
Learning Distributed Representations of Graphs with Geo2DR

Paul Scherer¹ Pietro Liò¹

Abstract

We present Geo2DR (*Geometric to Distributed Representations*), a GPU ready Python library for unsupervised learning on graph-structured data using discrete substructure patterns and neural language models. It contains efficient implementations of popular graph decomposition algorithms and neural language models in PyTorch which can be combined to learn representations of graphs using the distributive hypothesis. Furthermore, Geo2DR comes with general data processing and loading methods to bring substantial speed-up in the training of the neural language models. Through this we provide a modular set of tools and methods to quickly construct systems capable of learning distributed representations of graphs. This is useful for replication of existing methods, modification, or development of completely new methods. This paper serves to present the Geo2DR library and perform a comprehensive comparative analysis of existing methods re-implemented using Geo2DR across widely used graph classification benchmarks. Geo2DR displays a high reproducibility of results in published methods and interoperability with other libraries useful for distributive language modelling.

1. Introduction

Representation learning of graphs using neural networks has turned into a large and exciting hub of research driven by successive proposals of graph representation learning methods and datasets to apply them onto. A significant part of the activity has focused on *Graph Convolutional Neural Networks* (GCNN). Such neural networks are characterised by *graph convolutional* operators (Belkin & Niyogi, 2001; Defferrard et al., 2016; Kipf & Welling, 2017) that serve as useful inductive biases for learning representations of nodes and other graph substructures. Gilmer et al. (2017)

generalised the convolution operator over irregular domains as a message passing scheme, allowing the specification of a full spectrum of methods as variants of this equation. Representations of entire graphs are then created through the successive application of message passing operations followed by different *pooling* methods (Defferrard et al., 2016; Ying et al., 2018; Luzhnica et al., 2019) which aggregate node representations towards a single vector representation for the entire graph.

The difficulty of reliably constructing GCNN models has driven the need for toolkits and libraries to facilitate their development for replication, extension and creation of new models. Several such libraries have been made including: *Graph Nets* introduced by Battaglia et al. (2018), *DGL* by Wang et al. (2019), *GEM* by Goyal et al. (2018), and most recently *PyTorch Geometric* by Fey and Lenssen (2019). These libraries have greatly contributed to lowering the barrier of entry into GCNN research, fueling the development of novel methods and libraries supporting them in a healthy feedback cycle.

Alongside ongoing research into GCNNs and its variants, another approach has focused on extending graph kernel methods with neural language embedding methods (Yanardag & Vishwanathan, 2015; Narayanan et al., 2017; Ivanov & Burnaev, 2018) that exploit the distributive hypothesis to learn *distributed representations* of graphs. This is a useful alternative inductive bias to model the vector space embeddings of graphs over the distribution of the discrete substructure patterns *contextualising* them. Much like how the semantic meaning of words is similar to words that have similar context words around them (Harris, 1954), distributed representations of graphs are inductively biased to be similar if they contain similar substructure patterns, and dissimilar otherwise. This perspective enables the construction of a powerful class of unsupervised representation learning methods.

However, to our knowledge, no toolkit currently exists for rapidly composing methods capable of learning distributed representations of graphs. This project, Geo2DR, aims to fill this gap. The library along with links to documentation, example methods, experiment replication, and supporting material can be found on the GitHub repository¹.

¹Department of Computer Science and Technology, University of Cambridge, Cambridge, United Kingdom. Correspondence to: Paul Scherer <paul.scherer@cl.cam.ac.uk>.

¹<https://github.com/paulmorio/geo2dr>

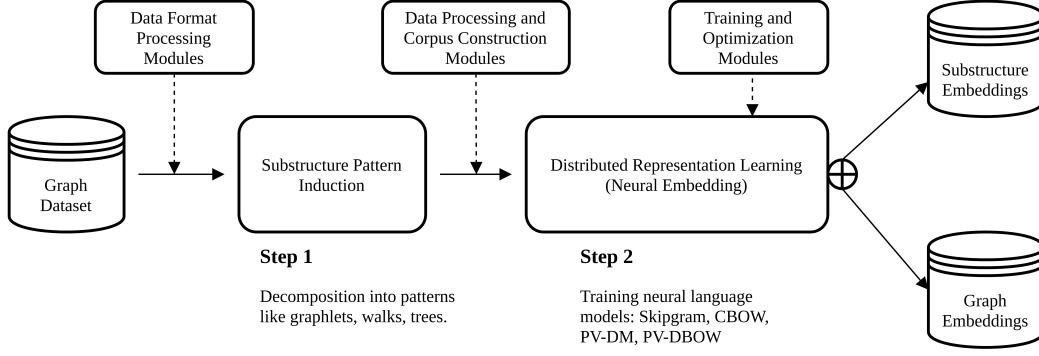


Figure 1. The two-stage design methodology for creating distributed representations of graphs and the various modules (in rectangles) included in Geo2DR to support this process. Each module can also be used independently for other tasks as mentioned in Section 3 and 4.

2. Background and Related Work

The approach towards distributive modelling of graphs was pioneered by Yanardag and Vishwanathan (2015). They observed that many graph kernel methods can be formulated as instances of the R-Convolutional framework. Herein, the similarity between different graphs is computed by decomposing graphs into discrete substructure patterns such as graphlets, shortest paths, and rooted subgraphs. This produces a $|\mathbf{V}|$ -dimensional bag-of-words or pattern frequency vectors for each graph where \mathbf{V} is the set of the unique patterns induced over all the graphs in a dataset. The graphs and their induced substructure patterns are input to a kernel function, such as counting the common substructures across pattern frequency vectors. This defines the relation or similarity measure between the graphs to construct the kernel matrix for use with kernel methods such as SVMs.

Yanardag and Vishwanathan (2015) further observed that as the size of graphs and the specificity of substructure patterns to be induced from graphs increases (via lengthening walks/paths, increasing the number of nodes in graphlet patterns) graphs are represented by extremely high dimensional pattern frequency vectors. As a result, only few substructure patterns are common across any given set of graphs producing sparse solutions where each graph is more similar to itself, a phenomenon known as *diagonal dominance*. To tackle this issue the authors proposed the use of neural language models which exploit the distributive hypothesis (Harris, 1954) to learn smooth low dimensional *distributed representations* of the substructures and construct graph kernel matrices. This was quickly followed up by works such as the aptly named Graph2Vec (Narayanan et al., 2017) and Anonymous Walk Embeddings (Ivanov & Burnaev, 2018) (AWE). These proposed different substructure patterns graphs could be reduced to and the use of Doc2Vec variants (Le & Mikolov, 2014) to build distributed representations of whole graphs directly.

Geo2DR provides various modules that can be used as “building blocks” to rapidly construct systems capable of learning such distributed representations of both substructure patterns and whole graphs of arbitrary size. Existing libraries for GNNs (Battaglia et al., 2018; Wang et al., 2019; Goyal & Ferrara, 2018; Fey & Lenssen, 2019) would require a substantial shift in philosophical focus from constructing message passing schemes and pooling methods. Hence Geo2DR is a complementary library alongside existing toolkits enabling researchers a broader range of options and tools for graph representation learning.

3. Overview of Geo2DR

Geo2DR contains various “building blocks” for rapid construction of methods capable of learning distributed representations of graphs. The conceptual framework for unsupervised learning of the representations for substructures and entire graphs is based around a simple two stage design methodology summarised in Figure 1.

Induction of descriptive substructure patterns: The first step consists of inducing discrete substructure patterns such as graphlets, rooted subgraphs, or anonymous walks within and across the dataset of graphs to construct a shared vocabulary and *corpus* dataset contextualizing the patterns and graphs. One may also use the output pattern distributions at this stage to construct a variety of graph kernels.

Learning distributed vector representations: The second stage consists of utilising the distributive hypothesis (Harris, 1954) to learn distributed representations of graphs contextualised by the induced substructure patterns. Embedding methods which exploit the distributive hypothesis such as skipgram (Mikolov et al., 2014) can be used to learn fixed-size vector embeddings of substructure patterns or whole graph in an unsupervised manner.

Combination of Geo2DR’s modules for decomposition and

Table 1. Table characterising each of the existing published methods by the substructure patterns induced and associated embedding method to create the graph kernel matrix (for DGK models) or graph embeddings.

Method	Induced substructure pattern	Embedding method	Object embedded
DGK-WL	WL rooted subgraphs	Skipgram or CBOW	Substructure patterns
DGK-SP	Shortest paths	Skipgram or CBOW	Substructure patterns
DGK-GK	Graphlets	Skipgram or CBOW	Substructure patterns
Graph2Vec	WL rooted subgraphs	PV-DBOW	Whole graphs
AWE-DD	Anonymous walks	PV-DM	Whole graphs

distributed representation learning methods can be used to quickly replicate existing models such as those shown in Table 1. Consistent input/output interfaces were implemented across modules to encourage exploration of novel methods. For example, one could create a “novel” unpublished method combining existing modules on inducing shortest path patterns and learning graph-level embeddings with PV-DBOW. This sort of experimentation fosters understanding and better control of the inductive biases involved in a graph learning task. We hope it would also encourage the creation of custom modules that can plug and play with the rest of the framework to create truly novel methods.

Practically, the library is centered around three subpackages under Geo2DR. The `data` subpackage, contains modules for transforming data formats used by popular dataset repositories such as Kersting et al. (2016) into consistent formats used by the decomposition algorithms implemented in Geo2DR. In Geo2DR, we chose to use the GEXF (Graph Exchange XML Format) as permanent storage format for individual instances of the graphs. This is because the format is compatible with network analysis software such as Gephi and NetworkX for detailed inspection.

The modules within the `decomposition` subpackage contain algorithms for inducing the substructure patterns in the graphs and forming vocabularies. The outputs of these algorithms are directly compatible with our PyTorch implementations of neural language models to utilize GPUs as well as those in Gensim (Řehůřek & Sojka, 2010). This essentially describes the packages and modules necessary for Step 1 of the process. The final subpackage `embedding_methods` contains modules for constructing corpus datasets and neural language models to build the distributed representation learning methods of Step 2. Several `Trainer` classes are also included which serve as battery-included corpus and neural net combinations that can be used to construct common architecture setups.

Existing methods for learning distributed representations as in Table 1 and several graph kernels can all be implemented using the modules and conceptual framework presented. We have included all methods as examples within the repository to get users started on creating their own variations. A brief code example is included in Appendix A.

4. Empirical Evaluation

As a form of validation for the various implemented modules, we empirically evaluate re-implementations of existing models using Geo2DR. Table 1 describes the induced substructure pattern and neural language model driving each method. We performed a series of common benchmark graph classification tasks under homogeneous data and evaluation scenarios giving a fairer picture of how they compare.

All datasets were downloaded from the benchmark dataset repository by Kersting et al. (2016) and processed into the format used by Geo2DR with the included data formatter. In each of the datasets the discrete node labels are exposed, but not the edge labels. For unlabelled datasets such as REDDIT-B, the node was labelled by their degree following practice of Shervashidze et al. (2011) to enable methods such as the WL rooted subgraph decomposition to induce patterns in the graphs; this was also applied to methods which can directly handle unlabelled graphs for conformity. As these datasets are standard benchmarks we have left specific descriptive details in Appendix B.

For all experiments, attempts were made to follow the hyperparameter setups described in the published papers of the original methods, with best-guess settings where details were not published. As we look at several kernels and embedding models specific hyperparameter ranges can be found in Appendix C. In all cases, the same off-the-shelf SVM implemented in SciKit-Learn (Pedregosa et al., 2011) was used with an RBF kernel trick for the supervised classification task on the graph embeddings learned. C values were estimated over the set (0.001, 0.01, 0.1, 1, 10, 100). We report the average score of 10 iterations of training and applying 10 fold cross-validation using the SVM over random data splits with individual training restarts in all cases. The exact setups of the experiments can be replicated using the experiment replication code provided within the Github repository².

Graph kernels: We start with an experiment suite based on the substructure patterns alone, using the decomposition algorithms to construct normalised bag-of-words frequency vectors for each of the graphs. Table 2 records the mean and

²<https://github.com/paulmorio/geo2dr/tree/master/replication>

Table 2. Random-split 10 fold cross-validation performance of SVM using RBF kernel on bag-of-words vectors of normalised frequencies of induced substructure patterns. Best scores or those within error of best are bolded. OOM denotes out-of-memory.

Substructure pattern	MUTAG	ENZYMES	PROTEINS	NCI1	REDDIT-B	IMDB-M
WL Rooted Subgraphs	88.95 ± 7.96	56.33 ± 6.18	74.29 ± 2.55	83.94 ± 1.99	77.35 ± 4.35	48.60 ± 4.33
Shortest Paths	83.68 ± 7.24	41.67 ± 4.83	74.73 ± 2.04	70.95 ± 1.95	OOM	50.20 ± 3.84
Graphlets	83.16 ± 6.16	25.33 ± 3.48	70.36 ± 3.59	54.09 ± 7.61	78.25 ± 2.71	44.40 ± 4.17
Anonymous Walks	80.53 ± 6.68	27.33 ± 6.23	71.87 ± 2.05	66.08 ± 2.21	81.30 ± 2.49	38.20 ± 3.91

Table 3. Graph classification performance over random-split 10 fold cross-validation in each of the re-implemented systems with standard deviation. Best scores or those within error of best are bolded. OOM denotes out-of-memory.

Method	MUTAG	ENZYMES	PROTEINS	NCI1	REDDIT-B	IMDB-M
DGK-WL	88.42 ± 8.42	41.00 ± 1.83	72.08 ± 0.74	77.54 ± 3.91	OOM	47.82 ± 0.79
DGK-SP	84.03 ± 7.16	44.27 ± 2.26	76.93 ± 2.56	69.22 ± 5.29	OOM	49.71 ± 1.18
DGK-GK	84.21 ± 6.74	23.61 ± 3.14	69.77 ± 3.13	53.92 ± 4.81	78.32 ± 1.92	44.40 ± 4.18
Graph2Vec	84.91 ± 2.79	51.77 ± 1.75	74.05 ± 2.28	71.34 ± 2.12	81.25 ± 2.64	47.11 ± 1.42
AWE-DD	79.29 ± 2.92	23.76 ± 1.74	69.70 ± 1.29	63.54 ± 1.82	81.46 ± 1.75	40.53 ± 6.42

standard deviation of randomly split 10 fold cross-validation using the SVM described above. The results closely match that of the published methods in (Yanardag & Vishwanathan, 2015; Shervashidze et al., 2011; Borgwardt & Kriegel, 2005; Ivanov & Burnaev, 2018).

Deep graph kernels and graph embeddings: Most of our experiments in Table 3 show a high reproducibility of the results published by the original proposers. Some discrepancies are to be expected due to the homogenised data setup, unpublished hyperparameter settings, and standardised neural architectures, but best effort was made by consulting original source code and communications with the authors. In particular, for AWE-DD, we do not use edge-labels for homogeneity of the experiment evaluation whilst the original paper used them if they gave better scores.

Runtime experiments and improvements in Geo2DR: Table 4 contains the average total training times incurred over 100 epochs, performed ten times with one standard deviation on a single quad-core Intel i5-4690 CPU. Comparison is drawn between the original reference implementation made available by each of the original papers and its re-implemented counterpart in Geo2DR. All methods were trained and compared on the MUTAG dataset as this was the only common dataset included in the reference implementations. None of the original reference implementations have scripts or tools to transform the publicly available datasets they used into the proprietary formats used by their own implementations, making reproduction difficult. This is why we have included data processing tools directly into the Geo2DR library to address this common limitation for the future.

Table 4. Total training run time (seconds) over 100 epochs on MUTAG. Bold text refers to lowest time taken for training or are within error bounds of being the fastest.

Method	Original reference implementation	Only Geo2DR PyTorch modules	Geo2DR with compatible libraries Gensim/TensorFlow
DGK-WL	3.06 ± 0.15	3.33 ± 0.07	3.19 ± 0.08
DGK-SP	6.95 ± 0.23	6.86 ± 0.27	7.39 ± 0.08
DGK-GK	9.46 ± 0.69	19.41 ± 0.49	9.89 ± 0.74
Graph2Vec	8.86 ± 0.05	10.64 ± 0.11	8.88 ± 0.06
AWE-DD	1231.75 ± 21.81	314.84 ± 8.91	—

5. Conclusion

Through the characterisation of existing methods, and the reproduction of their results in Geo2DR, we have shown that the library is a successful amalgamation of the various components that enable learning distributed representations of graphs. Using the simple design methodology, one can quickly re-implement existing models, an increasingly important part of reproducible research and designing novel architectures. By exploiting the modular structure and compatibility with other software and libraries the set of tools for constructing learning methods is broadened without having to deal with different data formats, language paradigms and workflows used by individual implementations. Using a host of re-implemented methods also allows for more homogenised experiment suites that can be used to more fairly compare existing and new methods in future research efforts. Geo2DR is available now with numerous examples and documentation as a starting point. The library will continue to evolve to add new components, compatibility with other libraries, tutorials, and accommodate new developments in the field.

Acknowledgements

Foremost, we would like to thank Dr. Yanardag (Yanardag & Vishwanathan, 2015), Dr. Narayanan (Narayanan et al., 2017) and Dr. Ivanov (Ivanov & Burnaev, 2018) for their correspondence, and making reference code available publicly. Furthermore we would like to thank the members of the AI Group at the Computer Laboratory for their patience and numerous proof readings of this work.

References

- Battaglia, P. W., Hamrick, J. B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V. F., Malinowski, M., and et al. Relational inductive biases, deep learning, and graph networks. *CoRR*, abs/1806.01261, 2018. URL <http://arxiv.org/abs/1806.01261>.
- Belkin, M. and Niyogi, P. Laplacian eigenmaps and spectral techniques for embedding and clustering. In *Proceedings of the 14th International Conference on Neural Information Processing Systems: Natural and Synthetic*, NeurIPS’01, pp. 585–591, Cambridge, MA, USA, 2001. MIT Press. URL <http://dl.acm.org/citation.cfm?id=2980539.2980616>.
- Borgwardt, K. M. and Kriegel, H.-P. Shortest-path kernels on graphs. In *Proceedings of the Fifth IEEE International Conference on Data Mining*, ICDM’05, pp. 74–81, Washington, DC, USA, 2005. IEEE Computer Society. ISBN 0-7695-2278-5. doi: 10.1109/ICDM.2005.132.
- Borgwardt, K. M., Ong, C. S., Schönauer, S., Vishwanathan, S. V. N., Smola, A. J., and Kriegel, H.-P. Protein function prediction via graph kernels. *Bioinformatics*, 21(1):47–56, 2005. ISSN 1367-4803.
- Debnath, A. K., Lopez de Compadre, R. L., Debnath, G., Shusterman, A. J., and Hansch, C. Structure-activity relationship of mutagenic aromatic and heteroaromatic nitro compounds. correlation with molecular orbital energies and hydrophobicity. *Journal of Medicinal Chemistry*, 34(2):786–797, Feb 1991. ISSN 0022-2623. doi: 10.1021/jm00106a046.
- Defferrard, M., Bresson, X., and Vandergheynst, P. Convolutional neural networks on graphs with fast localized spectral filtering. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, NeurIPS’16, pp. 3844–3852, USA, 2016. Curran Associates Inc. ISBN 978-1-5108-3881-9. URL <http://dl.acm.org/citation.cfm?id=3157382.3157527>.
- Fey, M. and Lenssen, J. E. Fast graph representation learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.
- Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and Dahl, G. E. Neural message passing for quantum chemistry. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ICML’17, pp. 1263–1272. JMLR.org, 2017. URL <http://dl.acm.org/citation.cfm?id=3305381.3305512>.
- Goyal, P. and Ferrara, E. Graph embedding techniques, applications, and performance: A survey. *Knowledge-Based Systems*, 151:78 – 94, 2018. ISSN 0950-7051. doi: <https://doi.org/10.1016/j.knosys.2018.03.022>. URL <http://www.sciencedirect.com/science/article/pii/S0950705118301540>.
- Harris, Z. S. Distributional structure. *WORD*, 10(2-3): 146–162, 1954. doi: 10.1080/00437956.1954.11659520.
- Ivanov, S. and Burnaev, E. Anonymous walk embeddings. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 2191–2200, Stockholmsmässan, Stockholm Sweden, 2018. PMLR. URL <http://proceedings.mlr.press/v80/ivanov18a.html>.
- Kersting, K., Kriege, N. M., Morris, C., Mutzel, P., and Neumann, M. Benchmark data sets for graph kernels, 2016. Datasets available at <http://graphkernels.cs.tu-dortmund.de>.
- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. In *Proceedings of the 5th International Conference on Learning Representations*, ICLR’17, 2017.
- Le, Q. and Mikolov, T. Distributed representations of sentences and documents. In *Proceedings of the 31st International Conference on International Conference on Machine Learning - Volume 32*, ICML’14, pp. II–1188–II–1196. JMLR.org, 2014. URL <http://dl.acm.org/citation.cfm?id=3044805.3045025>.
- Luzhnica, E., Day, B., and Liò, P. On graph classification networks, datasets and baselines. In *36th International Conference on Machine Learning*, ICML’19, 2019.
- Narayanan, A., Chandramohan, M., Venkatesan, R., Chen, L., Liu, Y., and Jaiswal, S. graph2vec: Learning distributed representations of graphs. *CoRR*, abs/1707.05005, 2017.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., and et al. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

Řehůřek, R. and Sojka, P. Software Framework for Topic Modelling with Large Corpora. In *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, pp. 45–50, Valletta, Malta, 2010. ELRA.

Shervashidze, N., Schweitzer, P., van Leeuwen, E. J., Mehlhorn, K., and Borgwardt, K. M. Weisfeiler-lehman graph kernels. *J. Mach. Learn. Res.*, 12:2539–2561, 2011. ISSN 1532-4435.

Wale, N., Watson, I. A., and Karypis, G. Comparison of descriptor spaces for chemical compound retrieval and classification. *Knowl. Inf. Syst.*, 14(3):347–375, March 2008. ISSN 0219-1377. doi: 10.1007/s10115-007-0103-5.

Wang, M., Yu, L., Zheng, D., Gan, Q., Gai, Y., Ye, Z., and et al. Deep graph library: Towards efficient and scalable deep learning on graphs. *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019. URL <https://arxiv.org/abs/1909.01315>.

Yanardag, P. and Vishwanathan, S. Deep graph kernels. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD’15, pp. 1365–1374, New York, NY, USA, 2015. ACM. ISBN 978-1-4503-3664-2. doi: 10.1145/2783258.2783417. URL <http://doi.acm.org/10.1145/2783258.2783417>.

Ying, R., You, J., Morris, C., Ren, X., Hamilton, W. L., and Leskovec, J. Hierarchical graph representation learning with differentiable pooling. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NeurIPS’18, pp. 4805–4815, USA, 2018. Curran Associates Inc. URL <http://dl.acm.org/citation.cfm?id=3327345.3327389>.

A. Code Example

We present a construction of a simplified Graph2Vec model and training it to produce 32 dimensional distributed vector embeddings of the MUTAG graphs using Geo2DR modules. To start we need to download a dataset to study. We will use the well known MUTAG (Debnath et al., 1991) downloaded from the TU Dortmund Graph Kernel Benchmark website (Kersting et al., 2016). Assume we have unpacked and saved the data into a directory called `org_data/` so the dataset as downloaded will be within the directory as `org_data/MUTAG/`.

Geo2DR uses the GEXF (Graph Exchange XML Format) as the permanent storage format for the graphs in a dataset. This is because it is compatible with network analysis software such as Gephi and NetworkX, and it is often useful to be able to study each graph individually; identified by a single file. Due to this design choice we need to transform

the format of the downloaded dataset using tools available within the `data` subpackage as in the code sample below.

```
1 from geometric2dr.data import DortmundGexf
2
3 gexifier = DortmundGexf("MUTAG", "org_data/", "data/")
4 gexifier.format_dataset()
```

Listing 1. Formatting the downloaded dataset into GEXF format

This will result in the following dataset format:

- `data/MUTAG/` : a directory containing individual `.gexf` files of each graph. A graph will be denoted by the graph IDs used in the original data. In this case graph 0 would be `data/MUTAG/0.gexf`
- `data/MUTAG.Labels` : a plain-text file with each line containing a graph’s file path and its classification label.

Given the preprocessed data we can now induce substructure patterns across the graph files. Here we will induce rooted subgraphs up to depth 2 using the Weisfeiler-Lehman node relabeling algorithm outlined in Shervashidze et al. (2011).

```
1 from geometric2dr.decomposition.
   weisfeiler_lehman_patterns import wl_corpus
2 import geometric2dr.embedding_methods.utils as utils
3
4 dataset_path = "data/MUTAG"
5 graph_files = utils.get_files(dataset_path, ".gexf")
6
7 wl_depth = 2
8 wl_corpus(graph_files, wl_depth)
```

Listing 2. Inducing rooted subgraphs across the graphs of the dataset

The `wl_corpus()` function induces rooted subgraph patterns across the list of `.gexf` files in `graph_files`, and builds a document for each graph describing the induced patterns within. These documents will have a special extension specific to each decomposition algorithm or can be set by the user. In this example the extension will be `.d2wl` to denote a Weisfeiler-Lehman decomposition to depth 2. Generating permanent files as a side effect of the graph decomposition process is useful for later study and also if we want to use the same induced patterns in the upcoming step of learning distributed representations of the graphs.

To learn distributed representations we need to construct a new target-context dataset. In Graph2Vec a graph is contextualised by the substructure patterns within it, and uses the PV-DBOW architecture with negative sampling to directly learn graph-level embeddings. Hence we use the `PVDBOWInMemoryCorpus` which is an extension of a standard `torch.utils.data.dataset` class. This can interface with a standard PyTorch dataloader to load the data into a `embedding_methods.skipgram` class that we train in a loop using a simple and recognizable `torch.nn` workflow.

```

1 import torch
2 import torch.optim as optim
3 from torch.utils.data import DataLoader
4 from geometric2dr.embedding_methods.pvdbow_data_reader
5     import PVDBOWInMemoryCorpus
6 from geometric2dr.embedding_methods.skipgram import
7     Skipgram
8
9 # Instantiate corpus dataset, dataloader and skipgram
10 # architecture
11 corpus = PVDBOWInMemoryCorpus(dataset_path, ".d2wl")
12 dataloader = DataLoader(corpus, batch_size=1000, shuffle=
13     False, collate_fn=corpus.collate)
14 skipgram = Skipgram(num_targets=corpus.num_graphs,
15     vocab_size=corpus.num_subgraphs, emb_dimension=32)
16
17 optimizer = optim.SGD(skipgram.parameters(), lr=0.1)
18 for epoch in range(100):
19     for i, sample_batched in enumerate(dataloader):
20         if len(sample_batched[0]) > 1:
21             pos_target = sample_batched[0]
22             pos_context = sample_batched[1]
23             neg_context = sample_batched[2]
24
25             optimizer.zero_grad()
26             loss = skipgram.forward(pos_target, pos_context,
27                 neg_context)
28             loss.backward()
29             optimizer.step()
30
31 final_graph_embeddings = skipgram.target_embeddings.
32     weight

```

Listing 3. Creating a target-context dataset then attaching a dataloader that feeds the corpus data into a skipgram model and training it.

The completion of the training provides the final graph embeddings. As this is such a common process, Geo2DR also comes with a number of `Trainer` classes which build corpus datasets, loaders, train neural language models, and save their outputs. All of the code above can be replaced with this short trainer.

```

1 from geometric2dr.embedding_methods.pvdbow_trainer
2     import InMemoryTrainer
3
4 trainer = InMemoryTrainer(corpus_dir=dataset_path,
5     extension=".d2wl", output_fh="graph_embeddings.json",
6     emb_dimension=32, batch_size=1000, epochs=100,
7     initial_lr=0.1, min_count=0)
8 trainer.train()
9 final_graph_embeddings = trainer.skipgram.
10     give_target_embeddings()

```

Listing 4. Trainer example of performing all of listing 1.3

B. Supplementary: Dataset Details

Table 5 contains descriptive information about each of the datasets as they were used within the empirical evaluation described in Section 4 of the main paper. All of the datasets are commonly used benchmark datasets downloaded from Kersting et al.’s (2016) repository³. After downloading the datasets they were processed into the format used by Geo2DR with the included data formatter. In each of the datasets the discrete node labels are exposed, but not the

edge labels. For unlabelled datasets such as REDDIT-B and IMDB-M, the nodes are labelled by their degree as in Shervashidze et al. (2011) to enable methods such as the WL rooted subgraph decomposition to induce patterns in the graphs. This was also applied to methods which can directly handle unlabelled graphs for conformity.

The graphs come from a variety of contexts and domains. MUTAG, ENZYMES and PROTEINS are datasets which have their roots in bioinformatics research. The graphs within them represent molecules with nodes representing atoms and edges denoting chemical bonds or spatial proximity between different atoms. Graph labels describe different properties of the molecules such as mutagenicity or whether a protein is an enzyme. NCII is a chemoinformatics dataset describing compounds screened for their ability to suppress or inhibit the growth of a panel of human tumor cell lines. REDDIT-B and IMDB-M are social network based datasets. In REDDIT-B each graph corresponds to an online discussion thread where nodes correspond to users, and there is an edge between the nodes if at least one responded to another’s comment. IMDB-M is a movie collaboration dataset where each graph corresponds to an ego-network of an actor or actress.

C. Supplementary: Hyperparameter Selections of Re-implemented Methods

For each of the methods described in Section 4 we prescribed a grid search over the following hyper-parameter settings inspired by the settings of the original papers:

C.1. Graph Kernels

- **WL Rooted Subgraphs:** Rooted subgraphs up to depth 2 induced.
- **Shortest Paths:** Shortest paths of all pairs of nodes induced.
- **Graphlets:** Graphlets of size 7 induced, sampling 100 graphlets per graph.
- **Anonymous Walks:** Anonymous walks of length 10 induced exhaustively from each node in the graph.

C.2. Deep Graph Kernels and Graph Embeddings

- **DGK-WL:** Rooted subgraphs of up to depth 2 induced. Trained Skipgram model with negative sampling using 10 negative samples with an Adam optimiser for 5 and 100 epochs using batch sizes of 10000 and 1000 with an initial learning rate of 0.1 and 0.01 adjusted by a cosine annealing scheme. Substructure embedding sizes of 2, 5, 10, 25, 50 dimensions were generated. Graph kernels were constructed using the formulation described in Yanardag and Vishwanathan (2015).

³ls11-www.cs.tu-dortmund.de/staff/morris/graphkerneldatasets

Table 5. Descriptive information about datasets used in the experimental evaluation. N refers the number of graphs in the datasets. C is the number of graph classification labels. Avg. Nodes and Avg. Edges denote the average number of nodes and edges found in the graphs of the dataset respectively. Finally Node Labels indicates whether the nodes are discretely labelled. The * refers to datasets which originally do not have node labels, but are subsequently labelled by their degree as described in Shervashidze et al. (2011)

Dataset	N	C	Avg. Nodes	Avg. Edges	Node Labels
MUTAG (Debnath et al., 1991)	188	2	17.93	19.79	Yes
ENZYMES (Borgwardt et al., 2005)	600	6	32.63	62.14	Yes
PROTEINS (Borgwardt et al., 2005)	1113	2	39.06	72.82	Yes
NCI1 (Wale et al., 2008)	4110	2	29.87	32.3	Yes
REDDIT-B (Yanardag & Vishwanathan, 2015)	2000	2	429.63	497.75	No*
IMDB-M (Yanardag & Vishwanathan, 2015)	1500	3	13	65.94	No*

- DGK-SP:** Shortest paths of all pairs of nodes induced. Trained Skipgram model with negative sampling using 10 negative samples with an Adam optimiser for 5 and 100 epochs using batch sizes of 10000 and 1000 with an initial learning rate of 0.1 and 0.01 adjusted by a cosine annealing scheme. Substructure embedding sizes of 2, 5, 10, 25, 50 dimensions were generated. Graph kernels were constructed using the formulation described in Yanardag and Vishwanathan (2015).
- DGK-GK:** Graphlets of size 7 induced, sampling 2, 5, 10, 25, and 50 graphlets for each graph. Trained Skipgram model with negative sampling using 10 negative samples with an Adam optimiser for 5 and 100 epochs using batch sizes of 10000 and 1000 with an initial learning rate of 0.1 and 0.01 adjusted by a cosine annealing scheme. Substructure embeddings of 2, 5, 10, 25, 50 dimensions were generated. Graph kernels were constructed using the formulation described in Yanardag and Vishwanathan (2015).
- Graph2Vec:** Rooted subgraphs of up to depth 2 induced. Trained over PV-DBOW (Skipgram) model with negative sampling using 10 negative samples with an Adam optimiser for 25, 50, 100 epochs and batch sizes of 512, 1024, 2048, 10000 with an initial learning rate of 0.1 adjusted by a cosine annealing scheme. Graph embeddings of 128 and 1024 dimensions were learned.
- AWE-DD:** Anonymous walks of length 10 induced exhaustively. Trained over PV-DM architecture with negative sampling using 10 negative samples with an Adagrad optimiser (as in reference implementation) for 100 epochs with batch sizes 100, 500, 1000, 5000, 10000 with an initial learning rate of 0.1. Window-sizes of 4, 8, 16 were used to extract context anonymous walks around the target anonymous walk in the PV-DM architecture.