

Deep Learning for the Web

Kyomin Jung
Seoul National University
1 Gwanak-ro, Gwanak-gu
Seoul, Korea
kjung@snu.ac.kr

Byoung-Tak Zhang
Seoul National University
1 Gwanak-ro, Gwanak-gu
Seoul, Korea
btzhang@snu.ac.kr

Prasenjit Mitra
Qatar Computing Research Institute
Tornado Tower
Doha, Qatar
pmitra@qf.org.qa

ABSTRACT

Deep learning is a machine learning technology that automatically extracts higher-level representations from raw data by stacking multiple layers of neuron-like units. The stacking allows for extracting representations of increasingly-complex features without time-consuming, offline feature engineering. Recent success of deep learning has shown that it outperforms state-of-the-art systems in image processing, voice recognition, web search, recommendation systems, etc [1]. A lot of industrial-scale big data processing systems including IBM Watson's Jeopardy Contest 2011, Google Now, Facebook's face recognition system, and the voice recognition systems by Google and Microsoft use deep learning [2][3][6]. Deep learning has a huge potential to improve the intelligence of the web and the web service systems by efficiently and effectively mining big data on the Web[4][5].

This tutorial provides the basics of deep learning as well as its key applications. We give the motivation and underlying ideas of deep learning and describe the architectures and learning algorithms for various deep learning models. We also cover applications of deep learning for image and video processing, natural language and text data analysis, social data analytics, and wearable IoT sensor data with an emphasis in the domain of Web systems. We will deliver the key insight and understanding of these techniques, using graphical illustrations and examples that could be important in analyzing a large amount of Web data. The tutorial is prepared to attract general audience at the WWW Conference, who are interested in machine learning and big data analysis for Web data.

The tutorial consists of five parts. The first part presents the basics of neural networks, and their structures. Then we explain the training algorithm via backpropagation, which is a common method of training artificial neural networks including deep neural networks. We will emphasize how each of these concepts can be used in various Web data

analysis. In the second part of the tutorial, we describe the learning algorithms for deep neural networks and related ideas, such as contrastive divergence, wake-sleep algorithms, and Monte Carlo simulation. We then describe various kinds of deep architectures, including stacked autoencoders, deep belief networks [7], convolutional neural networks [8], and deep hypernetworks [9].

In the third part, we present more details of the recursive neural networks, which can learn structured tree outputs as well as vector representations for phrases and sentences. We first show how training the recursive neural network can be achieved by a modified version of the back-propagation algorithm introduced before. These modifications allow the algorithm to work on tree structures. Then we will present its applications to sentence analysis including POS tagging, and sentiment analysis. The fourth part discusses the neural networks used to generate word embeddings, such as Word2Vec [10], DSSM for deep semantic similarity [11], and object detection in images [12], such as GoogLeNet, and AlexNet. We will explain in detail the applications of these deep learning techniques in the analysis of various social network data. By this point, the audience should have a clear understanding of how to build a deep learning system for word, sentence and document level tasks. The fifth part of the tutorial will cover other application examples of deep learning. These include object segmentation and action recognition from videos [9], web data analytics, and wearable/IoT sensor data modeling for smart services.

References

1. Ten breakthrough technologies 2013, MIT Technology Review, Apr 23, 2013.
2. Building Watson: An overview of the DeepQA project, AI Magazine, Fall 2010.
3. Y. Taigman, M. Yang, M. Ranzato, and L. Wolf, DeepFace: Closing the gap to human-level performance in face verification, Computer Vision and Pattern Recognition (CVPR) 2014.
4. P. Huang, X. He, J. Gao, L. Deng, A. Acero, and L. Heck, Learning deep structured semantic models for web search using clickthrough data, ACM International Conference on Information and Knowledge Management (CIKM) 2013.

5. X. Glorot, A. Bordes, and Y. Bengio, Domain adaptation for large-scale sentiment classification: A deep learning approach, International Conference on Machine Learning (ICML) 2011.
6. A. Graves, A. Mohamed, and G. E. Hinton, Speech recognition with deep recurrent neural networks, IEEE International Conference on Acoustic Speech and Signal Processing (ICASSP 2013), Vancouver, 2013.
7. G. E. Hinton, S. Osindero, and Y. The, A fast learning algorithm for deep belief nets, Neural Computation, 18: 1527-1544.
8. P. Sermanet, D. Eigen, X. Zhang, M. Mathieu, R. Fergus, and Y. LeCun, Overfeat: Integrated recognition, localization and detection using convolutional networks, International Conference on Learning Representations (ICLR2014), 2014.
9. J.-W. Ha, K.-M. Kim, and B.-T. Zhang, Automated construction of visual-linguistic knowledge via concept learning from cartoon videos, In Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI 2015), Austin, 2015.
10. Available at: <https://code.google.com/p/word2vec/> on Mar 15, 2015.
11. J. Gao, X. He, and L. Deng, L. Deep Learning for Web Search and Natural Language Processing, Microsoft Research Technical Report no. MSR-TR-2015-7, 2015.
12. A. Krizhevsky, I. Sutskever and G. Hinton, Imagenet classification with deep convolutional neural networks, Advances in neural information processing systems, 2012.