Deep-Learning-Roadmap Documentation

Release 1.0

Amirsina Torfi

Apr 21, 2020

Foreword

CHAPTER 1

Introduction

The purpose of this project is to introduce a shortcut to developers and researcher for finding useful resources about Deep Learning for Natural Language Processing.

1.1 Motivation

There are different motivations for this open source project.

1.1.1 What's the point of this open source project?

There other similar repositories similar to this repository and are very comprehensive and useful and to be honest they made me ponder if there is a necessity for this repository!

The point of this repository is that the resources are being targeted. The organization of the resources is such that the user can easily find the things he/she is looking for. We divided the resources to a large number of categories that in the beginning one may have a headache!!! However, if someone knows what is being located, it is very easy to find the most related resources. Even if someone doesn't know what to look for, in the beginning, the general resources have been provided.

CHAPTER 2

Papers

This chapter is associated with the papers published in deep learning.

2.1 Models

2.1.1 Convolutional Networks

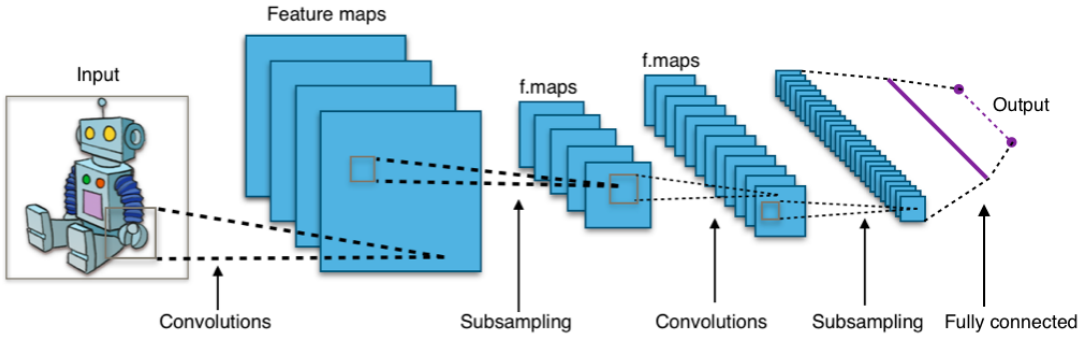

- **Imagenet classification with deep convolutional neural networks** : [Paper]

- **Convolutional Neural Networks for Sentence Classification** : [Paper]

- **Large-scale Video Classification with Convolutional Neural Networks** : [Paper]

- **Learning and Transferring Mid-Level Image Representations using Convolutional Neural Networks** : [Paper]

- **Deep convolutional neural networks for LVCSR** : [Paper]

- **Face recognition: a convolutional neural-network approach** : [Paper]

2.1.2 Recurrent Networks

- **An empirical exploration of recurrent network architectures** : [Paper]

- **LSTM: A search space odyssey** : [Paper]

- **On the difficulty of training recurrent neural networks** : [Paper]

- **Learning to forget: Continual prediction with LSTM** : [Paper]

2.1.3 Autoencoders

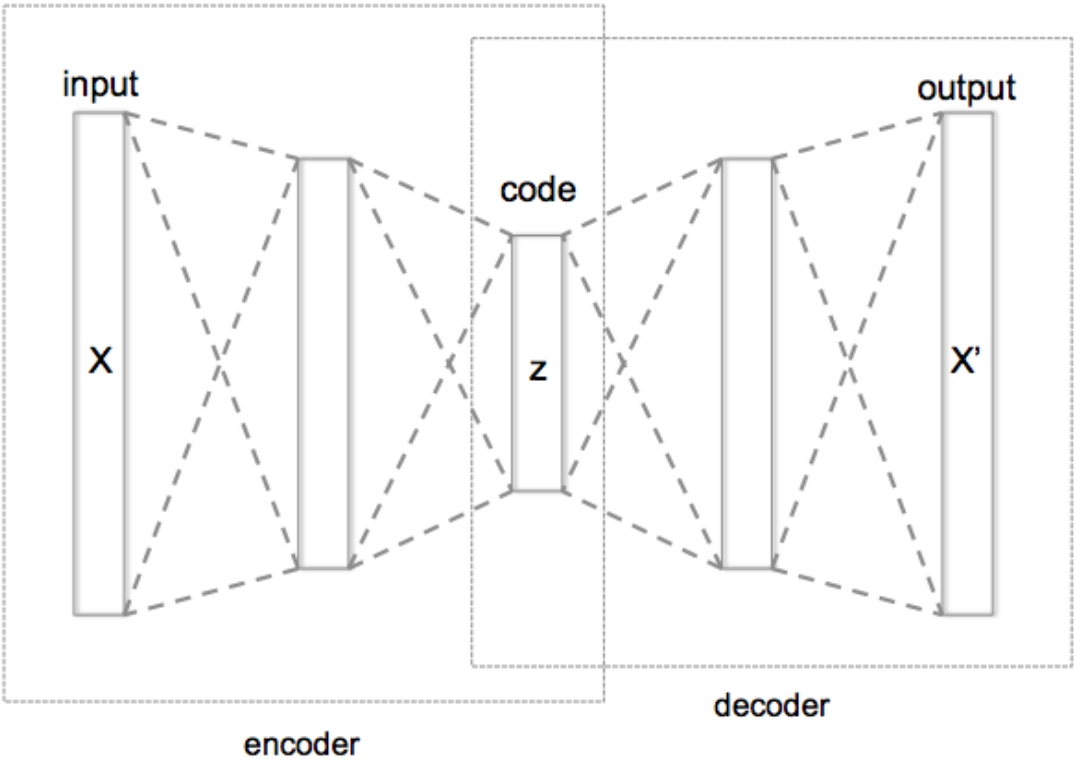

- **Extracting and composing robust features with denoising autoencoders** : [Paper]

- **Stacked Denoising Autoencoders: Learning Useful Representations in a Deep Network with a Local Denoising Criterion** : [Paper]

- **Adversarial Autoencoders** : [Paper]

- **Autoencoders, Unsupervised Learning, and Deep Architectures** : [Paper]

- **Reducing the Dimensionality of Data with Neural Networks** : [Paper]

2.1.4 Generative Models

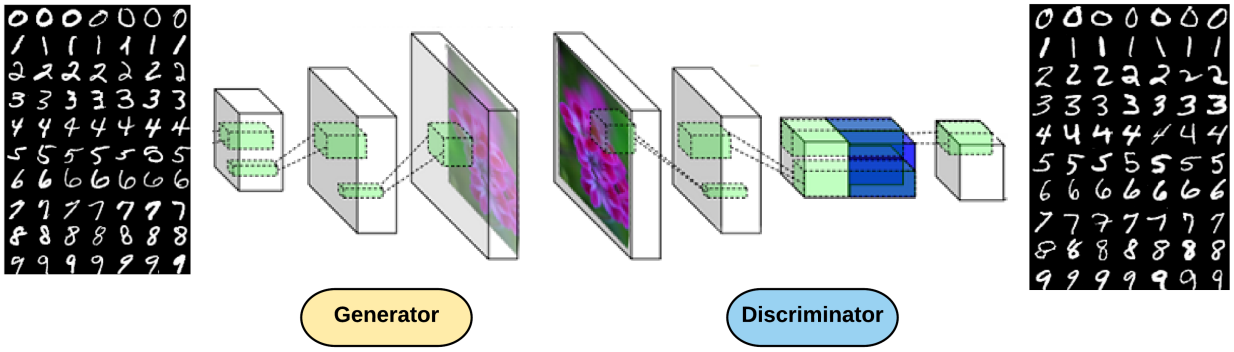

- **Exploiting generative models discriminative classifiers** : [Paper]

- **Semi-supervised Learning with Deep Generative Models** : [Paper]

- **Generative Adversarial Nets** : [Paper]

- **Generalized Denoising Auto-Encoders as Generative Models** : [Paper]

2.1.5 Probabilistic Models

- **Stochastic Backpropagation and Approximate Inference in Deep Generative Models** : [Paper]

- **Probabilistic models of cognition: exploring representations and inductive biases** : [Paper]

- **On deep generative models with applications to recognition** : [Paper]

2.2 Core

2.2.1 Optimization

- **Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift** : [Paper]

- **Dropout: A Simple Way to Prevent Neural Networks from Overfitting** : [Paper]

- **Training Very Deep Networks** : [Paper]

- **Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification** : [Paper]

- **Large Scale Distributed Deep Networks** : [Paper]

2.2.2 Representation Learning

- **Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks** : [Paper]

- **Representation Learning: A Review and New Perspectives** : [Paper]

- **InfoGAN: Interpretable Representation Learning by Information Maximizing Generative Adversarial Nets** : [Paper]

2.2.3 Understanding and Transfer Learning

- **Learning and Transferring Mid-Level Image Representations using Convolutional Neural Networks** : [Paper]

- **Distilling the Knowledge in a Neural Network** : [Paper]

- **DeCAF: A Deep Convolutional Activation Feature for Generic Visual Recognition** : [Paper]

- **How transferable are features in deep neural networks?** : [Paper]

2.2.4 Reinforcement Learning

- **Human-level control through deep reinforcement learning** : [Paper]

- **Playing Atari with Deep Reinforcement Learning** : [Paper]

- **Continuous control with deep reinforcement learning** : [Paper]

- **Deep Reinforcement Learning with Double Q-Learning** : [Paper]

- **Dueling Network Architectures for Deep Reinforcement Learning** : [Paper]

2.3 Applications

2.3.1 Image Recognition

- **Deep Residual Learning for Image Recognition** : [Paper]

- **Very Deep Convolutional Networks for Large-Scale Image Recognition** : [Paper]

- **Multi-column Deep Neural Networks for Image Classification** : [Paper]

- **DeepID3: Face Recognition with Very Deep Neural Networks** : [Paper]

- **Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps** : [Paper]

- **Deep Image: Scaling up Image Recognition** : [Paper]

- **Long-Term Recurrent Convolutional Networks for Visual Recognition and Description** : [Paper]

2.3.2 Object Recognition

- **ImageNet Classification with Deep Convolutional Neural Networks** : [Paper]

- **Learning Deep Features for Scene Recognition using Places Database** : [Paper]

- **Scalable Object Detection using Deep Neural Networks** : [Paper]

- **Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks** : [Paper]

- **OverFeat: Integrated Recognition, Localization and Detection using Convolutional Networks** : [Paper]

- **CNN Features Off-the-Shelf: An Astounding Baseline for Recognition** : [Paper]

- **What is the best multi-stage architecture for object recognition?** : [Paper]

2.3.3 Action Recognition

- **Long-Term Recurrent Convolutional Networks for Visual Recognition and Description** : [Paper]

- **Learning Spatiotemporal Features With 3D Convolutional Networks** : [Paper]

- **Describing Videos by Exploiting Temporal Structure** : [Paper]

- **Convolutional Two-Stream Network Fusion for Video Action Recognition** : [Paper]

- **Temporal segment networks: Towards good practices for deep action recognition** : [Paper]

2.3.4 Caption Generation

- **Show, Attend and Tell: Neural Image Caption Generation with Visual Attention** : [Paper]

- **Mind's Eye: A Recurrent Visual Representation for Image Caption Generation** : [Paper]

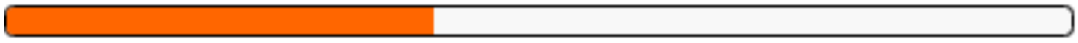

- **Generative Adversarial Text to Image Synthesis** : [Paper]

- **Deep Visual-Semantic Al60ignments for Generating Image Descriptions** : [Paper]

- **Show and Tell: A Neural Image Caption Generator** : [Paper]

2.3.5 Natural Language Processing

- **Distributed Representations of Words and Phrases and their Compositionality** : [Paper]

- **Efficient Estimation of Word Representations in Vector Space** : [Paper]

- **Sequence to Sequence Learning with Neural Networks** : [Paper]

- **Neural Machine Translation by Jointly Learning to Align and Translate** : [Paper]

- **Get To The Point: Summarization with Pointer-Generator Networks** : [Paper]

- **Attention Is All You Need** : [Paper]

- **Convolutional Neural Networks for Sentence Classification** : [Paper]

2.3.6 Speech Technology

- **Deep Neural Networks for Acoustic Modeling in Speech Recognition: The Shared Views of Four Research Groups** : [Paper]

- **Towards End-to-End Speech Recognition with Recurrent Neural Networks** : [Paper]

- **Speech recognition with deep recurrent neural networks** : [Paper]

- **Fast and Accurate Recurrent Neural Network Acoustic Models for Speech Recognition** : [Paper]

- **Deep Speech 2 : End-to-End Speech Recognition in English and Mandarin** : [Paper]

- **Deep Speech 2 : End-to-End Speech Recognition in English and Mandarin** : [Paper]

- **A novel scheme for speaker recognition using a phonetically-aware deep neural network** : [Paper]

CHAPTER 3

Courses

- **Machine Learning** by Stanford on Coursera : [Link]
- **Neural Networks and Deep Learning** Specialization by Coursera: [Link]
- **Intro to Deep Learning** by Google: [Link]
- **NVIDIA Deep Learning Institute** by NVIDIA: [Link]
- **Convolutional Neural Networks for Visual Recognition** by Standford: [Link]
- **Deep Learning for Natural Language Processing** by Standford: [Link]
- **Deep Learning** by fast.ai: [Link]

CHAPTER 4

Books

- **Deep Learning** by Ian Goodfellow: [Link]
- **Neural Networks and Deep Learning** : [Link]
- **Deep Learning with Python**: [Link]
- **Hands-On Machine Learning with Scikit-Learn and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems**: [Link]

CHAPTER 5

Blogs

- **Colah's blog**: [Link]
- **Andrej Karpathy blog**: [Link]
- **The Spectator** Shakir's Machine Learning Blog: [Link]
- **WILDML**: [Link]
- **Distill blog**: [Link]
- **BAIR** Berkeley Artificial Intelligent Research: [Link]
- **Sebastian Ruder's blog**: [Link]
- **inFERENCe**: [Link]
- **i am trask** A Machine Learning Craftsmanship Blog: [Link]

CHAPTER 6

Tutorials

- **Deep Learning Tutorials**: [Link]
- **Deep Learning for NLP with Pytorch** by Pytorch: [Link]
- **Deep Learning for Natural Language Processing: Tutorials with Jupyter Notebooks** by Jon Krohn: [Link]

CHAPTER 7

Contributing

For typos, please do not create a pull request. Instead, declare them in issues or email the repository owner. Please note we have a code of conduct, please follow it in all your interactions with the project.

7.1 Pull Request Process

Please consider the following criterions in order to help us in a better way:

1. The pull request is mainly expected to be a link suggestion.
2. Please make sure your suggested resources are not obsolete or broken.
3. Ensure any install or build dependencies are removed before the end of the layer when doing a build and creating a pull request.
4. Add comments with details of changes to the interface, this includes new environment variables, exposed ports, useful file locations and container parameters.
5. You may merge the Pull Request in once you have the sign-off of at least one other developer, or if you do not have permission to do that, you may request the owner to merge it for you if you believe all checks are passed.

7.2 Final Note

We are looking forward to your kind feedback. Please help us to improve this open source project and make our work better. For contribution, please create a pull request and we will investigate it promptly. Once again, we appreciate your kind feedback and elaborate code inspections.

CHAPTER 8

Contributor Covenant Code of Conduct

8.1 Our Pledge

In the interest of fostering an open and welcoming environment, we as contributors and maintainers pledge to making participation in our project and our community a harassment-free experience for everyone, regardless of age, body size, disability, ethnicity, gender identity and expression, level of experience, nationality, personal appearance, race, religion, or sexual identity and orientation.

8.2 Our Standards

Examples of behavior that contributes to creating a positive environment include:

- Using welcoming and inclusive language
- Being respectful of differing viewpoints and experiences
- Gracefully accepting constructive criticism
- Focusing on what is best for the community
- Showing empathy towards other community members

Examples of unacceptable behavior by participants include:

- The use of sexualized language or imagery and unwelcome sexual attention or advances
- Trolling, insulting/derogatory comments, and personal or political attacks
- Public or private harassment
- Publishing others' private information, such as a physical or electronic address, without explicit permission
- Other conduct which could reasonably be considered inappropriate in a professional setting

8.3 Our Responsibilities

Project maintainers are responsible for clarifying the standards of acceptable behavior and are expected to take appropriate and fair corrective action in response to any instances of unacceptable behavior.

Project maintainers have the right and responsibility to remove, edit, or reject comments, commits, code, wiki edits, issues, and other contributions that are not aligned to this Code of Conduct, or to ban temporarily or permanently any contributor for other behaviors that they deem inappropriate, threatening, offensive, or harmful.

8.4 Scope

This Code of Conduct applies both within project spaces and in public spaces when an individual is representing the project or its community. Examples of representing a project or community include using an official project e-mail address, posting via an official social media account, or acting as an appointed representative at an online or offline event. Representation of a project may be further defined and clarified by project maintainers.

8.5 Enforcement

Instances of abusive, harassing, or otherwise unacceptable behavior may be reported by contacting the project team at amirsina.torfi@gmail.com. The project team will review and investigate all complaints, and will respond in a way that it deems appropriate to the circumstances. The project team is obligated to maintain confidentiality with regard to the reporter of an incident. Further details of specific enforcement policies may be posted separately.

Project maintainers who do not follow or enforce the Code of Conduct in good faith may face temporary or permanent repercussions as determined by other members of the project's leadership.

8.6 Attribution

This Code of Conduct is adapted from the [Contributor Covenant][homepage], version 1.4, available at [http://contributor-covenant.org/version/1/4{]}{[}version]

[homepage]: http://contributor-covenant.org [version]: http://contributor-covenant.org/version/1/4/

CHAPTER 9

LICENSE