

A Comparative Accuracy Analysis of Node2Vec and FastRP

Ade Hodijah¹, Urip Teguh Setijohatmo², Sri Ratna Wulan³, Jonner Hutahaean⁴, Rahil Jumiyan⁵, Danu Mahesa⁶, Regi Purnama⁷

^{1,2,3,4,5,6,7}Department of Informatics and Computer Engineering, Bandung State Polytechnic, Indonesia

¹ade.hodijah@polban.ac.id, ²urip@jtk.polban.ac.id, ³sri.ratna@polban.ac.id, ⁴jonnerh@jtk.polban.ac.id,

⁵rahil.jumiyan@polban.ac.id, ⁶danu.mahesa1933@gmail.com, ⁷regiprnm432@gmail.com

Abstract

The rapid increase in scientific article publications has created a need for automated methods to support various applications, such as recommendation systems that help researchers discover relevant literature by measuring similarity between scientific publications using graph and text embeddings. This study mainly compares the accuracy of two graph embedding methods, Node2Vec and FastRP, where similarity is measured based on title, abstract, and citation attributes using a hybrid approach. In this approach, relational attributes (citations) are processed with graph embeddings, while textual attributes (titles and abstracts) are processed with text embeddings (S-BERT). The resulting vectors are then combined using a weighting scheme across attributes of title, abstract, and citation, with similarity values calculated using Cosine Similarity. The ordinal ranking results of these similarity scores are evaluated using Kendall Tau and Normalized Discounted Cumulative Gain (NDCG). The evaluation results indicate that Node2Vec achieves higher agreement with the similarity rankings generated by SPECTER than FastRP, as measured using NDCG and Kendall Tau. This study contributes to understanding the impact of graph structures in Node2Vec and FastRP on similarity scores derived from embeddings of title, abstract, and citation attributes. However, the current results are limited to a dataset consisting of only three attributes (title, abstract, and citation) from scientific publications, suggesting that further research must incorporate additional attributes (introduction, methods, results & discussion, and conclusion). The findings of this study are expected to provide practical guidance for researchers and developers of recommendation systems in selecting the most appropriate embedding method for their needs.

Keywords: Graph Embedding, Text Embedding, Node2Vec, FastRP, S-BERT.

1. Introduction

The rapid expansion of digital data has made similarity measurement technologies essential for many applications, including scientific publication retrieval, where the surge in article volume has complicated the discovery of relevant literature (Martín-Martín et al., 2021). Existing solutions often rely on citation counts (Sakib et al., 2020) or bibliographic coupling (Sjögårde & Ahlgren, 2024). Citation network analysis can reveal relationships and relevance among publications. Still, such approaches are biased toward highly cited works, prioritizing popular research while overlooking new or innovative studies with few citations (Sterling & Montemore, 2022). They are also vulnerable to citation cartel practices (Kojaku et al., 2021). Similarly, bibliographic coupling may connect documents that share references but differ in topical focus (Sjögårde & Ahlgren, 2024).

To overcome these limitations, embedding-based approaches represent attributes such as titles, abstracts, and citations as vectors (scientific paper

embeddings) using models like SPECTER (Cohan et al., 2020). Co-cited papers, for instance, often exhibit topical similarity reflected in their embedding proximity. DeepWalk, which treats random walks as a corpus, was later improved by Node2Vec to balance local and global structural exploration (Hiraoka et al., 2024). In Neo4j Graph Data Science (GDS), Node2Vec and FastRP are widely used algorithms: Node2Vec captures local and global structures, while FastRP prioritizes computational efficiency (Gupta et al., 2024).

Despite their popularity, comprehensive evaluations of Node2Vec and FastRP for multi-attribute scientific publication similarity remain limited. Furthermore, textual attributes (title, abstract) and relational attributes (citation) require different processing strategies. Sentence-transformer models such as S-BERT have been applied for automated journal recommendation (Colangelo et al., 2025), enabling semantic comparisons via cosine similarity and effectively capturing topical similarity across documents (Cohan et al., 2020). Meanwhile, graph-based approaches provide richer representations of citation re-

lations by modeling publications as interconnected nodes.

This study addresses these gaps by investigating two factors in order to offer clearer guidance on effective embedding methods for an alternative solution to capturing similarity among scientific publications.

1) The comparative accuracy of Node2Vec and FastRP in measuring similarity between scientific publications using Normalized Discounted Cumulative Gain (NDCG) (Järvelin & Kekäläinen, 2002) and Kendall Tau (Cabrera-Diego et al., 2019); and

2) The effect of incorporating text embeddings (title, abstract) on the accuracy of Node2Vec and FastRP.

2. Methods

This quantitative study uses an experimental approach to compare the effectiveness of graph embedding methods (Node2Vec, FastRP) in supporting similarity matching among scientific publication articles based on title, abstract, and paper's citation. The workflow as follows:

- Load the paper dataset containing the following attributes: paperId, title, abstract, embedding, citationIn, and citationOut.
- Construct a citation graph in Neo4j, where each paper is represented as a node and citation relationships (citationIn and citationOut) are represented as edges between nodes.
- Generate text embeddings from the title and abstract of each paper using the Sentence Transformer model (S-BERT). The resulting vectors are stored as the text_embedding property.
- Based on the constructed citation graph, generate graph embeddings using two different methods: Node2Vec and FastRP. The resulting embeddings are stored as embedding_node2vec and embedding_fastrp, respectively.
- Use the SPECTER embedding vectors already provided in the dataset as the reference embeddings, and store them in each paper node as ground_truth_embeddings.
- For both Node2Vec and FastRP, similarity scores were computed by combining graph embedding similarity and text embedding similarity using a weighted linear combination. Rather than fixing the weights at 50:50, the study evaluated multiple weighting schemes, where the graph embedding weight (w_1) varied from 0.0 to 1.0 in increments of 0.1 and the text embedding weight (w_2) was set to $1 - w_1$. This experiment was conducted to identify the weighting combination that produced the highest agreement with the SPECTER reference ranking. The weighting parameters were implemented

in the function `save_similarity_json` (emb_matrix, text_matrix, filename, w_1 , w_2), where w_1 denotes the graph embedding weight and w_2 denotes the text embedding weight.

- For each source paper, sort the similarity scores in descending order to produce a ranking of the most similar papers.
- Compare the rankings and similarity scores generated by Node2Vec and Text Embedding, FastRP and Text Embedding against those generated using the ground-truth embeddings.
- Evaluate the performance of each method using the NDCG and Kendall Tau metrics.

Node2Vec is a graph-based learning algorithm designed to generate low-dimensional vector representations (embeddings) of each node in a graph (Grover & Leskovec, 2016) by adapting the skip-gram approach from Word2Vec. Meanwhile, FastRP (Fast Random Projection) is a graph embedding algorithm designed to generate efficient node representations in large-scale graphs by combining the generation of a node similarity matrix and dimensionality reduction using very sparse random projection (Chen et al., 2019). In Node2Vec, the varying parameters include embedding dimension, walk length, in-out factor, and return factor. Meanwhile, FastRP parameters that varied include embedding dimension, iteration weights, and normalization strength.

The dataset of scientific publication articles used as the ground-truth consists of 57 scientific publications collected from a publicly available benchmark. However, the dataset was expanded from the initial 57 scientific publications to 9000 scientific publications to provide a larger evaluation sample and improve the robustness of the experimental results. The expanded dataset contains the required title, abstract, citation, and SPECTER embedding attributes. The larger number of publications enables the Node2Vec and FastRP methods to be evaluated on a more extensive citation graph and reduces the limitation associated with the relatively small sample size of the initial dataset. In addition, evaluate the performance of each method using four evaluation metrics: Normalized Discounted Cumulative Gain (NDCG), Kendall Tau, Mean Squared Error (MSE), and Mean Absolute Error (MAE). NDCG and Kendall Tau are used to evaluate the agreement and ranking consistency between the generated similarity rankings and the SPECTER reference ranking, while MSE and MAE are used to measure the numerical difference between the similarity scores generated by the proposed methods and the corresponding SPECTER similarity scores. Lower MSE and MAE values indicate smaller prediction errors, whereas higher NDCG and Kendall Tau values indicate better ranking agreement.

Table 1 presents an example of the metadata of scientific publication articles. The research methodology, as illustrated in Figure 1, outlines five stages of the study with two main focuses, with various attribute weighting of title and abstract have been evaluated based on the best NDCG and Kendall Tau scores to support the experiments on parameters combination of Node2Vec and FastRP methods. The following steps outline the process of determining the most effective embedding method (Node2Vec, FastRP) based on experiments with attribute variations (title, abstract, citation) and embedding types (text or graph) applied to the title and abstract attributes.




Figure 1. Research Methods

2.1 Experiment on Attribute Weighting (Title, Abstract, Citation)

This experiment examines the optimal weighting scheme for titles, abstracts, and citations, from weight=0 to weight=1. Accordingly, the weighting experiment replaces the initial equal-weight assumption by systematically evaluating all combinations of w_1 and w_2 before the parameter optimization experiments are conducted. Attribute weights are varied because each attribute contributes differently to the overall similarity of scientific publications. Adjusting these weights reveals the influence of each attribute on similarity measurement. Transforming the attributes into vector representations through embedding methods enables comparing their vectors with cosine similarity (Januzaj, Y. & Luma, A., 2022) to evaluate content similarity (Colangelo et al., 2025). The accuracy of among weight attributes (title, abstract, citation) is then assessed using the Kendall Tau (Cabrera-Diego et al., 2019) and NDCG (Järvelin & Kekäläinen, 2002) metrics. The Kendall Tau measures ranking consistency (ranking correlation) between the system-generated order and the ground-truth. At the same time, NDCG evaluates how effectively the system ranks relevant scientific publications at the top (top-k). To ensure the experiment remains focused on the influence of weight variations, the graph embedding algorithms (Node2Vec and FastRP) were run with the default parameter configurations from the Neo4j Graph Data Science library to maintain consistency.

The default configuration for Node2Vec includes Embedding Dimension=128, Walk Length=80, Return Factor=1.0, and In-Out Factor=1.0. As for FastRP, the parameters used are Embedding Dimension=128, Iteration Weights=[0.0,

Table 1. Metadata of Scientific Article Publications (Tran & Kim, 2024)

Attributes	Type	Example
Title	String	Design of 5G Architecture Enhancements for Supporting Serverless Computing.
Abstract	String	Serverless computing is a recent emerged trending technology in... and the edge resource consumption amount from 10 to 41%.
Citation	List of String	[Edge computing for Internet of Everything: A survey, A survey on approximate edge AI for energy efficient autonomous driving services, ..., Design of computing-aware traffic steering architecture for 5G mobile user plane]

Table 2. Node2Vec Parameters and Justification for Selection

Name	Range	Selection Justification
Embedding Dimension	[64, 128]	Determines the size (length) of the representation vector. Larger dimensions can capture more complex information but risk overfitting and impose higher computational costs. Dimension 128 was tested as the default/baseline, while dimension 64 was evaluated for efficiency in assessing richer representations.
Walk Length	[40, 80, 120]	Sets the length of the random walk to capture context. A value of 80 was tested as the default, 40 to focus on local structure (nearby neighbors), and 120 to examine whether broader exploration could better capture global context.
Return Factor (p)	[0.25, 0.5, 1.0, 2.0, 4.0]	Controls the probability of returning to the previous node. A value of $p < 1$ increases the likelihood of returning, encouraging highly local exploration. A value of $p > 1$ decreases the likelihood of returning, pushing the walk to move forward and explore new nodes.
In-out Factor (q)	[0.25, 0.5, 1.0, 2.0, 4.0]	Controls the balance between local exploration (similar to BFS) and global exploration (similar to DFS). A value of $q > 1$ prioritizes nearby neighbors (useful for detecting communities), whereas a value of $q < 1$ encourages traversal to more distant nodes (useful for learning structural roles).

1.0, 1.0], and Normalization Strength=0.0. The details of each parameter, along with the justification for its selection, are presented in Table 2 (Node2Vec) and Table 3 (FastRP).

Table 3. FastRP Parameters and Justification for Selection

Name	Range	Selection Justification
Iteration Weights	[1,0,0], [0,1,0], [0,0,1], [1,1,0], [1,0,1], [0,1,1], [1,1,1]	Determines the contribution of information from neighbors at different distances (hops). Variations were tested to examine the effect of exploration depth. For example, [1,0,0] focuses only on direct neighbors (1-hop), whereas [1,1,1] combines information from neighbors up to 3-hops away to produce a more comprehensive representation.
Normalization strength	[0.0, 0.25, 0.5, 0.75, 1.0]	Controls the strength of normalization applied to the embedding vectors. Tested to evaluate the impact of normalization on the stability and quality of the vectors, ranging from no normalization (0.0) to full normalization (1.0), in order to identify the optimal balance.

2.2 Experiment on Parameters Combination for Node2Vec and FastRP

This experiment uses different graph structures to identify the most effective embedding method for title and abstract attributes. After determining the optimal weighting scheme in Experiment 1, the second experiment evaluates different parameter combinations for Node2Vec and FastRP using the best-performing weights obtained for each method. The study adopts a combination of citation-based and text-based similarity (Yu et al., 2017), with the following scenarios.

- Textual attributes (title, abstract). These are first processed with text embedding (S-BERT), producing vector representations incorporated into graph embedding (Node2Vec, FastRP). The goal is to determine whether preprocessing with S-BERT provides more accurate embeddings than raw data within the graph embedding process.
- Relational attribute (citation). Citation data, consisting of references cited by a given publication, are processed using graph embedding methods (Node2Vec, FastRP).

3. Results and Discussion

There was a Paper node labels with CITED_BY (incoming paper's citation or citationIn) and CITES (outgoing paper's citation or citationOut) relationship types (Cohan et al., 2020) used in this research. The data model as seen in Figure 2 and the properties of the Paper Node as seen in Figure 3. The main properties (Figure 3) used in the experiment are text_embedding (title, abstract) and embedding_node2vec and embedding_fastp for citation (graph embedding) based on the relationship of CITED-BY and CITES (Figure 2).




Figure 2. Data Model of Paper




Figure 3. Properties of Paper

In the first experiment, the study determined the optimal weight settings (Table 4, Table 5) by repeating the experiment ten times to obtain the average values of Kendall Tau and NDCG. Based on the best results from this weight tuning, both Node2Vec and FastRP achieved its highest average performance with weights $w_1=0.3$ and $w_2=0.7$. The w_1 was for the citation attribute and w_2 was for the title and abstract attributes. These optimal weights were subsequently fixed and used throughout the parameter optimization experiments reported in Section 2.2.

Table 4. The Optimal Weighting Scheme of Title, Abstract, and Citation for FastRP

W1	W2	NDCG	Kendall Tau	Average
0	1	0.6061	0.2281	0.4171
0.1	0.9	0.5970	0.3130	0.4550
0.2	0.8	0.5923	0.3328	0.4626
0.3	0.7	0.5834	0.3697	0.4765
0.4	0.6	0.5716	0.3283	0.4500
0.5	0.5	0.5536	0.2739	0.4138
0.6	0.4	0.5242	0.2183	0.3712
0.7	0.3	0.4853	0.0648	0.2751
0.8	0.2	0.4308	0.0557	0.2433
0.9	0.1	0.3751	0.0240	0.1995
1	0	0.3395	0.0230	0.1812

Table 5. The Optimal Weighting Scheme of Title, Abstract, and Citation for Node2Vec

W1	W2	NDCG	Kendall Tau	Average
0	1	0.6061	0.2281	0.4171
0.1	0.9	0.5961	0.3314	0.4637
0.2	0.8	0.5950	0.3241	0.4595
0.3	0.7	0.5925	0.3098	0.4512
0.4	0.6	0.5891	0.2725	0.4308
0.5	0.5	0.5873	0.3023	0.4448
0.6	0.4	0.5726	0.1039	0.3383
0.7	0.3	0.5451	0.0614	0.3032
0.8	0.2	0.5023	0.0416	0.2720
0.9	0.1	0.4513	0.0241	0.2377
1	0	0.3923	0.0223	0.2073

For Node2Vec, the optimal configuration in-

Table 6. The Optimal Node2Vec Parameters (A: Embedding Dimension; B: Walk Length; C: Return Factor (p); D: In-Out Factor (q); E: NDCG; F: Kendall Tau; G: Average)

A	B	C	D	E	F	G
64	40	0.25	0.25	0.6038	0.6424	0.6231
64	40	0.25	0.50	0.6070	0.5963	0.6017
64	40	0.25	1.00	0.6103	0.3213	0.4658
...	...	...	...	...	...	...
128	80	0.50	1.00	0.6126	0.3308	0.4717
128	80	0.50	2.00	0.6219	0.6293	0.6256
...	...	...	...	...	...	...
128	120	4.00	2.00	0.6085	0.2744	0.4415
128	120	4.00	4.00	0.6034	0.4098	0.5066
64	40	0.25	0.25	0.6038	0.6424	0.6231
64	40	0.25	0.50	0.6070	0.5963	0.6017

cluded an embedding dimension of 128, a walk length of 80, an in-out factor ($q = 2$), and a return factor ($p = 0.5$), as seen in Table 6 based on the optimal weight scheme ($w_1=0.1$ and $w_2=0.9$) with average: 0.625648 from NDCG's average: 0.621908 and Kendall's average: 0.629388. Meanwhile, FastRP (Table 7) demonstrated the best configuration at a dimension of 64 with iteration weights $[0.0, 1.0, 1.0]$ and a normalization strength of 0.75 based on the optimal weight scheme ($w_1=0.3$ and $w_2=0.7$) with average: 0.569895 from NDCG's average: 0.585709 and Kendall's average: 0.554082.

The Embedding Dimension parameter determines the size of the vector space for each node. Theoretically, higher dimensions are capable of storing more complex information. However, experiments show that increasing the dimension does not always lead to significant performance improvement. The difference is negligible between dimension 64 (Node2Vec = 0.603896, FastRP = 0.595969) and dimension 128 (Node2Vec = 0.603495, FastRP = 0.594287). This indicates that dimension 64 is more than sufficient to capture the essential structural information of the article publication graph. Increasing the dimension to 128 does not provide any substantial benefits; instead, it raises computational costs. Moreover, excessively high dimensions risk overfitting, where the model captures too many irrelevant patterns without corresponding gains in performance. Therefore, higher dimensions do not necessarily yield better results.

The optimal average (0.625648) of Node2Vec was achieved for the Walk Length parameter at a value of 80, indicating that a step length of 80 strikes a balance between exploring the graph structure and maintaining computational efficiency. In addition, the In-Out Factor (q) and Return Factor (p) demonstrated a significant influence on performance. The configuration with $q = 2.00$ and $p = 0.50$ emerged as the most optimal combination. This pattern reflects a traversal strategy that emphasizes local exploration without

frequent returns to previously visited nodes. It further reinforces earlier findings that local homophily (e.g., the connectivity among article publications with similar titles, abstracts, or citations) is a crucial element in the underlying data structure.

Table 7. The Optimal FastRP Parameters (A: Embedding Dimension; B: Iteration Weight; C: Normalization Strength; D: NDCG; E: Kendall Tau; F: Average)

A	B	C	D	E	F
64	[1.0, 0.0, 0.0]	0.00	0.5761	0.1591	0.3676
64	[1.0, 0.0, 0.0]	0.25	0.5716	0.4809	0.5262
64	[1.0, 0.0, 0.0]	0.50	0.5959	0.1497	0.3728
...	...	...	...	...	...
64	[0.0, 1.0, 0.0]	0.25	0.5879	0.1370	0.3625
64	[0.0, 1.0, 1.0]	0.50	0.5954	0.4214	0.5084
64	[0.0, 1.0, 1.0]	0.75	0.5857	0.5540	0.5698
...	...	...	...	...	...
128	[1.0, 1.0, 1.0]	1.00	0.5942	0.3709	0.4826

The Iteration Weights parameter plays a critical role in determining how much information from neighbors at different distances (hops) contributes to the final embedding. This experiment tested several weighting schemes, including:

- [1.0, 0.0, 0.0]: relying solely on the node itself (0-hop).
- [0.0, 1.0, 0.0]: considering only direct neighbors (1-hop).
- [0.0, 0.0, 1.0]: focusing only on second-degree neighbors (2-hop).
- [1.0, 1.0, 0.0]: combining the node itself and direct neighbors (0-hop and 1-hop).
- [1.0, 0.0, 1.0]: combining the node itself and second-degree neighbors (0-hop and 2-hop).
- [0.0, 1.0, 1.0]: considering both direct and second-degree neighbors (1-hop and 2-hop).
- [1.0, 1.0, 1.0]: incorporating all layers equally: self, 1-hop, and 2-hop.

Based on the evaluation results in Table 7, the iteration weight configuration $[0.0, 1.0, 1.0]$, which utilizes information from both 1-hop and 2-hop neighbors, achieved the best performance with the highest average score of 0.569895. This finding demonstrates that multi-hop combinations are necessary to capture inter-entity relationship patterns in the graph. When

relying only on 1-hop information [0.0, 1.0, 0.0], performance declined considerably (0.362535). Similarly, depending solely on the node itself ([1.0, 0.0, 0.0]) with Normalization Strength = 0.50 produced the lowest average (0.367658). However, at Normalization Strength = 0.25, the average score increased significantly (0.526278), though the Kendall's average (0.480952) remained lower than that at Normalization Strength = 0.75 (the optimal FastRP parameters, 0.554082). A high Kendall Tau score indicates that the ranking generated by the similarity measure closely resembles the actual ranking (SPECTER reference ranking), showing that the similarity measure effectively captures semantic similarity.

Normalization Strength is a parameter used to regulate how much structural information—particularly the degree or number of neighbors—affects the vector representation (embedding) of nodes in the graph. A higher Normalization Strength amplifies the normalization effect, making vectors less dominated by highly connected nodes. Conversely, a lower value preserves the original scale of connectivity, including highly connected nodes. In this study, Normalization Strength = 0.75 yielded the optimal FastRP parameters. This suggests that normalization improved the quality of the embeddings. One plausible explanation is that the embedding distribution in the graph was not naturally balanced (e.g., varying citation counts across papers), and normalization provided additional benefits.

However, when Normalization Strength = 1.0, the average NDCG increased (0.594287), but the average Kendall Tau decreased (0.370975), causing the overall average to drop (0.482631). This may be attributed to the relatively small size and moderate complexity of the graph, where scaling effects across vectors were less critical, rendering normalization less impactful. These findings highlight that in certain cases—particularly in relatively homogeneous graphs—excessive normalization can hinder performance rather than enhance it. Although normalization did not improve performance in this dataset, the experiments were still necessary to provide empirical evidence, rather than assumptions, that this feature was not required for article publication graph data.

In summary, these findings confirm two important insights. First, for this dataset, the raw connectivity information of the graph is more valuable than its normalized version. Second, they show that no single parameter setting is universally optimal; its effectiveness depends heavily on the unique structural characteristics of the data being analyzed.

The results have shown Node2vec keeps traditionally better precision than FastRP where on average Node2vec got 0.625648 and FastRP 0.569895. As this work first objective to find how far influence of paper citation to just text similarity based on title and abstract, take a look at the Table 4 that says the varies portion w_1 and w_2 have emerged to at least two find-

ings:

- Title and abstract (w_2) were more important than citation (w_1), for each their solely role represented in the first row and the last row of the table. Surprisingly, according to (Cohan et al., 2020), SPECTER did not agree with this work. Just using solely title plus abstract produce low precision although it prepared with pretrained SciBert. The table shows with $w_1=0$ and $w_2=1$ which tells using solely title plus abstract got NDCG 0.606170 in contrast to $w_1=1$ and $w_2=0$ which says using citation solely got NDCG 0.339508.
- The combination of the two factors got the best result. We then have three choices to get the best value depend on what metric we choose: Average NDCG, Average Kendall Tau, or Average of the previous two. This work chose the best Average NDCG and Average Kendall Tau mentioned as highlighted row which compose of the 0.3 portion of w_1 and 0.7 of w_2 , got:

- FastRP with average of NDCG: 0.583468, average of Kendall Tau: 0.369728, and average: 0.476598.

- Node2vec, average of NDCG: 0.592583, average of Kendall Tau: 0.309864, and average: 0.451223.

From another overall objective, the results have shown Node2Vec keeps traditionally better precision than FastRP where on average Node2Vec got 0.625648 and FastRP 0.569895. This means the semantic context will get better by appropriate walk given the characteristic of problem domain in the form of dataset. It seems to us that graph structure of the dataset is less uniform indicated by facts where papers have sharply fluctuative referencing relation to other papers, in which random walk is more appropriate than projection. Another factor, also from the characteristic of dataset, tells us the majority citation is not primed subject of the source paper so the text embedding plays the major role compare to citation one.

To further examine the contribution of textual information, an additional comparison was conducted between the best-performing hybrid configuration, which combines citation-based graph embedding with title and abstract text embedding, and a citation-only configuration using 9000 data of scientific publications. In the hybrid configuration, the citation weight was 0.3 and the title-and-abstract weight was 0.7, as seen in Table 8. In the citation-only configuration, the citation weight was 1.0 and the text weight was 0.0, as seen in Table 9. The comparison was performed for both Node2Vec and FastRP using NDCG, Kendall Tau, MSE, and MAE.

Table 8 shows that Combining citation-based graph embedding with text embedding using the optimal weighting scheme (0.3 for citation and 0.7 for

Table 8. The Hybrid Optimal Weighting Scheme of Title, Abstract, and Citation for Node2Vec using 9000 data

Method	NDCG	Kendall Tau	MSE	MAE
Node2Vec	0.9562	0.469	0.00378	0.04907
FastRP	0.9105	0.412	0.01522	0.09843

title and abstract) improves recommendation performance compared with using graph embedding alone. Table 9 shows the result that Node2Vec were still better than FastRP when the experiment only used citation without text embeddings of title and abstract.

Table 9. The Citation Only Experimental Configuration Without Title and Abstract for Node2Vec and FastRP Using 9000 Data

Method	NDCG	Kendall Tau	MSE	MAE
Node2Vec	0.9075	0.141	0.00552	0.05912
FastRP	0.8123	0.113	0.04077	0.18544

For Node2Vec, the improvement in Kendall Tau was particularly pronounced, increasing by 0.3283 points (from 0.141 to 0.469). This result suggests that the textual information from titles and abstracts contributes substantially to ranking consistency when combined with citation-based graph representations. However, these findings should be interpreted as results for the evaluated dataset and weighting configurations rather than as evidence of universal superiority.

The comparison between the hybrid (Table 8) and citation-only configurations (Table 9) further indicates the contribution of title and abstract text embeddings. For Node2Vec, adding text information increased NDCG from 0.9075 to 0.9562 and Kendall Tau from 0.141 to 0.4693, while reducing MSE from 0.00552 to 0.00378 and MAE from 0.05912 to 0.04907. For FastRP, the corresponding NDCG increased from 0.8123 to 0.9105 and Kendall Tau from 0.113 to 0.412, while MSE decreased from 0.04077 to 0.01522 and MAE from 0.18544 to 0.09843. Thus, within these experiments, combining citation-based graph embeddings with title and abstract text embeddings produced better agreement than using citation information alone. These results indicate that Node2Vec generated rankings that were more consistent with the SPECTER reference ranking while also producing smaller prediction errors in this comparison.

4. Conclusion

This study found that Node2Vec was better than FastRP when the optimal parameters applied, but FastRP had the best result at the first experiment (weighting scheme), where FastRP (0.369728) demonstrated more consistent similar to SPECTER ranking based on the Kendall Tau score than Node2Vec, which only reached 0.331451. The best performance was ob-

tained using an appropriate a weighting combination of text embeddings (title, abstract) and citation attributes (citation-based graph embeddings). The experimental results indicate that the text embeddings derived from the title and abstract contributed more to similarity measurement than citation-based graph embeddings. Therefore, further analysis is required to investigate the contribution of other textual attributes in different sections of a scientific article, such as the introduction, methods, and conclusion, which share characteristics similar to those of the abstract and title. However, this study has several limitations. First, the experiments were conducted on only 57 scientific publications for the optimal scheme of citation weight was 0.3 and the title-and-abstract weight was 0.7 and 9000 data for the optimal scheme and the citation weight was 1.0 and the text weight was 0.0, which limits the generalizability of the findings. Future studies should evaluate larger citation graphs covering multiple research domains to validate whether the observed trends remain consistent. Second, although Node2Vec and FastRP can directly utilize raw title and abstract text as node properties, this scenario was not evaluated in the present study and should be investigated in future work. Finally, computational efficiency was not evaluated in this study. Future work should compare runtime because FastRP is primarily designed for scalable graph embedding.

References

- Cabrera-Diego, L. A., El-Bèze, M., Torres-Moreno, J.-M., & Durette, B. (2019): *Ranking résumés automatically using only résumés: A method free of job offers*, Expert Systems with Applications, 123, 91–107. <https://doi.org/10.1016/j.eswa.2018.12.054>
- Chen, H., Sultan, S. F., Tian, Y., Chen, M., & Skiena, S. (2019): *Fast and accurate network embeddings via very sparse random projection*, In Proceedings of the 28th ACM International Conference on Information and Knowledge Management (pp. 399–408). Association for Computing Machinery. <https://doi.org/10.1145/3357384.3357879>
- Cohan, A., Feldman, S., Beltagy, I., Downey, D., & Weld, D. S. (2020): *SPECTER: Document-level representation learning using citation-informed transformers*, In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (pp. 2270–2282). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.acl-main.207>
- Colangelo, M. T., Meleti, M., Guizzardi, S., Calciolari, E., & Galli, C. (2025): *A comparative analysis of sentence transformer models for automated journal recommendation using PubMed metadata*, Biology and Life Sciences. <https://doi.org/10.20944/preprints202501.1334.v1>

- Grover, A., & Leskovec, J. (2016): *node2vec: Scalable feature learning for networks*, In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 855–864). Association for Computing Machinery. <https://doi.org/10.1145/2939672.2939754>
- Gupta, R., Krishnamoorthy, M., & Vangala, V. (2024): *Better graph embeddings for enterprise graphs*, In Proceedings of the 7th Joint International Conference on Data Science & Management of Data (11th ACM IKDD CODS and 29th COMAD) (pp. 368–374). Association for Computing Machinery. <https://doi.org/10.1145/3632410.3632412>
- Hiraoka, Y., Imoto, Y., Lacombe, T., Meehan, K., & Yachimura, T. (2024): *Topological Node2vec: Enhanced graph embedding via persistent homology*, Journal of Machine Learning Research, 25, 1–26. <https://jmlr.org/papers/volume25/23-1185/23-1185.pdf>
- Januzaj, Y., & Luma, A. (2022): *Cosine similarity – A computing approach to match similarity between higher education programs and job market demands based on maximum number of common words*, International Journal of Emerging Technologies in Learning (iJET), 17(12), 258–268. <https://doi.org/10.3991/ijet.v17i12.30375>
- Järvelin, K., & Kekäläinen, J. (2002): *Cumulated gain-based evaluation of IR techniques*, ACM Transactions on Information Systems, 20(4), 422–446. <https://doi.org/10.1145/582415.582418>
- Kojaku, S., Livan, G., & Masuda, N. (2021): *Detecting anomalous citation groups in journal networks*, Scientific Reports, 11(1), 14524. <https://doi.org/10.1038/s41598-021-93572-3>
- Martín-Martín, A., Thelwall, M., Orduna-Malea, E., & Delgado López-Cózar, E. (2021): *Google Scholar, Microsoft Academic, Scopus, Dimensions, Web of Science, and OpenCitations' COCI: A multidisciplinary comparison of coverage via citations*, Scientometrics, 126(1), 871–906. <https://doi.org/10.1007/s11192-020-03690-4>
- Sakib, N., Ahmad, R. B., & Haruna, K. (2020): *A collaborative approach toward scientific paper recommendation using citation context*, IEEE Access, 8, 51246–51255. <https://doi.org/10.1109/ACCESS.2020.2980589>
- Sjögårde, P., & Ahlgren, P. (2024): *Seed-based information retrieval in networks of research publications: Evaluation of direct citations, bibliographic coupling, co-citations, and PUBMED-related article score*, Journal of the Association for Information Science and Technology, 75(13), 1453–1465. <https://doi.org/10.1002/asi.24951>
- Sterling, J. A., & Montemore, M. M. (2022): *Combining citation network information and text similarity for research article recommender systems*, IEEE Access, 10, 16–23. <https://doi.org/10.1109/ACCESS.2021.3137960>
- Tran, M.-N., & Kim, Y. (2024): *Design of 5G architecture enhancements for supporting serverless computing*, IEEE Access, 12, 163382–163395. <https://doi.org/10.1109/ACCESS.2024.3490671>
- Yu, D., Wang, W., Zhang, S., Zhang, W., & Liu, R. (2017): *Hybrid self-optimized clustering model based on citation links and textual features to detect research topics*, PLOS ONE, 12(10), e0187164. <https://doi.org/10.1371/journal.pone.0187164>