



So what is this thing called “Deep Learning” anyway?

Mark Huckvale


Speech, Hearing and Phonetic Sciences

Deep Learning





10.9m



45.7m



68.8m




78.8m



80.4m




Deep Learning Perception
Distance Detection



Speech and Language


- Image captioning
 - <https://www.captionbot.ai/>
- Machine translation
 - <https://translate.google.com/>
- Spoken dialogue
 - <https://ai.googleblog.com/2018/05/duplex-ai-system-for-natural-conversation.html>

1950s: Perceptron




$$y = \varphi(\vec{w} \cdot \vec{x} + b)$$

1980s: Multi-Layer Perceptron



“Back propagation of error”
Rumelhart, Hinton, Williams, 1986

2010s: Deep Neural Networks



Christian Szegedy, et al., Going Deeper with Convolutions, 2014

DNN Synergy

- improved optimisation algorithms
- more machine-readable data
- faster computers (parallel computing)
- open-source data, tools, models & recipes
- open competitions with standard training data, test data and evaluation criteria
 - <https://www.kaggle.com/competitions>


Deep Learning Characteristics

- Tackles real-world problems
- Relies less on human domain expertise
- “End-to-end” processing
- Exploits a hierarchy of abstraction
- State of the art performance on many tasks


DL Problem Types

- Classification
 - image recognition, sentiment analysis
- Regression
 - house prices, power consumption
- Planning
 - autonomous driving, game playing
- Modelling
 - word semantics, language modelling, time series prediction
- Transformation
 - language translation, speech recognition, speech synthesis



Dense Network



Convolutional Network




Recurrent Network



Gated Recurrent Unit

Adversarial Network




DL Tools

- Python programming language
- TensorFlow machine learning toolkit
- Jupyter project notebooks



TensorFlow

$e = (a+b) * (b-1)$



```
a = tf.placeholder(tf.float32)
b = tf.placeholder(tf.float32)
c = tf.add(a, b)
d = tf.sub(b, 1)
e = tf.mul(c, d)
```

Functions (i.e. networks) are expressed as computational graphs

TensorFlow

- Computational instructions as graph
- Matches neural network computation
- Allows for distributed and parallel computation
- Automatic differentiation

$$y = f(x) \rightarrow dy/dx = f'(x)$$

- High-level toolkits




DL Example

- Classification of hand-written digits (MNIST)
- [Google Colaboratory](#)
 - Jupyter notebook hosted by Google backed by GPU



DL Getting Started

- Python
 - Lynda.com
- Machine Learning
 - Géron, Hands-on machine learning with Scikit-Learn and TensorFlow, O'Reilly 2017.
- Tutorials
 - Hvass-Labs TensorFlow-Tutorials



Artificial General Intelligence

- Does deep-learning solve the problem of AGI?
- Is deep learning a step in the direction of AGI?
- Deep learning is inspiring advances in:
 - multi-task learning
 - long-term memory
 - logical inference
 - on-line learning

"The role of FAIR is to advance the science and the technology of AI and do experiments that demonstrate that technology for new applications like computer vision, dialogue systems, virtual assistants, speech recognition, natural language understanding, translation, things like that," Yann Lecun, Chief AI Scientist, Facebook, 2016.

Summary

- Deep learning is a practical technology for machine learning providing state of the art performance on many tasks
- Deep learning does not of itself solve AI
- But its success encourages further research into the challenges of AI