Deep learning for graphs

Davide Bacciu (bacciu@di.unipi.it)

Department of Computer Science, University of Pisa, Italy

Abstract

Graphs are a rich and powerful abstraction to represent complex compound information made of atomic information pieces, i.e. the vertices and their labels, enriched by complex relational information, represented by the edges and associated labels/types. Much of the data produced by artificial and natural processes can be naturally and effectively represented as either a single graph (a network) of interactions or as a collection of graph-structured samples, with generally varying size and connectivity. Examples of network-type information are social graphs, biological networks of proteomics or genomic interactions, etc. Examples of structure varying graphs are, for instance, molecular compounds, point-clouds in 3D scans, visual regions in machine vision. Being able to process such data in this rich graph structured form provides a fundamental advantage when it comes to identifying data patterns useful for achieving predictive and/or explorative analyses. This has motivated a recent increasing interest of the machine learning community into the development of deep learning models for the adaptive processing of graph structured information, which is witnessed by the increasing amounts of works on the topic being presented at the major international Machine Learning (ML) conferences in over the last year.

Dealing with graph data from a ML perspective requires designing learning models capable of adapting to structured samples of varying size and topology, that can identify, represent and store those structural patterns that are most relevant to perform predictive and explorative tasks. This has to be achieved while maintaining the efficiency and scalability necessary to process (also) large scale networks. The lecture will introduce the emerging field of deep learning for graphs and its applications to network data analysis, to bioinformatics, chemistry, physics and vision. We will start by introducing early works [2, 5] that, almost coincidentally, proposed an extension of the convolutional neural network (CNN) to graphs, using spectral properties of convolutions on the graph adjacency matrix. We will highlight how such approaches are limited in scalability and in the capability of dealing with populations of topology-varying graphs. Then, we will focus on works that have proposed proper generalizations of the convolutional operators to convolutions on graphs, enabling to tackle structures of varying topology [3, 6, 4]. We will then consider some recent and advanced

research topics which include unsupervised learning for graphs [1], the definition of general and differentiable pooling functions on non-constrained structures [7] as well as the design of generative models which can learn distributions and sampling processes on spaces of graphs [8]. We will conclude the overview by survey impacting applications proposed in this context, ranging from the automatic synthesis of molecular compounds with specified chemo-physical properties, to large scale biological network analysis for drug repositioning.

References

[1] Davide Bacciu, Federico Errica, and Alessio Micheli. Contextual graph markov model: A deep and generative approach to graph processing. In *International Conference on Machine Learning*, 2018.

[2] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, NIPS'16, pages 3844–3852, USA, 2016. Curran Associates Inc.

[3] David K Duvenaud, Dougal Maclaurin, Jorge Iparraguirre, Rafael Bombarell, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan P Adams. Convolutional networks on graphs for learning molecular fingerprints. In *Advances in neural information processing systems*, pages 2224–2232, 2015.

[4] William L. Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *NIPS*, 2017.

[5] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *Proceedings of ICLR 2017*, 2017.

[6] Mathias Niepert, Mohamed Ahmed, and Konstantin Kutzkov. Learning convolutional neural networks for graphs. In *International Conference on Machine Learning*, pages 2014–2023, 2016.

[7] Rex Ying, Jiaxuan You, Christopher Morris, Xiang Ren, William L Hamilton, and Jure Leskovec. Hierarchical graph representation learning with differentiable pooling. In *Advances in neural information processing systems*, 2018.

[8] Jiaxuan You, Rex Ying, Xiang Ren, William L. Hamilton, and Jure Leskovec. Graphrnn: A deep generative model for graphs. In *International Conference on Machine Learning*, 2018.