

---

# Graph Neural Networks in TensorFlow and Keras with Spektral

---

Daniele Grattarola<sup>1</sup> Cesare Alippi<sup>1,2</sup>

## Abstract

In this paper we present Spektral, an open-source Python library for building graph neural networks with TensorFlow and the Keras application programming interface. Spektral implements a large set of methods for deep learning on graphs, including message-passing and pooling operators, as well as utilities for processing graphs and loading popular benchmark datasets. The purpose of this library is to provide the essential building blocks for creating graph neural networks, focusing on the guiding principles of user-friendliness and quick prototyping on which Keras is based. Spektral is, therefore, suitable for absolute beginners and expert deep learning practitioners alike. In this work, we present an overview of Spektral’s features and report the performance of the methods implemented by the library in scenarios of node classification, graph classification, and graph regression.

## 1. Introduction

Graph Neural Networks (GNNs) are a class of deep learning methods designed to perform inference on data described by graphs (Battaglia et al., 2018). Due to the different possibilities offered by graph machine learning and the large number of applications where graphs are naturally found, GNNs have been successfully applied to a diverse spectrum of fields to solve a variety of tasks. In physics, GNNs have been used to learn physical models of complex systems of interacting particles (Battaglia et al., 2016; Kipf et al., 2018; Sanchez-Gonzalez et al., 2018; Farrell et al., 2018). In recommender systems, the interactions between users and items can be represented as a bipartite graph and the goal is to predict new potential edges (*i.e.*, which items could a user be interested in), which can be achieved with GNNs (Berg et al., 2017; Ying et al.,

2018a). GNNs have also been largely applied to the biological sciences, with applications ranging from the recommendation of medications (Shang et al., 2019), to the prediction of protein-protein and protein-ligand interactions (Gainza et al., 2020), and in chemistry, for the prediction of quantum molecular properties as well as the generation of novel compounds and drugs (Do et al., 2019; You et al., 2018). Finally, GNNs have been successfully applied in fields like natural language processing (Fernandes et al., 2018; De Cao et al., 2018) and even more complex tasks like abstract reasoning (Santoro et al., 2017; Allamanis et al., 2017; Schlichtkrull et al., 2018) and decision making with reinforcement learning (Zambaldi et al., 2018; Hamrick et al., 2018).

At the core of GNNs there are two main types of operations, which can be interpreted as a generalisation of the convolution and pooling operators in convolutional neural networks: *message passing* and *graph pooling* (Fig. 1). The former is used to learn a non-linear transformation of the input graphs and the latter to reduce their size. When combined, these two operations enable graph representation learning as a general tool to predict node-level, edge-level, and global properties of graphs. Several works in recent literature have introduced models for either message passing (Gilmer et al., 2017; Scarselli et al., 2009; Defferrard et al., 2016; Kipf & Welling, 2016; Simonovsky & Komodakis, 2017; Velickovic et al., 2017; Hamilton et al., 2017; Bianchi et al., 2019; Xu et al., 2019; Klicpera et al., 2019) or graph pooling (Ying et al., 2018b; Gao & Ji, 2019; Bianchi et al., 2020; Cangea et al., 2018).

Thanks to the increasing popularity of GNNs, many software libraries implementing the building blocks of graph representation learning have been developed in recent years, paving the way for the adoption of GNNs in other fields of science. One of the major challenges faced by researchers and software developers who wish to contribute to the larger scientific community is to make software both accessible and intuitive, so that even non-technical audiences can benefit from the advances carried by intelligent systems. In this spirit, Keras is an application programming interface (API) for creating neural networks, developed according to the guiding principle that “being able to go from idea to result with the least possible delay is key to doing good research” (Chollet et al., 2015). Keras

---

<sup>1</sup>Università della Svizzera italiana, Lugano, Switzerland  
<sup>2</sup>Politecnico di Milano, Milan, Italy. Correspondence to: Daniele Grattarola <grattd@usi.ch>.




Figure 1. 1(a) Schematic view of a graph neural network with message-passing, pooling, and global pooling layers. The role of message-passing layers is to compute a representation of each node in the graph, leveraging *local* information (*messages*) from its neighbours. The role of pooling layers is to reduce the size of the graph by aggregating or discarding redundant information, so that the GNN can learn a hierarchical representation of the input data. Finally, global pooling layers reduce the graph to a single vector, usually to feed it as input to a multi-layer perceptron for classification or regression. 1(b) The stylised ghost logo of Spektral.

is designed to reduce the cognitive load of end users, shifting the focus away from the boilerplate implementation details and allowing instead to focus on the creation of models. As such, Keras is extremely beginner-friendly and, for many, an entry point to machine learning itself. At the same time, Keras integrates smoothly with its TensorFlow (Abadi et al., 2016) backend and enables users to build any model that they could have implemented in pure TensorFlow. This flexibility makes Keras an excellent tool even for expert deep learning practitioners and has recently led to TensorFlow’s adoption of Keras as the official interface to the framework.

In this paper we present Spektral, a Python library for building graph neural networks using TensorFlow and the Keras API. Spektral implements some of the most important papers from the GNN literature as Keras layers, and it integrates seamlessly within Keras models and with the most important features of Keras like the training loop, callbacks, distributed training, and automatic support for GPUs and TPUs. As such, Spektral inherits the philosophy of ease of use and flexibility that characterises Keras. The components of Spektral act as standard TensorFlow operations and can be easily used even in more advanced settings, integrating tightly with all the features of TensorFlow and allowing for an easy deployment to production systems. For these reasons, Spektral is the ideal library to implement GNNs in the TensorFlow ecosystem, both for total beginners and experts alike.

All features of Spektral are documented in detail<sup>1</sup> and a collection of examples is provided with the source code. The project is released on GitHub<sup>2</sup> under MIT license.

<sup>1</sup><https://graphneural.network>

<sup>2</sup><https://github.com/danielegrattarola/spektral>

## 2. Library Overview

### 2.1. Representing graphs

Let  $\mathcal{G} = \{\mathcal{X}, \mathcal{E}\}$  be a graph where  $\mathcal{X} = \{\mathbf{x}_i \in \mathbb{R}^F \mid i = 1, \dots, N\}$  is the set of nodes with  $F$ -dimensional real attributes, and  $\mathcal{E} = \{\mathbf{e}_{ij} \in \mathbb{R}^S \mid \mathbf{x}_i, \mathbf{x}_j \in \mathcal{X}\}$  the set of edges with  $S$  dimensional real attributes. In Spektral, we represent  $\mathcal{G}$  by its binary adjacency matrix  $\mathbf{A} \in \{0, 1\}^{N \times N}$ , node features  $\mathbf{X} \in \mathbb{R}^{N \times F}$ , and edge features  $\mathbf{E} \in \mathbb{R}^{N \times N \times S}$ . Any format accepted by Keras to represent the above matrices is also supported by Spektral, which means that it also natively supports the NumPy stack of scientific computing libraries for Python. Most of the layers and utilities implemented in Spektral also support the sparse matrices of the SciPy library, making them computationally efficient both in time and memory. Additionally, Spektral makes very few assumptions on how a user may want to represent graphs, and transparently deals with batches of graphs represented as higher-order tensors or disjoint unions (see Appendix A for more details).

### 2.2. Message passing

Message-passing networks are a general paradigm introduced by Gilmer et al. (2017) that unifies most GNN methods found in the literature as a combination of *message*, *aggregation*, and *update* functions. Message-passing layers are equivalent in role to the convolutional operators in convolutional neural networks, and are the essential component of graph representation learning. Message-passing layers in Spektral are available in the `layers.convolutional` module.<sup>3</sup> Currently, Spektral implements fifteen different message-passing layers including Graph Convolutional Networks (GCN) (Kipf & Welling,

<sup>3</sup>The name *convolutional* derives from the homonymous module in Keras, as well as message-passing layers being originally derived as a generalisation of convolutional operators.

2016), ChebNets (Defferrard et al., 2016), GraphSAGE (Hamilton et al., 2017), ARMA convolutions (Bianchi et al., 2019), Edge-Conditioned Convolutions (ECC) (Simonovsky & Komodakis, 2017), Graph Attention Networks (GAT) (Velickovic et al., 2017), APPNP (Klicpera et al., 2019), and Graph Isomorphism Networks (GIN) (Xu et al., 2019), as well as the methods proposed by Li et al. (2017; 2015); Thekumparampil et al. (2018); Du et al. (2017); Xie & Grossman (2018); Wang et al. (2018) and a general interface that can be extended to implement message-passing layers. The available methods are sufficient to deal with all kinds of graphs, including those with attributed edges.

### 2.3. Graph pooling

Graph pooling refers to any operation to reduce the number of nodes in a graph and has a similar role to pooling in traditional convolutional networks for learning hierarchical representations. Because pooling computes a coarser version of the graph at each step, ultimately resulting in a single vector representation, it is usually applied to problems of graph-level inference. Graph pooling layers are available in `layers.pooling` and include: DiffPool (Ying et al., 2018b), MinCut pooling (Bianchi et al., 2020), Top- $K$  pooling (Cangea et al., 2018; Gao & Ji, 2019), and Self-Attention Graph Pooling (SAGPool) (Lee et al., 2019). Spektral also implements *global* graph pooling methods, which can be seen as a limit case of graph pooling where a graph is reduced to a single node, *i.e.*, its node features are reduced to a single vector. Spektral implements six different global pooling strategies: sum, average, max, gated attention (GAP) (Li et al., 2015), SortPool (Zhang et al., 2018), and attention-weighted sum (AWSP).<sup>4</sup>

### 2.4. Datasets

Spektral comes with a large variety of popular graph datasets accessible from the `datasets` module. The datasets available from Spektral provide benchmarks for transductive and inductive node classification, graph signal classification, graph classification, and graph regression. In particular, the following datasets can be loaded with Spektral: the citation networks, Cora, CiteSeer, and Pubmed (Sen et al., 2008); the protein-protein interaction dataset (PPI) (Stark et al., 2006; Zitnik & Leskovec, 2017; Hamilton et al., 2017) and the Reddit communities network dataset (Hamilton et al., 2017) from the GraphSAGE paper (Hamilton et al., 2017); the QM9 chemical dataset of small molecules (Ramakrishnan et al., 2014); the MNIST 8-NN

<sup>4</sup>While never published in the literature, attention-weighted sum is a straightforward concept that consists of computing a weighted sum of the node features, where the weights are computed through a simple attentional mechanism

Table 1. Comparison of different GNN libraries. The *Framework* column indicates the backend framework supported by the library, while the *MP* and *Pooling* columns indicate the number of different message-passing and pooling layers implemented by the library, respectively.

Library	Framework	MP	Pooling
<b>Spektral</b>	TensorFlow	15	10
<b>PyG</b>	PyTorch	28	14
<b>DGL</b>	PyTorch, others	15	7
<b>StellarGraph</b>	TensorFlow	6	N/A

graph for graph signal classification as proposed by Defferrard et al. (2016); the Benchmark Data Sets for Graph Kernels (Kersting et al., 2016). Each dataset is automatically downloaded and stored locally when necessary.

### 2.5. Other tools

Some of the secondary functionalities implemented by Spektral include: the `utils` module, which exposes some useful utilities for graph deep learning (*e.g.*, methods for computing the characteristic graph matrices or manipulating the data); the `chem` module, which offers tools for loading and processing molecular graphs; the `layers.ops` module, which offers a set of common operations that can be used by advanced users to create new GNN layers, like wrappers for common matrix operations that automatically handle sparse inputs and batches of graphs.

## 3. Comparison to other libraries

Given the growing popularity of the field, several libraries for GNNs have appeared in recent years. Among the most notable, we cite PyTorch Geometric<sup>5</sup> (PyG) (Fey & Lenssen, 2019) and the Deep Graph Library<sup>6</sup> (DGL) (Wang et al., 2019), both of which are based on the PyTorch deep learning library.<sup>7</sup> Instead, Spektral is specifically developed for the TensorFlow ecosystem, which to this date is estimated to support the majority of deep learning applications both in research and industry (Keras, 2019). The features offered by Spektral, summarized in Table 1, are largely similar to those offered both by PyG (which however implements a much larger variety of message-passing methods and other algorithms from GNN literature) and by DGL. The computational performance of Spektral’s layers is also comparable to that of PyG, with small differences due to implementation details and differences between the

<sup>5</sup><https://pytorch-geometric.readthedocs.io/>

<sup>6</sup><https://docs.dgl.ai/>

<sup>7</sup>Note that DGL also supports MXNet and TensorFlow as backends, albeit with a reduced set of features. In this paper, we consider DGL to be a PyTorch library.

two respective backend frameworks. We also mention the StellarGraph library for GNNs which, like Spektral, is based on Keras. This library implements six message-passing layers, four of which are available in Spektral (GCN, GraphSAGE, GAT and APPNP), but does not offer pooling layers and relies on a custom format for graph data, which limits flexibility.

## 4. Applications

In this section, we report some experimental results on several well-known benchmark tasks, in order to provide a high-level overview of how the different methods implemented by Spektral perform in a standard research use case scenario. We report the results for three main settings: a node classification task and two tasks of graph-level property prediction, one of classification and one of regression. All experimental details are reported in Appendix C.

**Node classification** In our first experiment, we consider a task of semi-supervised node classification on the Cora, CiteSeer, and Pubmed citation networks. We evaluate GCN (Kipf & Welling, 2016), ChebNets (Defferrard et al., 2016), ARMA (Bianchi et al., 2019), GAT (Velickovic et al., 2017) and APPNP (Klicpera et al., 2019). We reproduce the same experimental settings described in the original papers, but we use the random data splits suggested by Shchur et al. (2018) for a fairer evaluation.

**Graph classification** To evaluate the pooling layers, DiffPool, MinCut, Top- $K$ , and SAGPool, we consider a task of graph-level classification, where each graph represents an individual sample to be classified. We use four datasets from the Benchmark Data Sets for Graph Kernels: Proteins, IMDB-Binary, Mutag and NCI1. Here, we adopt a fixed GNN architecture (described in Appendix C) where we only change the pooling method. To assess whether each pooling layer is actually beneficial for learning a representation of the data, we also evaluate the same GNN without pooling (*Flat*).

**Graph regression** To evaluate the global pooling methods, we use the QM9 molecular database and train a GNN on four different regression targets: dipole moment (Mu), isotropic polarizability (Alpha), energy of HOMO (Homo), and internal energy at OK (U0). Because the molecules in QM9 have attributed edges, we adopt a GNN based on ECC, which is designed to integrate edge attributes in the message-passing operation.

### 4.1. Results

The results for each experiment are reported in Tables 2, 3, and 4. In the first experiment, results are compatible with what reported in the literature, although some differences in performance are present due to using a random data split

Table 2. Classification accuracy on the node classification tasks.

Dataset	ChebNets	GCN	GAT	ARMA	APPNP
<b>Cora</b>	77.4 $\pm$ 1.5	79.4 $\pm$ 1.3	82.0 $\pm$ 1.2	80.5 $\pm$ 1.2	<b>82.8</b> $\pm$ 0.9
<b>CiteSeer</b>	68.2 $\pm$ 1.6	68.8 $\pm$ 1.4	70.0 $\pm$ 1.0	<b>70.6</b> $\pm$ 0.9	70.0 $\pm$ 1.0
<b>Pubmed</b>	74.0 $\pm$ 2.7	76.6 $\pm$ 2.5	73.8 $\pm$ 3.3	77.2 $\pm$ 1.6	<b>78.2</b> $\pm$ 2.1

Table 3. Classification accuracy on the graph classification tasks.

Dataset	Flat	MinCut	DiffPool	Top- $K$	SAGPool
<b>Proteins</b>	74.3 $\pm$ 4.5	<b>75.5</b> $\pm$ 2.0	74.1 $\pm$ 3.9	70.5 $\pm$ 3.4	71.3 $\pm$ 5.0
<b>IMDB-B</b>	72.8 $\pm$ 7.2	<b>73.6</b> $\pm$ 5.4	70.6 $\pm$ 6.6	67.7 $\pm$ 8.2	69.3 $\pm$ 5.7
<b>Mutag</b>	72.5 $\pm$ 14.0	81.4 $\pm$ 10.7	<b>83.5</b> $\pm$ 9.7	79.2 $\pm$ 8.0	78.5 $\pm$ 8.3
<b>NCI1</b>	<b>77.3</b> $\pm$ 2.6	74.4 $\pm$ 1.9	71.1 $\pm$ 3.0	72.0 $\pm$ 3.0	69.4 $\pm$ 8.4

rather than the pre-defined one used in the original experiments. The APPNP operator consistently achieves good results on the citation networks, outperforming the other methods on Cora and Pubmed, and coming close to ARMA on CiteSeer. For graph classification, the results are sometimes different than what is reported in the literature, due to the standardized architecture that we used in this experiment. MinCut generally achieves the best performance followed by DiffPool. We also note that the Flat baseline often achieves better results than the equivalent GNNs equipped with pooling. For graph regression, results show that the choice of global pooling method can have a significant impact on performance. In particular, the GAP operator performs best on Alpha and U0, while the best results on Mu and Homo are obtained with max pooling and AWSP, respectively. However, we note how these latter two operators have largely unstable performances depending on the datasets, as both fail on Alpha and U0.

## 5. Conclusion

We presented Spektral, a library for building graph neural networks using the Keras API. Spektral implements several state-of-the-art methods for GNNs, including message-passing and pooling layers, a wide set of utilities, and comes with many popular graph datasets. The library is designed for providing a streamlined user experience and is currently the most mature library for GNNs in the TensorFlow ecosystem. In the future, we will keep Spektral up to date with the ever-growing field of GNN research, and we will focus on improving the performance of its core components.

Table 4. Mean-squared error on the graph regression tasks. Results for Homo are in scale of  $10^{-5}$ .

Dataset	Average	Sum	Max	GAP	AWSP
<b>Mu</b>	1.12 $\pm$ 0.03	1.02 $\pm$ 0.02	<b>0.90</b> $\pm$ 0.04	1.04 $\pm$ 0.05	0.99 $\pm$ 0.03
<b>Alpha</b>	3.15 $\pm$ 0.65	2.38 $\pm$ 0.64	6.20 $\pm$ 0.33	<b>1.89</b> $\pm$ 0.59	31.1 $\pm$ 0.37
<b>Homo</b>	9.24 $\pm$ 0.41	9.22 $\pm$ 0.51	8.90 $\pm$ 0.36	9.04 $\pm$ 0.29	<b>8.05</b> $\pm$ 0.29
<b>U0</b>	0.42 $\pm$ 0.14	0.50 $\pm$ 0.13	110.7 $\pm$ 4.5	<b>0.22</b> $\pm$ 0.13	624.0 $\pm$ 19.0

## References

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al. Tensorflow: A system for large-scale machine learning. In *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)*, pp. 265–283, 2016.
- Allamanis, M., Brockschmidt, M., and Khademi, M. Learning to represent programs with graphs. *arXiv preprint arXiv:1711.00740*, 2017.
- Battaglia, P., Pascanu, R., Lai, M., Rezendes, D. J., et al. Interaction networks for learning about objects, relations and physics. In *Advances in neural information processing systems*, pp. 4502–4510, 2016.
- Battaglia, P. W., Hamrick, J. B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V., Malinowski, M., Tacchetti, A., Raposo, D., Santoro, A., Faulkner, R., et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018.
- Berg, R. v. d., Kipf, T. N., and Welling, M. Graph convolutional matrix completion. *arXiv preprint arXiv:1706.02263*, 2017.
- Bianchi, F. M., Grattarola, D., Livi, L., and Alippi, C. Graph neural networks with convolutional ARMA filters. *arXiv preprint arXiv:1901.01343*, 2019.
- Bianchi, F. M., Grattarola, D., and Alippi, C. Spectral clustering with graph neural networks for graph pooling. In *Proceedings of the 37th international conference on Machine learning*. ACM, 2020.
- Cangea, C., Velicković, P., Jovanović, N., Kipf, T., and Liò, P. Towards sparse hierarchical graph classifiers. *arXiv preprint arXiv:1811.01287*, 2018.
- Chollet, F. et al. Keras. <https://keras.io>, 2015.
- De Cao, N., Aziz, W., and Titov, I. Question answering by reasoning across documents with graph convolutional networks. *arXiv preprint arXiv:1808.09920*, 2018.
- Defferrard, M., Bresson, X., and Vandergheynst, P. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in Neural Information Processing Systems*, pp. 3844–3852, 2016.
- Do, K., Tran, T., and Venkatesh, S. Graph transformation policy network for chemical reaction prediction. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 750–760, 2019.
- Du, J., Zhang, S., Wu, G., Moura, J. M., and Kar, S. Topology adaptive graph convolutional networks. *arXiv preprint arXiv:1710.10370*, 2017.
- Farrell, S., Calafiura, P., Mudigonda, M., Prabhat, Anderson, D., Vlimant, J.-R., Zheng, S., Bendavid, J., Spiropulu, M., Cerati, G., Gray, L., Kowalkowski, J., Spentzouris, P., and Tsaris, A. Novel deep learning methods for track reconstruction, 2018.
- Fernandes, P., Allamanis, M., and Brockschmidt, M. Structured neural summarization. *arXiv preprint arXiv:1811.01824*, 2018.
- Fey, M. and Lenssen, J. E. Fast graph representation learning with pytorch geometric. *arXiv preprint arXiv:1903.02428*, 2019.
- Gainza, P., Sverrisson, F., Monti, F., Rodolà, E., Boscai, D., Bronstein, M., and Correia, B. Deciphering interaction fingerprints from protein molecular surfaces using geometric deep learning. *Nature Methods*, 17(2):184–192, 2020.
- Gao, H. and Ji, S. Graph u-nets. *CoRR*, abs/1905.05178, 2019. URL <http://arxiv.org/abs/1905.05178>.
- Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and Dahl, G. E. Neural message passing for quantum chemistry. *arXiv preprint arXiv:1704.01212*, 2017.
- Hamilton, W., Ying, Z., and Leskovec, J. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*, pp. 1024–1034, 2017.
- Hamrick, J. B., Allen, K. R., Bapst, V., Zhu, T., McKee, K. R., Tenenbaum, J. B., and Battaglia, P. W. Relational inductive bias for physical construction in humans and machines. *arXiv preprint arXiv:1806.01203*, 2018.
- Ivanov, S., Sviridov, S., and Burnaev, E. Understanding isomorphism bias in graph data sets, 2019.
- Keras. Why use Keras - Keras Documentation, 2019. URL <https://keras.io/why-use-keras/>.
- Kersting, K., Kriege, N. M., Morris, C., Mutzel, P., and Neumann, M. Benchmark data sets for graph kernels, 2016. URL <https://ls11-www.cs.tu-dortmund.de/staff/morris/graphkerneldatasets>.
- Kipf, T., Fetaya, E., Wang, K.-C., Welling, M., and Zemel, R. Neural relational inference for interacting systems. *arXiv preprint arXiv:1802.04687*, 2018.

- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*, 2016.
- Klicpera, J., Bojchevski, A., and Günnemann, S. Predict then propagate: Graph neural networks meet personalized pagerank. In *International Conference on Learning Representations (ICLR)*, 2019.
- Lee, J., Lee, I., and Kang, J. Self-attention graph pooling. *CoRR*, abs/1904.08082, 2019. URL <http://arxiv.org/abs/1904.08082>.
- Li, Y., Tarlow, D., Brockschmidt, M., and Zemel, R. Gated graph sequence neural networks. *arXiv preprint arXiv:1511.05493*, 2015.
- Li, Y., Yu, R., Shahabi, C., and Liu, Y. Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. *arXiv preprint arXiv:1707.01926*, 2017.
- Ramakrishnan, R., Dral, P. O., Rupp, M., and Von Lilienfeld, O. A. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific data*, 1:140022, 2014.
- Sanchez-Gonzalez, A., Heess, N., Springenberg, J. T., Merel, J., Riedmiller, M., Hadsell, R., and Battaglia, P. Graph networks as learnable physics engines for inference and control. *arXiv preprint arXiv:1806.01242*, 2018.
- Santoro, A., Raposo, D., Barrett, D. G., Malinowski, M., Pascanu, R., Battaglia, P., and Lillicrap, T. A simple neural network module for relational reasoning. In *Advances in neural information processing systems*, pp. 4967–4976, 2017.
- Scarselli, F., Gori, M., Tsoi, A. C., Hagenbuchner, M., and Monfardini, G. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2009.
- Schlichtkrull, M., Kipf, T. N., Bloem, P., Van Den Berg, R., Titov, I., and Welling, M. Modeling relational data with graph convolutional networks. In *European Semantic Web Conference*, pp. 593–607. Springer, 2018.
- Sen, P., Namata, G., Bilgic, M., Getoor, L., Galligher, B., and Eliassi-Rad, T. Collective classification in network data. *AI magazine*, 29(3):93–93, 2008.
- Shang, J., Xiao, C., Ma, T., Li, H., and Sun, J. Gamenet: graph augmented memory networks for recommending medication combination. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 1126–1133, 2019.
- Shchur, O., Mumme, M., Bojchevski, A., and Günnemann, S. Pitfalls of graph neural network evaluation. *arXiv preprint arXiv:1811.05868*, 2018.
- Simonovsky, M. and Komodakis, N. Dynamic edge-conditioned filters in convolutional neural networks on graphs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- Stark, C., Breikreutz, B.-J., Reguly, T., Boucher, L., Breikreutz, A., and Tyers, M. Biogrid: a general repository for interaction datasets. *Nucleic acids research*, 34 (suppl\_1):D535–D539, 2006.
- Thekumparampil, K. K., Wang, C., Oh, S., and Li, L.-J. Attention-based graph neural network for semi-supervised learning. *arXiv preprint arXiv:1803.03735*, 2018.
- Velickovic, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., and Bengio, Y. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.
- Wang, M., Yu, L., Zheng, D., Gan, Q., Gai, Y., Ye, Z., Li, M., Zhou, J., Huang, Q., Ma, C., et al. Deep graph library: Towards efficient and scalable deep learning on graphs. *arXiv preprint arXiv:1909.01315*, 2019.
- Wang, Y., Sun, Y., Liu, Z., Sarma, S. E., Bronstein, M. M., and Solomon, J. M. Dynamic graph cnn for learning on point clouds.(2018). *arXiv preprint arXiv:1801.07829*, 2018.
- Xie, T. and Grossman, J. C. Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Physical review letters*, 120 (14):145301, 2018.
- Xu, K., Hu, W., Leskovec, J., and Jegelka, S. How powerful are graph neural networks? In *International Conference on Learning Representations (ICLR)*, 2019.
- Ying, R., He, R., Chen, K., Eksombatchai, P., Hamilton, W. L., and Leskovec, J. Graph convolutional neural networks for web-scale recommender systems. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 974–983, 2018a.
- Ying, R., You, J., Morris, C., Ren, X., Hamilton, W. L., and Leskovec, J. Hierarchical graph representation learning with differentiable pooling. *arXiv preprint arXiv:1806.08804*, 2018b.
- You, J., Ying, R., Ren, X., Hamilton, W. L., and Leskovec, J. Graphrnn: Generating realistic graphs with deep autoregressive models. *arXiv preprint arXiv:1802.08773*, 2018.
- Zambaldi, V., Raposo, D., Santoro, A., Bapst, V., Li, Y., Babuschkin, I., Tuyls, K., Reichert, D., Lillicrap, T., Lockhart, E., et al. Relational deep reinforcement learning. *arXiv preprint arXiv:1806.01830*, 2018.

Zhang, M., Cui, Z., Neumann, M., and Chen, Y. An end-to-end deep learning architecture for graph classification. In *Proceedings of AAAI Conference on Artificial Intelligence*, 2018.

Zitnik, M. and Leskovec, J. Predicting multicellular function through multi-layer tissue networks. *Bioinformatics*, 33(14):i190–i198, 2017.

## A. Data modes

Spektral supports four different ways of representing graphs (or batches thereof), which we refer to as *data modes*.

In *single mode*, the data describes a single graph with its adjacency matrix and attributes, and inference usually happens at the level of individual nodes. *Disjoint mode* is a special case of single mode, where the graph is obtained as the disjoint union of a set of smaller graphs. In this case the node attributes of the graphs are stacked in a single matrix and their adjacency matrices are combined in a block-diagonal matrix. This is a practical way of representing batches of variable-order graphs, although it requires an additional data structure to keep track of the different graphs in a batch. Alternatively, in *batch mode*, a set of graphs is represented by stacking their adjacency and attributes matrices in higher-order tensors of shape  $B \times N \times \dots$ . This mode is akin to traditional batch processing in machine learning and can be more naturally adopted in deep learning architectures. However, it requires the graphs in a batch to have the same number of nodes. Finally, *mixed mode* is the one most often found in traditional machine learning literature and consists of a graph with fixed support but variable attributes. Common examples of this mode are found in computer vision, where images have a fixed 2-dimensional grid support and only differ in the pixel values (*i.e.*, the node attributes), and in traditional graph signal processing applications.

In Spektral, all layers implementing message-passing operations are compatible with single and disjoint mode, and more than half of the layers also support mixed and batch mode. A similar consideration holds for pooling methods where however, due to structural limits of the methods themselves, some data modes cannot be supported directly.

## B. Technical notes

Spektral is distributed through the Python Package Index (package name: `spektral`), supports all UNIX-like platforms,<sup>8</sup> and has no proprietary dependencies. The library is compatible with Python version 3.5 and above. Starting from version 0.2, Spektral is developed for TensorFlow 2 and its integrated implementation of Keras. Version 0.1 of Spektral, which is based on TensorFlow 1 and the standalone version of Keras, will be maintained until TensorFlow 1 is officially discontinued, although new features will only be added to the newer versions of Spektral.

<sup>8</sup>It is also largely compatible with Windows.

## C. Experimental details

This section summarises the architectures and hyperparameters used in the experiments of Section 4. All experiments were run on a single NVIDIA Titan V GPU with 12GB of video memory.

### C.1. Node classification

Hyperparameters:

- Learning rate: see original papers;
- Weight decay: see original papers;
- Epochs: see original papers;
- Patience: see original papers;
- Repetitions per method and per dataset: 100;
- Data: we used Cora, Citeseer and Pubmed. As suggested in (Shchur et al., 2018), we use random splits with 20 labels per class for training, 30 labels per class for early stopping, all the remaining labels for testing.

### C.2. Graph classification

We configure a GNN with the following structure: GCS - POOLING - GCS - POOLING - GCS - GLOBAL-SUMPOOLING - DENSE, where GCS indicates a *Graph Convolutional Skip* layer as described in (Bianchi et al., 2019), POOLING indicates the graph pooling layer being tested, GLOBALSUMPOOLING represents a global sum pooling layer, and DENSE represents the fully-connected output layer. GCS layers have 32 units each, ReLU activation, and L2 regularisation applied to both weight matrices. The same L2 regularization is applied to pooling layers when possible. Top- $K$  and SAGPool layers are configured to output half of the nodes for each input graph. DiffPool and MinCut are configured to output  $K = \frac{\bar{N}}{2}$  nodes at the first layer, and  $K = \frac{\bar{N}}{4}$  nodes at the second layer, where  $\bar{N}$  is the average order of the graphs in the dataset. When using DiffPool, we remove the first two GCS layers, because DiffPool has an internal message-passing layer for the input features. DiffPool and MinCut were trained in batch mode by zero-padding the adjacency and node attributes matrices. All networks were trained using Adam with the default parameters of Keras, except for the learning rate.

Hyperparameters:

- Batch size: 8;
- Learning rate: 0.001;
- Weight decay: 0.00001;



- Epochs: models trained to convergence;
- Patience: 50 epochs;
- Repetitions per method and per dataset: 10;
- Data: we used the Benchmark Datasets for Graph Kernels as described in (Ivanov et al., 2019), that were modified to contain no isomorphic graphs. For each run, we randomly split the dataset and use 80% of the data for training, 10% for early stopping, and 10% for testing.

### C.3. Graph regression

We configure a GNN with the following structure: ECC - ECC - GLOBALPOOLING - DENSE, where ECC indicates an Edge-Conditioned Convolutional layer (Simonovsky & Komodakis, 2017) and GLOBALPOOLING indicates the global pooling layer being tested. ECC layers have 32 units each, and ReLU activation. No regularization is applied to the GNN. GAP is configured to use 32 units. All networks were trained using Adam with the default parameters of Keras, except for the learning rate. We use the mean squared error as loss.

Node features are one-hot encodings of the atomic number of each atom. Edge features are one-hot encodings of the bond type. The units of measurement for the target variables are: debye units (D) for  $\mu$ ,  $a_0^3$  ( $a_0$  is the Bohr radius) for Alpha, and Hartree (Ha) for Homo and U0 (Ramakrishnan et al., 2014).

Hyperparameters:

- Batch size: 32;
- Learning rate: 0.0005;
- Epochs: models trained to convergence;
- Patience: 10 epochs;
- Repetitions per method and per dataset: 5;
- Data: for each run, we randomly split the dataset and use 80% of the molecules for training, 10% for early stopping, and 10% for testing.