

Guided Tour to Deep Learning Using MATLAB

© Oge Marques, PhD – 2016-2017

Introduction

- This document guides you through several tutorials, papers, and resources related to Deep Learning (for computer vision).
- It assumes no prior exposure to Deep Learning, but it is targeted at readers with basic knowledge of image processing / computer vision / machine learning and prior experience with MATLAB and some of its toolboxes.
- It is structured as a step-by-step guide. It is best that you follow it in the intended sequence.

Part 1- The big picture

1. Read the paper: LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436-444¹.
2. Read Chapter 1 of the book "Deep Learning" by Ian Goodfellow and Yoshua Bengio and Aaron Courville: <http://www.deeplearningbook.org/contents/intro.html>

Part 2- Fundamentals

1. Read Chapter 1 of "*Neural Networks and Deep Learning*" by Michael Nielsen². Ignore (for now) the exercises and don't worry about downloading and running the example code (in Python).
2. Read Chapter 6 of "*Neural Networks and Deep Learning*" by Michael Nielsen. Ignore (for now) the exercises and don't worry about downloading and running the example code (in Python).

¹ <http://www.nature.com/nature/journal/v521/n7553/full/nature14539.html>

² <http://neuralnetworksanddeeplearning.com/>

Part 3- Deep Learning and MATLAB

1. Go to: <http://www.mathworks.com/campaigns/products/artifacts/deep-learning.html> for an overview of the steps within this part of the guided tour.
2. (OPTIONAL) Watch the (first) video(s) from the series at the following YouTube playlist: <https://www.youtube.com/playlist?list=PLjJh1vISEYgvGod9wWiydumYl8hOXixNu>
3. Skip (for now, at least) Andrew Ng's "Unsupervised Feature Learning and Deep Learning" tutorial: http://ufldl.stanford.edu/wiki/index.php/UFLDL_Tutorial
4. (OPTIONAL) Watch the 27-min "Object Recognition: Deep Learning and Machine Learning for Computer Vision" video (archived webinar): <https://goo.gl/yiakuj>

Don't forget to download the associated source code: <https://goo.gl/d4oeGR>

5. (OPTIONAL) Read the article (and try the code³) at: <http://www.mathworks.com/company/newsletters/articles/deep-learning-for-computer-vision-with-matlab.html>
6. Try the "Image Category Classification Using Deep Learning" example at: <https://goo.gl/tkHRSP>
7. Try the "Training a Deep Neural Network for Digit Classification" example at: <https://goo.gl/EpJucE>
8. (OPTIONAL) Try the "Construct Deep Network Using Autoencoders" example at: <https://goo.gl/GRnkXe>

³ The code (available at <http://www.mathworks.com/matlabcentral/fileexchange/57116-deep-learning-for-computer-vision-demo-code>) is clunky, has dependencies on local files/folders, and relies on images and videos that you'd have to find/create separately.

Useful links⁴:

- Deep Learning Portal: <http://deeplearning.net/>
- A curated list of deep learning resources for computer vision:
<https://github.com/kjw0612/awesome-deep-vision>
- [Deep Learning](#) by Ian Goodfellow and Yoshua Bengio and Aaron Courville (an MIT Press book): <http://www.deeplearningbook.org/>
- UFLDL Tutorial: <http://ufldl.stanford.edu/tutorial/>
- UFLDL Tutorial: http://deeplearning.stanford.edu/wiki/index.php/UFLDL_Tutorial
- "Deep Learning for Computer Vision Barcelona" Summer seminar UPC TelecomBCN (July 4-8, 2016): <https://imatge-upc.github.io/telecombcn-2016-dlcv/>
- Hands-on Deep Learning using the [MatConvNet](#) framework:
<http://www.cvc.uab.es/~gros/index.php/hands-on-deep-learning-with-matconvnet/>
- ML/DL resources: <http://goo.gl/lzkDEa>
- Deep Learning talks, lectures, tutorials (video recordings):
<http://www.computervisiontalks.com/tag/deepUlearning/>
- Overview article (MIT Technology Review, 2013):
<http://www.technologyreview.com/featuredstory/513696/deepUlearning/>
- Overview paper (Nature, May 2015):
<http://www.nature.com/nature/journal/v521/n7553/full/nature14539.html>
- Tutorial (Intl. Conference in Computer Vision (CVPR) 2014):
<https://sites.google.com/site/deeplearningcvpr2014/>
- Deep Vision (*Deep Learning in Computer Vision* workshop, June 2015):
<http://www.deepUvision.net/>
- ImageNet (<http://imageUnet.org/>) and its Large Scale Visual Recognition Challenge 2015 (ILSVRC2015) (<http://imageUnet.org/challenges/LSVRC/2015/index>)
- Project Adam (Microsoft Research): <http://research.microsoft.com/enUs/news/features/dnnvisionU071414.aspx>
- Andrej Karpathy's blog: <https://karpathy.github.io/>
- Tom Stone's Computer Vision Blog: <http://www.computervisionblog.com/>
- Facebook AI research: <https://research.facebook.com/ai>
- "The Rapid Rise of Deep Learning Computer Vision Technology": recent (June 2015) news story with several useful links and (links to) demos and APIs: <http://goo.gl/G7gkbb>

⁴ In no particular order.

- ConvJS (a JavaScript library for training Deep Learning models entirely in your browser!): <https://cs.stanford.edu/people/karpathy/convnetjs/>
- Caffe: <http://caffe.berkeleyvision.org/>
- Torch: <http://torch.ch/>
- "Training a Deep Neural Network for Digit Classification" (Mathworks): <https://goo.gl/IIQBkS>
- "Deep Learning Toolbox" for MATLAB by [Rasmus Berg Palm](#): <https://www.mathworks.com/matlabcentral/fileexchange/38310-deep-learning-toolbox>
- "Neural Networks and Deep Learning" by Michael Nielsen (free online book, examples in Python): <http://neuralnetworksanddeeplearning.com/index.html>