

Topics and Reading Material

Note most of the papers can be downloaded from <http://deeplearning.net/reading-list/>. Others are freely available on arxiv.org or similar websites. If you can not access a paper, please let me know.

Note that the suggested papers and topics are only a recommendation. You are neither expected to read all of the papers in detail nor do you need to include all the information contained in the papers / all the suggested topics in your talk. I am aware that you won't be able to explain all the details in 25 minutes. You should focus your talk on few aspects which you find important. Of course you are free to use additional material.

1. Deep Learning Overview

Suggested Papers

- Arel et al.: "Deep Machine Learning—A New Frontier in Artificial Intelligence Research"
- Schmidhuber: "Deep Learning in Neural Networks: An Overview"
- Bengio et al.: "Representation Learning: A Review and New Perspectives"

Suggested Topics

- History of Neural Networks / Deep Learning
- Mainstream Deep Learning approaches (ConvNets, DBNs, Stacked Autoencoders etc.)
- What applications benefit from deep learning / success stories of deep learning
- Future of Deep Learning

2. Deep Learning vs. Shallow Learning

Suggested Papers

- Bengio: "Learning Deep Architectures for AI"
- Ba & Caruana: "Do Deep Nets Really Need to be Deep"
- Bengio et al.: "Representation Learning: A Review and New Perspectives"

Suggested Topics

- Theoretical Advantages of Deep Architectures
- Comparison with Kernel Methods
- (Dis)Advantages in practice (when do deep architectures work, when do they fail ?)
- Do Deep Nets Really Need to be Deep ?

3. How to train a DNNs (tips & tricks)

Suggested Papers

- Erhan et al.: "Why Does Unsupervised Pre-training Help Deep Learning?"
- Glorot & Bengio: "Understanding the difficulty of training deep feedforward neural networks"
- Bengio et al.: "Practical Recommendations for Gradient-Based Training of Deep Architectures"
- Hinton et al.: "Improving neural networks by preventing co-adaptation of feature detectors"

Suggested Topics

- Training a DNN (Basics)
- Pre-training, Dropout and other tricks
- Practical Recommendations, why are GPUs so important

4. Understanding / Interpreting DNNs

Suggested Papers

- Zeiler & Fergus: "Visualizing and Understanding Convolutional Networks"
- Mahendran & Vedaldi: Understanding Deep Image Representations by Inverting Them
- Simonyan et al.: "Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps"

Suggested Topics

- Interpreting DNNs
- Saliency Maps
- Possible Applications

5. Deep Learning in image classification

Suggested Papers

- Szegedy et al.: "Going deeper with convolutions"
- Krizhevsky et al.: "Imagenet classification with deep convolutional neural networks"
- Howard: "Some Improvements on Deep Convolutional Neural Network Based Image Classification"
- Serre et al.: "Robust object recognition with cortex-like mechanisms"

Suggested Topics

- Why convolutional nets work well for image classification
- How many layers do we need, what types of layers are used ?
- How to train a DNN for image classification in practice ?
- Can we related DNNs to cortex-like mechanisms ?

6. Deep Learning for object detection

Suggested Papers

- Sermanet et al.: "Overfeat: Integrated recognition, localization and detection using convolutional networks"
- Szegedy et al.: "Deep neural networks for object detection."
- Erhan et al.: "Scalable object detection using deep neural networks"

Suggested Topics

- Object Detection / Localization
- Sliding Window Efficiency
- How to train a DNN for object detection in practice ?

7. Deep Learning in speech recognition

Suggested Papers

- Deng et al.: "Ensemble Deep Learning for Speech Recognition"
- Hannun et al.: "Deep Speech: Scaling up end-to-end speech recognition"
- Hinton et al.: "Deep neural networks for acoustic modeling in speech recognition"

Suggested Topics

- Different models for speech recognition
- Practical Recommendations

8. Deep Learning in text analysis

Suggested Papers

- Mikolov et al.: "Distributed Representations of Words and Phrases and their Compositionality"
- Sutskever et al.: "Sequence to Sequence Learning with Neural Networks"
- Bengio et al.: "A neural probabilistic language model"
- Denil et al.: "Modelling, Visualising and Summarising Documents with a Single Convolutional Neural Network"

Suggested Topics

- Word Representation
- Document Classification
- Practical Recommendations

9. Deep Learning in video analysis

Suggested Papers

- Simonyan & Zisserman: "Two-Stream Convolutional Networks for Action Recognition in Videos"
- Ji et al.: "3D convolutional neural networks for human action recognition"
- Karpathy et al.: "Large-scale video classification with convolutional neural networks"

Suggested Topics

- Human action recognition
- Different types of architectures
- How to use temporal information in video
- How to train a DNN for video analysis in practice ?

10. Deep Learning in fMRI analysis

Suggested Papers

- Hatakeyama et al.: "Multi-Voxel Pattern Analysis of fMRI Based on Deep Learning Methods"
- Plis et al.: "Deep learning for neuroimaging: a validation study"
- Koyamada et al.: "Construction of Subject-independent Brain Decoders for Human FMRI with Deep Learning"

Suggested Topics

- Feature Learning in fMRI
- Decoding with Deep Learning

- (Dis)Advantages compared to linear methods ? Interpretability ?
- Is Deep Learning also applicable to EEG?

11. Deep Boltzmann Machines

Suggested Papers

- Salakhutdinov & Hinton: "Deep Boltzmann Machines"
- Le Roux & Bengio: "Representational Power of Restricted Boltzmann Machines and Deep Belief Networks"
- Montavon & Müller: "Deep Boltzmann Machines and the Centering Trick"
- Srivastava et al.: "Multimodal learning with deep boltzmann machines"

Suggested Topics

- Application scenarios of DBM
- Training a DBM

12. Recurrent Nets

Suggested Papers

- Hochreiter et al.: "Gradient Flow in Recurrent Nets: the difficulty of learning long-term dependencies"
- Pascanu et al.: "How to Construct Deep Recurrent Neural Networks"
- Graves et al.: "Speech recognition with deep recurrent neural networks"
- Sutskever et al.: "Generating text with recurrent neural networks"

Suggested Topics

- When do we need recurrent nets
- How to train recurrent nets
- Application example