

# GraPE: fast and scalable Graph Processing and Embedding

**Luca Cappelletti**

Universita' degli Studi di Milano

**Tommaso Fontana**

Universita' degli Studi di Milano

**Elena Casiraghi**

Universita' degli Studi di Milano

**Vida Ravanmehr**

The Jackson Laboratory for Genomic Medicine

**Tiffany Callahan**

Lawrence Berkeley National Laboratory <https://orcid.org/0000-0002-8169-9049>

**Marcin Joachimiak**

Lawrence Berkeley National Laboratory

**Christopher Mungall**

Lawrence Berkeley National Laboratory

**Justin Reese**

Lawrence Berkeley National Laboratory

**Peter Robinson**

The Jackson Laboratory for Genomic Medicine

**Giorgio Valentini** (✉ [valentini@di.unimi.it](mailto:valentini@di.unimi.it))

AnacletoLab, Dipartimento di Informatica, Università degli Studi di Milano, Italy <https://orcid.org/0000-0002-5694-3919>

---

## Resource

## Keywords:

**Posted Date:** August 26th, 2022

**DOI:** <https://doi.org/10.21203/rs.3.rs-967145/v1>

**License:**   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

---

# GRAPE: Fast and Scalable Graph Processing and Embedding

Luca Cappelletti<sup>1,7</sup>, Tommaso Fontana<sup>1,7</sup>, Elena Casiraghi<sup>1,7</sup>, Vida Ravanmehr<sup>2,3</sup>,  
Tiffany J. Callahan<sup>4</sup>, Marcin P. Joachimiak<sup>5</sup>, Christopher J. Mungall<sup>5</sup>, Peter N.  
Robinson<sup>2</sup>, Justin Reese<sup>5</sup>, and Giorgio Valentini<sup>1,6,7,8</sup>

<sup>1</sup>AnacletoLab, Dipartimento di Informatica, Università degli Studi di Milano, Italy

<sup>2</sup>The Jackson Laboratory for Genomic Medicine, Farmington, CT, USA

<sup>3</sup>Department of Lymphoma and Myeloma, The University of Texas MD Anderson  
Cancer Center, Houston, TX, USA

<sup>4</sup>Department of Biomedical Informatics, Columbia University, New York, USA

<sup>5</sup>Division of Environmental Genomics and Systems Biology, Lawrence Berkeley  
National Laboratory, Berkeley, USA

<sup>6</sup>European Laboratory for Learning and Intelligent Systems (ELLIS)

<sup>7</sup>CINI, National Laboratory in Artificial Intelligence and Intelligent Systems—AIIS,  
Roma, Italy

<sup>8</sup>Data Science Research Center, Università degli Studi di Milano, Italy

## Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
<b>2</b>	<b>Results</b>	<b>4</b>
2.1	Overview of the <i>GRAPE</i> resource: <i>Embiggen</i> and <i>Ensmallen</i>	4
2.2	Fast error-resilient graph loading	6
2.3	<i>GRAPE</i> outperforms state-of-the-art libraries on random walk generation	6
2.3.1	Experimental comparison of graph processing libraries	6
2.3.2	Approximated random walks to process graphs with high-degree nodes	8
2.4	Node and edge embedding models	9
2.5	GraPE enables a fair and reproducible comparison of graph embedding and graph-based prediction methods	9
2.5.1	FAIR graph retrieval	9
	Findability and Accessibility	10
	Interoperability	10
	Reusability	10
2.5.2	FAIR evaluation pipelines	10
2.5.3	Experimental comparison of node embedding methods	11
2.6	Scaling with big real-world graphs	13
2.6.1	Experimental setup	13
	English Wikipedia	13
	Comparative Toxicogenomic Database (CTD)	13
	PheKnowLator biomedical data	13
	Graph libraries compared in the experiments	13
	Evaluation of the results	14

2.6.2	Results	15
	<i>GRAPE</i> is able to scale with big graphs when the other competing libraries fail.	16
	<i>GRAPE</i> improves the empirical time complexity of state-of-the-art libraries.	16
	<i>GRAPE</i> boosts edge prediction performance.	16
<b>3</b>	<b>Discussion</b>	<b>16</b>
<b>4</b>	<b>Methods</b>	<b>18</b>
4.1	Succinct data structures for adjacency matrices	18
4.1.1	Memory Complexity	19
4.1.2	Edge Encoding	19
4.1.3	Operations on Elias-Fano.	19
4.1.4	Efficient implementation of Elias-Fano.	20
4.1.5	Optional memory-time trade-offs	21
	Explicit destinations vector	21
	Explicit out-bound degrees vector	21
	Explicit sources vector	21
4.2	Spectral and matrix factorization embedding methods	21
4.3	First and second-order random walk methods.	22
4.3.1	Implementation of second-order random walks	23
	Efficient computation of the in-out and return bias.	24
4.3.2	Efficient sampling for Node2Vec random walks	24
4.3.3	Approximated random walks	25
4.4	Triple-sampling methods.	25
4.5	Corrupted triple-sampling methods.	26
4.6	Edge embedding methods	26
4.7	Node-label, edge-label, and edge-label prediction models	27
4.8	Pipelines for the evaluation of graph-prediction tasks	28

## Abstract

Graph Representation Learning methods opened new possibilities for addressing complex, real-world problems represented by graphs. However, many graphs used in these applications comprise millions of nodes and billions of edges and are beyond the capabilities of current methods and software implementations. We present *GRAPE*, a software resource for graph processing and representation learning that is able to scale with big graphs by using specialized and smart data structures, algorithms, and a fast parallel implementation. When compared with state of the art software resources, *GRAPE* shows an improvement of orders of magnitude in empirical space and time complexity, as well as a substantial and statistically significant improvement in edge prediction and node label prediction performance. Furthermore, *GRAPE* provides over 80,000 graphs from the literature and other sources, standardized interfaces allowing a straightforward integration of third-party libraries, 61 node embedding methods, 25 inference models, and 3 modular pipelines to allow a FAIR and reproducible comparison of methods and libraries for graph processing and embedding.

## 1 Introduction

In various fields such as biology, medicine, data and network science, graphs can naturally model available knowledge as interrelated concepts, represented by a network of (possibly attributed) nodes connected by edges. The wide range of graph applications has motivated the development of a rich literature on Graph Representation Learning (GRL) and inference models.

GRL models compute embeddings, i.e. vector representations (usually in a metric space) of the graph and its constituent elements, capturing their topological, structural, and semantic relationships. Graph inference models can use such embeddings and available additional features for several tasks, e.g., visualization, clustering, node-label, edge-label, and edge prediction problems [1, 2].

GRL methods available in the literature include, among others, matrix factorization-based methods (section 4.2), random walk-based methods (section 4.3) and triple-sampling methods (sections 4.4 and 4.5) [1]. They have shown their effectiveness in the analysis of networks from sociology, biology, medicine, and many other disciplines [2].

Although a great deal of research has been devoted to the development of libraries to process and analyze graphs (e.g., iGraph [3], GraphLab [4], NetworkX [5], GraphX [6] and SNAP [7]), as well as GRL software resources (e.g., PecanPy [8], PyKeen [9], DGL [10], PytorchGeometric [11], Spektral [12]), real-world networks often include millions of nodes and billions of edges, thus raising the problem of the scalability of existing software resources implementing Graph Representation Learning algorithms [2, 13].

Another issue with random walk-based graph embeddings computed by current state-of-the-art libraries is the reduced capability to efficiently generate enough data to accurately represent the topology of the underlying graph. This limitation affects the performance of node and edge label prediction methods, since they strongly depend on the informativeness of the underlying embedded graph representation. The efficient generation of billions of sampled random walks can lead to more accurate embedded representations of graphs and boosts the performance of machine learning methods that learn from the embedded vector representation of nodes and edges.

The fair and reproducible comparison of different graph-based methods under different experimental set-ups is a relevant open issue usually not considered by current state-of-the-art libraries. To this end, standardized interfaces to easily embed methods from other libraries and standardized experimental pipelines could facilitate and make easy the comparison of different methods and libraries for the analysis of graph-based data.

*GRAPE* offers a modular and flexible solution to the above problems through well-documented, scalable, and fast software libraries that can run on general-purpose desktop and laptop computers, as

well as on high performance computing clusters.

## 2 Results

### 2.1 Overview of the *GRAPE* resource: *Embiggen* and *Ensmallen*

*GRAPE* is a fast graph processing and embedding library; it extensively uses parallel computation and efficient data structures to scale with big graphs. The library’s high-level structure, overall functionalities, and its two core modules, *Ensmallen* (ENabler of SMALL computational resources for Large Networks) and *Embiggen* (EMBeddInG GENerator), are depicted in Fig. 1a.

*Ensmallen* efficiently loads big graphs and executes graph processing operations, owing to its Rust [14] implementation, and use of traditional map-reduce thread-based parallelism and branch-less Single Instruction Multiple Data (SIMD) parallelism. It also provides Python bindings for ease of use. Designed to leverage succinct data structures [15], *GRAPE* requires only a fraction of the memory required by other libraries, and guarantees average constant-time rank and select operations. This makes it possible to execute many graph processing tasks, e.g. accessing node neighbours and running first- and second-order random walks, with memory usage close to the theoretical minimum. Among the many high-performance algorithms it provides, the library implements efficient and effective *approximated* weighted DeepWalk & Node2Vec embedding models. They can process graphs containing high-degree nodes (degree  $> 10^6$ ), an otherwise unmanageable task when using the analogous exact algorithms, and allow one to obtain edge-prediction performance comparable to those achieved by using the *exact* version (Section 2.3).

*Ensmallen* also provides many other methods and utilities, such as multiple holdout techniques, Bader and Kruskal algorithms for computing random and minimum spanning arborescence and connected components, stress and betweenness centrality [16], node and edge filtering methods, and algebraic set operations on graphs. *Ensmallen* allows graphs to be loaded from a wide variety of node and edge list formats (section 2.2). In addition, users can automatically load data from an ever-increasing list of over 80,000 graphs from the literature and elsewhere (Fig. 1b, detailed in section 2.5.1).

*Embiggen* provides GRL and inference models (sections 2.4), including an exhaustive set of node embedding methods, e.g., spectral and matrix factorization models such as HOPE [17], NetMF [18] and their variations (GLEE [19], SocioDim [20] - section 4.2), CBOW, SkipGram and GloVe embedding methods [21, 22] exploiting random walk-based methods such as DeepWalk, Node2Vec and Walklets [23, 24] (section 4.3), triple sampling methods such as LINE [25] (section 4.4), and corrupted-triple sampling methods such as TransE [26] (section 4.5), and, more generally, a wide range of inference methods (sections 4.7).

*GRAPE* provides three modular pipelines to compare and evaluate node-label, edge-label and edge prediction performance under different experimental settings (section 2.5.2, fig. 1b), as well as utilities for graph visualization (fig. 1c). These pipelines allow non-expert users to tailor their desired experimental setup and easily obtain actionable and reproducible results (Fig. 1b). Furthermore, *GRAPE* provides interfaces to integrate third-party models and libraries (e.g., KarateClub [27] and PyKeen [9] libraries). In this way, the evaluation pipelines can be used to obtain a fair comparison between models implemented or integrated in *GRAPE*.

The possibility to integrate external models and the availability of graphs for testing them on the same datasets allows to answer a still open and crucial issue in literature, which regards the FAIR (Findable, Accessible, Interoperable, and Reusable), objective, reproducible, and efficient comparison of graph-based methods and software implementations (Section 2.5.1).

We used the evaluation pipelines to compare the edge prediction and node-label prediction performance of 16 node embedding models, 12 reimplemented in *GRAPE* and 4 integrated from the KarateClub

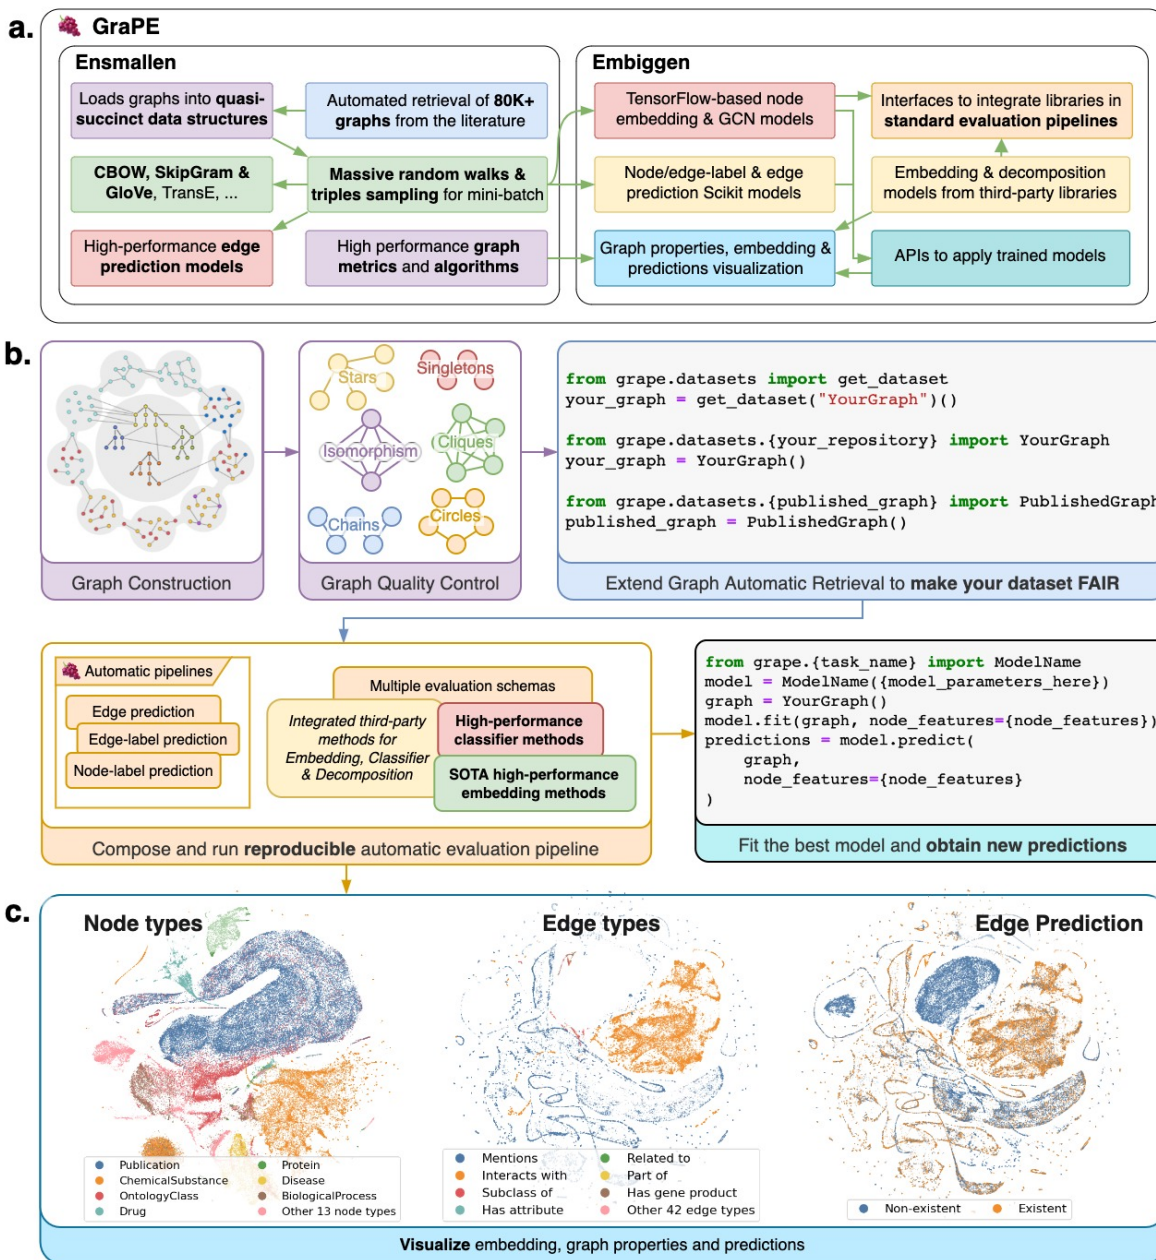


Figure 1: **Schematic diagram of GRAPE (Ensmallen + Embiggen) functionalities.** **a.** High level structure of the GRAPE software resource. **b.** Pipelines for an easy, fair, and reproducible comparison of graph embedding techniques, graph-processing methods, and libraries. **c.** Visualization of KGCOVID19 graph, obtained by displaying the first two components of the t-SNEdecomposition of the embeddings computed by using a Node2Vec SkipGram model that ignores the node and edge type during the computation. The clusters colors indicate: (left) the Biolink category for each node; (center) the Biolink category for each edge; (right) the predicted edge existence.

library [27]. Moreover, we compared GRAPE with state-of-the-art graph-processing libraries across several types of graphs having different size and characteristics, including big real-world graphs such as Wikipedia, the CTD, Comparative Toxicogenomic Database [28] and biomedical Knowledge Graphs

generated through *PhKnowLator* [29], showing that *GRAPE* achieves state-of-the-art results in processing big real-world graphs both in terms of empirical time and space complexity and prediction performance.

## 2.2 Fast error-resilient graph loading

*GRAPE* can process many graph formats and simultaneously check for common format errors. Fig. 2 shows the empirical space (a) and time complexity (b) required by *GRAPE* and by state-of-the-art graph processing libraries, including *NetworkX* [30], *iGraph* [3], *CSRGraph*, *PecanPy* [8], when loading 44 real-world graphs. Results show that *GRAPE* is faster and requires less memory as compared to state-of-the-art libraries. For instance, *GRAPE* loads the *ClueWeb09* graph (1.7B nodes and 8B undirected edges) in less than 10 minutes and requires about 60GB of memory, whereas the other libraries were not able to load this graph. All graphs and libraries used in these experiments are directly available from *GRAPE*'s and are detailed in the Supplementary Information ( S1 and S2).

## 2.3 *GRAPE* outperforms state-of-the-art libraries on random walk generation

Through extensive use of thread and SIMD parallelism and specialized quasi-succinct data structures, *GRAPE* outperforms state-of-the-art libraries by one to four orders of magnitude in the computation of random walks, both in terms of empirical computational time and space requirements (Figure 2-c, d, e, f and Section 2.3.1).

Further speed-up of second-order random walk computation is obtained by dispatching one of the 8 optimized implementations of *Node2Vec* sampling [23]. The dispatching is based on the values of the *return* and *in-out* parameters and on the type of the graph (weighted or unweighted). *GRAPE* automatically provides the version best suited to the requested task, with minimal code redundancy (Section 4.3.1). The time performance difference between the least and the most computationally expensive implementations is around two orders of magnitude (Supplementary Section S7.2 and Tables S50 and S51).

### 2.3.1 Experimental comparison of graph processing libraries.

We compared *GRAPE* with a set of state-of-the-art libraries including *GraphEmbedding*, *Node2Vec*, *CSRGraph* and *PecanPy* [8], on a large set of first and second-order random walk tasks. The random walk procedures in the *GraphEmbedding* and *Node2Vec* libraries use the alias method (Supplementary Information S7.2.3). The *PecanPy* library also employs the alias method for small graphs use-cases (less than 10,000 nodes). *CSRGraph*, on the other hand, computes the random walks lazily using *Numba* [31]. Similarly, *PecanPy* leverages *Numba* lazy generation for graphs having more than 10,000 nodes. All libraries are further detailed in Supplementary Information S1.

Figure 2 shows the experimental results of a complete iteration of one-hundred step random walks on all the nodes across 44 graphs having a number of edges ranging from some thousands to several billions (Section 2.2). *GRAPE* greatly outperforms all the compared graph libraries on both first and second-order random walks in terms of both space and time complexity. Note that *GRAPE* scales well with the biggest graphs considered in the experiments, while the other libraries either crash when exceeding 200GB of memory, or take more than 4 hours to execute the task (Figure 2 c, d, e, f).




Figure 2: **Experimental comparison of *GRAPE* with state-of-the-art graph processing libraries across 44 graphs.** **Top row - graph loading:** **a.** Empirical execution time. **b.** Peak memory usage. The horizontal axis shows the number of edges, vertical axis peak memory usage. **Middle row - first-order random walk:** **c.** Empirical execution time. **d.** Peak memory usage. **Bottom row - second order random walk:** **e.** Empirical execution time. **f.** Peak memory usage. The horizontal axis shows the number of nodes, and the vertical axis respectively execution time (c,e) and memory usage (d,f). All axes are in logarithmic scale. The  $\times$  represent when either a library crashes, exceeds 200GB of memory or takes more than 4 hours to execute the task. Each line corresponds to a graph resource/library, and points on the lines refer to the 44 graphs used in the experimental comparison. The method used to properly measure execution time and peak memory usage is presented in the Supplementary section S.6.3. Note that the blue line representing *GRAPE* is always below all the other lines.



### 2.3.2 Approximated random walks to process graphs with high-degree nodes

Random walks on graphs containing high-degree nodes is challenging, since multiple paths from the same node need to be processed. To overcome this computational burden, *GRAPE* provides an approximated implementation of weighted random walks that undersamples the neighbors to scale with graphs containing nodes with high-degree, e.g. with millions of neighbors (Figure 3 a, b, c, Section 4.3.3). We compared exact and approximated random walk samples for Node2Vec-based SkipGram for edge prediction problem on the (unfiltered) *H. sapiens* STRING PPI network [32], achieving statistically equivalent performance (Two-sided Wilcoxon rank-sum  $p$ -value  $> 0.2$ , Fig. 3,d), by running 30 holdouts, and setting a (deliberately low) degree threshold equal to 10 for the approximated random walk, while the maximum degree in the training set ranged between 3325 and 4184 across the holdouts. These results show no relevant performance decay, even when using a relatively stringent degree threshold.

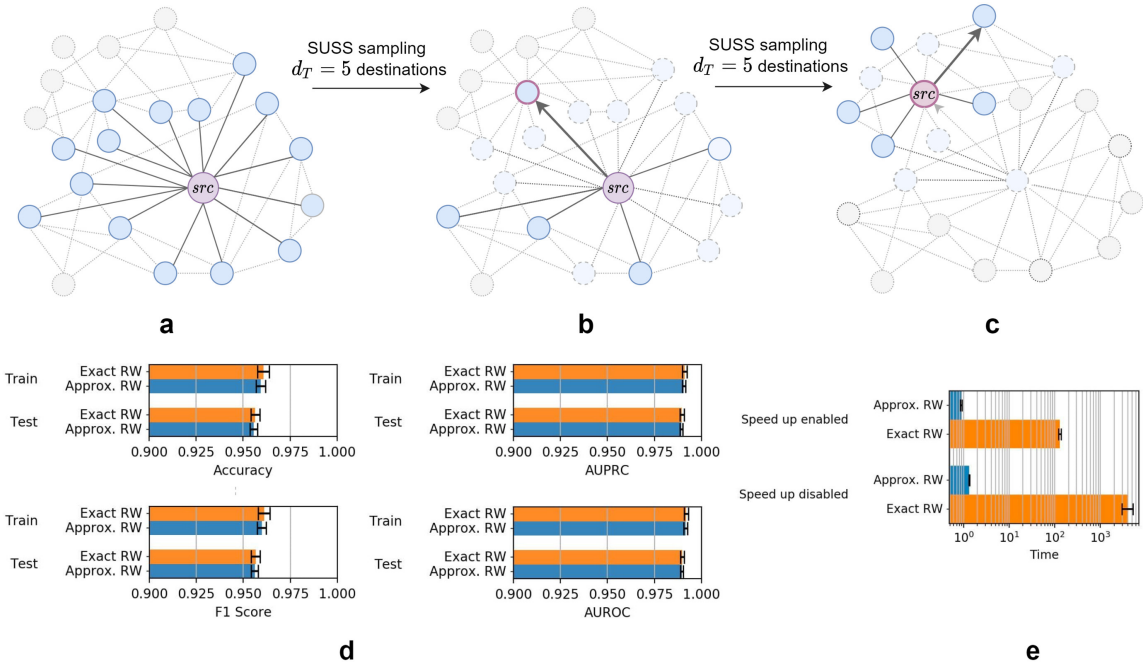


Figure 3: **Approximated random walk.** **a.** The random walk starts at node *src*; its 15 neighbourhood nodes are highlighted in cyan. **b.** We sample  $d_T = 5$  destination nodes ( $d_T$  being the degree threshold) from the available 15 destinations, using our novel Sorted Unique Sub-Sampling algorithm (SUSS, see Supplementary Information S7.2.4), and performs a random step (edge highlighted with an arrow). **c.** A further step will be performed on the successor node (that now becomes the novel source node *src*, and the same process is repeated until the end of the walk). **d.** Edge prediction performance comparison (Accuracy, AUPRC, F1 Score, and AUROC) using Skipgram-based embeddings and RW samples obtained with exact and approximated RWs for both the training and the test set with the STRING-PPI dataset. Bar plots are zoomed-in at 0.9 to 1.0, with error bars representing the standard deviation, computed over 30 holdouts. **e.** Empirical time comparison (in msec) of the approximated and exact second-order random walk algorithm on the graph *sk-2005* [33]: 100 – steps random walks are run on 100 randomly selected nodes. Error bars represent the standard deviation across 10 repetitions. Time is on a logarithmic scale.

## 2.4 Node and edge embedding models

*GRAPE* provides GRL methods (sections 4.2, 4.3, 4.4 and 4.5) and prediction models (section 4.7) with high-performance implementations for kernel preprocessing as well as generation of random-walk and triple mini-batches.

*GRAPE* provides both its own implementations and Keras-based implementations for all shallow neural network models (e.g. CBOW, SkipGram, TransE). Nevertheless, since shallow models allow for particularly efficient data-race aware and synchronization-free implementations [26], the “by-scratch” *GRAPE* implementations significantly outperform the Keras-based ones, as TensorFlow APIs are too coarse and high-level for such fine-grained optimizations. While GPU training is available for the TensorFlow models, their overhead with shallow models tends to be so relevant that “by-scratch” CPU implementations outperform those based on GPU. Moreover, the embedding of large graphs (such as Wikipedia) do not fit in most GPU hardware memory. dtill, Keras-based models allow users to experiment with the open-software available in the literature for Keras, including, e.g., advanced optimizer and learning rate scheduling methods.

The provided spectral and matrix-factorization-based models, including HOPE [17], NetMF [18] and their variations (GLEE [19], SocioDim [20]), compute the (generally sparse) weighted adjacency matrix in parallel and then rely on LAPACK [34] routines for the singular values decomposition and eigenvectors computations. Lastly, *GRAPE* also provide Keras/TensorFlow-based first-order and second-order LINE models, as they generally suffer from noisy gradients and therefore require higher-order optimisers such as Adam. Notable third parties libraries integrated are Sciki-Learn [35], PyKeen [9] and KarateClub [27].

Furthermore, *GRAPE* provides many methods to compute edge embeddings given some computed node embedding, ranging from concatenation, Hadamard (element-wise multiplication), the element-wise difference in  $L_1$  or  $L_2$  norm, and element-wise mean, subtraction or sum between the node embedding vectors. The library also comes equipped with tools to visualize the computed node and edge embedding and their properties, including edge weights, node degrees, connected components, node types and edge types. For example, in figure 1 c we display the node (left) and edge types (center) of the KG-COVID19 graph and whether sampled edges exist (right) by using the first two components of the t-SNE decomposition of the node/edge embeddings [36].

## 2.5 GraPE enables a fair and reproducible comparison of graph embedding and graph-based prediction methods

*GRAPE* provides both a large set of ready-to-use graphs that can be used in the experiments and standardized pipelines to fairly compare different models and graph libraries ensuring reproducibility of the results (Fig. 1 b). Graph embedding are efficiently implemented in Rust by scratch (with a Python interface) or are integrated from other libraries by implementing the interface methods of an abstract *GRAPE* class. *GRAPE* users can compare different embedding methods and prediction models and can also add their own methods to the standardized pipelines. Our experiments (Section 2.5.3) show how to use the standardized pipelines to fairly compare a large set of methods and different implementations by using only a few lines of Python code.

### 2.5.1 FAIR graph retrieval

*GRAPE* facilitates *FAIR* access to an extensive set of graphs and related datasets, including both commonly used benchmark datasets and graphs actively used in biomedical research. Any of the available graphs can be retrieved and loaded with a single line of Python code (Fig. 1 b.), and their list is constantly expanding, thanks to the generous contributions of *GRAPE* users. The list of resources currently supported can be found at Supplementary Information S3.1.

**Findability and Accessibility.** Datasets may change locations, versions may appear in more than one location, and file formats may change. Using an ensemble of custom web scrapers, we collect, curate and normalize the most up-to-date datasets from an extensive resources list (currently over 80,000 graphs). The collected metadata is shipped with each *GRAPE* release, ensuring end-users can always find and immediately access any available versions of the provided datasets.

**Interoperability.** The graph retrieval phase contains steps that robustly convert data from (even malformed) datasets into general-use TSV documents that, while primarily used as graph data, can be used for any desired application case.

**Reusability.** Once loaded, the graphs can be arbitrarily processed and combined, used with any of the many embedding and classifier models from either the *GRAPE* library or any third-party model integrated in *GRAPE* by implementing the interface described in section 2.5.2.

### 2.5.2 FAIR evaluation pipelines

The fair and objective comparative evaluation of datasets, graph embedding and prediction models is fundamental for scientific research. To provide actionable results, this comparison requires specifically designed and real-world benchmark datasets [37], as well as pipelines that could allow non-expert users to easily test and compare graphs and inference algorithms on the desired graphs.

Beside FAIR graphs, *GRAPE* allows even users with minimal Python language experience to implement experimental designs through pipelines for running node-label, edge-label and edge prediction experiments with task-specific evaluation schemas.

More precisely, each pipeline allows users to tailor the experiment by choosing: (a) the set of graphs to be used in the experiments, (b) the functions to be called for graph filtering (if needed), (c) the set of (embedding and prediction) algorithms to be applied, and (d) the evaluation schema (Section 4.8).

In particular, the evaluation schemas for edge prediction models are K-fold cross-validations, Monte Carlo, and Connected Monte Carlo (Monte Carlo designed to avoid the introduction of new connected components in the training graph) holdouts. All of the edge prediction evaluation schemas may sample the edges in a uniform or stratified way, with respect to a provided list of edge-types. Sampling of negative (non-existing) edges may be executed by either following a uniform or a scale-free distribution. Furthermore, the edge-prediction evaluation may be performed by using varying unbalance ratios (between existent and non-existent edges) to better gauge the true-negative rate (specificity) and false-positive rate (fall-out). Stratified Kfold and stratified Monte Carlo holdouts are also provided for node and edge-label prediction models.

For all tasks, an exhaustive set of evaluation metrics are computed, including AUROC, AUPRC, Balanced Accuracy, Miss-rate, Diagnostic odds ratio, Markedness, Matthews correlation coefficient and many others.

Each pipeline can receive as input any model that implements the specific task’s interface (a Python abstract class). All interfaces follow the familiar scikit-learn style, with embedding models required to implement the *fit\_transform* method, and prediction models required to implement the *fit*, *predict*, *predict\_proba* methods, plus some additional metadata necessary to check for biases in the considered task. Interfaces are made available for [embedding models](#), [node-label prediction](#), [edge-label prediction](#), and [edge prediction](#). All models available in *GRAPE* implement these interfaces, and they can be used as starting points for custom integrations. Many usage examples are available in the library tutorials: <https://github.com/AnacletoLAB/grape/tree/main/tutorials>.

### 2.5.3 Experimental comparison of node embedding methods

We selected 16 among the 61 node embedding methods available in *GRAPE*, and we used the edge prediction and node-label standardized prediction pipelines to compare the prediction results obtained by Perceptrons, Decision Trees, and Random Forests classifiers (Fig. 4). For the edge prediction tasks we used the Hadamard product to construct edge embeddings from node embeddings. We applied a “connected Monte Carlo” evaluation schema for edge prediction and a stratified Monte Carlo evaluation schema for node-label prediction (see Supplementary S9.2 for more details).

The models have been tested on 3 graphs for edge prediction (Fig. 4-a,b) and 3 graphs for node-label prediction (Fig. 4-c,d). The graph reports, describing the characteristics of the analyzed graphs, automatically generated with *GRAPE*, are available in the Supplementary information S3.2, S3.3. Since they are homogeneous graphs<sup>1</sup> we considered only homogeneous node embedding methods. Moreover, we discarded non-scalable models, e.g. models based on the factorization of dense adjacency matrices.

Among the 16 methods, 12 are implemented in *GRAPE* (purple in Fig. 4) and 4 have been integrated from the Karate Club library [27] (cyan in Fig. 4). They can be grouped into three broad classes:

- a. Spectral and matrix factorization methods:** Geometric Laplacian Eigenmap Embedding (GLEE) [19], Alternating Direction Method of Multipliers for Non-Negative Matrix Factorization (NMFADMM) [38], High-Order Proximity preserved Embedding (HOPE) [17], Iterative Random Projection Network Embedding (RandNE) [39], Network Matrix Factorization (NetMF) [18], and Graph Representations (GraRep) [40].
- b. First-order random-walk methods:** DeepWalk-based GloVe, CBOW, and SkipGram, Walklets SkipGram [24, 22, 21], and Role2Vec with Weisfeiler-Lehman Hashing [41, 42, 27].
- c. Second-order random-walk methods:** Node2Vec-based GloVe, CBOW, and SkipGram [22, 21, 23].
- d. Triple-sampling methods:** first and second order LINE [25].

All the embedding methods and classifiers are described in more detail in sections 4.2, 4.3, 4.4, and 4.7.

Results show that no model is consistently better with respect to the others across the types of task and the data sets used in the experiments (Figure 4). These results are analogous to those obtained by Kadlec et al. [43] for TransE model family, and those obtained by Errica et al [44] for GNN models, highlighting the need for objective pipelines to systematically compare a wide array of possible methods for a desired task. The standardized pipelines implementing the experiments are available from the online *GRAPE* tutorials and allow the full reproducibility of the results summarized in Fig. 4. Full results relative to other evaluation metrics are available in Supplementary Information S5.1 and S5.2.

---

<sup>1</sup>Hereafter, graph homogeneity/heterogeneity refers to the homogeneity/heterogeneity of node and edge types.

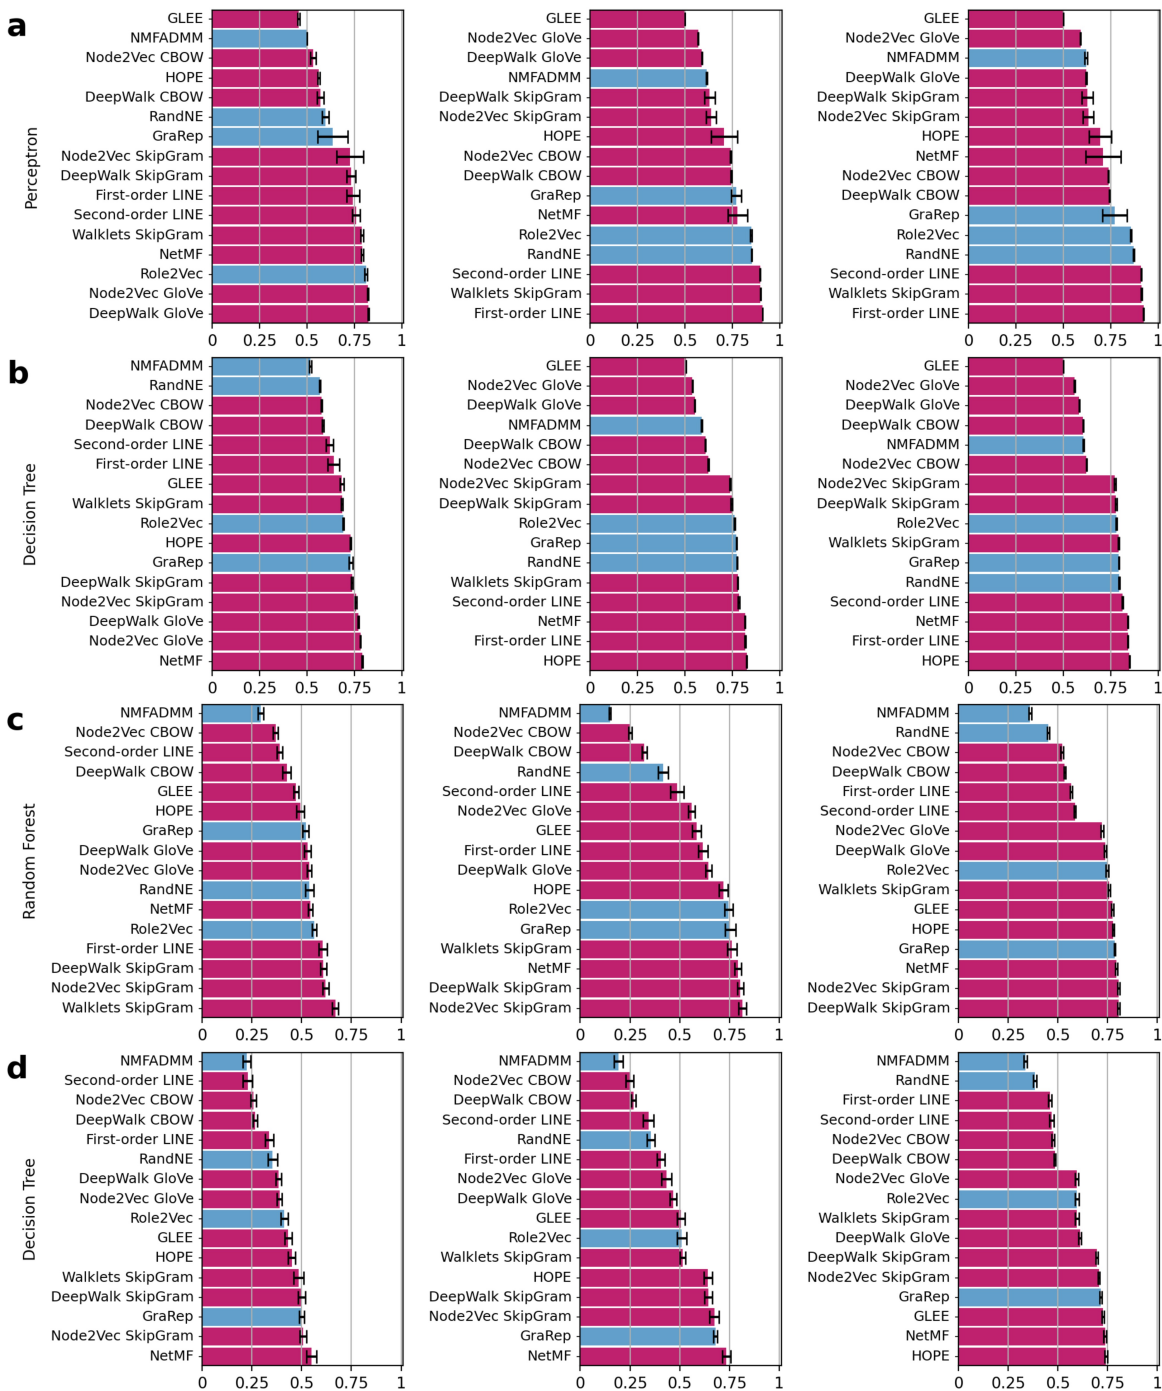


Figure 4: Comparison of embedding methods through the *GRAPE* pipelines: edge and node label prediction results. Results represent the balanced accuracy averaged across ten holdouts (results relative to other evaluation metrics are available in the Supplementary Information S5.). We sorted the embedding models by performance for each task; methods directly implemented in *GRAPE* are in purple, while integrated methods are in cyan. (a, b): Edge prediction results obtained through a Perceptron (a) and a Decision tree (b). Barplots from left to right, show the balanced accuracy results obtained with the *Human Phenotype Ontology* (left), *STRING Homo sapiens* (center) and *STRING Mus musculus* (right). (c, d): Node-label prediction results obtained through a Random Forest (c) and a Decision Tree (d). Barplots from left to right show the balanced accuracy respectively achieved with *CiteSeer* (left), *Cora* (center) and *PubMed Diabetes* (right) datasets.

## 2.6 Scaling with big real-world graphs

To show that *GRAPE* can scale and boost edge prediction in big real-world graphs, we compared its Node2Vec-based models with state-of-the-art implementations on three big graphs: 1) English Wikipedia; 2) Comparative Toxicogenomic Database (CTD [28]); 3) A biomedical graph generated through PheKnowLator [29] (Supplementary Information S6).

### 2.6.1 Experimental setup

**English Wikipedia.** Wikipedia graphs are web graphs with nodes representing either Wiki sites pages or related websites; edges represent the links between the pages. In the experiments, we used the English Wikipedia graph having 17 million nodes and 130 million (undirected) edges (2021-11-01 version). The task for the English Wikipedia graph is a whole-graph edge prediction, that is, predicting whether two given nodes in the entire graph are connected by an edge. The set of positive edges is defined as the undirected edges present in the entire graph (about 130 million). The set of negative edges, is instead defined as the edges that are not present in the graph, that is around 150 trillion undirected edges.

**Comparative Toxicogenomic Database (CTD).** CTD is a publicly available database that aims to advance understanding of how environmental exposures affect human health. It provides manually curated information about chemical–gene/protein interactions, chemical–disease, and gene–disease relationships, and includes also information about phenotypes, pathways, ontologies, and their relations with genes, chemicals, and diseases, including about 45 million edges and more than 100K nodes. The CTD edge-prediction task consists in predicting *gene–disease* associations. The set of positive edges is defined as the set of all the existing (undirected) relationships (edges) between gene and disease nodes, that amounts to about 29 million. Negative edges were defined by pairs gene–disease being unrelated in the CTD dataset (about 362 million “negative” edges).

**PheKnowLator biomedical data.** PheKnowLator is a software resource designed to facilitate the construction of large-scale biomedical knowledge graphs using a number of knowledge models (instance-based and subclass-based). PheKnowLator currently integrates 12 Open Biomedical Ontologies and 31 linked open-data sources. In our experiments we used a 2022-04-11 build including about 7 millions of (undirected) edges and about 800K nodes. The PheKnowLator task consists in the prediction of *genetic variant–disease* associations, using 44K known “positive” associations and a set of “negative” edges, including about 3 billions of variant–disease pairs having no known associations. Detailed information about the source data and scripts used to generate the above three big real-world graphs are available in Supplementary Section S6.

**Graph libraries compared in the experiments.** In the experiments we used two *GRAPE* implementations of embedding algorithms: CBOV and SkipGram. We compared them with the following state-of-the-art embedding libraries, widely used by the scientific community:

- *PecanPy* [8] is a Python library implementing a Numba-based version of *node2vec*, leveraging Numba’s just-in-time Python compilation [31] to generate the random walks on the input graph, and forwarding them to an embedding model provided by Gensim natural language processing library [45].
- *NodeVectors*<sup>2</sup> is a Python package that enables fast and scalable node embedding algorithms. It leverages CSR matrix storage for graphs, but it also support NetworkX [30] graph loading. Besides *node2vec*, the library also implements several kinds of first and second-order random walks.

---

<sup>2</sup><https://github.com/VHRanger/nodevectors>

- *SNAP* [7], Stanford Network Analysis Platform, is a general-purpose system for the manipulation and analysis of large networks, written in C++. Once compiled, it becomes an executable to analyze and compute different statistics about the graphs; it also implements different kinds of graph-processing algorithms and allows computing node embeddings, by using a pre-processing phase for pre-computation of transition probabilities through the Alias method [46].
- *Node2Vec*<sup>3</sup> is a Python package for embedding networks through random walk-based algorithms like *node2vec*. Similar to SNAP, it employs the Alias method [46] to pre-compute transition probabilities. It also handles the graph loading through the NetworkX library [30].
- *GraphEmbedding*<sup>4</sup> is a Python package that handles network embeddings with random walk-based methods. Again, the transition probability is pre-computed via the Alias method and employs NetworkX library to handle the loading of a graph.
- *FastNode2Vec*<sup>5</sup> implements the *node2vec* algorithm, leveraging both Numba and Gensim. This implementation scales linearly, in time and memory, with respect to the dimension of the input graph.

**Evaluation of the results** For all of the considered tasks, we firstly computed the embedded graphs using graph libraries and then the resulting embeddings have been used to train machine learning methods for an edge prediction problem. To evaluate the ML models we adopted a connected Monte Carlo (Supplementary Informations S9.2) repeated ten times, with a train:test ratio equal to 80% : 20% of the data. As evaluation metrics we applied precision, recall, accuracy, balanced accuracy, F1, AUROC, and AUPRC. In the experimental set-up we imposed the following memory and time constraints, using a Google Cloud VM with 64 cores<sup>6</sup>:

- A maximum time of 48 hours for each holdout to produce the embedding;
- The maximum memory usage allowed during the embedding phase is 64GB.
- The maximum memory usage allowed during the prediction phase is 256GB.

To keep track of memory and time requirements and of possible stops for exceptions and system-related errors (out of memory, core dumps), the Python library `memory_time_tracker` was used<sup>7</sup>.

---

<sup>3</sup><https://github.com/eliorc/node2vec>

<sup>4</sup><https://github.com/shenweichen/GraphEmbedding>

<sup>5</sup><https://github.com/louisabraham/fastnode2vec>

<sup>6</sup>N1 Cpus with Intel Haswell micro-architecture

<sup>7</sup>[https://github.com/LucaCappelletti94/memory\\_time\\_tracker](https://github.com/LucaCappelletti94/memory_time_tracker)

## 2.6.2 Results

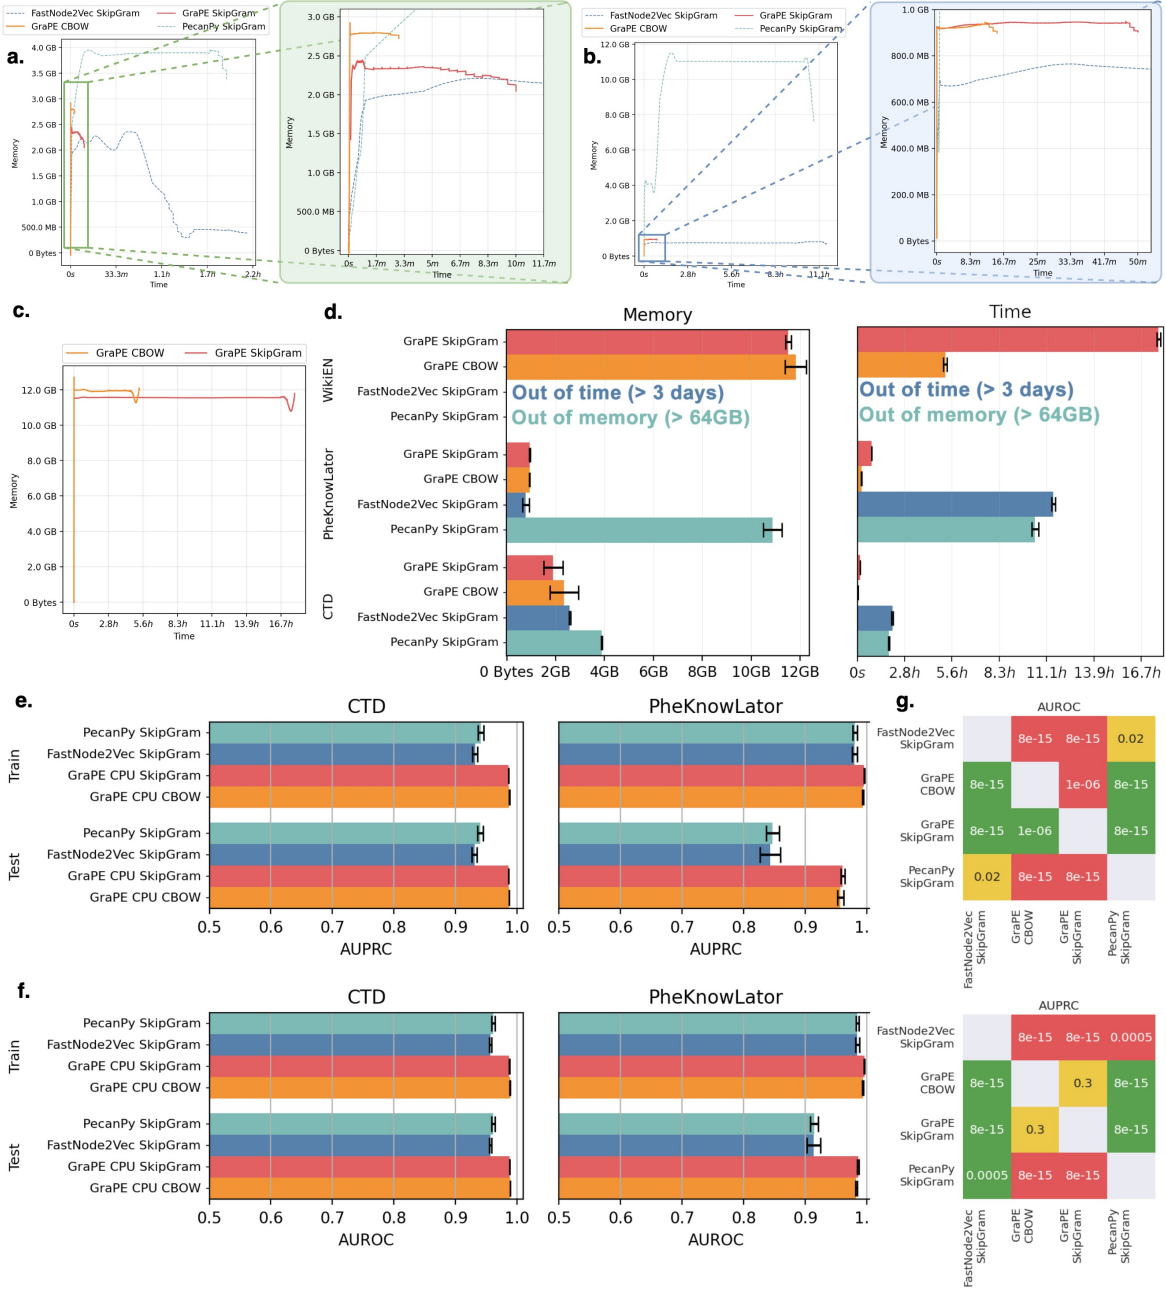


Figure 5: Performance comparison between *GRAPE* and state-of-the-art implementations of Node2Vec on real-world big graphs. *GRAPE* implementations achieve significantly better empirical time complexity: (a.), (b.) and (c.) show the worst performance (maximum time and memory, denoised using a Savitzky–Golay filter) over 10 holdouts on *CTD*, *PheKnowLator* and *Wikipedia*, respectively. In textbfa. and textbfb. the rectangles in the left figure are magnified in the right figure to highlight *GRAPE* performances. In the *Wikipedia* plot (c.) only *GRAPE* results are available as the others either go out-of-time or out-of-memory. (d.) Average memory and computational time across the holdouts; error bars represent standard deviation. (e.) AUPRC and (f.) AUROC results of Decision Trees trained with different graph embedding libraries: *GRAPE* embedding achieve better edge prediction performance than those obtained by the other libraries. (g.) Wilcoxon signed-rank tests results (p-values) between *GRAPE* and the other state-of-the-art libraries, where the win of a row against a columns is in green, the tie in yellow, and the loss in red). Top: AUROC, bottom: AUPRC.



***GRAPE* is able to scale with big graphs when the other competing libraries fail.** Most of the competing libraries were not able to complete the embedding and prediction tasks on big real-world graphs. Indeed *NodeVectors* exceeded the time computation limit, while *SNAP*, *Node2Vec*, and *GraphEmbedding* went out of memory in the embedding phase due to the high memory complexity required by the Alias method they use for pre-computing the transition probabilities (Supplementary Information S7.2.3)<sup>8</sup>. *FastNode2Vec* and *PecanPy* went out of time (more than 48h of computation) on the biggest Wikipedia graph. In practice only *GRAPE* was able to successfully terminate the embedding and prediction tasks with all the three big real-world graphs considered here.

***GRAPE* improves the empirical time complexity of state-of-the-art libraries.** Fig. 5 a, b and c show the memory and time requirements of *GRAPE*, *FastNode2Vec* and *PecanPy* (note that the other state-of-the-art libraries ran out of time or memory on these real-world graph prediction tasks. With *CTD* and *PheKnowLator* biomedical graphs we can observe a speed-up of about one order of magnitude (Fig. 5 a, b) of *GRAPE* with respect to both *FastNode2Vec* and *PecanPy* with also a significant gain in memory usage with respect to *PecanPy* and a comparable memory footprint with *FastNode2Vec*. These results are confirmed by the average memory and time requirements across ten holdouts (Fig. 5 d). Note that both *FastNode2Vec* and *PecanPy* fail with the Wikipedia task, while *GRAPE* was able to terminate the computation in a few hours using a reasonable amount of memory (Fig. 5 c and d).

***GRAPE* boosts edge prediction performance.** *GRAPE* not only enables big graph embedding and speed-up computation, but can boost prediction performance on big real world graphs. Fig. 5-e and f show that *GRAPE* achieves better results on edge prediction tasks with both *CTD* and *PheKnowLator* biomedical graphs. *GRAPE* outperforms the other competing libraries at 0.001 significance level, according to the Wilcoxon rank sum test (Fig. 5 g). The edge embeddings have been used to train a decision tree to allow a safe comparison between the embedding libraries.

Supplementary Information S6 reports AUROC, accuracy, and F1-score performances and other more detailed results about the experimental comparison of *GRAPE* with state of the art libraries.

### 3 Discussion

*GRAPE* is a software resource with specialized data structures, algorithms, and fast parallel implementations of graph processing coupled to implementations of algorithms for graph machine learning. The experimental results demonstrate that *GRAPE* significantly outperforms state of the art graph processing libraries in terms of empirical space and time complexity (Wilcoxon rank-sum test at 0.05 significance level) and can even load and process massive graphs (e.g. the Twitter graph) when the other libraries fail (Figure 2 a and b).

Moreover, for both first and second-order random walks, we observed a speedup ranging from two to four orders of magnitude compared with state of the art libraries (Figure 2 c and e). We registered a statistically significant improvement for all the considered random walk tasks (Wilcoxon rank-sum test at 0.01 significance level) for both the memory use and empirical time requirements (Figure 2 c, d, e and f).

When used to compare 16 representative node embedding methods through the comparative pipelines, results show that no model is consistently better with respect to the others across the types of task and the data sets used in the experiments Figure 4. This highlights the need for objective evaluation pipelines to systematically compare of a wide array of possible methods for a desired task.

---

<sup>8</sup>the Alias method has quadratic complexity with respect to the number of nodes in the graph, therefore becoming quickly too expensive on big graphs.

The scaling properties of *GRAPE* with respect to state-of-the-art libraries have been analyzed through extensive experiments with three real-world big graphs, i.e. *Wikipedia*, *CTD* and a biomedical Knowledge Graph generated by *PheKnowLator*. *GRAPE* significantly outperforms state-of-the-art libraries both in terms of empirical time and space complexity and in edge prediction performance, and is able to process big graphs (e.g. *Wikipedia*) even when the other competing libraries fail (Fig. 5).

Quasi-succinct data structures and lazy loading are crucial to efficiently generate large sets of random-walk training data and enable the processing of large graphs in off-the-shelf computers. In fact, when we analyze a graph of the size of KG-COVID-19 [47], the static generation of a huge amount of data fed into the models would easily occupy tens of terabytes of memory, making data processing unfeasible even in well-equipped computers.

*GRAPE* can also run an approximated version of first and second-order random walks on big and highly connected graphs where the exact algorithm is extremely inefficient. For instance, on the graph *sk-2005* [48], that includes about 50 millions of nodes and 1.8 billions of edges and some nodes with degrees over 8 million, by linearly extrapolating the results reported in Fig. 3-e to the entire graph, the exact algorithm requires about 23 days, while the approximate one about 11 minutes, both running on a PC with two AMD EPYC 7662 64-core processors, 256 CPU threads, and 1TB RAM.

The synergy between parallel computation and efficient data structures enables a fast generation of training batches (random-walk and triple samples), which are used to obtain more refined and accurate graph embeddings [2]. Moreover, the training speed obtained using *Ensmellen* opens the door to model selection, through, e.g., Bayesian optimization [49], that would have previously required a prohibitive amount of computation time. Indeed, the sheer amount of training samples that can be quickly generated meets the well-known greedy nature of neural network training sample requirements, leading to boosted performance in node and edge predictions tasks.

*GRAPE* also implements optional time-space trade-offs (Section 4.1.5), thanks to which the algorithms can be executed on computing environments ranging from notebooks to HPC clusters.

Finally, *GRAPE* provides facilities for graph analysis, human readable reports in natural language to summarize the graph characteristics, and, most importantly, it allows a straightforward integration of methods from other libraries, thus facilitating a fair and reproducible comparison between methods and implementations.

Though *GRAPE* allowed to compare different experimental setups by composing experiments on different graphs, and by using several embedding methods and prediction models, no method systematically outranked other models 2.5.3. To close this knowledge gap, in future work we plan to use *GRAPE* to execute a large scale grid-search to identify task-specific trends for the various combinations of models and their parameters.

## 4 Methods

*GRAPE* provides a wide spectrum of graph processing methods, implemented within the *Ensmallen* module, including node embedding methods (sections 4.2, 4.3, 4.4, 4.5), methods to combine the node embeddings for obtaining edge embeddings (section 4.6), and models for node-label, edge-label and edge prediction (section 4.7), implemented within the *Embiggen* module. The graph processing methods include fast graph loading, multiple graph holdouts, efficient first and second-order random walks, triple and corrupted-triple sampling, plus a wide range of graph processing algorithms that nicely scale with big graphs, using parallel computation and efficient data structures to speed up the computation.

*Ensmallen* is implemented using Rust, with fully documented Python bindings. Rust is a compiled language gaining importance in the scientific community [14] thanks to its robustness, reliability, and speed. Rust allows threads and data parallelism to be exploited robustly and safely. To further improve efficiency, some core functionalities of the library, such as the generation of pseudo-random numbers and sampling procedures from a discrete distribution, are written in assembly (see Supplementary Sections S7.2.1 and S7.2.2).

Graph representation learning methods encode nodes as vectors that capture the graph topological and structural characteristics. In other words, node embeddings map structurally and/or topologically similar nodes into vectors that are close in the mapped metric space. Once computed, node embeddings can be used for graph-prediction tasks [2], including node-label prediction, edge and edge-label prediction, unsupervised graph analysis (e.g. node clustering) or to visualize graphs and their properties (fig. 1 c) in combination with dimensionality reduction techniques (e.g. t-SNE [36]). *GRAPE* currently provides 49 unique node embedding models (61 considering redundant implementations, important for benchmarks), with 19 being “by-scratch” implementations and 30 integrated from third-party libraries. The list of available node embedding methods is constantly growing, with the ultimate goal to provide a complete set of efficient node embedding models. The input for the various models (e.g. random walks and triples) are provided by *Ensmallen* in a scalable, highly efficient, and parallel way (Fig. 1a). All models were designed according to the “composition over inheritance” paradigm, to ensure a better user experience through increased modularity and polymorphic behaviour [50]. More specifically, *Embiggen* provides interfaces, specific for either the embedding or each of the prediction-tasks, that must be implemented by all models; third-party models, such as PyKeen [9], KarateClub [27] and Scikit-Learn [35] libraries, are already integrated within *GRAPE* by implementing these interfaces. *GRAPE* users can straightforwardly create their models and wrap them by implementing the appropriate interface.

*GRAPE* has a comprehensive test suite. However, to thoroughly test it against many scenarios, we also employed fuzzers, that is tools that iteratively generate inputs to find corner cases in the library.

In the next section we describe the succinct data structures used in the library and detail their efficient *GRAPE* implementation (Section 4.1). We then summarize the spectral and matrix factorization (Section 4.2), the random walk-based (Section 4.3), the triple and corrupted triples-based (Section 4.4 and 4.5) embedding methods and their *GRAPE* implementation. In section 4.6 we describe the edge embedding methods and in Section 4.7 the node and edge label prediction methods available in *GRAPE*. Finally in Section 4.8 we detail the *GRAPE* standardized pipelines to evaluate and compare models for graph prediction tasks.

### 4.1 Succinct data structures for adjacency matrices

Besides heavy exploitation of parallelism, the second pillar of our efficient implementation is the careful design of the data structures for using as little memory as possible and quickly performing operations on them. The naive representation of graphs explicitly stores its adjacency matrix, with a  $\mathcal{O}(|V|^2)$  time and memory complexity, being  $|V|$  the number of nodes, which leads to intractable memory costs on large graphs. However, since most large graphs are highly sparse, this problem can be mitigated

by storing only the existing edges. Often, the adopted data structure is a Compressed Sparse Rows matrix (CSR [51]), which stores the source and destination indices of existing edges into two sorted vectors. In *Ensmallen* we further compressed the graph adjacency matrix by adopting the Elias-Fano succinct data scheme, to efficiently store the edges (Supplementary section S7.1). Since Elias-Fano representation stores a sorted set of integers using memory close to the information-theoretical limit, we defined a bijective map from the graph-edge set and a sorted integer set. To define such encoding, we firstly assigned a numerical id from a dense set to each node, and then we defined the encoding of an edge as the concatenation of the binary representations of the numerical ids of the source and destination nodes. This edge encoding has the appealing property of representing the neighbours of a node as a sequential and sorted set of numeric values, and can therefore be employed in the Elias-Fano data structure. Elias-Fano has faster sequential access than random access (Supplementary section S7.1.1) and is well suited for graph processing tasks such as retrieving neighbours during random walk computation and executing negative sampling using the outbound or inbound node degrees scale-free distributions.

#### 4.1.1 Memory Complexity

Elias-Fano is a quasi-succinct data representation scheme, which provides a memory efficient storage of a monotone list of  $n$  sorted integers, bounded by  $u$ , by using at most  $\mathcal{EF}(n, u) = 2n + n \lceil \log_2 \frac{u}{n} \rceil$  bits, which was proven to be less than half a bit per element away from optimality [15] and assures random access to data in average constant-time. Thus, when Elias-Fano is paired with the previously presented encoding, the final memory complexity to represent a graph  $G(V, E)$  is  $\mathcal{EF}_\phi(|V|, |E|) = \mathcal{O}\left(|E| \log \frac{|V|^2}{|E|}\right)$ ; this is asymptotically better than the  $\mathcal{O}(|E| \log |V|^2)$  complexity of the CSR scheme.

#### 4.1.2 Edge Encoding

*Ensmallen* converts all the edges of a graph  $G(V, E)$  into a sorted list of integers. Considering an edge  $e = (v, x) \in E$  connecting nodes  $v$  and  $x$  represented with, respectively, integers  $a$  and  $b$ , the binary representation of  $a$  and  $b$  are concatenated through the function  $\phi_k(a, b)$  to generate an integer index uniquely representing the edge  $e$  itself:

$$\phi_k(a, b) = a 2^k + b, \text{ where } k = \lceil \log_2 |V| \rceil \quad \Rightarrow \quad a = \left\lfloor \frac{\phi_k(a, b) - b}{2^k} \right\rfloor, \quad b = \phi_k(a, b) - a 2^k$$

This implementation is particularly fast because it requires only few bit-wise instructions:

$$\phi_k(a, b) = a \ll k | b \quad \Rightarrow \quad a = \phi_k(a, b) \gg k, \quad b = \phi_k(a, b) \& (2^k - 1)$$

where  $\ll$  is the left bit-shift,  $|$  is the bit-wise OR and  $\&$  is the bit-wise AND (see Supplementary S7.1.1 for an example and an implementation of the encoding). Since the encoding uses  $2k$  bits, it has the best performances when it fits into a CPU word, which is usually 64-bits on modern computers, meaning that the graph must have less than  $2^{32}$  nodes and less than  $2^{64}$  edges. However, by using multi-word integers it can be easily extended to even larger graphs.

#### 4.1.3 Operations on Elias-Fano.

The aforementioned encoding, when paired with Elias-Fano representation, allows an even more efficient computation of random-walk samples. Indeed, Elias-Fano representation allows performing **rank** and **select** operations by requiring on average constant time. These two operations were initially introduced by Jacobson to simulate operations on general trees, and were subsequently proven fundamental to support operations on data structures encoded through efficient schemes. In particular, given a set of integers  $S$ , Jacobson defined the **rank** and **select** operations as follows [52, 53]:

**rank**( $S, m$ )      returns the number of elements in  $S$  less or equal than  $m$   
**select**( $S, i$ )      returns the  $i$ -th smallest value in  $S$

As explained below, to speed up computation, we deviate from this definition by defining the rank operation as the number of elements strictly lower than  $m$ . To compute the neighbours of a node using the rank and select operations, we observe that for every pair of nodes  $\alpha, \beta$  with numerical ids  $a, b$  respectively, it holds that:

$$a 2^k \leq a 2^k + b < (a + 1) 2^k \quad \Rightarrow \quad \phi_k(a, 0) \leq \phi_k(a, b) < \phi_k(a + 1, 0)$$

Thus, the encoding of all the edges with source  $\alpha$  will fall in the discrete range

$$\left[ \phi_k(a, 0), \phi_k(a + 1, 0) \right) = \left[ a 2^k, (a + 1) 2^k \right)$$

Thanks to our definition of the **rank** operation and the aforementioned property of the encoding, we can easily derive the computation of the degree  $d(a)$  of any node  $v$  with numerical id  $a$  for the set of encoded edges  $\Gamma$  of a given graph, which is equivalent to the number of outgoing edges from that node:

$$d(a) = \mathbf{rank}(\Gamma, \phi_k(a + 1, 0)) - \mathbf{rank}(\Gamma, \phi_k(a, 0))$$

Moreover, we can retrieve the encoding of all the edges  $\Gamma_a$  starting from  $v$  encoded as  $a$ , by selecting every index value  $i$  falling in in the range  $[\phi_k(a, 0), \phi_k(a + 1, 0))$ :

$$\Gamma_a = \left\{ \mathbf{select}(\Gamma, i) \mid \mathbf{rank}(\Gamma, \phi_k(a, 0)) \leq i < \mathbf{rank}(\Gamma, \phi_k(a + 1, 0)) \right\}$$

We can then decode the numerical id of the destination nodes from  $\Gamma_a$ , thus finally obtaining the set of numerical ids of the neighbours nodes  $N(a)$ :

$$N(a) = \left\{ \mathbf{select}(\Gamma, i) \ \&\mathcal{L}(2^k - 1) \mid \mathbf{rank}(\Gamma, \phi_k(a, 0)) \leq i < \mathbf{rank}(\Gamma, \phi_k(a + 1, 0)) \right\}$$

In this way, by exploiting the above integer encoding of the graph and the Elias-Fano data scheme, we can efficiently compute the degree and neighbours of a node using rank and select operations.

#### 4.1.4 Efficient implementation of Elias-Fano.

The performance and complexity of Elias-Fano heavily relies on the details of its implementation. In this section our implementation is sketched, to show how we obtain an average constant time complexity for rank and select operations. A more detailed explanation can be found in the Supplementary Section S7.1.

Elias-Fano is essentially aimed at the efficient representation of a sorted list of integers  $y_0, \dots, y_n$  bounded by  $u$ , i.e.  $\forall i \in \{1, \dots, n - 1\}$  it represents  $0 \leq y_{i-1} \leq y_i \leq y_{i+1} \leq u$ .

To this aim, it firstly splits each value,  $y_i$ , into a low-bits,  $l_i$ , and a high-bits part,  $h_i$ , where it can be proven that the optimal split between the high and low bits requires  $\lceil \log_2 \frac{u}{n} \rceil$  bits [53].

The lower-bits are consecutively stored into a low-bits array  $L = [l_1, \dots, l_n]$ , while the high-bits are stored in a bit-vector  $H = [h_1, \dots, h_n]$ , by concatenating the inverted unary encoding<sup>9</sup>,  $\mathcal{U}$ , of the differences (gaps) between consecutive high-bits parts:  $H = [\mathcal{U}(h_1 - 0), \mathcal{U}(h_2 - h_1), \dots, \mathcal{U}(h_n - h_{n-1})]$  (see supplementary Figures S21 and S23 for a more detailed illustration of this scheme).

The rank and select operations on the Elias-Fano representation require two fundamental operations that consist in finding the  $i$ -th 1 or 0 on a bit-vector. To perform them in an average constant time, having preset a quantum  $q$ , we build an index for the zeros,  $O_0 = [o_1, \dots, o_k]$ , that stores the position of every  $q$  zeros, and an index for the ones,  $O_1 = [o_1, \dots, o_k]$ , that similarly stores the position of every  $q$  ones.

<sup>9</sup>The inverted unary encoding represents a non-negative integer,  $n$ , with  $n$  zeros followed by a one. As an example, 5 is represented a 000001.

Thanks to the constructed index, when the  $i$ -th value  $v$  must be found, the scan can be started from a position,  $o_j$ , for  $j = \lfloor \frac{i}{q} \rfloor$  that is already near to the  $i$ -th  $v$ . Therefore, instead of scanning the whole high-bits array for each search, we only need to scan the high-bits array from position  $o_j$  to position  $o_{j+1}$ .

It can be shown that such scans take an average constant time  $\mathcal{O}(q)$  at a low expense of the memory complexity, since we need  $\mathcal{O}\left(\frac{n}{q} \log_2 n\right)$  bits for storing the two indexes (Supplementary section S7.1). Indeed, in our implementation we chose  $q = 1024$  which provides good performance at the cost of a low memory overhead of 3.125% over the high-bits and, on average, for every select operation we need to scan 16 words of memory.

#### 4.1.5 Optional memory-time trade-offs

Even if the Elias-Fano quasi-succinct data structure enables efficient operations on graphs, *Ensmallen* provides the following three options that may be set to further speed-up the computation at the expense of a more expensive usage of the main memory.

**Explicit destinations vector** The first and most important option is to explicitly create the vector of the destination nodes, avoiding to execute a select from the Elias-Fano data structure each time a given destination node must be chosen. This allows achieving a speedup during the random walks (on average an x3-4 speedup) while spending twice as much memory.

**Explicit out-bound degrees vector** The second most important option is to create the vector of the out-bound node degree, which avoids extracting the degree of a source node from the Elias-Fano data structure. While spending a relatively limited amount of RAM, this grants on average an additional 10% speedup in the computation of random walks. When combined with the explicit destinations vector, it can achieve a combined speedup of up to x5-6 of random walks' computation. We suggest enabling this option when computing a random walk-based model, such as CBOW or SkipGram.

**Explicit sources vector** In the context of the generation of edge-prediction batches, the explicit creation of both the vector of sources and the vector of destinations avoids accessing the Elias-Fano data-structure at all. In this way we spend around three times more memory, but we can achieve between three to four times speedup for the generation of edge-prediction batches.

## 4.2 Spectral and matrix factorization embedding methods

Spectral and matrix factorization methods start by computing weighted adjacency matrices, and may include one or more factorization steps. Secondly, given a target embedding dimensionality  $k$ , these models generally use as embeddings the  $k$  eigenvectors or singular vectors corresponding to spectral or singular values of interest.

Laplacian Eigenmap (LE) computes the symmetrically normalized Laplacian and then computes as embeddings the eigenvectors corresponding to the  $k$  *smallest* eigenvalues. Building on top of LE, Geometric Laplacian Eigenmap Embedding (GLEE) [19], also computes the symmetrically normalized Laplacian and then computes the eigenvectors corresponding to the  $k$  *largest* eigenvalues. High-Order Proximity preserved Embedding (HOPE) [17] starts by computing a node-proximity matrix, where the proximity between two nodes may be defined in different ways, e.g. by using the number of common neighbors, or the Adamic-Adar index. Then HOPE computes the singular vectors corresponding to the  $k$  most significant singular values of the proximity matrix and uses the left and right product of the singular values with the singular vectors as the embeddings of, respectively, the source and destination nodes. Similar to HOPE, the Social Dimensions (SocioDim) approach computes the **dense** modularity

matrix and uses as node embedding the eigenvectors corresponding to the *largest*  $k$  eigenvalues [20]. Alternating Direction Method of Multipliers for Non-Negative Matrix Factorization (NMFADMM) [38] leverages NMF to factorize the left Laplacian matrix into two non-negative matrices, which correspond to the embeddings of the source and destination nodes. Similar to NMFADMM, Iterative Random Projection Network Embedding (RandNE) [39] applies an iterative procedure to factorize the dot product of the left Laplacian and an (initially) random matrix  $R$ ; after a user defined number of factorization the matrix  $R$  is used as the node embeddings. Graph Representations (GraRep) [40] analogously factorizes the left Laplacian matrix and, at every iteration, computes the singular vectors corresponding to the  $k$  most significant singular values, hence producing several embeddings equal to the number of iterations. The Network Matrix Factorization (NetMF), given a window size, first computes a sparse log co-occurrence matrix by using first-order random walks and then proceeds to compute the singular vectors corresponding to the  $k$  largest singular values [18].

*GRAPE* provides efficient parallel methods to compute the initial weighted adjacency matrix of the various implemented methods, which are computed either as dense or sparse matrices depending on how many non-zero values the metrics are expected to generate. The computation of the singular vectors and eigenvectors are currently computed using the state-of-the-art LAPACK library [34], though more scalable methods that compute the vectors using an implicit representation of the weighted matrices are currently under investigation.

### 4.3 First and second-order random walk methods.

First- and second-order random-walk embedding models are shallow neural networks generally composed by two layers and trained on batches of random-walk samples. Given a window size, these models learn some properties of the sliding windows on the random walks, such as the co-occurrence of two nodes in each window, the window central node given the other nodes in the window, or vice-versa the nodes in the window from the window central node. The optimal window size value may vary considerably depending on the graph diameter and overall topology. Once the shallow model has been optimized, the weights in either the first or the second layer can be used as node embeddings.

DeepWalk and its Walklets [24] extension to a multiscale random-walk representation (detailed below) are first-order random-walk sampling methods.

Node2Vec [23], is a second-order random walk method that uses weights to bias the walk towards breadth-first search or depth-first search (section 4.3.1). Node2Vec random walks are more computationally expensive than first-order random walks (see figure 2 c and e), since they require to tune two parameters, and our experimental results showed that models trained on Node2Vec walks do not necessarily outperform models trained on first-order walks, when a sufficient amount of training samples is made available (see figure 4). This of course depends also on the characteristics of the graph, since it is well-known that by tuning the return and in-out parameters of Node2vec we can capture different topological and structural features of the underlying graph [23].

GloVe [22] trains a two-layer neural network to predict the logarithm of the co-occurrence frequency of two nodes within the contextual window of size  $w$  in random walks. CBOW [21] also trains a two-layer neural network to predict the central node of a random walk sequence given the other contextual nodes. SkipGram [21] resembles a transposed version of CBOW: it predicts the contextual nodes of a sequence given its central node. SkipGram has a computational complexity  $w$  times higher than CBOW, as it executes  $w$  times more weights updates for each training sample. As a result, SkipGram models often achieve better performance than CBOW models. Glove, CBOW and SkipGram may be trained with sequences sampled using either DeepWalk or Node2Vec.

Walklets-based SkipGram (or CBOW) computes  $w$  times a DeepWalk-based SkipGram (or CBOW) embedding with window size 1. For each  $0 \dots i \dots w$  embedding, Walklets filters the random-walks by keeping only nodes whose position within the random walk belongs to the congruence class in module

$i$  [24]. This is done to obtain node embedding that learn multi-scale random walk representations.

Role2Vec with Weisfeiler-Lehman Hashing [41, 42, 27] uses first-order random walks to approximate the point-wise mutual information matrix obtained by multiplying the pooled adjacency power matrix with a structural feature matrix (in this case, Weisfeiler-Lehman features) to obtain a structural node embedding.

SkipGram and CBOW models are trained using scale-free negative sampling, which is efficiently implemented using the Elias-Fano data structure rank method.

To obtain reliable embeddings, the training phase of the shallow model would need an exhaustive set of random-walk samples to be provided for each source node, so as to fully represent the source-node context. When dealing with big graphs, the computation of a proper amount of random-walk samples needs efficient routines to represent the graph into memory, retrieve and access the neighbors of each node, randomly sample an integer, and, in case of (Node2Vec) second-order random walks [23], compute the transition probabilities, which must be recomputed at each step of the walk.

The first-order random walk is implemented using a SIMD routine for sampling integers (Supplementary Information S7.2.1). When the graph is weighted, another SIMD routine is used to compute the cumulative sum of the unnormalized probability distribution (Supplementary Information S7.2.2). The implementation of the second-order random walk requires more sophisticated routines described in sections 4.3.1, and 4.3.2. Moreover, in section 4.3.3 we present an approximated weighted and second-order random walk that allows to deal with high-degree nodes.

#### 4.3.1 Implementation of second-order random walks

Node2Vec is a second-order random-walk sampling method [23], whose peculiarity relies in the fact that the probability of stepping from one node  $v$  to its neighbours considers the preceding step of the walk (Supplementary Fig. S27). More precisely, Node2Vec defines the un-normalized transition probability  $\pi_{vx}$  of moving from  $v$  to any direct neighbor  $x$ , starting at a previous step from node  $t$ , as a function of the weight  $w_{vx}$  on the edge connecting  $v$  and  $x$  ( $v, x$ ), and a search bias  $\alpha_{pq}(t, x)$ :

$$\pi_{vx} = \alpha_{pq}(t, x) w_{vx}$$

The search bias  $\alpha_{pq}(t, x)$  is defined as a function of the distance  $d(t, x)$  between  $t$  and  $x$ , and two parameters  $p$  and  $q$ , called, respectively, the *return* and *in-out* parameters:

$$\alpha_{pq}(t, x) = \begin{cases} \frac{1}{p} & \text{if } d(t, x) = 0 \\ 1 & \text{if } d(t, x) = 1 \\ \frac{1}{q} & \text{if } d(t, x) = 2 \end{cases} \quad (1)$$

If the return parameter  $p$  is small, the walk will be enforced to return to the preceding node; if  $p$  is large, the walk will otherwise be encouraged to visit new nodes. The in-out parameter  $q$  allows to vary smoothly between Breadth First Search (BFS) and Depth First Search (DFS) behaviour. Indeed, when  $q$  is small the walk will prefer outward nodes, thus mimicking DFS; it will otherwise prefer inward nodes emulating in this case BFS. Since  $\alpha$  must be recomputed at each step of the walk, the algorithm to compute it must be carefully designed to guarantee scalability.

In *GRAPE* we sped up its computation by decomposing the search bias  $\alpha_{pq}(t, x)$  into the in-out bias  $\beta_q(t, x)$ , related to the  $q$  parameter, and the return bias  $\gamma_p(t, x)$ , related to  $p$ :

$$\alpha_{pq}(t, x) = \beta_q(t, x) \gamma_p(t, x) \quad (2)$$

where the two new biases are defined as:

$$\beta_q(t, x) = \begin{cases} 1 & \text{if } d(t, x) \leq 1 \\ \frac{1}{q} & \text{if } d(t, x) = 2 \end{cases} \quad \gamma_p(t, x) = \begin{cases} \frac{1}{p} & \text{if } d(t, x) = 0 \\ 1 & \text{if } d(t, x) > 0 \end{cases} \quad (3)$$



It is easy to see that eq. 2 is equivalent to eq. 1.

**Efficient computation of the in-out and return bias.** The in-out bias can be re-formulated to allow an efficient implementation: starting from an edge  $(t, v)$  we need to compute  $\beta_q(t, x)$  for each  $x \in N(v)$ , where  $N(v)$  is the set of nodes adjacent to  $v$  including node  $v$  itself.

$$\beta_q(t, x) = \begin{cases} 1 & \text{if } d(t, x) \leq 1 \\ \frac{1}{q} & \text{otherwise} \end{cases} \quad \Rightarrow \quad \beta_q(t, x) = \begin{cases} 1 & \text{if } x \in N(t) \\ \frac{1}{q} & \text{otherwise} \end{cases}$$

This formulation (Supplementary Fig. S26) allows us to compute in batch the set of nodes  $X_\beta$  affected by the in-out parameter  $q$ :

$$X_\beta = \left\{ x \mid \beta_q(t, x) = \frac{1}{q}, q \neq 1 \right\} = N(v) \setminus N(t)$$

where  $N(v)$  are the direct neighbors of node  $v$ . In this way, the selection of the nodes  $X_\beta$  affected by  $\beta_q$  simply require to compute the difference of the two sets  $N(v) \setminus N(t)$ . We efficiently compute  $X_\beta$  by using a SIMD algorithm implemented in assembly, leveraging AVX2 instructions that work on node-set representations as sorted vectors of the indices of the nodes (see Supplementary sections S7.2.1 and S7.2.2 for more details). The return bias  $\gamma_p$  can be simplified as:

$$\gamma_p(t, x) = \begin{cases} \frac{1}{p} & \text{if } d(t, x) = 0 \\ 1 & \text{otherwise} \end{cases} \quad \Rightarrow \quad \gamma_p(t, x) = \begin{cases} \frac{1}{p} & \text{if } t = x \\ 1 & \text{otherwise} \end{cases}$$

It can be efficiently computed using a binary search for the node  $t$  in the sorted vector of neighbours. Summarizing, we re-formulated the transition probability  $\pi_{vx}$  of a second-order random walk in the following way:

$$\pi_{vx} = \beta_q(t, x) \gamma_p(t, v, x) w_{vx} \quad \beta_q(t, x) = \begin{cases} 1 & \text{if } x \in N(t) \\ \frac{1}{q} & \text{otherwise} \end{cases} \quad \gamma_p(t, v, x) = \begin{cases} \frac{1}{p} & \text{if } t = x \\ 1 & \text{otherwise} \end{cases}$$

If  $p, q$  are equal to one, the biases can be simplified, so that we can avoid computing them. In general, depending on the values of  $p, q$  and on the type of the graph (weighted or unweighted), *GRAPE* provides eight specialized implementation of the Node2Vec algorithm, to significantly speed-up the computation (Supplementary Tables S50 and S51). *GRAPE* automatically selects and runs the specialized algorithm that corresponds to the choice of the parameters  $p, q$  and the graph type. This strategy allows a significant speed-up. For instance, in the base case ( $p = q = 1$  and an unweighted graph) the specialized algorithm runs more than 100 times faster than the most complex one ( $p \neq 1, q \neq 1$ , weighted graph). Moreover, as expected, we observe that the major bottleneck is the computation of the in-out bias (Supplementary Table S51).

### 4.3.2 Efficient sampling for Node2Vec random walks

Sampling from a discrete probability distribution is a fundamental step for computing a random walk and can be a significant bottleneck. Many graph libraries implementing the Node2Vec algorithm speed up sampling by using the Alias method (see Supplementary section S.7.2.3), which allows sampling in constant time from a discrete probability distribution with support of cardinality  $n$ , with a pre-processing phase that scales linearly with  $n$ .

The use of the Alias Method for Node2Vec incurs the “memory explosion problem” since the preprocessing phase for a second-order random walk on a graph with  $|E|$  edges has a support whose cardinality is  $\mathcal{O}\left(\sum_{e_{ij} \in E} \deg(j)\right)$ , where  $\deg(j)$  is the degree of the destination node of the edge  $e_{ij} \in E$ .

Therefore, the time and memory complexities needed for preprocessing make the Alias method impractical even on relatively small graphs. For instance, on the unfiltered Human STRING PPI graph (19.354 nodes and 5.879.727 edges) it would require 777 GB of RAM.

To avoid this problem, we compute the distributions on the fly. For a given source node  $v$ , our sampling algorithm applies the following steps: 1) computation of the un-normalized transition probabilities to each neighbour of  $v$  according to the provided *in-out* and *return* biases; 2) computation of the un-normalized cumulative distribution, which is equivalent to a cumulative sum; 3) uniform sampling of a random value between 0 and the maximum value in the un-normalized cumulative distribution; 4) identification of the corresponding index through either a linear scan or a binary search, according to the degree of the node  $v$ .

To compute the cumulative sum efficiently, we implemented a SIMD routine that processes at once in CPU batches of 24 values. Moreover, when the length of the vector is smaller than 128, we apply a linear scan instead of a binary search because it is faster thanks to lower branching and better cache locality. Further details are available in the Supplementary section S7.2.2.

### 4.3.3 Approximated random walks

Since the computational time complexity of the sampling algorithm for either weighted or second-order random walks scales linearly with the degree of the considered source node, computing an exact random walk on graphs with high degree nodes (where "high" refers to nodes having an outbound degree larger than 10000) would be impractical, also considering that such nodes have a higher probability to be visited.

To cope with this problem, we designed an approximated random walk algorithm, where each step of the walk considers only a sub-sampled set of  $k$  neighbors, where the parameter  $k$  is set to a value significantly lower than the maximum node degree.

An efficient neighborhood sub-sampling for nodes with degree greater than  $k$  requires to uniformly sample unique neighbors whose original order must be maintained. To uniformly sample distinct neighbors in a discrete range  $[0, n]$  we developed an algorithm (Sorted Unique Sub-Sampling - SUSS, see Supplementary S7.2.4) that divides the range  $[0, n]$  into  $k$  uniformly spaced buckets and then randomly samples a value in each bucket. This allows to efficiently extract  $k$  distinct and sorted values with complexity  $\Theta(k)$ . The disadvantage of this sub-sampling approach is that two consecutive neighbors will never be selected in the same sub-sampled neighborhood.

Nevertheless considering that the sub-sampling is repeated at each step of the walk, consecutive neighbors have the same probability of being selected in different sub-samplings.

## 4.4 Triple-sampling methods.

Triple sampling methods are shallow neural networks trained on triples,  $(v, \ell, s)$ , where  $\{v, s\}$  is a node-pair composed of a source ( $v$ ) and a destination node ( $s$ ), and  $\ell$  is a property of the edge  $(v, s)$  connecting them; this property may, for instance, tell whether the edge exists, or can define the edge type, or its edge weight. They often include a single layer representing the node embedding, and in some models a second layer representing the contextual node embedding or the considered edge properties.

First-order LINE [25] samples node tuples and evaluates whether such tuples correspond to a known edge in the graph. The model optimizes a single layer corresponding to the node embedding; in this way, connected nodes will obtain a *parallel* embedding, while disconnected nodes will obtain an *orthogonal* embedding. Second-order LINE [25] differ from first-order models as they optimize two layers. The first layer corresponds to the source-node embeddings, the second layer corresponds to the destination-node (or contextual-node) embeddings. Analogous to the first-order model, connected

source and destination nodes will obtain a parallel embedding, while disconnected nodes will obtain an orthogonal embedding.

## 4.5 Corrupted triple-sampling methods.

Similar to triple sampling methods, *corrupted*-triple sampling methods are shallow neural networks trained on the (*true*) triples  $(v, \ell, s)$  defined by the existing edges in the graph (where  $v$  is the source node,  $\ell$  is the property of the edge  $(v, s)$ , and  $s$  is the destination node, see section 4.4), but also on *corrupted triples*, that are obtained by corrupting the original triples by substituting the source and/or destination nodes  $\{v, s\}$  with randomly sampled nodes  $\{v', s'\}$ , while maintaining the attribute unchanged  $(v', \ell, s')$ .

The shallow neural network models used on corrupted-triple sampling batches include a weight matrix representing the node embedding, plus one or more matrices for representing the edge attributes, which are composed to capture the attribute meaning as algebraic operations (e.g. *woman+is\_royal = queen*). For this reason, they are particularly well suited to compute node and edge properties embedding of attributed graphs where the edge properties represent meaningful directed transitions (e.g. *is\_royal*), while being out of scope when dealing with local undirected properties (e.g. *interacts with*). Given a distance metric defined for the triples the shallow models are generally optimized to minimize the distance of *true* triples while maximizing the distance of the *corrupted* ones.

The distance defined for triples is often a feature-wise distance, whose advantage is that the computation of the gradient of each feature is independent from any other feature. This allows for particularly effective data-racing-aware and synchronization-free parallel implementations [26].

TransE [26] is among the first and possibly one of the most commonly used of the corrupted-triple sampling methods presented in the literature, from which a large family of variations has been defined. The model trains a shallow neural network composed of two weight matrices representing the node embedding and the edge type embedding. It generally uses as distance metric a feature-wise euclidean distance (though any element-wise distance metric may be used) and defines its energy loss as:

$$\mathcal{L}_{TransE} = \sum_{(v,v',\ell,s,s')} \text{ReLU} [\text{constant} + (v + \ell - s)^2 - (v' + \ell - s')^2]$$

As per the aforementioned properties of feature-wise distances, our Rust implementation of the TransE model is a synchronization-free parallel implementation.

A large set of corrupted-triple sampling models is integrated from the PyKeen library. The integrated models include TransH, DistMult, HolE, AutoSF, TransF, TorusE, DistMA, ProjE, ConvE, RESCAL, QuatE, TransD, ERMLP, CrossE, TuckER, TransR, PairRE, RotatE, ComplEx, and BoxE [9]. We refer to each of the original papers for the extensive explanation. The parameters used for the evaluation of node embedding models in *GRAPE* pipelines are available in the Supplementary Information S4.1.

## 4.6 Edge embedding methods

*GRAPE* offers an extensive set of methods to compute edge embeddings from node embeddings (e.g. concatenation, average, cosine distance, L1, L2 and Hadamard operators [23])<sup>10</sup>. To meet the various model requirements, the library provides three implementations of the edge embedding. In the first one, all edge embedding methods are implemented as Keras/TensorFlow layers and may be employed in any Keras model. In the second one, all methods are also provided in a NumPy implementation. Finally, a third one uses Rust for models where performance is particularly relevant. For instance, the cosine similarity computation in the Rust implementation is over 250× faster than the analogous

<sup>10</sup>the choice of the specific edge-embedding operator is left to the user, who can set it through a parameter

NumPy implementation. Whenever possible, the computation of edge embeddings is executed lazily for a given subset of the edges at a time since the amount of RAM required to explicitly rasterize the edge embedding can be prohibitive on most systems, depending on the edge set cardinality of the considered graph. More specifically, while the lazy generation of edge embeddings is possible during training for only a subset of the supported edge and edge-label prediction models, it is supported for all models during inference.

## 4.7 Node-label, edge-label, and edge-label prediction models

*GRAPE* provides implementations to perform node-label prediction, edge-label prediction and edge prediction tasks.

All the models devoted to any of the three prediction tasks share the following implementation similarities. Firstly, they all implement the *abstract classifier interface* and therefore provide straightforward methods for training (*fit*) and inference (*predict* and *predict\_proba*).

Secondly, all models are multi-modal, that is, they not only can receive the (user-defined) node/edge embedded representation, but also other embeddings computed in multiple ways and therefore carrying different semantics (e.g., topological node/edge embeddings or BERT embeddings). For edge prediction and edge-label prediction models, this also generalizes to multiple node-type features, which, if available, are concatenated to the considered node features, and to the possibility of computing traditional edge metrics (e.g. Jaccard, Adamic-Adar, and so on).

For each task, we make available at least eight models from the literature, adapted to the considered task: 5 are Scikit-learn-based models, namely Random Forest, Extra Trees, Decision Tree, Multi-Layer Perceptron (MLP), and Gradient Boosting. The remaining 3 are TensorFlow-based models, namely GraphSAGE [1], Kipf GCN [54] and a baseline GNN.

As per the node embedding models, custom and third party models can be integrated through task-specific Python abstract classes (Section 4.8).

Scikit-learn-based models make available all the parameters that are available in the Scikit version. TensorFlow-based models make available parameters to set the number of layers in each provided feature’s sub-module and head module.

Visualizations of the Kipf GCN model for node-label, edge-label and edge prediction tasks are also available (see Supplementary Information S8).

All edge prediction models can be trained by sampling the graph negative edges by either following a uniform or a scale-free distribution; by default we set a scale-free distribution because it generally produces more informative negative-training sets, characterized by a smaller covariate-shift with respect to the positive-set. This approach still guarantees a negligible number of false negatives edges. The unbalance between positive and negative edges is also a free parameter which may be arbitrarily set: by default the models are trained using a balanced approach, that is we sample a number of negative edges equal to the number of positive edges.

In addition to the eight models presented in section 4.7, we also make available a multi-modal perceptron model implemented in Rust. This model, analogously to all other models, supports lazy computation of edge embedding and edge features, but does this in an extensively parallel manner with no additional memory requirement over the model weights. The model optimizer is Nadam. The Perceptron is a great baseline for comparison, given its rapid convergence, minimal hardware requirements (no GPUs nor significant RAM requirement), and competitive performance in many considered tasks. Such a model is essential to put into perspective the improvements achieved by more complex and often significantly more expensive models.

Parameters used for the evaluation of edge prediction models in *GRAPE* pipelines are available in the Supplementary Information S4.2.

All of the provided edge-label prediction models support binary and multi-class classification tasks. We currently lack support for multi-label classification tasks, which is being addressed.

All of the provided node-label prediction models support binary, multi-class and multi-label classification tasks. Parameters used for the evaluation of node-label prediction models in *GRAPE* pipelines are available in the Supplementary Information S4.3.

## 4.8 Pipelines for the evaluation of graph-prediction tasks

To provide actionable and reliable results, the fair and objective comparative evaluation of datasets, graph embedding, and prediction models is crucial and not only requires specifically designed and real-world benchmark datasets [37], but also pipelines that could allow non-expert users to easily test and compare graphs and inference algorithms on the desired graphs.

*GRAPE* provides pipelines for evaluating node-label, edge-label and edge prediction experiments trained on user-defined embedding features and by using task-specific evaluation schemas (Section 2.5.2).

All the implemented pipelines have integrated support for differential caching, storing the results of every step of the specific experiment, and for “smoke tests”, i.e. for running a lightweight version of the experimental setup with minimal requirements to ensure execution until completion before running the full experiment.

The pipelines can use any model implementing a standard interface we developed. The interface requires the model to implement methods for training (*fit* or *fit\_transform*), inference (*predict* and *predict\_proba*) plus additional metadata methods (e.g., whether to use node types, edge types, and others) which are used to identify experimental flaws and biases. As an example, in an edge-label prediction task using node embeddings, *GRAPE* will use the provided metadata to check whether the selected node embedding method also uses edge labels. If so, the node embedding will be recomputed during each holdout. Conversely, if the edge labels are not used in the node embedding method, it may be computed only once. The choice to recompute the node embedding for each holdout, which may be helpful to gauge how much different random seeds change the performance, is left to the user in this latter case.

To configure one of the comparative pipelines, users have to import the desired pipeline from the *GRAPE* library and specify the following modular elements:

**Graphs** The graphs to evaluate, which can be either graph objects or strings matching the names from graphs retrieval.

**Graph normalization callback** For some graphs it is necessary to execute normalization steps and filtering, such as the STRING graphs which can e.g. to be filtered at 700 minimum edge weight. For this reason, users can provide this optional normalization callback.

**Classifier models** The classifier models to evaluate, which can either be a model implemented in *GRAPE* or custom models implementing the proper interface.

**Node, node type, and edge features** The features to be used to train the provided classifier models. These features can be node embedding models, either implemented in *GRAPE* or custom embedding models implementing the node embedding interface.

**Evaluation schema** The evaluation schema to follow for the evaluation.

Given any input graph, each pipeline starts by retrieving it (if the name of the graph was provided) and validating the provided features (checking for NaNs, constant columns, compatibility with the

provided graphs); next, and if requested by the user, it computes all the node-embeddings to be used as additional features for the prediction task. Once this preliminary phase is completed, the pipeline starts to iterate and generate holdouts following the provided evaluation schema.

For each holdout, *GRAPE* then computes the node embeddings required to perform the prediction task (such as topological node embeddings for a node-label prediction task, or topological node embeddings followed by their combination through a user-defined edge embedding operator - see Section 4.6 - to obtain the edge embedding in an edge-prediction task), so that a new instance of the provided classifier models can be fitted and evaluated (by using both the required embedding and, eventually, the additional, label-independent, features computed in the preliminary phase). The classifier evaluation is finally performed by computing an exhaustive set of metrics including AUROC, AUPRC, Balanced Accuracy, Miss-rate, Diagnostic odds ratio, Markedness, Matthews correlation coefficient and many others.

## Data availability

*GRAPE* graph retrieval includes all the graphs used in the *Ensmallen* benchmarks and the pipeline experiments and are all available from <https://github.com/AnacletoLAB/grape>. The real world graphs used in Section 2.6 are downloadable from [https://archive.org/download/ctd\\_20220404/CTD.tar](https://archive.org/download/ctd_20220404/CTD.tar) (Pre-built CTD), [https://archive.org/download/pheknowlator\\_20220411/PheKnowLator.tar](https://archive.org/download/pheknowlator_20220411/PheKnowLator.tar) (Pre-built biomedical PheKnowLator data), and [https://archive.org/download/wikipedia\\_edge\\_list.npy/wikipedia\\_edge\\_list.npy.gz](https://archive.org/download/wikipedia_edge_list.npy/wikipedia_edge_list.npy.gz) (Pre-built English Wikipedia). More details are available in Supplementary Information S6.

## Code availability

All the code of the experiments presented within the manuscript are publicly available from GitHub repositories. *GRAPE* can be installed from PyPI: <https://pypi.org/project/grape>. The source code, reference manual and tutorials for its usage, alongside several application examples, are available on GitHub: <https://github.com/AnacletoLAB/grape>. In particular more than 50 tutorials to learn how to use the main functionalities of *GRAPE* are available from <https://github.com/AnacletoLAB/grape/tree/main/tutorials>.

All the scripts to reproduce the experiments showed in the paper are available from GitHub:

- *Ensmallen* benchmarks: loading graphs, executing first and second-order random walks: [https://github.com/LucaCappelletti94/ensmallen\\_experiments](https://github.com/LucaCappelletti94/ensmallen_experiments)
- Approximated random walk experiments: <https://github.com/AnacletoLAB/grape/blob/main/tutorials/Comparing%20DeepWalk%20and%20Node2Vec%20running%20on%20exact%20and%20approximated%20random%20walks.ipynb>
- Experimental comparison of node embedding methods
  - (a) Edge prediction experiments: <https://github.com/AnacletoLAB/grape/blob/main/tutorials/Using%20the%20edge%20prediction%20pipeline.ipynb>
  - (b) Node-label prediction experiments: <https://github.com/AnacletoLAB/grape/blob/main/tutorials/Using%20the%20node-label%20prediction%20pipeline.ipynb>
- Comparison of *GRAPE* with state-of-the-art libraries on big real-world graphs: [https://github.com/LucaCappelletti94/embiggen\\_experiments/tree/master/node2vec\\_comparisons](https://github.com/LucaCappelletti94/embiggen_experiments/tree/master/node2vec_comparisons)

The software is delivered under the MIT license.

## Competing interests

The authors declare that they have no competing interests.

## Authors' Contributions

Conceptualization and Methodology: L.C., T.F., G.V., E.C., J.R. and P.N.R. Software - Design and implementation: L.C., T.F. with the contribution of V.R., T.C. and J.R. Software - documentation: L.C., T.F. Data Curation and Investigation: J.R., P.N.R., V.R. and T.C. Supervision: G.V., P.N.R., J.R. Funding Acquisition: C.M., P.N.R., G.V. Writing - Original Draft Preparation: G.V., E.C., J.R., P.N.R. Writing - Review & Editing: all authors.

## Funding

This research was supported in part by Transition Grant Line 1A Project "UNIMI PARTENARIATI H2020" [PSR2015-1720GVALE\_01], the Common Fund, Office of the Director, National Institutes of Health [U01-CA239108-02], the Monarch Initiative, National Institute of Health [1R24OD011883-01], and from the Director, Office of Science, Office of Basic Energy Sciences of the U.S. Department of Energy under U.S. Department of Energy Contract No. DE-AC02-05CH11231.

## References

- [1] William L Hamilton. Graph representation learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 14(3):1–159, 2020.
- [2] V. Martinez, F. Berzal, and JC Cubero. A survey of link prediction in complex networks. *ACM Computing Surveys*, 49(6), 2017.
- [3] Gabor Csardi and Tamas Nepusz. The igraph software package for complex network research. *InterJournal*, Complex Systems:1695, 2006.
- [4] Yucheng Low, Joseph Gonzalez, Aapo Kyrola, Danny Bickson, Carlos Guestrin, and Joseph Hellerstein. Graphlab: A new framework for parallel machine learning. In *Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence*, UAI'10, page 340–349, Arlington, Virginia, USA, 2010. AUAI Press.
- [5] A. Hagberg, D. Schult, and P. Swart. Exploring network structure, dynamics, and function using networkx. In *Proceedings of the 7th Python in Science Conference, Pasadena, CA USA*, 2008.
- [6] J. Gonzalez, R. Xin, A. Dave, D. Crankshaw, MJ Franklin, and I. Stoica. GraphX: Graph processing in a distributed dataflow framework. In *11th USENIX Symposium on Operating Systems Design and Implementation (OSDI 14)*, pages 599–613, Broomfield, CO, 2014. USENIX Association.
- [7] Jure Leskovec and Rok Sosič. Snap: A general-purpose network analysis and graph-mining library. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 8(1):1, 2016.
- [8] Renming Liu and Arjun Krishnan. Pecanpy: a fast, efficient and parallelized python implementation of node2vec. *Bioinformatics*, 37(19):3377–3379, 2021.
- [9] Mehdi Ali, Max Berrendorf, Charles Tapley Hoyt, Laurent Vermue, Sahand Sharifzadeh, Volker Tresp, and Jens Lehmann. PyKEEN 1.0: A Python Library for Training and Evaluating Knowledge Graph Embeddings. *Journal of Machine Learning Research*, 22(82):1–6, 2021.

- [10] Minjie Wang, Da Zheng, Zihao Ye, Quan Gan, Mufei Li, Xiang Song, Jinjing Zhou, Chao Ma, Lingfan Yu, Yu Gai, Tianjun Xiao, Tong He, George Karypis, Jinyang Li, and Zheng Zhang. Deep graph library: Towards efficient and scalable deep learning on graphs. In *ICLR workshop on representation learning on graphs and manifolds*, 2019.
- [11] Matthias Fey and Jan E. Lenssen. Fast graph representation learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.
- [12] Daniele Grattarola and Cesare Alippi. Graph neural networks in tensorflow and keras with spektral. *IEEE Computational Intelligence Magazine*, 16(1):99–106, 2021.
- [13] D Zhang, J Yin, X Zhu, and Zhang C. Network representation learning: A survey. *IEEE Transactions on Big Data*, 1:3–28, 2020.
- [14] Jeffrey M Perkel. Why scientists are turning to Rust. *Nature*, 588(7836):185–186, 2020.
- [15] Peter Elias. Efficient storage and retrieval by content and address of static files. *Journal of the ACM (JACM)*, 21(2):246–260, 1974.
- [16] David A Bader and Kamesh Madduri. Parallel algorithms for evaluating centrality indices in real-world networks. In *2006 International Conference on Parallel Processing (ICPP’06)*, pages 539–550. IEEE, 2006.
- [17] Mingdong Ou, Peng Cui, Jian Pei, Ziwei Zhang, and Wenwu Zhu. Asymmetric transitivity preserving graph embedding. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1105–1114, 2016.
- [18] Jiezhong Qiu, Yuxiao Dong, Hao Ma, Jian Li, Kuansan Wang, and Jie Tang. Network embedding as matrix factorization: Unifying deepwalk, line, pte, and node2vec. In *Proceedings of the eleventh ACM international conference on web search and data mining*, pages 459–467, 2018.
- [19] Leo Torres, Kevin S Chan, and Tina Eliassi-Rad. Glee: Geometric laplacian eigenmap embedding. *Journal of Complex Networks*, 8(2):cnaa007, 2020.
- [20] Lei Tang and Huan Liu. Relational learning via latent social dimensions. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 817–826, 2009.
- [21] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeff Dean. Efficient estimation of word representations in vector space. In Yoshua Bengio and Yann LeCun, editors, *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*, 2013.
- [22] Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543, 2014.
- [23] Aditya Grover and Jure Leskovec. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 855–864, 2016.
- [24] Bryan Perozzi, Vivek Kulkarni, Haochen Chen, and Steven Skiena. Don’t walk, skip! online learning of multi-scale network embeddings. In *Proceedings of the 2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2017*, pages 258–265, 2017.
- [25] Jian Tang, Meng Qu, Mingzhe Wang, Ming Zhang, Jun Yan, and Qiaozhu Mei. Line: Large-scale information network embedding. In *Proceedings of the 24th International Conference on World Wide Web*, pages 1067–1077, 2015.



- [26] Denghui Zhang, Manling Li, Yantao Jia, Yuanzhuo Wang, and Xueqi Cheng. Efficient parallel translating embedding for knowledge graphs. In *Proceedings of the International Conference on Web Intelligence*, pages 460–468, 2017.
- [27] Benedek Rozemberczki, Oliver Kiss, and Rik Sarkar. Karate Club: An API Oriented Open-source Python Framework for Unsupervised Learning on Graphs. In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management (CIKM '20)*, page 3125–3132. ACM, 2020.
- [28] Allan Peter Davis, Cynthia J Grondin, Robin J Johnson, Daniela Sciaky, Jolene Wieggers, Thomas C Wieggers, and Carolyn J Mattingly. Comparative Toxicogenomics Database (CTD): update 2021. *Nucleic Acids Research*, 49(D1):D1138–D1143, 10 2020.
- [29] TJ Callahan, IJ Tripodi, LE Hunter, and WA Baumgartner. A framework for automated construction of heterogeneous large-scale biomedical knowledge graphs. In *Bio-Ontologies COSI - Intelligent Systems for Molecular Biology*, 2020.
- [30] Aric Hagberg, Pieter Swart, and Daniel S Chult. Exploring network structure, dynamics, and function using networkx. Technical report, Los Alamos National Lab.(LANL), Los Alamos, NM (United States), 2008.
- [31] Siu Kwan Lam, Antoine Pitrou, and Stanley Seibert. Numba: A llvm-based python jit compiler. In *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC*, pages 1–6, 2015.
- [32] Damian Szklarczyk, Annika L Gable, David Lyon, Alexander Junge, Stefan Wyder, Jaime Huerta-Cepas, Milan Simonovic, Nadezhda T Doncheva, John H Morris, Peer Bork, Lars J Jensen, and Christian von Mering. STRING v11: protein–protein association networks with increased coverage, supporting functional discovery in genome-wide experimental datasets. *Nucleic Acids Research*, 47(D1):D607–D613, 2018.
- [33] Paolo Boldi, Marco Rosa, Massimo Santini, and Sebastiano Vigna. Layered label propagation: A multiresolution coordinate-free ordering for compressing social networks. In Sadagopan Srinivasan, Krithi Ramamritham, Arun Kumar, M. P. Ravindra, Elisa Bertino, and Ravi Kumar, editors, *Proceedings of the 20th international conference on World Wide Web*, pages 587–596. ACM Press, 2011.
- [34] Edward Anderson, Zhaojun Bai, Christian Bischof, L Susan Blackford, James Demmel, Jack Dongarra, Jeremy Du Croz, Anne Greenbaum, Sven Hammarling, Alan McKenney, and Danny C. Sorensen. *LAPACK users' guide*, 1999.
- [35] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [36] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- [37] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. Open graph benchmark: Datasets for machine learning on graphs. *Advances in neural information processing systems*, 33:22118–22133, 2020.
- [38] Dennis L Sun and Cedric Fevotte. Alternating direction method of multipliers for non-negative matrix factorization with the beta-divergence. In *2014 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 6201–6205. IEEE, 2014.

- [39] Ziwei Zhang, Peng Cui, Haoyang Li, Xiao Wang, and Wenwu Zhu. Billion-scale network embedding with iterative random projection. In *2018 IEEE International Conference on Data Mining (ICDM)*, pages 787–796. IEEE, 2018.
- [40] Shaosheng Cao, Wei Lu, and Qionгкаi Xu. Grarep: Learning graph representations with global structural information. In *Proceedings of the 24th ACM international on conference on information and knowledge management*, pages 891–900, 2015.
- [41] Nesreen K Ahmed, Ryan A Rossi, John Boaz Lee, Theodore L Willke, Rong Zhou, Xiangnan Kong, and Hoda Eldardiry. role2vec: Role-based network embeddings. *Proc. DLG KDD*, pages 1–7, 2019.
- [42] Nino Shervashidze, Pascal Schweitzer, Erik Jan Van Leeuwen, Kurt Mehlhorn, and Karsten M Borgwardt. Weisfeiler-lehman graph kernels. *Journal of Machine Learning Research*, 12(9), 2011.
- [43] Rudolf Kadlec, Ondrej Bajgar, and Jan Kleindienst. Knowledge base completion: Baselines strike back. In *Proceedings of the 2nd Workshop on Representation Learning for NLP*, pages 69–74, Vancouver, Canada, August 2017. Association for Computational Linguistics.
- [44] Federico Errica, Marco Podda, Davide Bacciu, and Alessio Micheli. A fair comparison of graph neural networks for graph classification. In *Proceedings of the 8th International Conference on Learning Representations (ICLR)*, 2020.
- [45] Radim Řehůřek and Petr Sojka. Software Framework for Topic Modelling with Large Corpora. In *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, pages 45–50, Valletta, Malta, May 2010. ELRA. <http://is.muni.cz/publication/884893/en>.
- [46] Richard A Kronmal and Arthur V Peterson Jr. On the alias method for generating random variables from a discrete distribution. *The American Statistician*, 33(4):214–218, 1979.
- [47] Justin T. Reese, Deepak Unni, Tiffany J. Callahan, Luca Cappelletti, Vida Ravanmehr, Seth Carbon, Kent A. Shefchek, Benjamin M. Good, James P. Balhoff, Tommaso Fontana, Hannah Blau, Nicolas Matentzoglou, Nomi L. Harris, Monica C. Munoz-Torres, Melissa A. Haendel, Peter N. Robinson, Marcin P. Joachimiak, and Christopher J. Mungall. Kg-covid-19: A framework to produce customized knowledge graphs for covid-19 response. *Patterns*, 2(1):100155, Jan 2021.
- [48] Paolo Boldi, Marco Rosa, Massimo Santini, and Sebastiano Vigna. Layered label propagation: A multiresolution coordinate-free ordering for compressing social networks. In Sadagopan Srinivasan, Krithi Ramamritham, Arun Kumar, M. P. Ravindra, Elisa Bertino, and Ravi Kumar, editors, *Proceedings of the 20th international conference on World Wide Web*, pages 587–596. ACM Press, 2011.
- [49] Gustavo Malkomes, Charles Schaff, and Roman Garnett. Bayesian optimization for automated model selection. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29, Barcelona, Spain, 2016. Curran Associates, Inc.
- [50] Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides. *Design Patterns: Elements of reusable object-oriented software*, 1995.
- [51] Yousef Saad. Parallel iterative methods for sparse linear systems. In *Studies in Computational Mathematics*, volume 8, pages 423–440. Elsevier, 2001.
- [52] Gonzalo Navarro and Eliana Provedel. Fast, small, simple rank/select on bitmaps. In *International Symposium on Experimental Algorithms*, pages 295–306. Springer, Springer, 2012.
- [53] Giulio Ermanno Pibiri and Rossano Venturini. Dynamic elias-fano representation. In *28th Annual symposium on combinatorial pattern matching (CPM 2017)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.

- [54] Max Welling and Thomas N Kipf. Semi-supervised classification with graph convolutional networks. In *J. International Conference on Learning Representations (ICLR 2017)*, 2016.

## Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [GraPEsupplementaryreview.pdf](#)