

Deep learning style transfer: architecture

Deep learning allows easy combination architecture tasks.




Figure 2. Style transfer algorithm. First content and style features are extracted and stored. The style image $\vec{a}$ is passed through the network and its style representation A^l on all layers included are computed and stored (left). The content image $\vec{p}$ is passed through the network and its style features G^l and content features F^l are computed. On each layer included in the style representation, the element-wise mean squared difference between G^l and A^l is computed to give the style loss $\mathcal{L}_{style}$ (left). Also the mean squared difference between F^l and P^l is computed to give the content loss $\mathcal{L}_{content}$ (right). The total loss $\mathcal{L}_{total}$ is then a linear combination between the content and the style loss. Its derivative with respect to the pixel values can be computed using error back-propagation (middle). This gradient is used to iteratively update the image $\vec{x}$ until it simultaneously matches the style features of the style image $\vec{a}$ and the content features of the content image $\vec{p}$ (middle, bottom).

More complex architectures can allow learning and transferring characteristics of objects.

https://www.cv-foundation.org/openaccess/content_cvpr_2016/papers/Gatys_Image_Style_Transfer_CVPR_2016_paper.pdf

Deep learning style transfer: examples

Applications such as these seem impossible without DL



Photos can be transitioned to paintings with prescribed styles

https://www.cv-foundation.org/openaccess/content_cvpr_2016/papers/Gatys_Image_Style_Transfer_CVPR_2016_paper.pdf

Deep learning sound generation from video

DL is able to generate realistic synthetic data.




Figure 1: We train a model to synthesize plausible impact sounds from silent videos, a task that requires implicit knowledge of material properties and physical interactions. In each video, someone probes the scene with a drumstick, hitting and scratching different objects. We show frames from two videos and below them the predicted audio tracks. The locations of these sampled frames are indicated by the dotted lines on the audio track. The predicted audio tracks show seven seconds of sound, corresponding to multiple hits in the videos.

Expected sound characteristics can be learned from video with sound and then generated and added to video lacking sound.

<https://arxiv.org/pdf/1512.08512.pdf>

<https://www.youtube.com/watch?v=0FW99AQmMc8>

Many similar examples exist, see e.g. *DeepFakes* and automatic text generations. <https://en.wikipedia.org/wiki/GPT-3>

DeepLens for detecting rare events

Some tasks are not scalable without machine learning.

For many science applications the quantity of data is beyond human inspection, use DL: Large Synoptic Survey Telescope
<https://academic.oup.com/mnras/article/473/3/3895/3930852>

One approach to scale the visual inspection effort to the size of these surveys is to use crowdsourcing. This is the idea behind the Space Warps project (Marshall et al. 2015; More et al. 2015), which crowdsourced the visual inspection of a sample of 430 000 images from the CHFTLS to a crowd of 37 000 citizen scientists, yielding a new sample of gravitational lens candidates. The authors further estimate that a similar crowdsourcing effort can be scaled up to LSST sizes, where a considerable crowd of 10^6 volunteers could visually inspect 10^6 LSST targets in a matter of weeks.

Is your area of interest immune from deep learning?

Deep deep trouble: A case study by Michael Elad for image denoising:

Thousands of papers addressing this fundamental task [image denoising] were written over the years. Researchers developed beautiful and deep mathematical ideas with tools from partial differential equations, such as anisotropic diffusion and total variation, energy minimization viewpoint, adoption of a geometric interpretation of images as manifolds, use of the Beltrami flow, and more.... We have hence gained vast knowledge in image processing over the past three decades.

In 2012, Harold Burger, Christian Schuler, and Stefan Harmeling decided to throw deep learning into this problem. The idea was conceptually quite simple: take a huge set of clean images, add synthetic noise, and then feed them to the learning process that aims to turn a noisy image into its clean version. While the process was tedious, frustrating, and lengthy — the end result was a network that performed better than any known image denoising algorithm at that time.

<https://sinews.siam.org/Details-Page/deep-deep-trouble>