Curriculum INF5860/INF9860 2018

Lecture notes, **exercises and mandatory exercises** define the curriculum.

Lecture 1: Introduction to convolution

Know and be able to compute by hand the simple operation that convolution of a given filter with an image gives. Know mean filters and edge detection filters.

Lecture 2: KNN, regression, cost functions and simple gradient descent

Overview:

- KNN-classification, distance and how to selct k
- The linear regression problem and cost function
- Gradient descent
- Logistic regression including cost function

Readings:

Text:

- http://cs229.stanford.edu/notes/cs229-notes1.pdf (pages 1-7 and 16-19)
- http://cs229.stanford.edu/notes/cs229-notes1.pdf (section 1-3)
- Deep Learning book Chapter 4.3

Video:

- https://www.youtube.com/watch?v=OoUX-nOEjG0&t=0s&list=PL3FW7Lu3i5JvHM8ljYj-zLfQRF3EO8sYv&index=2

Lecture 3: Softmax classification and regularization

Remark: SVM is not part of the curriculum.

Overview:

- Multiclass logistic regression and softmax classification
- Cost functions
- The need for regularization and simple L2 regularization
- Know equations behind these classifiers.

Readings:

Text:

- http://cs231n.github.io/linear-classify/

Video:

- https://www.youtube.com/watch?v=OoUX-nOEjG0&list=PL3FW7Lu3i5JvHM8ljYj-zLfQRF3EO8sYv
- https://www.youtube.com/watch?v=h7iBpEHGVNc&list=PL3FW7Lu3i5JvHM8ljYj-zLfQRF3EO8sYv&index=2

Lecture 4: Dense neural network classifiers

Everything covered in the lecture is considered curriculum, this includes the derivation of the various results.

Overview:

- Dense neural network structure
- Forward propagation
- ReLU and sigmoid activation functions (and their derivative)
- The softmax function and its derivative
- Cross entropy cost function
- Stochastic gradient descent optimization function
- Backward propagation
- Vectorization over multiple nodes and examples.

Readings

The lecture slides should be self-contained, but more information can be found in

Deep Learning Book:

- Chapter 5.9
- Chapter 6

Lecture 5: Introduction to TensorFlow

All material covering in the lecture is considered curriculum. It is not expected from you to remember detailed syntax, but you should be able to reason about the structure of TensorFlow.

Overview:

- TensorFlow graphs
- TensorFlow session
- TensorFlow constants
- TensorFlow variables
- TensorFlow feeding data to the graph
- Tensorboard
- TensorFlow save/restore models (weights)

Readings:

- https://www.tensorflow.org/
- https://web.stanford.edu/class/cs20si/2017/syllabus.html [lecture notes/slides 1-5]
- https://www.youtube.com/channel/UCMq6IdbXar_KtYixMS_wHcQ/videos

Lecture 6: Convolutional neural networks

Pages from the lecture notes which are NOT curriculum: 80

Overview:

- Challenges with image classification
- Fully connected neural network on images
- Convolutional layer
- Convolutional layer hyperparameters
- Convolutional layer example
- Receptive field (Field of View)
- Dilated convolutions
- Pooling
- Visualizing and Understanding CNN
- Applications were CNN are used

Readings:

Text:

- http://cs231n.github.io/convolutional-networks/

Video:

- https://www.youtube.com/watch?v=bNb2fEVKeEo&index=5&list=PLC1qU-LWwrF64f4QKQTVg5Wr4qEE1Zxk

Optional text:

- Receptive field: http://www.cs.toronto.edu/~wenjie/papers/nips16/top.pdf
- Visualizing and Understanding CNN: https://cs.nyu.edu/~fergus/papers/zeilerECCV2014.pdf
- Dilated convolutions: https://arxiv.org/abs/1511.07122

Optional videos:

- https://www.youtube.com/watch?v=ghEmQSxT6tw
- https://www.youtube.com/watch?v=SQ67NBCLV98

Lecture 7: Training a neural net

Overview:

- Weight initialization
- Data input normalization
- Batchnorm
- Activation functions
- Monitoring the training process

- Variations of gradient descent, including momentm, RMSprop og ADAM

Readings:

Text:

- http://cs231n.github.io/neural-networks-2/ (Data scaling, weight initialization, batchnorm)
- http://cs231n.github.io/neural-networks-3/ (Monitoring the loss, parameter update schemes)
- Deep Learning 6.2.2 and 6.3 on activation functions
- Deep Learning 8.7.1 on Batchnorm

Video:

- https://www.youtube.com/watch?v=wEoyxE0GP2M&index=7&list=PL3FW7Lu3i5JvHM8ljYj-zLfQRF3EO8sYv&t=0s

Lecture 8: Generalization

All material covering in the lecture is considered curriculum except "Growth function" and "The Vapnik-Chervonenkis Inequality" (Pages 37-42).

Overview:

- Is learning feasible?
- Model complexity
- Overfitting
- Evaluating performance
- Learning from small datasets
- Rethinking generalization
- Capacity of dense neural networks

Readings:

Learning theory (caltech course):

- https://work.caltech.edu/lectures.html
- Lecture (Videos): 2,5,6,7,8,11

Read: CS231n: section "Dropouts"

- http://cs231n.github.io/neural-networks-2/

Optional:

- Read: The Curse of Dimensionality in classification
    - http://www.visiondummy.com/2014/04/curse-dimensionality-affectclassification/
- Read: Rethinking generalization
    - https://arxiv.org/pdf/1611.03530.pdf

Lecture 9: CNN Architectures

It is expected that you know the main gist of the covered networks, what ideas they introduced, and why they are successful. It is not expected that you should know every small detail about the network or training procedure.

Overview

- LeNet: General architecture
- AlexNet: General architecture
- VGG: General architecture
- GoogLeNet: How inception modules are built up, and how they are used in the network.
- ResNet: General architecture. Residual learning, residual blocks and the motivation. Familiarize yourself with the theories on why they are effective.

Readings

Links to every relevant publication is found in the lecture notes and can be useful to look at, but they will cover more detail than necessary. It is not expected to know more than what is covered by the lecture slides.

Lecture 10: Segmentation and object detection

This lecture also covers some specific network models, and the same remarks on lecture 9 also applies to this. We are not interested in to much network specific details, but more of a general overview.

Overview

- Evaluation metrics
- Object localization vs object detection
- R-CNN: General architecture, pros and cons
- Fast R-CNN: General architecture, pros and cons
- Faster R-CNN: General architecture, pros and cons
- YOLO: General architecture, pros and cons
- Semantic segmentation vs instance segmentation
- General idea of segmentation networks
- Upsampling with transposed convolution
- Upsampling with dilated convolution
- Why we use conditional random fields (but not details of how they work)

Readings

References to the relevant architecture is found in the lecture slides. But any detail extending the lecture slides is not curriculum.

Lecture 11: Visualization and understanding

Tiny details in the methods are not required, but should know the main purpose and main idea behind the different methods coverd.

Overview:

- Visualization of filters
- Guided backprop
- Occlusion experiments
- Saliency maps, gradcam
- Visualizing filters by gradient ascent
- Briefly about fooling the network
- Feature inversion
- Neural style transfer

Reading:

Text:

- https://blog.keras.io/how-convolutional-neural-networks-see-the-world.html
- A link to selected research papers is given in the lecture slides. These help understanding more, but are not part of the curriculum, but useful when you apply visualization.

Video:

- https://www.youtube.com/watch?v=6wcs6szJWMY&list=PL3FW7Lu3i5JvHM8ljYj-zLfQRF3EO8sYv&index=12

Lecture 12: Recurrent neural networks

Pages from the lecture notes which are NOT curriculum: 41, 48, 78, 81, 82, 84, 85

Overview:

- Overview (Feed forward and convolution neural networks)
- Vanilla Recurrent Neural Network (RNN)
- Input-output structure of RNN's
- Training Recurrent Neural Networks
- Simple examples
- Advanced models
- Advanced examples

Readings:

Video:

- Cs231n: lecture 10 | Recurrent Neural Networks.
  - https://www.youtube.com/watch?v=6niqTuYFZLQ&list=PLzUTmXVwsnXod6WNdg57Yc3zFx_f-RYsq&index=10
- Cs231n (2016): Lecture 14: First part "ConvNet for Video" to slides 38
  - https://www.youtube.com/watch?v=ekyBklxwQMU&list=PLkt2uSq6rBVctENoVBg1TpCC7OQi31AlC&index=13

Read:

- The Unreasonable Effectiveness of Recurrent Neural Networks
  - http://karpathy.github.io/2015/05/21/rnn-effectiveness/

Optional:

- Video: cs224d L8:
  - https://www.youtube.com/watch?v=Keqep_PKrY8&list=PL3FW7Lu3i5Jsnh1rnUwq_TcylNr7EkRe6&index=8
- Video: cs224d L9:
  - https://www.youtube.com/watch?v=QuELiw8tbx8&index=9&list=PL3FW7Lu3i5Jsnh1rnUwq_TcylNr7EkRe6
- Read: Deep learning book: CH10 (10.1-10.2)
  - http://www.deeplearningbook.org/contents/rnn.html

Lecture 13: Unsupervised learning

The most important subjects are t-SNE, autoencoders, and variational autoencoders. K-means clustering and PCA are assumed known and considered background knowledge.

Overview

- K-means clustering
- Principal Component Analysis (PCA)
- SNE and t-SNE
- Autoencoders
- Variational autoencoders

Readings

Visualizing data using t-SNE

Laurens van der Maaten, Geoffrey Hinton

JMLR (2008)

http://www.jmlr.org/papers/volume9/vandermaaten08a/vandermaaten08a.pdf

- Section 1, 2, and 3.

Deep Learning Book

- Chapter 14.1, 14.2, and 14.3

Tutorial on Variational Autoencoders, Carl Doersch

`https://arxiv.org/pdf/1606.05908.pdf`

- Chapter 1 (introduction) and chapter 2.

Lecture 14: Generative adversarial networks

It is important to understand the basic idea behind GANs. Variants such as DCGAN and Wasserstein GAN were only briefly mentioned, and details about them are not part of the curriculum

Overview

- Understanding the basic idea behind GANs
- The role of the discriminator and the generator
- The discriminator and generator cost function, and how they are used together
- Challenges with GAN, especially mode collapse
- Adversarial domain adaption: the general idea

Readings

NIPS 2016 Tutorial: Generative Adversarial Networks

Ian Goodfellow

https://arxiv.org/pdf/1701.00160.pdf

Selected sections with varying importance (only the sections listed below are curriculum)

- Introduction, chapter 1, and chapter 2 is considered background
- Chapter 3 up to (including) 3.2 is important
- Chapter 4, background
- Chapter 5.1 is important
- Rest of chapter 5 is considered background

Lecture 15: Reinforcement learning

Pages from the lecture notes which are NOT curriculum: 59-62, 77-78, 81-82.

Overview:

- The RL notation and terminology.
- Value function based methods (Q-learning)
- Policy based methods (policy gradients)

Readings:

Video:

- CS231n: Lecture 14 | Deep Reinforcement Learning: https://www.youtube.com/watch?v=lvoHnicueoE&index=14&list=PLC1qULWwrF64f4QKQT-Vg5Wr4qEE1Zxk&t=3s

Text:

- Karpathy blog: (Reinforcement learning/Policy learning) http://karpathy.github.io/2016/05/31/rl/

Optional:

- RL Course by David Silver: https://www.youtube.com/watch?v=2pWv7GOvuf0&list=PL7-jPKtc4r78- wCZcQn5IqyuWhBZ8fOxT&index=0