Homework 2
CS584: Deep Learning
Spring 2020

This assignment is due on **Wednesday, January 22, 2020 at 4:00pm ET** on the course website.

Problem 1. *(30 points)* Each student will present a paper in class. Each presentation will be 20 minutes long with 5 minutes for questions, and each paper will be about deep learning or somehow related to it. Some papers are relatively classical papers in the field while others are more recent.

Please list at least 10 papers (from most to least preferred) that you would like to present, e.g., 4(a), 4(b), 1(a), 1(g), 1(f), 1(a), 1(b), 1(c), 1(d), 3(b). If you list more papers, then there is a higher chance that you will be able to present one of your preferences. You may also include papers not on this list.

Pick a paper that you can tell a story about. Most likely, you will be unable to present every detail from your paper, so you will want to focus on the important parts of the paper. For example, you should try to answer the following questions in your presentation: What problem is the paper addressing? Why is the problem important, and why is the approach in the paper interesting or successful? How did the authors address the problem, and how is this approach related to earlier and/or later approaches?

For this homework, please write why you chose the papers that you chose. Are you interested in a particular issue or approach? Your responses will help me choose who presents which papers.

Note: Some links or otherwise may have errors. Please let me know if you see any mistakes.

1. Theory

   (a) Rumelhart, Hinton, Williams. *Learning representations by back-propagating errors.* Nature (1997).
   `https://www.nature.com/articles/323533a0`

   (b) Rahimi, Recht. *Random Features for Large-Scale Kernel Machines.* NIPS (2007).
   `https://papers.nips.cc/paper/3182-random-features-for-large-scale-kernel-machi`

   (c) Glorot, Bengio. *Understanding the difficulty of training deep feedforward neural networks.* AISTATS (2010).
   `http://proceedings.mlr.press/v9/glorot10a/glorot10a.pdf`

   (d) Ge, Huang, Jin, Yuan. *Escaping From Saddle Points – Online Stochastic Gradient for Tensor Decomposition.* JMLR (2015).
   `http://proceedings.mlr.press/v40/Ge15.pdf`

   (e) Belkin, Ma, Mandal. *To Understand Deep Learning We Need to Understand Kernel Learning*. PMLR (2018).
       `http://proceedings.mlr.press/v80/belkin18a/belkin18a.pdf`
   (f) Belkin, Hsu, Ma, Mandal. *Reconciling modern machine learning practice and the bias-variance trade-off*. PNAS (2019).
       `https://www.pnas.org/content/116/32/15849`
   (g) Nakkiran, Kaplun, Bansal, Yang, Barak, Sutskever. *Deep Double Descent: Where Bigger Models and More Data Hurt*. arXiv (2019).
       `https://arxiv.org/abs/1912.02292`

2. Natural language processing

   (a) Collobert, Weston, Bottou, Karlen, Kavukcuoglu, Kuksa. *Natural Language Processing (almost) from Scratch*. JMLR (2011).
       `http://www.jmlr.org/papers/volume12/collobert11a/collobert11a.pdf`
   (b) Wu, Roberts, Datta, Du, Ji, Si, Soni, Wang, Wei, Xiang, Zhao, Xu. *Deep learning in clinical natural language processing: a methodical review*. JAMIA (2019).
       `https://academic.oup.com/jamia/advance-article-abstract/doi/10.1093/jamia/ocz200/5651084`

3. Object recognition

   (a) LeCun, Bottou, Bengio, Haner. *Gradient-Based Learning Applied to Document Recognition*. Proc. IEEE (1998).
       `http://vision.stanford.edu/cs598_spring07/papers/Lecun98.pdf`
   (b) Krizhevsky, Sutskever, Hinton. *ImageNet Classification with Deep Convolutional Neural Networks*. NIPS (2012).
       `https://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutio`
   (c) Goodfellow, Warde-Farley, Mirza, Courville, Bengio. *Maxout Networks*. JMLR (2013).
       `http://proceedings.mlr.press/v28/goodfellow13.pdf`
   (d) Girshick, Donahue, Darrell, Malik. *Rich feature hierarchies for accurate object detection and semantic segmentation*. CVPR (2014).
       `https://www.cv-foundation.org/openaccess/content_cvpr_2014/papers/Girshick_Rich_Feature_Hierarchies_2014_CVPR_paper.pdf`
   (e) Zeiler, Fergus. *Visualizing and Understanding Convolutional Networks*. ECCV (2014).
       `https://link.springer.com/chapter/10.1007/978-3-319-10590-1_53`
   (f) Szegedy, Liu, Jia, Sermanet, Reed, Anguelov, Erhan, Vanhoucke, Rabinovich. *Going Deeper With Convolutions*. CVPR (2015).

https://www.cv-foundation.org/openaccess/content_cvpr_2015/html/Szegedy_Going_Deeper_With_2015_CVPR_paper.html

   (g) He, Zhang, Ren, Sun. *Deep Residual Learning for Image Recognition.* CVPR (2016).
   https://www.cv-foundation.org/openaccess/content_cvpr_2016/papers/He_Deep_Residual_Learning_CVPR_2016_paper.pdf

   (h) Chao, Changpinyo, Gong, Sha. *An Empirical Study and Analysis of Generalized Zero-Shot Learning for Object Recognition in the Wild.* ECCV (2016).
   https://link.springer.com/chapter/10.1007/978-3-319-46475-6_4

4. Reinforcement learning

   (a) Mnih, Kavukcuoglu, Silver, Rusu, Veness, Bellemare, Graves, Riedmiller, Fidjeland, Ostrovski, Petersen, Beattie, Sadik, Antonoglou, King, Kumaran, Wierstra Legg, Hassabis. *Human-level control through deep reinforcement learning.* Nature (2015).
   https://www.nature.com/articles/nature14236

   (b) Silver, Schrittwieser, Simonyan, Antonoglou, Huang, Guez, Hubert, Baker, Lai, Bolton, Chen, Lillicrap, Hui, Sifre, van den Driessche, Graepel, Hassabis. *Mastering the game of Go without human knowledge.* Nature (2017).
   https://www.nature.com/articles/nature24270

5. Speech recognition

   (a) Deng, Seltzer, Yu, Acero, Mohamed, Hinton. *Binary Coding of Speech Spectrograms Using a Deep Auto-encoder.* Interspeech (2010).
   https://www.microsoft.com/en-us/research/wp-content/uploads/2016/02/IS100232.pdf

   (b) Hinton, Deng, Yu, Dahl, Mohamed, Jaitly, Senior, Vanhoucke, Nguyen, Sainath, Kingsbury. *Deep Neural Networks for Acoustic Modeling in Speech Recognition.* IEEE (2012).
   https://ieeexplore.ieee.org/document/6296526

6. Other

   (a) Hochreiter, Schmidhuber. *Long Short-Term Memory.* Neural Computation (1997).
   https://www.bioinf.jku.at/publications/older/2604.pdf

   (b) Goodfellow, Pouget-Abadie, Mirza, Xu, Warde-Farley, Ozair, Courville, Bengio. *Generative Adversarial Nets.* NIPS (2014).
   https://papers.nips.cc/paper/5423-generative-adversarial-nets

   (c) Greff, Srivastava, Koutník, Steunebrink, Schmidhuber. *LSTM: A Space Search Odyssey.* IEEE (2016).

https://ieeexplore.ieee.org/document/7508408

(d) Poplin, Varadarajan, Blumer, Liu, McConnell, Corrado, Peng, Webster. *Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning*. Nature Biomedical Engineering (2018).
https://www.nature.com/articles/s41551-018-0195-0

(e) Obermeyer, Powers, Vogeli, Mullainathan. *Dissecting racial bias in an algorithm used to manage the health of populations*. Science (2019).
https://science.sciencemag.org/content/366/6464/447.abstract