

# RAFEN – Regularized Alignment Framework for Embeddings of Nodes

Kamil Tagowski<sup>1</sup>[0000–0003–4809–3587], Piotr Bielak<sup>1</sup>[0000–0002–1487–2569], Jakub Binkowski<sup>1</sup>[0000–0001–7386–5150], and Tomasz Kajdanowicz<sup>1</sup>[0000–0002–8417–1012]

Department of Artificial Intelligence,  
Wrocław University of Science and Technology,  
Wrocław, Poland  
`kamil.tagowski@pwr.edu.pl`

**Abstract.** Learning representations of nodes has been a crucial area of the graph machine learning research area. A well-defined node embedding model should reflect both node features and the graph structure in the final embedding. In the case of dynamic graphs, this problem becomes even more complex as both features and structure may change over time. The embeddings of particular nodes should remain comparable during the evolution of the graph, what can be achieved by applying an alignment procedure. This step was often applied in existing works after the node embedding was already computed. In this paper, we introduce a framework – **RAFEN** – that allows to enrich any existing node embedding method using the aforementioned alignment term and learning aligned node embedding during training time. We propose several variants of our framework and demonstrate its performance on six real-world datasets. **RAFEN** achieves on-par or better performance than existing approaches without requiring additional processing steps.

**Keywords:** Dynamic Graphs · Graph Embedding · Embedding Alignment · Graph Neural Networks · Link prediction · Machine Learning.

## 1 Introduction

In recent years, representation learning on structured objects, like graphs, gained much attention in the community. The wide application range, which includes recommender systems, molecular biology, or social networks, motivated a rapid development of new methods for learning adequate representations of such data structures. Especially real-world data often involves accounting for temporal phenomena, which requires handling structural changes over time (i.e., dynamic graphs).

Researchers widely use the embedding of dynamic graphs in a discrete-time manner, i.e., the graph data is aggregated into snapshots and embeddings are computed for each one of them. In opposite to the continuous-time approach, one does not require to provide an online update mechanism, thus allowing one to utilize any classical static embedding approaches or temporal embedding. On

the other hand, the discrete-time approach requires an embedding aggregation method that transforms the sequence of node snapshot embeddings into a single node embedding matrix which captures the whole graph evolution. Unfortunately, due to the stochastic properties of node embedding methods, such embeddings are algebraically incomparable. To solve this problem, most existing approaches apply a post-hoc correction [1,14,15,16].

In this work, we consider a different scenario where we build aligned node embeddings by introducing an additional regularization term to a model’s loss function. This term explicitly aligns embeddings of corresponding nodes from consecutive timesteps, which eventually helps to improve results on a downstream task. In particular, we summarize the contributions of the paper as follows:

- C1. We propose a novel framework **RAFEN** for learning aligned node embeddings in dynamic networks. It can be used with any existing node embedding method (both for static and dynamic networks) and, contrary to existing post-hoc methods, it does not rely on a time-consuming matrix factorization approach.
- C2. We propose three different versions of the alignment function, including one that treats all nodes equally and one that utilizes temporal network measures for weighting the alignment process.
- C3. We evaluate the proposed framework variants in a link prediction task over six real-world datasets and show that our approach allows to improve results on the downstream task compared to other existing methods.

## 2 Related Work

**Static Node Embedding.** To date, plethora of node embedding methods have been introduced. The seminal work of Grover et al. [4] introduced Node2Vec method, which enables learning structural embeddings of nodes in an unsupervised manner. It pushes node embeddings apart when nodes are away from each other but makes embeddings closer for nearby nodes, leveraging nodes co-occurrences on common random walks as a measure of node "distance". Recently, the main body of literature was focused on Graph Neural Networks, among which GCN [8], GIN [18], and GAT [17] are the most notable ones. In particular, the Graph AutoEncoder (GAE) [7] architecture leverages GCNs to learn unsupervised node structural embeddings, using reconstruction of the adjacency matrix as the training task. As static node embeddings can be adopted for time-discrete dynamic graphs, we utilize Node2Vec and GAE as base for computing dynamic node representations in this work.

**Dynamic Node Embedding.** The problem of dynamic node embedding has recently gained more focus from researchers. We can distinguish meta approaches that could utilize any static embedding method to produce dynamic snapshot embeddings [1,14,15,16]. However, most of them utilize only the Node2Vec [4] embeddings method. Another group of models are the ones dedicated for dynamic node embedding, such as Dyngraph2vec [3], where the authors utilize the Autoencoder (AE) architecture to capture the temporal network’s evolving structure

and provide three variants of the model which differ in the representation method of a neighbor vector.

**Learning of compatible representation.** This problem has already been studied in literature [5,10,13]. In [5], the authors present the BC-Aligner method that allows obtaining previous embedding models based on a learnable transformation matrix. [10] achieves model compatibility with an alignment loss that aligns class centers between models with a boundary loss to constraint new features to be more centralized to class centers. [13] proposes backward compatibility training for image classification by adding an influence loss that can capture dataset changes along with new classes.

**Post-hoc alignment for node embeddings** There are already approaches for alignment of node embeddings based on the Orthogonal Procrustes method [1,14,15,16]. In [14,16], the authors use all common nodes between snapshots to perform the alignment. In [1], it has been found that some nodes change their behavior too significantly. Therefore, they use only a subset of common nodes, which is selected based on activity and selection method. This approach is further explored in [15] by introducing a wide range of activity functions that can be used for the alignment process.

### 3 The proposed RAFEN framework

**Definition 1.** A dynamic network (graph) is a tuple  $G_{0,T} = (V_{0,T}, E_{0,T})$ , where  $V_{0,T}$  denotes a set of all nodes (vertices) observed between timestamp 0 and  $T$ , and  $E_{0,T}$  is a set of edges in the same timestamp range. We model such a network as a sequence of graph snapshots  $G_{0,1}, G_{1,2}, G_{2,3}, \dots, G_{T-1,T}$ .

**Definition 2.** Node embedding is a function  $f : V \rightarrow \mathbb{R}^{|V| \times d}$  that maps a set of nodes  $V$  into low-dimensional (i.e.,  $d \ll |V|$ ) embedding matrix  $F$ , where each row represents the embedding of a single node.

**Definition 3.** Node activity scoring function  $s : V_{com} \rightarrow \mathbb{R}^{|V_{com}|}$  (with  $V_{com} = V_{i,j} \cap V_{k,l}$  and  $i < j < k < l$ ) assigns a scalar score to each node from the set of common nodes  $V_{com}$ , which reflects the change in a node’s activity between two snapshots  $G_{i,j}$  and  $G_{k,l}$ .

**Definition 4.** Reference nodes  $V_{ref}$  is a subset of common nodes between two snapshots retrieved by a selection function. We use the percent selection function as defined in [1], which chooses the top  $p\%$  of nodes – i.e.,  $select(S, V) = V_{ref} \subseteq sort_S(V)$ , s.t.  $|V_{ref}| = p|V|$ .

**Problem statement.** Given a dynamic graph  $G_{0,T}$  in the form of discrete graph snapshots (see Definition 1), the objective is to find at any given timestamp  $t \in \{1, 2, \dots, T\}$ , the node embedding  $F_{t-1,t}$  of the current snapshot  $G_{t-1,t}$ , such that  $F_{t-1,t}$  will be algebraically compatible (aligned) with the node embedding

$F_{t-2,t-1}$  of the previous snapshot  $G_{t-2,t-1}$ . Such compatibility is required to properly aggregate snapshot embeddings  $F_{0,1}, F_{1,2}, \dots, F_{t-1,t}$  into a single node embedding matrix  $F_{0,t}$  that summarizes the whole graph evolution (see [1] for more details about the aggregation mechanism).

**RAFEN.** As a solution for the alignment problem, we propose a novel framework RAFEN that enhances the loss function of any existing node embedding method  $L_{\text{model}}$  by means of an alignment regularization term  $L_{\text{alignment}}$ . It allows to learn node embeddings that are aligned with a given anchor embedding (in our case: the previous graph snapshot). Our framework can be used along with expert knowledge by setting a model hyperparameter  $\alpha$  or using temporal network measures to automatically determine the alignment coefficient of each node. We will discuss all framework details in the sections below. Moreover, we present a general overview of our pipeline in Figure 1.

**How to incorporate the alignment term?** We enhance the loss function of the node embedding model  $L_{\text{model}}$  by adding an alignment term  $L_{\text{alignment}}$ . In order to balance between the information from the previous graph snapshot (through the optimization of the  $L_{\text{alignment}}$  loss) and the information from the current snapshot (through the optimization of the  $L_{\text{model}}$  loss), we propose to combine both terms in the following way:

$$L = (1 - \alpha) \cdot L_{\text{model}} + \alpha \cdot L_{\text{alignment}} \quad (1)$$

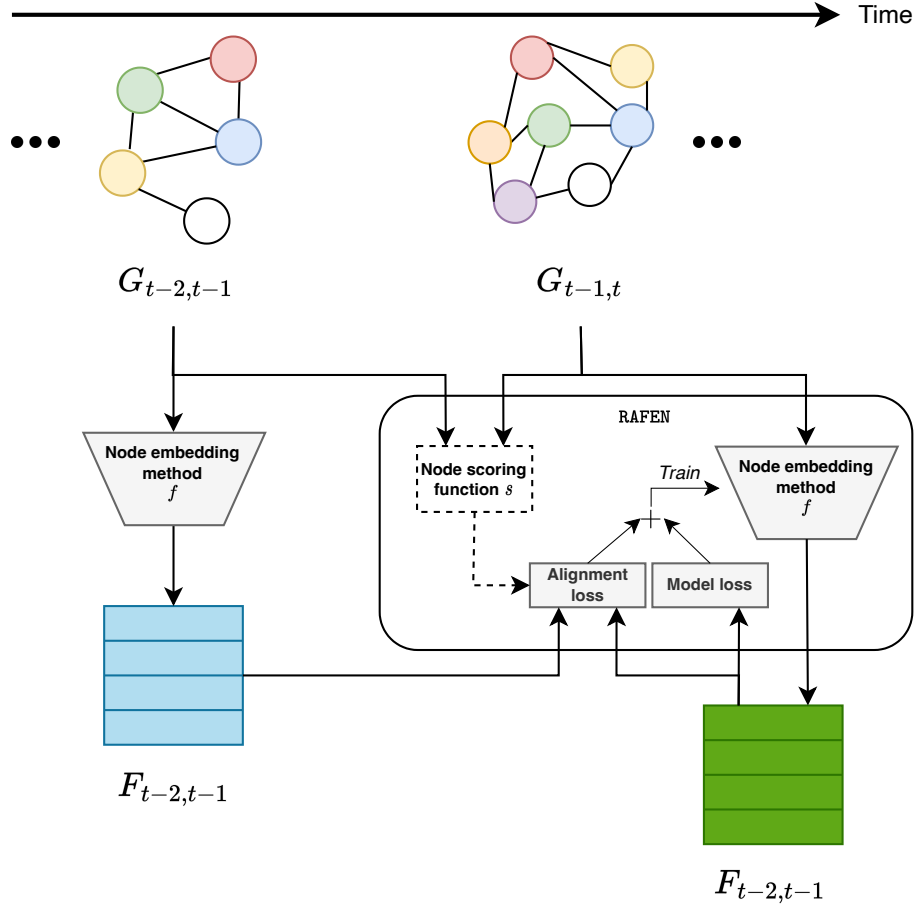
The  $\alpha$  parameter controls the trade-off between the model’s loss and the alignment term, which are both critical factors in the model’s learning behavior. We can distinguish two boundary conditions of the  $\alpha$  parameter:

- $\alpha = 0$  – the model uses only model loss,
- $\alpha = 1$  – the model uses only alignment term.

Choosing the right  $\alpha$  value is essential for the model’s performance. It can be done by a simple grid search or through expert knowledge. However, performing an  $\alpha$  grid search might be time-consuming and computationally expensive. To address this, we propose a simplified version of our framework, where we set equal importance to both loss terms, which is equivalent to setting  $\alpha = 0.5$

$$L = L_{\text{model}} + L_{\text{alignment}} \quad (2)$$

**Alignment loss term ( $L_{\text{alignment}}$ )** The foundation of our proposed RAFEN framework is the alignment loss term  $L_{\text{alignment}}$ . We build upon the mean squared error (MSE) between node  $v$ ’s embedding from the current snapshot  $F_{t-1,t}^{(v)}$  and its representation from the previous snapshot  $F_{t-2,t-1}^{(v)}$ . We apply such regularization to a subset of the common nodes  $V_{\text{com}} \in V_{t-2,t-1} \cap V_{t-1,t}$ .



**Fig. 1.** Overview of the RAFEN pipeline. The pipeline can be decomposed into two steps: (1) For the snapshot  $G_{t-2,t-1}$  that we want align to, we compute node embeddings  $F_{t-2,t-1}$  using the vanilla node embedding method  $f$ . (2) For the next snapshot  $G_{t-1,t}$ , we additionally compute the alignment loss term, which we add to the node embedding method model loss – Eq. (2). The alignment loss term is calculated as an MSE (mean squared error) – Eq. (3) between the representation of common nodes of currently trained node embeddings  $F_{t-1,t}$  and embeddings from the previous snapshot  $F_{t-2,t}$ . There is no need to take all common nodes to calculate the alignment loss. Instead, we apply node scoring function  $s$  that computes the difference in node’s activity between snapshots and selects a subset of common nodes – Eq. (5) or use the output as weights in alignment loss term – Eq. (4)

Let us now define three RAFEN variants:

- **RAFEN\_ALL** – we use all common nodes between consecutive snapshots:

$$L_{\text{alignment}} = \frac{1}{|V_{\text{com}}|} \sum_{v \in V_{\text{com}}} \left( F_{t-1,t}^{(v)} - F_{t-2,t-1}^{(v)} \right)^2 \quad (3)$$

- **RAFEN\_Weighted** – similarly to the previous case, we use all common nodes, but instead of treating each one of them equally, we assign weights that reflect the node activity change (node activity scoring function  $s(\cdot)$  - see Definition 3) of a particular node between the graph snapshots;

$$L_{\text{alignment}} = \frac{1}{|V_{\text{com}}|} \sum_{v \in V_{\text{com}}} \left( \left( F_{t-1,t}^{(v)} - F_{t-2,t-1}^{(v)} \right)^2 \cdot s(v) \right) \quad (4)$$

- **RAFEN\_REF** – in [1,15], the authors show that utilizing all common nodes in the alignment process may lead to degraded performance in the downstream tasks due to the fact that some nodes change their behavior too much. They use  $V_{\text{ref}} \subset V_{\text{com}}$  subset of common nodes, which is built according to the node activity scores (the top  $p\%$  of the common nodes are selected).

$$L_{\text{alignment}} = \frac{1}{|V_{\text{ref}}|} \sum_{v \in V_{\text{ref}}} \left( F_{t-1,t}^{(v)} - F_{t-2,t-1}^{(v)} \right)^2 \quad (5)$$

## 4 Experimental Setup

We evaluate all methods for the link prediction task. Following, we discuss the details and hyperparameters of the experiments.

### 4.1 Datasets

We conducted experiments on six real-world datasets. We followed the graph snapshots split procedure defined in [15], such that we split graphs based on the timestamp frequency (monthly or yearly). Moreover, the first four snapshots in *ppi*, and *ogbl-collab* datasets were ignored, as merging them would result in a too broad timespan of such merged snapshots. Also, the two last snapshots of *bitcoin-alpha* and *bitcoin-otc* were merged to provide more data for validation. We present an overview of datasets in Table 1. Despite the similarities with the experimental setup in [15], we reproduced all experiments due to a refined evaluation protocol on link prediction. Also, our method differs in alignment procedure, i.e., it aligns embedding to the previous snapshot instead of the first.

**Table 1.** Statistics of graph datasets.  $|\mathcal{V}|$  - number of nodes,  $|\mathcal{E}|$  - number of edges, Dir - whether the graph is directed or not, T - number of snapshots

Dataset	$ \mathcal{V} $	$ \mathcal{E} $	Dir	Timespan	T	Snapshot timespan	Network domain
fb-forum [12]	899	33 720	×	5.5 months	5	1 month	social
fb-messages [12]	1 899	61 734	×	7.2 months	7	1 month	social
bitcoin-alpha [12]	3 783	24 186	✓	5.2 years	5	1 year	crypto
bitcoin-otc [12]	5 881	35 592	✓	5.2 years	5	1 year	crypto
ppi [12]	16 386	141 836	×	24 years	5	5 years	protein
ogbl-collab [6]	233 513	1 171 947	×	34 years	7	5 years	citation

## 4.2 Node embeddings

We selected two different methods as our base models for computing node embeddings. First, we used Node2Vec, a widely renowned node structural embedding method that relies on random walks. To show our framework’s universality, we included a second method – GAE, an encoder-decoder model that utilizes the graph convolution operation. Both methods are capable of capturing meaningful structural features of the graph and are trained in an unsupervised manner. Nonetheless, due to the flexibility of our framework, other node embedding models could also be utilized. For both models, each snapshot representation was recomputed 25 times to account for randomness in these methods (e.g., random walks in Node2Vec, weight initialization in the encoder of the GAE model).

**Node2Vec** We trained the Node2Vec model, reusing the hyperparameters configuration from [15], as these were obtained from a hyperparameter search procedure. In particular, we use 128-dimensional embeddings, except for the *bitcoin-alpha* dataset where we use 32 as the embedding size. Further, we use the same configurations in our *RAFEN* models, as we desire that our models perform well without additional hyperparameter search. We utilize the Node2Vec implementation from the PyTorch Geometric library [2].

**GAE** We use a Graph Autoencoder with two Graph Convolutional layers (GCN) as an encoder and the Inner-Product decoder. Since datasets do not come with node features required by the GCN layers, we use an additional trainable lookup embedding layer, which serves as a node feature matrix for the model (we evaluated different strategies in preliminary experiment – this one turned out to work the best). We train these models for 100 epochs using the Adam optimizer with a learning rate of 0.01, and a hidden layer size of 128.

### 4.3 Aligned models

We tested variants of our RAFEN framework with loss functions introduced in section 3, namely `RAFEN_ALL`, `RAFEN_WEIGHTED`, and `RAFEN_REF`. In addition, `RAFEN_WEIGHTED` involves scoring function, denoted as  $s(\cdot)$  in Eq. (4). For scoring function, we selected *Edge Jaccard* (EJ) [15], and *Temporal Betweenness* (TB) [19] due to their good performance in post-hoc alignment methods in previous work. Also, these two functions provide desired variability, such that they are from a different family of methods: *Edge Jaccard* is based only on edge changes in the static scheme, while *Temporal Betweenness* takes into account the time aspect. We denote the two variants as `RAFEN_WEIGHTED_EJ` and `RAFEN_WEIGHTED_TB`, for *Edge Jaccard* and *Temporal Betweenness* respectively.

It is worth noting that  $L_{model}$  and  $L_{alignment}$  are in different scales. Therefore, we additionally scale them with the loss value of the model from the first batch.

### 4.4 Link Prediction

We define the link prediction task as edge existence prediction on the last snapshot of the graph  $G_{T-1,T}$ , based on the previously learned representations from previous snapshots  $F_{0,1}, \dots, F_{T-2,T-1}$ . We take existing links in the graph  $G_{T-1,T}$  as positive examples. To avoid out-of-distribution nodes in evaluation, we filter edges wherein one of the nodes was not observed previously. The number of negative edges (non-existing edges) is equal in size to the positive ones. In the negative sampling process, we also applied a simple edge reject criterion that prevents adding already existing edges as negative ones. Having edges sampled, we split the datasets into train, val and test splits in the proportion: 60%, 20%, 20%, respectively. This split is both leveraged during grid-search hyperparameter optimization and the method evaluation. Finally, using optimal hyperparameters, we performed the final evaluation with training and validation subsets merged. The results of the evaluation are reported in Section 5.

To fuse the node embeddings from all snapshots into a single one that reflects the graph evolution, we evaluate 4 different aggregation schemes:

- `mean` – node embeddings are averaged across snapshots,
- `last` – only the node embedding from the last snapshot is taken,
- `FILDNE` – weighted incremental combination of previous snapshot embeddings as in [1],
- `k-FILDNE` – weighted incremental combination of previous snapshot embeddings with automated weight estimation as in [1].

Further, we use a logistic regression classifier, which takes the Hadamard product of the two aggregated node embeddings. We utilize the implementation from the scikit-learn [11] library with *liblinear* solver, keeping other hyperparameters values default.

Summarizing, the experiments were performed on link prediction with embeddings trained with 3 variants of loss function and 4 embeddings aggregation schemes. We also report results for the two embedding methods (Node2Vec, GAE)



trained without alignment. In addition, we used 3 post-hoc alignment methods from [15] as our baselines. To sum up, we evaluate the following models:

- **(whole graph)** – embeddings are trained on whole graph  $G_{0,T-1}$ ,
- **(last snapshot only)** – embeddings are trained on last graph snapshot  $G_{T-2,T-1}$ ,
- **Posthoc-PA** – Procrustes Aligner (PA),
- **Posthoc-EJ** – Edge Jaccard Aligner (EJA),
- **Posthoc-TB** – Temporal Betweenness (TB),
- **RAFEN\_ALL** – RAFEN variant with all common nodes between consecutive snapshots and simplified loss importance term,
- **RAFEN\_Weighted\_EJ** – RAFEN\_Weighted variant with Edge Jaccard node activity scoring method and simplified loss importance term,
- **RAFEN\_Weighted\_TB** – RAFEN\_Weighted variant with Temporal Betweenness node activity scoring method and simplified loss importance term.

## 5 Results

This section presents the results of our experiments, divided into three subsections presenting different framework characteristics. First, we compare the embedding aggregation methods. Next, we show the results for the two formulations of RAFEN loss function, i.e., with and without hyperparameter  $\alpha$ . Finally, we present the results of several variants of RAFEN models compared to the baselines approaches.

### 5.1 Embeddings aggregation method comparison

Foremost, we evaluated the embedding aggregation methods, namely **last**, **mean**, **FILDNE**, **k-FILDNE**. We summarize these results in Table 2. We chose the best aggregation method that performs well on both GAE and Node2Vec-based models based on mean ranks for each dataset separately, i.e., bitcoin-alpha: **k-FILDNE**; bitcoin-otc: **last**; fb-forum: **FILDNE**; fb-messages: **k-FILDNE**; ogbl-collab: **last**, ppi: **last**. Please note that for the bitcoin-alpha dataset, the best aggregation method for Node2Vec and GAE differs. However, we chose **k-FILDNE** due to a better combination of mean ranks (i.e., 1.22 and 1.56 vs 3.44 and 1.44).

Counterintuitively to previous dynamic graph representation learning papers [1] for four out of six datasets, the best aggregation method was **last**, which does not aggregate all historical representations. That behavior led us to an additional study that we conducted to further explore both Posthoc and RAFEN methods. We contrasted AUC on the link-prediction of **RAFEN\_ALL** and **Posthoc-PA** with plain Node2Vec embeddings and computed their relative performance. In particular, we evaluated two scenarios:

- **Prev** – we take each node’s  $v$  embedding  $F_{t-1,t}^{(v)}$  at time step  $t$  and evaluate it on link-prediction using the previous snapshot’s graph structure  $G_{t-2,t-1}$  as a target,

**Table 2.** Mean ranks of aggregation methods on different datasets. Ranks were calculated on the link prediction task (evaluated on last snapshot), using the AUC metric over 25 runs. Then we took mean ranks of all the RAFEN and Posthoc models for Node2Vec and GAE models, separately.

Dataset	Model	Aggregation method			
		FILDNE	k-FILDNE	last	mean
bitcoin-alpha	Node2Vec	2.11	<b>1.22</b>	3.22	3.44
	GAE	3.00	1.56	4.00	<b>1.44</b>
bitcoin-otc	Node2Vec	1.89	3.00	<b>1.11</b>	4.00
	GAE	1.67	2.89	<b>1.44</b>	4.00
fb-forum	Node2Vec	<b>1.22</b>	2.11	3.11	3.56
	GAE	<b>1.33</b>	2.22	2.78	3.67
fb-messages	Node2Vec	2.00	3.00	<b>1.00</b>	4.00
	GAE	1.83	2.50	<b>1.67</b>	4.00
ogbl-collab	Node2Vec	2.00	3.00	<b>1.00</b>	4.00
	GAE	2.00	3.00	<b>1.00</b>	4.00
ppi	Node2Vec	2.33	2.67	<b>1.00</b>	4.00
	GAE	2.89	2.06	<b>1.11</b>	3.94

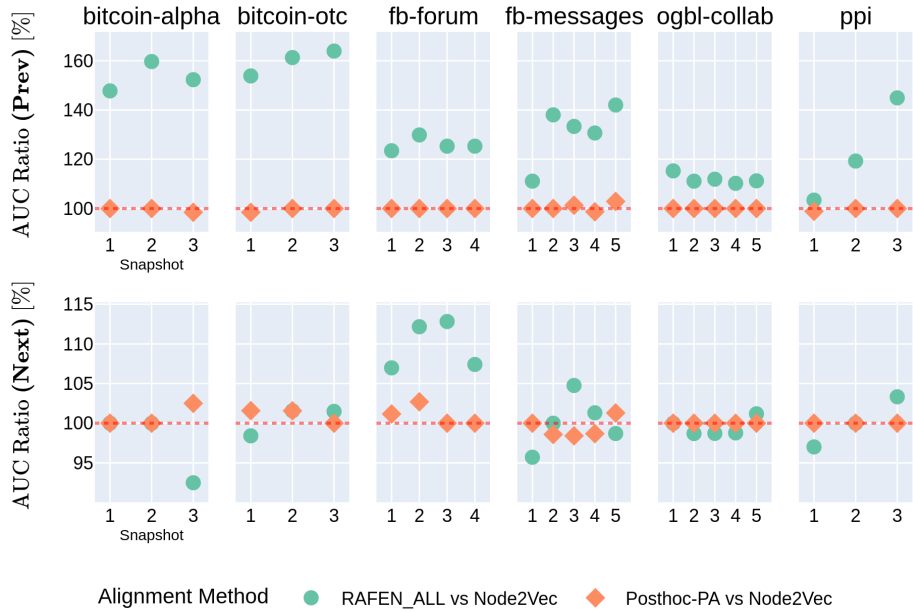
- **Next** – we take each node’s  $v$  embedding  $F_{t-1,t}^{(v)}$  at time step  $t$  and evaluate it on link-prediction using next snapshot’s graph structure  $G_{t,t+1}$  as a target.

We report the link prediction AUC ratio between alignment methods and the non-aligned Node2Vec version. Results are shown in Figure 2.

Our proposed RAFEN method increased the evaluation performance on previous snapshots for all of the benchmarked datasets. The gain ranges from 3.5% on the first snapshot of the ppi dataset up to a 63% difference on the third snapshot of bitcoin-otc. In the case of the **Next** scenario, we observe that the last snapshot evaluation performance increased on four out of six datasets – ranging from 1.2% in the case of the ogbl-collab dataset up to 7.4% in fb-forum.

## 5.2 Comparison of RAFEN’s $\alpha$ -based variants

In section 3, we proposed two formulations of the loss function for our RAFEN framework. Since the first variant defined in Eq. (1) involves an  $\alpha$  hyperparameter, which requires additional domain knowledge or hyperparameter search procedure (e.g., grid search), we investigate whether its simplified variant with  $L_{model}$  and  $L_{alignment}$  treated equally important provides competitive performance. Likewise, RAFEN\_REF models require defining a reference node selection function with hyperparameters, like top  $p\%$  in our case. Therefore, we compare it to RAFEN\_Weighted models that rely only on scoring function and do not involve additional hyperparameters. Results of the comparisons are shown in Table 3.



**Fig. 2.** Single snapshot representation evaluation study. In the first figure row, the previous snapshot evaluation is shown (denoted as **Prev**), and in the second row next snapshot evaluation (denoted as **Next**). Each column represents a different dataset. Y-Axis contains a link prediction AUC metric Ratio calculated between Alignment Method and vanilla Node2Vec. X-Axis presents snapshot IDs on which node representation was trained.

Across almost all the shown comparisons, we can observe that percentage AUC difference in *alpha*-based models to simplified/weighted versions is in the range (-1.61% to 1.71%). We can observe outlier values in both **bitcoin-alpha** and **ppi** favoring  $\alpha$  hyperparameter. Conversely, there is an outlier in **fb-forum** favoring the weighted method. As in most cases, the differences are negligible, we decided to use simplified/weighted versions for the comparison with baselines. In the aftermath, **RAFEN** requires only one additional hyperparameter compared to plain GAE or Node2Vec models, which is the node activity scoring function.

### 5.3 RAFEN comparison to baselines

The results of the comparison between our proposed **RAFEN** framework and baseline models are presented in Table 4. Compared to snapshot-based Node2Vec and GAE baselines, **RAFEN** was better on four datasets: **bitcoin-otc**, **fb-forum**, **ppi**, and **ogbl-collab** – only losing to the variant trained on the full graph data. For the remaining two datasets, namely **fb-messages**, and **bitcoin-alpha**, we observe the superior performance of post-hoc methods. However, as the previous study

**Table 3.** Comparison of  $\alpha$ -based RAFEN with simplified/weighted versions. We present the percentage difference of averaged link prediction AUC metric on the last snapshot over 25 runs. Values below zero denote better performance of the simplified/weighted variant, and values above zero denote better  $\alpha$  variant performance. We mark outlier values with **bold**.

Model	Selector	RAFEN variant	bitcoin alpha	bitcoin otc	fb forum	fb messages	ogbl collab	ppi
Node2Vec	All	$\alpha$ vs Simplified	-0.55%	0.42%	-0.58%	1.71%	0.59%	0.16%
	TB	REF_ $\alpha$ vs Weighted	<b>4.84%</b>	0.00%	<b>-2.59%</b>	1.34%	0.35%	-0.49%
	EJ	REF_ $\alpha$ vs Weighted	<b>5.01%</b>	-0.42%	1.04%	<b>2.79%</b>	0.47%	-0.49%
GAE	All	$\alpha$ vs Simplified	0.49%	1.16%	-0.12%	-0.54%	0.76%	0.50%
	TB	REF_ $\alpha$ vs Weighted	<b>2.07%</b>	0.91%	-1.53%	<b>2.12%</b>	0.25%	0.17%
	EJ	REF_ $\alpha$ vs Weighted	<b>3.96%</b>	0.52%	0.81%	-1.64%	1.27%	0.00%

(see Section 5.2) shows, these datasets provided the most significant differences between  $\alpha$  based and weighted. When the models were tuned, they would have also improved performance compared to `RAFEN_Weighted` on these datasets.

In the case of `RAFEN_Weighted`, the temporal betweenness scoring function was barely better than Edge Jaccard, showing notable differences only in the case of GAE on ogbl-collab and Node2Vec on fb-forum. When comparing `RAFEN_ALL` and `RAFEN_Weighted`, `RAFEN_ALL` approach served as a strong baseline, showing better performance and higher rank in most of the experiments. We hypothesize that the scoring functions we leveraged led to decreased performance. Hence, it requires further investigation and refinements, which we leave for future work.

## 6 Conclusion and future work

In this paper, we tackled the problem of learning aligned node embedding for dynamic graphs modeled as a series of discrete graph snapshots. Contrary to existing approaches, such as post-hoc methods, we perform the alignment step during the embedding training. We proposed a novel framework – `RAFEN` – which allows to enrich any existing node embedding method with the aforementioned alignment capabilities. We conducted experiments on six real-world datasets and showed that our approach, even in its simplified version (with minimum set of hyperparameters), achieves better or on-par performance compared to existing approaches. For future work, we want to explore more advanced mechanisms for the actual alignment step, including better investigation of the scoring function in `RAFEN_Weighted` setting and the applicability of `RAFEN` for continuous-time dynamic graphs.

Our code with experiments and data, all enclosed in DVC pipelines [9], is available publicly at <https://github.com/graphml-lab-pwr/rafen>.

**Table 4.** Last snapshot link prediction evaluation results. We report AUC metric values with mean and standard deviation over 25 model retrains. We mark the top three results with **bold**. Methods mean ranks represent the ranking established upon mean AUC values over retrains.

Model	Alignment method	bitcoin alpha	bitcoin otc	fb forum	fb messages	ogbl collab	ppi	Mean rank
Node2Vec								
	(whole graph)	70.60 ± 5.80	71.70 ± 2.60	<b>91.40</b> ± 0.80	71.70 ± 4.70	<b>89.00</b> ± 0.10	<b>63.80</b> ± 0.20	<b>7.00</b>
	(last snapshot only)	72.70 ± 4.40	69.30 ± 2.70	80.00 ± 2.40	74.30 ± 3.90	84.50 ± 0.10	60.30 ± 0.40	10.42
	Posthoc-PA	74.60 ± 7.00	69.70 ± 2.20	84.40 ± 2.10	73.60 ± 6.50	84.50 ± 0.20	60.30 ± 0.30	10.00
	Posthoc-EJ	76.50 ± 7.50	69.60 ± 1.80	84.20 ± 2.10	73.30 ± 6.00	84.50 ± 0.20	60.30 ± 0.30	10.33
	Posthoc-TB	72.70 ± 7.60	70.20 ± 2.80	84.10 ± 1.90	74.00 ± 5.40	84.50 ± 0.20	60.30 ± 0.30	10.00
	RAFEN_ALL	73.10 ± 5.70	70.80 ± 2.50	<b>87.20</b> ± 1.70	74.60 ± 5.40	<b>84.80</b> ± 0.10	<b>61.90</b> ± 0.40	<b>5.75</b>
	RAFEN_Weighted_EJ	70.20 ± 7.20	71.10 ± 2.60	85.80 ± 1.80	73.20 ± 2.90	84.70 ± 0.20	<b>61.90</b> ± 0.40	9.17
	RAFEN_Weighted_TB	70.80 ± 8.80	71.40 ± 3.00	<b>87.00</b> ± 1.70	73.70 ± 3.80	<b>84.80</b> ± 0.10	61.90 ± 0.30	7.25
	(whole graph)	72.10 ± 4.30	<b>77.80</b> ± 1.70	68.20 ± 6.30	65.60 ± 4.50	80.70 ± 0.20	61.30 ± 0.60	10.00
	(last snapshot only)	77.80 ± 2.00	74.50 ± 1.90	79.60 ± 1.70	<b>75.80</b> ± 3.80	78.20 ± 0.30	58.00 ± 1.00	9.92
GAE								
	Posthoc-PA	<b>82.70</b> ± 2.80	74.40 ± 2.30	84.80 ± 1.70	<b>75.70</b> ± 3.80	78.20 ± 0.30	58.00 ± 0.80	8.42
	Posthoc-EJ	<b>82.60</b> ± 3.00	74.50 ± 2.20	84.90 ± 1.80	75.40 ± 3.30	78.20 ± 0.30	58.10 ± 0.70	7.92
	Posthoc-TB	<b>82.30</b> ± 3.50	74.80 ± 2.00	84.80 ± 1.60	<b>75.90</b> ± 3.50	78.20 ± 0.30	58.10 ± 0.60	7.58
	RAFEN_ALL	81.50 ± 2.60	<b>76.80</b> ± 1.40	86.60 ± 1.10	74.00 ± 3.50	78.50 ± 0.40	59.30 ± 0.50	<b>6.75</b>
	RAFEN_Weighted_EJ	80.10 ± 2.60	76.30 ± 1.50	86.10 ± 1.20	74.50 ± 2.60	78.00 ± 0.30	59.40 ± 0.60	8.00
	RAFEN_Weighted_TB	80.50 ± 3.30	<b>76.40</b> ± 1.90	86.30 ± 1.20	73.90 ± 2.80	78.50 ± 0.30	59.30 ± 0.50	7.50

**Acknowledgments** This work was financed by (1) the Polish Ministry of Education and Science, CLARIN-PL; (2) the European Regional Development Fund as a part of the 2014–2020 Smart Growth Operational Programme, CLARIN — Common Language Resources and Technology Infrastructure, project no. POIR.04.02.00-00C002/19; (3) the statutory funds of the Department of Artificial Intelligence, Wrocław University of Science and Technology, Poland; (4) Horizon Europe Framework Programme MSCA Staff Exchanges grant no. 101086321 (OMINO).

## References

1. Bielak, P., Tagowski, K., Falkiewicz, M., Kajdanowicz, T., Chawla, N.V.: Fildne: a framework for incremental learning of dynamic networks embeddings. *Knowledge-Based Systems* **236**, 107453 (2022)
2. Fey, M., Lenssen, J.E.: Fast graph representation learning with PyTorch Geometric. In: *ICLR Workshop on Representation Learning on Graphs and Manifolds* (2019)
3. Goyal, P., Chhetri, S.R., Canedo, A.: dyngraph2vec: Capturing network dynamics using dynamic graph representation learning. *Knowledge-Based Systems* **187**, 104816 (2020)
4. Grover, A., Leskovec, J.: node2vec: Scalable feature learning for networks. In: *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*. pp. 855–864 (2016)
5. Hu, W., Bansal, R., Cao, K., Rao, N., Subbian, K., Leskovec, J.: Learning backward compatible embeddings. In: *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery & Data Mining* (2022)
6. Hu, W., Fey, M., Zitnik, M., Dong, Y., Ren, H., Liu, B., Catasta, M., Leskovec, J.: Open graph benchmark: Datasets for machine learning on graphs. *arXiv preprint arXiv:2005.00687* (2020)
7. Kipf, T.N., Welling, M.: Variational graph auto-encoders. *arXiv preprint arXiv:1611.07308* (2016)
8. Kipf, T.N., Welling, M.: Semi-Supervised Classification with Graph Convolutional Networks. *arXiv* (Feb 2017). <https://doi.org/10.48550/arXiv.1609.02907>, <http://arxiv.org/abs/1609.02907>, arXiv:1609.02907 [cs, stat]
9. Kuprieiev, R., skshetry, Petrov, D., Redzynski, P., Rowlands, P., da Costa-Luis, C., Schepanovski, A., Shcheklein, I., Taskaya, B., Gao, Orpinel, J., de la Iglesia Castro, D., Santos, F., Sharma, A., Berenbaum, D., Zhanibek, Hodovic, D., daniele, Kodenko, N., Grigorev, A., Earl, Dash, N., Vyshnya, G., Lamy, R., maykulkarni, Hora, M., Vera, Mangal, S.: Dvc: Data version control - git for data & models (Sep 2022). <https://doi.org/10.5281/zenodo.7093084>, <https://doi.org/10.5281/zenodo.7093084>
10. Meng, Q., Zhang, C., Xu, X., Zhou, F.: Learning compatible embeddings. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 9939–9948 (2021)
11. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* **12**, 2825–2830 (2011)

12. Rossi, R., Ahmed, N.: The network data repository with interactive graph analytics and visualization. In: Twenty-Ninth AAAI Conference on Artificial Intelligence (2015)
13. Shen, Y., Xiong, Y., Xia, W., Soatto, S.: Towards backward-compatible representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6368–6377 (2020)
14. Singer, U., Guy, I., Radinsky, K.: Node embedding over temporal graphs. In: Proceedings of the 28th International Joint Conference on Artificial Intelligence. pp. 4605–4612. AAAI Press (2019)
15. Tagowski, K., Bielak, P., Kajdanowicz, T.: Embedding alignment methods in dynamic networks. In: Paszynski, M., Kranzlmüller, D., Krzhizhanovskaya, V.V., Dongarra, J.J., Sloot, P.M.A. (eds.) Computational Science – ICCS 2021. pp. 599–613. Springer International Publishing, Cham (2021)
16. Trivedi, P., Büyükçakır, A., Lin, Y., Qian, Y., Jin, D., Koutra, D.: On structural vs. proximity-based temporal node embeddings (2020)
17. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., Bengio, Y.: Graph attention networks. arXiv preprint arXiv:1710.10903 (2017)
18. Xu, K., Hu, W., Leskovec, J., Jegelka, S.: How Powerful are Graph Neural Networks? (Feb 2019), <http://arxiv.org/abs/1810.00826>, arXiv:1810.00826 [cs, stat]
19. Yu, E.Y., Fu, Y., Chen, X., Xie, M., Chen, D.B.: Identifying critical nodes in temporal networks by network embedding. *Scientific Reports* **10**(1), 1–8 (2020)