Appraisal of Citation Reliability Using a Gan-Based Approach

Dvora Toledano Kitai[a], Renata Avros[b], Ilya Lev[c], Biran Fridman*[d] and Zeev Volkovich*[e]
Software Engineering Department, Braude College of Engineering, Snonit st., Karmiel, Israel

Keywords: Citation Manipulation, Network Perturbations, Complex Network Analysis, Link Prediction, Generative, Adversarial Network.

Abstract: This paper addresses the pressing issue of citation manipulation in academic publications. Traditional detection methods, which rely on expert manual review, struggle to keep pace with the ever-growing volume of research output. To overcome these limitations, this study introduces an automated, network-based approach for identifying unreliable citations using an Encoder-Decoder model. By learning regular citation patterns, the model detects anomalies through reconstruction errors. Citation reliability is assessed by systematically removing edges from a citation network and predicting their reinstatement using a modified GAN-based framework. Successful predictions indicate legitimate citations, while failures suggest potential manipulation. The proposed methodology is validated on the CORA dataset, demonstrating its effectiveness in distinguishing genuine references from manipulated ones. This approach provides a scalable and data-driven solution for enhancing research integrity and mitigating citation distortions in scholarly literature.

1 INTRODUCTION

In the academic world, scholarly publications are essential tools for advancing knowledge and fostering research development via appropriate citations and interactions, commonly used as an indicator of scientific career development. Indeed, such widespread practices inevitably lead to efforts to affect the citation process, encouraging potential authors to include unnecessary or only loosely relevant references to inflate the perceived importance of their work. A type of unethical practice in academic publishing is citation manipulation, where authors, editors, or journals deliberately alter citation behavior to artificially increase citation counts. Another kind appears when citation cartels, groups of authors, or journals collaborate to cite each other's work excessively, and coercive citations. Often, editors or reviewers request authors to add citations as a condition for publication. Additionally, reference padding involves adding unnecessary citations to reference lists without engaging with the cited work. Such manipulations aim to improve academic metrics such as the h-index, impact factor, or perceived influence of specific authors or journals.

Including irrelevant citations negatively impacts the quality and relevance of academic papers, undermining both scholarly integrity and the reliability of scientific literature, which are crucial for advancing research and knowledge. Consequently, it is essential for academic institutions and publishers to actively address this concern by promoting ethical citation practices and offering clear guidelines to authors on responsible citation.

The exploration of citation patterns is a subject of extensive academic study. The pioneering research (Garfield, 1979) established the foundation for understanding citation practices across different fields. Subsequent studies (Wang and White,1996), (Case and Higgins, 2000]), and (Bornmann and Daniel, 2008) delved into the underlying reasons for

[a] https://orcid.org/0009-0002-1923-3640
[b] https://orcid.org/0000-0001-9528-0636
[c] https://orcid.org/0009-0002-5222-8077
[d] https://orcid.org/0009-0007-3118-4980
[e] https://orcid.org/0000-0003-4636-9762
* Corresponding author

Kitai, D. T., Avros, R., Lev, I., Fridman, B., Volkovich and Z.
Appraisal of Citation Reliability Using a Gan-Based Approach.
DOI: 10.5220/0013586500003967
In *Proceedings of the 14th International Conference on Data Science, Technology and Applications (DATA 2025)*, pages 778-785
ISBN: 978-989-758-758-0; ISSN: 2184-285X

citations, shedding light on both scientific and non-scientific influences. Generally, these inquiries underscore the intricate nature of citation behavior and its crucial role in evaluating scholarly output. The research (Mammola, Piano, Doretto, Caprio, and Chamberlain, 2022) emphasizes that while scholarly content should be the primary basis for citing, other elements such as the length of the paper, the number of authors, their collaborative networks, and individual characteristics can also influence citation behaviors.

The paper (Prabha, 1983) suggests that more than two-thirds of references in academic papers are unnecessary, highlighting the prevalent issue of questionable citations. The research presented by (Wilhite and Fong, 2012), as well as by (Wren and Georgescu, 2022), has delved into various aspects of reference list manipulation, uncovering practices such as coercive citation and unusual referencing patterns as departures from established standards.

Traditional methods for identifying citation manipulation involve experts carefully examining citation patterns in scholarly articles. This process entails assessing the relevance and context of citations, detecting potential biases or inconsistencies, and exploring the relationships between cited and citing works. While manual review can provide valuable insights by leveraging the expertise of subject matter specialists, it is labor-intensive and challenging to implement on a large scale. With the increasing volume of academic publications, the shortcomings of manual detection methods have become increasingly evident. As a result, automated approaches have been developed to improve efficiency and consistency in identifying citation manipulation.

Several studies highlight the utility of network analysis in detecting citation manipulation. Research (Ding, Y., 2011) explores the connection between collaboration and citation patterns, while (Liu, J., Bai, X., Wang, M., Tuarob, S., & Xia, F, 2024) introduces ACTION, a framework for identifying anomalous citations in heterogeneous networks. A study [Isfandyari-Moghaddam, A., Saberi, M. K., Tahmasebi-Limoni, S., Mohammadian, S., & Naderbeigi, F., 2023) examines co-authorship networks among leading research nations.

Studies (Avros, Haim, Madar, Ravve, and Volkovich, 2023) and (Avros, Keshet, Kitai, Vexler, and Volkovich, 2023) have investigated the automation of detecting manipulated citations in academic papers using advanced graph-based techniques. These considerations have constructed robust frameworks that scrutinize citation networks' structural and contextual relationships by employing self-learning graph transformers, perturbation methods, and Graph embeddings.

The current paper addresses the challenge of assessing the reliability and consistency of citations within a citation network. Following the general standpoint outlined in the mentioned works, the aim is to investigate the stability of ideal ("genie") references under network distortions. This core problem can be reframed in the context of anomaly detection using an Encoder-Decoder model. Specifically, the methodology leverages the model's ability to learn the underlying structure of normal (i.e., consistent and reliable) citation patterns. Trained solely based on these normal citation examples, the model learns a compressed latent representation that facilitates an accurate reconstruction of such citations. While the model succeeds at reconstructing normal citation data with minimal error, it struggles with anomalous citations that are unreliable or inconsistent and thus deviate from the learned patterns. Critically, the difference between the original citation data and its reconstructed version, the reconstruction error, serves as the primary metric for identifying these anomalous citations.

The process presented in this study is inspired by the work outlined by (Jin, Xu, Cheng, Liu, and Wu, 2022). This paper addresses the limitations of traditional link prediction methods by proposing a novel approach utilizing Generative Adversarial Networks (GANs). The suggested method organizes the network into hierarchical layers, preserving local and global structural features. A GAN is employed to iteratively learn low-dimensional vector representations of vertices at each layer, using these representations to initialize the previous layer.

In our study, we utilize a modified version of this method. We randomly remove a fixed fraction of citations (edges) from the network through multiple trials. The described GAN-based approach is then employed to predict the missing citations, comparing them with the omitted ones. The reconstruction rate calculated within the trials indicates the reliability of the corresponding edges. So, successful predictions indicate the likely importance of the citation, while failed predictions suggest potential irrelevance or inclusion for non-scholarly reasons.

The subsequent sections of the paper are dedicated to presenting the necessary background concepts, describing the proposed model, and reporting numerical results. At this stage, we aim to validate the proposed model using just a single dataset, with plans to extend the study and evaluate its reliability across additional datasets in future work. Section 2 provides the mathematical foundations underlying the proposed approach. Section 3 details the proposed methodology and outlines the GAN-based framework for citation relevance prediction. Section 4 presents numerical results, demonstrating the model's effectiveness on the well-known CORA dataset. Section 5 is devoted to a conclusion.

2 PRELIMINARIES

The mathematical models forming the algorithmic framework of this research are discussed in this Section.

2.1 EmbedGAN: Embedding Generation with Generative Adversarial Networks

Generative Adversarial Networks (GANs) (see, e.g., Goodfellow, 2014) employ two competing neural networks: a Generator and a Discriminator. Acting as a data creator, the Generator produces samples meant to imitate actual data; conversely, the Discriminator assesses the authenticity of given samples. Through adversarial training, the Generator continuously improves its ability to generate realistic data while the Discriminator refines its ability to distinguish between authentic and artificial examples. This iterative process continues until the generator produces synthetic data practically indistinguishable from the actual dataset.

EmbedGAN (Zhao, Zhang and Zhang, 2021) is an innovative approach to network embedding by leveraging a GAN to generate high-quality node representations. At its core, following the general agenda, EmbedGAN utilizes an adversarial training process involving two mentioned neural networks: a Generator and a Discriminator. The Generator aims to create synthetic network embeddings that resemble real network-derived embeddings, while the Discriminator is trained to distinguish between authentic and generated embeddings.

A key component of EmbedGAN is its Builder Sampling Strategy, which optimizes the selection of training samples to enhance adversarial learning. Rather than exhaustively considering all node pairs, this strategy purposefully chooses samples that effectively capture structural and semantic relationships within the network. A particularly important aspect is hard negative sampling, which incorporates structurally similar but unconnected nodes to challenge the Discriminator, thereby improving its ability to distinguish between real and generated embeddings. Furthermore, hierarchical sampling ensures the preservation of local and global network structures in the resulting embeddings. An adaptive selection process prioritizes more challenging samples as training advances, leading to higher-quality embedding while reducing computational costs.

EmbedGAN employs a crucial two-stage Embedding Assignment and Refinement process to ensure its node representations accurately reflect the network's structural and relational properties. Initially, nodes are mapped to a lower-dimensional latent space, capturing local and global network characteristics. Subsequently, these initial embeddings undergo iterative refinement through adversarial GAN training. Training occurs over multiple epochs, regularly incorporating varying batch sizes to optimize convergence. K-Fold cross-validation is employed to enhance model generalization and prevent overfitting.

In addition to EmbedGAN, another graph embedding technique applied in the research is the famous Node2Vec approach (Grover and Leskovec, 2016), being a network embedding algorithm designed to generate vector representations of nodes while preserving structural properties. It achieves this by performing biased random walks guided by parameters p and q, which controls the balance between local and global exploration. Node sequences generated by these walks are fed into the Word2Vec algorithm, which learns informative embeddings. The ability to tune p and q allows Node2Vec to capture different structural aspects of the network, making it a flexible and scalable approach. Integrating biased random walks with the famous Word2Vec (Mikolov, Chen, Corrado, and Dean, 2013).

2.2 NetLay: Hierarchical Graph Representation for Link Prediction

NetLay (Jin, Xu, Cheng, Liu and Wu, 2022) is a hierarchical graph representation learning method designed to improve link prediction by capturing local and global network structures. Unlike traditional approaches that rely only on instant neighbors, NetLay constructs a multi-scale hierarchical representation, grouping nodes based on structural roles such as community membership or core-periphery relationships. This hierarchy provides deeper insights into network connectivity, enhancing prediction accuracy. The method involves several key components.

- *Graph coarsening* or *clustering* organizes nodes into progressively larger groups, forming a hierarchical structure.
- *Neighborhood aggregation* integrates information from different hierarchy levels using weighted aggregation or attention mechanisms.
- *Embedding learning* refines node representations at each level through graph neural networks (Node2Vec in our study). These embeddings are then used to compute link probabilities based on similarity measures, such as cosine similarity.

By incorporating hierarchical information, NetLay can identify connections extending beyond instant neighbourhoods, capture long-range dependencies, and uncover hidden relationships within the network, leading to more accurate link prediction than methods focusing solely on local structures.

3 APPROACH

This section introduces the proposed methodology, which aims to address the issue of irrelevant citations in scholarly articles through a GAN-based algorithm. The process involves constructing a citation graph from a given dataset, predicting citation relevance, and subsequently identifying potentially relevant or irrelevant citations.

As previously discussed, the methodology employed assesses the reliability of citations by examining their behavior under network perturbation. This assessment is performed by systematically and randomly removing edges from the citation network. Following each removal, the restoration of these connections is analyzed. Observing and quantifying the recovery of these links gives insights into the stability and importance of individual citations within the network's structure. A consistently and easily re-established citation after perturbation suggests a crucial structural role within the network, indicating its robustness and significance. Conversely, a citation that fails to reappear after removal implies a weaker or less vital connection, potentially signifying a less essential role in maintaining the network's integrity.

3.1 Initialization of Network Perturbation Sequential Process

- Parameters:
  - *N*: Number of iterations
  - *Fr*: Fraction of randomly omitted edges at each iteration
- Graph Loading:
  - Load the graph $G = \langle V, E \rangle$: V (nodes), E (edges)
- Initialization of an indication array:
  - Create a zero-filled Z array of size $|V|$

3.2 Network Perturbation Sequential Process

For each current iteration k within N iterations a modified graph $G_0^{(k)} = (V, E_k)$ is constructed by randomly removing a fraction of Fr edges in the edges E of the source graph.

3.2.1 Transformation of the Modified Graph into a Weighted Citation Network

Edge weights are determined based on both citation links and content resemblance, computed using cosine similarity between feature word vectors papers reduced to a manageable size using PCA.

3.2.2 MNL-Modified NetLay Algorithm

The suggested modified NetLay Algorithm (MNL) simplifies the complex citation network by recursively generating hierarchical coarsened graphs

$$\left(G_0^{(k)}, G_1^{(k)}, \dots, G_n^{(k)}\right).$$

MNL enhances the original NetLay method by incorporating the Infomap approach [32] to identify communities of densely connected nodes, thereby improving the graph coarsening process. This technique merges nodes within the same community into super nodes, effectively reducing graph complexity while maintaining essential connectivity patterns. To facilitate analysis, edge weights are normalized to a fixed range of [0,1]. Additionally, feature vectors for super nodes are determined by averaging the attributes of their constituent nodes. At the final stage of coarsening, $G_n^{(k)}$ represents the most simplified yet structurally representative version of the considered modified graph $G_0^{(k)}$.

3.2.3 Node Embeddings via a Hierarchical Graph Networks

This phase aims to generate informative node embeddings by leveraging hierarchical graph structures. The process begins with the Node2Vec procedure applied to the most refined hierarchical layer obtained in the previous step. Following the initialization at $G_n^{(k)}$ a recursive embedding refinement process is performed, propagating embeddings back through the hierarchical layers to the original graph $G_0^{(k)}$. In each intermediate layer $G_i^{(k)}$, $0 \leq i \leq (n-1)$ embeddings are adjusted by introducing a controlled noise factor α. For each node $V_i^{(k)}$ within a super node, an updated embedding is computed as

$$V_i^{new} = V_i^{(k)} + \alpha \cdot FV_i,$$

where FV_i is the node's feature vector derived from the preprocessing phase. Finally, the concluding embeddings at $G_0^{(k)}$ integrate hierarchical information from all preceding layers, providing a rich and context-aware representation of the citation network.

3.2.4 EmbedGAN

The approach begins with a pre-training phase. This phase is crucial in preparing the GAN model's Generator and Discriminator components. The pre-training of the generator starts with random noise as an input to train it to produce embeddings that resemble the simplest graph layer edges, as obtained from Node2Vec. Meanwhile, the Discriminator is trained to distinguish between real embeddings derived from actual edges in the graph and fake embeddings generated by the model.

For each fold in the hierarchical graph layer pyramid, positive and negative examples are defined as follows: positive examples correspond to actual edges, whereas negative examples are artificially generated by random walking among the graph nodes. The length of this random walk is set based on the average size of the strongly connected components within the graph.

The generated positive samples consist of two main groups:

- Existing edges that are already present in the graph.
- Connections formed between a randomly generated walk's first and last node.

On the other hand, negative (fake) samples are randomly generated edges that do not exist in the graph. If a randomly generated edge coincides with a positive edge, the process is repeated until a truly negative edge is obtained.

The training is structured as a recursive process across multiple hierarchical graph layers, progressively learning structural patterns from simplified network representations to the full citation network. A two-loop training procedure is employed: an outer loop using K-Fold cross-validation to enhance model generalization and an inner loop performing iterative training through adversarial learning. Genuine and synthetic edge embeddings are evaluated at each step, and model performance is assessed using precision, recall, and F1-score metrics. The final output consists of an optimized Generator and Discriminator, capable of accurately predicting citation relevance based on learned network structures.

3.2.5 Final Stage: Link Prediction

At this stage, the trained Discriminator model is utilized to evaluate the removed edges from the current iteration, assigning each a prediction score ranging from 0 to 1, with higher scores indicating a more substantial likelihood of scholarly significance. A classification threshold (commonly 0.5) is applied, categorizing edges as relevant or likely irrelevant

3.3 Process Summarization

The iterative computation of reconstruction rates yields a distribution that functions as a proxy measure for edge reliability. Specifically, diminished reconstruction rates indicate potentially unstable edges, suggesting a lack of consistent patterns in the network's connections. This instability, by extension, implies unreliable citations, as the model's difficulty in reconstructing these edges signifies a deviation from expected citation behaviors. Furthermore, the variance of this distribution affords insight into the network's structural dynamics, revealing the degree of heterogeneity in edge reliability. A high variance, for example, may suggest the presence of distinct clusters with varying citation practices, thereby elucidating potential anomalies within citation patterns. Such anomalies could indicate manipulative activities, evolving research trends, or inherent structural weaknesses within the network, all of which warrant further investigation to ensure the integrity of scholarly communication.

4 NUMERICAL EXPERIMENTS

The validation of the proposed model is conducted using the CORA dataset, a well-established benchmark in citation network analysis. This dataset comprises 2,708 scientific publications categorized into seven distinct disciplines and interconnected through a citation network comprising 5,429 links. Each publication is represented as a binary word vector, indicating the presence or absence of specific terms from a dictionary of 1,433 unique words commonly used within these fields. The dataset is particularly valuable for examining publication relationships, analyzing term distributions across disciplines, and predicting future citation patterns. Its structured representation enables a comprehensive assessment of the model's effectiveness in capturing structural and contextual patterns within citation networks.

This dataset, widely used in testing various approaches like clustering, link prediction, citation validation, etc. (see, e.g., McCallum, 2024). The data consists of nodes representing academic articles spanning various research fields, including Neural Networks, Probabilistic Methods, Rule Learning, Genetic Algorithms, Reinforcement Learning, Theory, and Case-Based Reasoning.

A citation graph underwent modification by randomly removing 25% and 50% of its edges. The algorithm is then applied with 50 iterations to this altered network, each time attempting to predict the presence of the removed edges. The success in correctly identifying each edge in the graph is reconstructed, which is exhibited by the proportion of successful predictions across all iterations.

The following Fig.1 represents a distribution of the reconstruction rate obtained for a 25% random removal repeated 50 times. The category borders are [9,21), [21,33), [33,45), [45,50], and the relative frequencies (0.4105 0.5262 0.0059 0.0575)

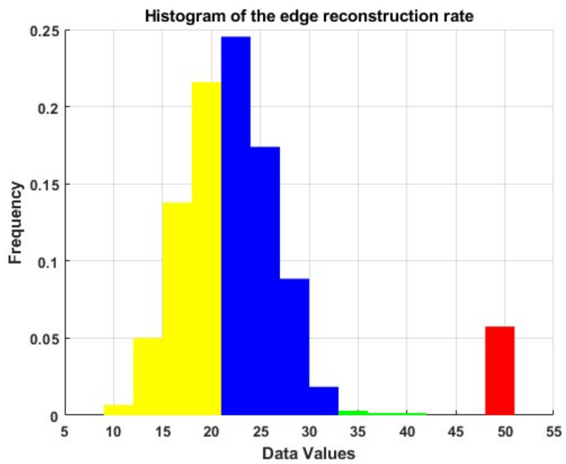

Figure 1: Histogram of the edge reconstruction rate obtained for random removing of 25%.

The distribution is positively skewed, resulting in an asymmetric form. This fact is well coordinated with results obtained in (Avros, Haim, Madar, Ravve and Volkovich, 2023) and (Avros, Keshet, Kitai, Vexler and Volkovich, 2023).

The second scenario being analyzed involves randomly removing 50% of the edges. Fig.2 exhibits the obtained histogram

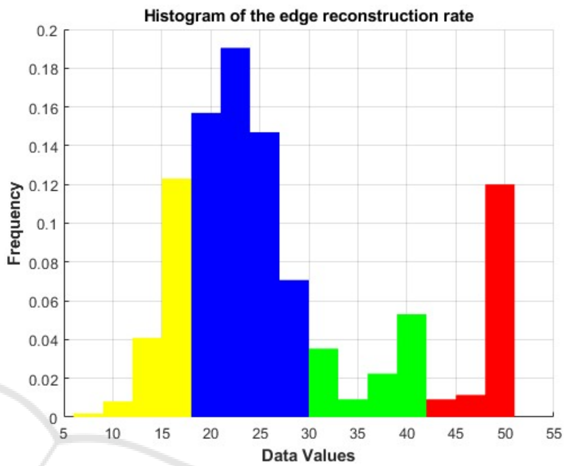

Figure 2: A histogram of the edge reconstruction rate was obtained for random removal of 50%.

Also, the data is unevenly distributed in this case, with a longer tail on the right, making it asymmetric. The edges of the categories are [6, 18), [18, 30), [30, 42), [42,48] with the relative frequencies (0.1738, 0.5650, 0.1201, 0.1407). Overall, this distribution is shifted to the left compared to the previous case.

The last considered case is a sanity check.

Sanity checks are basic, initial tests that confirm a system, model, or dataset is functioning as expected before more in-depth analysis. They prevent apparent errors and inconsistencies, ensuring the validity of later evaluations. These checks simplify debugging, enhance efficiency, and stop errors from spreading by catching fundamental issues early. Sanity checks are essential across diverse fields, verifying that inputs, outputs, and system behaviours meet predefined standards.

In our case, it is an experiment randomly added to the data connections. More in detail, a central fraction of edges is randomly added to the network aiming to take part in the testing procedure. It is natural to anticipate that most such edges must not be recognized as genuine ones. In our study, 10% of the overall source quantity of edges is randomly added. The result is presented in Fig.3.

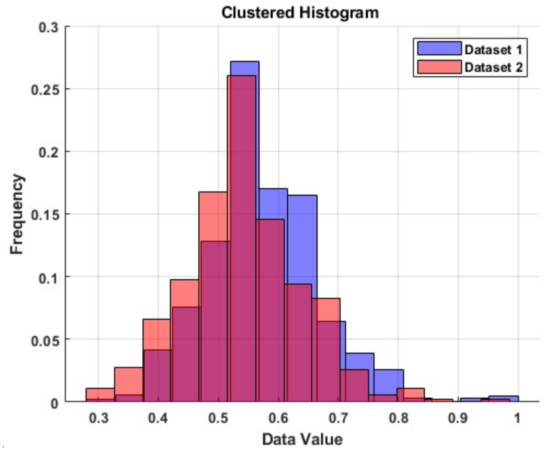

Figure 3: Histograms of the reconstructed rate of the source and 10% noised dataset.

The figure illustrates the difference between Dataset 1, the original dataset, and Dataset 2, which includes a 10% random edge addition. The noised dataset's histogram reveals an expected concentration towards the left, suggesting that the artificially added edges are less amenable to reconstruction. The skewness values of 0.6056 and 0.3373 confirm this observation.

Thus, the provided sanity check corroborates the suitability of the model.

5 CONCLUSIONS

This paper presents a novel, data-driven approach to uncovering and systematically analyzing the intricate internal structure of citation networks. At the heart of this methodology lies a Generative Adversarial Network (GAN)- based graph model designed to learn and internalize standard citation patterns that emerge naturally within academic literature. By capturing these normative relationships between citing and cited works, the model establishes a statistical baseline for expected citation behavior. Deviations from this learned baseline, measured through significant reconstruction errors, serve as strong indicators of potential citation anomalies.

A systematic perturbation strategy is employed to evaluate the reliability of individual citations. Citation links, represented as edges within the network, are selectively removed, and the trained GAN-based framework is then tasked with predicting their reinstatement. The underlying principle is intuitive: citations that align with established, legitimate patterns are more likely to be accurately reconstructed, while those exhibiting irregularities or inconsistencies remain unrecognized by the model. The inability to predict reinstatement serves as a potential marker of citation manipulation, irrelevance, or artificial inflation.

The effectiveness of this approach is rigorously validated using the CORA dataset, a widely recognized benchmark in citation network analysis. Experimental results demonstrate the model's ability to distinguish between genuine, contextually relevant citations and those potentially introduced to artificially enhance scholarly influence. This validation highlights the potential of the proposed methodology to provide a scalable, automated framework for preserving research integrity.

Beyond anomaly detection, this study addresses the broader issue of citation distortions within academic literature. By offering an objective, quantitative measure of citation reliability, this approach equips researchers, publishers, and academic institutions with a powerful tool for identifying and mitigating unethical citation practices. Moreover, the insights derived from structural anomalies in citation networks contribute to a deeper understanding of how citation behavior influences scholarly impact and knowledge dissemination. Ultimately, this research promotes a more transparent and trustworthy academic ecosystem by encouraging responsible citation practices and ensuring that scholarly recognition is grounded in genuine contributions.

In future studies, it's important to focus on preventing overfitting, which can cause the model to perform poorly on new data when trained on smaller datasets.

REFERENCES

Avros, R., Haim, M. B., Madar, A., Ravve, E., & Volkovich, Z. (2024). Spotting suspicious academic citations using self-learning graph transformers. *Mathematics, 12*(6), 814. https://doi.org/10.3390/math12060814

Avros, R., Keshet, S., Kitai, D. T., Vexler, E., & Volkovich, Z. (2023). Detecting manipulated citations through disturbed node2vec embedding. In *Proceedings of the 25th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC)*, Nancy, France, 2023 (pp. 274–278). IEEE. https://doi.org/10.1109/SYNASC61333.2023.00047

Avros, R., Keshet, S., Kitai, D. T., Vexler, E., & Volkovich, Z. (2023). Detecting pseudo-manipulated citations in scientific literature through perturbations of the citation graph. *Mathematics, 11*(18), 3820. https://doi.org/10.3390/math11123820

Bornmann, L., & Daniel, H.-D. (2008). What do citation counts measure? A review of studies on citing behavior. *Journal of Documentation, 64*(1), 45–80. https://doi.org/10.1108/00220410810844150

Case, D. O., & Higgins, G. M. (2000). How can we investigate citation behavior? A study of reasons for citing literature in communication. *Journal of the American Society for Information Science, 51*(7), 635–645. https://doi.org/10.1002/(SICI)1097-4571(2000)51:7<635::AID-ASI6>3.0.CO;2-H

Ding, Y. (2011). Scientific collaboration and endorsement: Network analysis of coauthorship and citation networks. *Journal of Informetrics*, *5*(1), 187–203.

Garfield, E. (1979). *Citation Indexing: Its Theory and Application in Science, Technology, and Humanities*. New York: Wiley.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., & Courville, A. (2014). Generative adversarial nets. *Advances in Neural Information Processing Systems, 27*, 2672–2680. Retrieved from https://arxiv.org/abs/1406.2661

Grover, A., & Leskovec, J. (2016). node2vec: Scalable feature learning for networks. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '16)*, 855–864. https://doi.org/10.1145/2939672.2939754

Isfandyari-Moghaddam, A., Saberi, M. K., Tahmasebi-Limoni, S., Mohammadian, S., & Naderbeigi, F. (2023). Global scientific collaboration: A social network analysis and data mining of the co-authorship networks. Journal of Information Science, 49(4), 1126–1141.

Jin, H., Xu, G., Cheng, K., Liu, J., & Wu, Z. (2022). A link prediction algorithm based on GAN. *Electronics, 11*(13), 2059. https://doi.org/10.3390/electronics11132059.

Liu, J., Bai, X., Wang, M., Tuarob, S., & Xia, F. (2024). Anomalous citations detection in academic networks. *Artificial Intelligence Review*, *57*

Mammola, S., Piano, E., Doretto, A., Caprio, E., & Chamberlain, D. (2022). Measuring the influence of nonscientific features on citations. *Scientometrics, 127*, 41123–4137. https://doi.org/10.1007/s11192-022-04421-7

McCallum, A. (2024). Cora. Available at: https://dx.doi.org/10.21227/jsg4-wp31

Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*. https://arxiv.org/abs/1301.3781

Prabha, C. G. (1983). Some aspects of citation behavior: A study in business administration. *Journal of the American Society for Information Science, 34*(3), 202–206.

Wang, P., & White, M. D. (1996). A qualitative study of scholars' citation behavior. In *Proceedings of ASIS Annual Meeting*, Baltimore, MD (pp. 255–261). ASI

Wilhite, A., & Fong, E. (2012). Coercive citation in academic publishing. *Science, 335*(6068), 542–543. https://doi.org/10.1126/science.1212540

Wren, J. D., & Georgescu, C. (2022). Detecting anomalous referencing patterns in PubMed papers suggestive of author-centric reference list manipulation. *Scientometrics, 127*, 5753–5771.

Zhao, Z., Zhang, T., & Zhang, Y. (2021). embedGAN: A method to embed images in GAN latent space. In *Proceedings of the International Conference on Artificial Intelligence and Robotics* (245-260). Springer. https://doi.org/10.1007/978-981-33-4400-6_20