

Learning Deep Representations of Time Series

Supervisor: Igor Kulev
igor.kulev@epfl.ch

June 12, 2017

1 Overview and Goal

Time series data are ubiquitous. The explosion of new sensing technologies (wearable sensors, satellites, mobile phones, etc.), combined with increasingly cheap and effective storage, is generating an unprecedented and growing amount of time series data in a variety of domains [1]. Understanding this time series data is important for many Artificial Intelligence tasks. Recently, recurrent neural networks using the Long Short Term Memory (LSTM) architecture have been used successfully to perform various supervised sequence learning tasks, such as speech recognition, machine translation, and caption generation for images [2]. However, most of the available data is "unlabeled". There is a need for methods that are able to describe the hidden structure from "unlabeled" data. Srivastava et al. used LSTM to learn representations of video sequences [2]. The goal of this project is to develop and evaluate methods based on LSTM for learning representations of time series. As a part of the project you will implement at least two methods and apply them on at least two different datasets. The aim of the next phase of the project will be to extend one of the methods and improve its performance on the datasets.

2 Required Skills

- Familiar with Machine Learning and Probability
- Experienced with Python

References

- [1] "Sigkdd workshop on mining and learning from time series," 2015. [Online]. Available: <http://www-bcf.usc.edu/~liu32/milets/>

- [2] N. Srivastava, E. Mansimov, and R. Salakhudinov, “Unsupervised learning of video representations using lstms,” in *International Conference on Machine Learning*, 2015, pp. 843–852.
- [3] A. Makhzani, J. Shlens, N. Jaitly, I. Goodfellow, and B. Frey, “Adversarial autoencoders,” *arXiv preprint arXiv:1511.05644*, 2015.
- [4] W. Harchaoui, P.-A. Mattei, and C. Bouveyron, “Deep adversarial gaussian mixture auto-encoder for clustering,” 2017.
- [5] Y. Zheng, H. Tan, B. Tang, H. Zhou *et al.*, “Variational deep embedding: A generative approach to clustering,” *arXiv preprint arXiv:1611.05148*, 2016.
- [6] J. Xie, R. Girshick, and A. Farhadi, “Unsupervised deep embedding for clustering analysis,” in *International Conference on Machine Learning (ICML)*, 2016.
- [7] I. Sutskever, O. Vinyals, and Q. V. Le, “Sequence to sequence learning with neural networks,” in *Advances in neural information processing systems*, 2014, pp. 3104–3112.
- [8] D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” *arXiv preprint arXiv:1312.6114*, 2013.
- [9] C. Doersch, “Tutorial on variational autoencoders,” *arXiv preprint arXiv:1606.05908*, 2016.
- [10] O. Fabius and J. R. van Amersfoort, “Variational recurrent auto-encoders,” *arXiv preprint arXiv:1412.6581*, 2014.
- [11] J. Chung, K. Kastner, L. Dinh, K. Goel, A. C. Courville, and Y. Bengio, “A recurrent latent variable model for sequential data,” in *Advances in neural information processing systems*, 2015, pp. 2980–2988.