

Conf-GNNRec: Quantifying and Calibrating the Prediction Confidence for GNN-based Recommendation Methods

Meng Yan
Xidian University
Xi'an, China
mengyan@stu.xidian.edu.cn

Cai Xu
Xidian University
Xi'an, China
cxu@xidian.edu.cn

Xujing Wang
Xidian University
Xi'an, China
xjwong@stu.xidian.edu.cn

Ziyu Guan*
Xidian University
Xi'an, China
zyguan@xidian.edu.cn

Wei Zhao
Xidian University
Xi'an, China
ywzhao@mail.xidian.edu.cn

Yuhang Zhou
Communication University Of China
Beijing, China
yhz@mails.cuc.edu.cn

Abstract

Recommender systems based on graph neural networks perform well in tasks such as rating and ranking. However, in real-world recommendation scenarios, noise such as user misuse and malicious advertisement gradually accumulates through the message propagation mechanism. Even if existing studies mitigate their effects by reducing the noise propagation weights, the severe sparsity of the recommender system still leads to the low-weighted noisy neighbors being mistaken as meaningful information, and the prediction result obtained based on the polluted nodes is not entirely trustworthy. Therefore, it is crucial to measure the confidence of the prediction results in this highly noisy framework. Furthermore, our evaluation of the existing representative GNN-based recommendation shows that it suffers from overconfidence. Based on the above considerations, we propose a new method to quantify and calibrate the prediction confidence of GNN-based recommendations (Conf-GNNRec). Specifically, we propose a rating calibration method that dynamically adjusts excessive ratings to mitigate overconfidence based on user personalization. We also design a confidence loss function to reduce the overconfidence of negative samples and effectively improve recommendation performance. Experiments on public datasets demonstrate the validity of Conf-GNNRec in prediction confidence and recommendation performance.

CCS Concepts

• **Information systems** → **Novelty in information retrieval; newty in information retrieval; Probabilistic retrieval models**; • **Applied computing** → **Online shopping; Online shopping.**

Keywords

Recommendation, Confidence, Trustworthy

*Corresponding author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
WWW Companion '25, April 28-May 2, 2025, Sydney, NSW, Australia

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-1331-6/2025/04
<https://doi.org/10.1145/3701716.3715475>

ACM Reference Format:

Meng Yan, Cai Xu, Xujing Wang, Ziyu Guan, Wei Zhao, and Yuhang Zhou. 2025. Conf-GNNRec: Quantifying and Calibrating the Prediction Confidence for GNN-based Recommendation Methods. In *Companion Proceedings of the ACM Web Conference 2025 (WWW Companion '25)*, April 28-May 2, 2025, Sydney, NSW, Australia. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3701716.3715475>

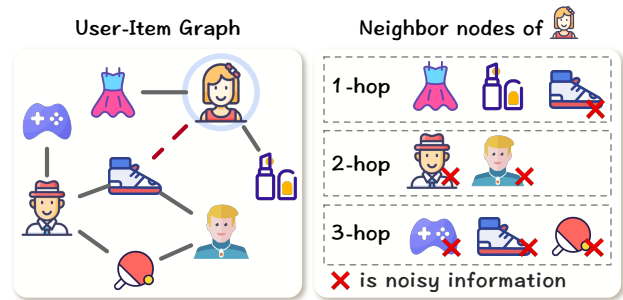


Figure 1: Illustration of the propagation of noise in the GNN-based recommendation. Sparse interactions result in 2-hop and 3-hop neighbor nodes being noisy information.

1 Introduction

Recommender systems play an important role in prediction making of various social media such as e-commerce [1], entertainment [2, 3], and social networks[4, 5]. Interactions between users and items can be naturally constructed as graphs, and graph structures can represent the higher-order connections that users establish with other users via common items. Graph Neural Network (GNN), by virtue of its ability to model interactions and capture higher-order dependencies, has become the powerful framework for recommender systems[6].

Most of the GNN-based recommendation methods make improvements in all or some of the following three steps: 1) **Graph Construction**. Transform the user-item interactions in the recommender system into a graph structure, which can introduce side information to add additional nodes and edges. 2) **Message Propagation**. Propagate and aggregate users and items information on the graph based on GNN. 3) **Recommendation**. Calculate interaction probabilities or ratings based on the final embeddings of users

and items. In practical recommendation scenarios, noise such as user misuse and malicious advertisements will gradually accumulate through the message propagation mechanism. Currently, there have been research attempts to reduce the weight of noise through attention [7–9], confidence [10–12], and other methods in the propagation process. Since node features are mainly complemented by the neighbor nodes, severe sparsity in recommendation can lead to low weighted neighbors also being recognized as meaningful information, which can lead to layer-by-layer aggregation of noise and pollution of node features [13, 14]. As shown in Figure 1, a female user accidentally clicked on a pair of shoes, and the shoe preference belongs to noise information. Meanwhile, other items that the user interacts with are very unpopular, the noise of this pair of shoes will be more aggregated during the message propagation process, leading to further propagation of its noise impact.

Recommendation results usually rely on similarity calculations between user nodes and item nodes, and the interference of noise can lead to predictions that are not fully trustworthy. Therefore, it is crucial to measure the confidence of predictions in this high noise framework. In other words, we need to assess the trustworthiness of the result through the probability, i.e., confidence level, associated with the prediction result. In general, the predicted probability at the output layer represents the confidence for the prediction. Therefore, it naturally raises the fundamental question: *does the confidence given directly by a GNN-based recommendation model match the accuracy of the prediction?*

To explore this question, we evaluate the relationship between confidence and accuracy for lightGCN [6] and KGCL [15] on publicly available datasets respectively, where confidence is the activation function output of the last layer of the model. Surprisingly, confidence in existing GNN-based recommendations is significantly flawed, with the confidence being much higher than the accuracy. These probabilities obtained from the activation layer only take into account the relative magnitude of the outputs and do not truly reflect the level of confidence the model has in the predicted results. More importantly, the overconfidence confidence that the model assigns to the prediction results can harm the user experience. In high-risk scenarios such as finance and healthcare, wrong recommendations can cause significant loss of life and property.

Based on the above challenges, we propose a new method to quantify and calibrate the prediction confidence of GNN-based recommendations (Conf-GNNRec), which adopts a post-calibration strategy without other additional inputs. Specifically, we design a rating calibration method that dynamically adjusts excessive ratings to mitigate overconfidence based on user personalization. In addition, we design a new confidence loss function that effectively reduces the high confidence of negative samples, further bridging the gap between recommendation accuracy and confidence. The main contributions of this work can be summarized as follows:

- To the best of our knowledge, we are the first to focus on prediction confidence in GNN-based recommendation and to point out that existing methods are overconfident in predicting results.
- We propose a new prediction confidence measure method that can improve recommendation performance while mitigating overconfidence.

- We compare the performance changes of representative GNN-based recommendation methods before and after applying Conf-GNNRec on a public dataset, and the experimental results demonstrate the effectiveness of Conf-GNNRec.

2 Method

2.1 Problem Statement

In recommender systems, user-item interactions can be naturally constructed as graphs $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$. The $\mathcal{V} = \mathcal{V}_U \cup \mathcal{V}_I$ is the node set with $(N + M)$ nodes, where $\mathcal{V}_U = \{v_1, v_2, \dots, v_N\}$ represents the user set, and $\mathcal{V}_I = \{v_{N+1}, v_{N+2}, \dots, v_{N+M}\}$ represents the item set. $\mathcal{E} = \{(v_u, v_i) | y_{ui} = 1\}$ is the edge set, where $y_{ui} = 1$ indicating that there is an implicit interaction between $v_u \in \mathcal{V}_U$ and $v_i \in \mathcal{V}_I$. Given a GNN-based recommendation model f_θ where θ is the learnable parameters.

2.2 GNN-based Recommendation Confidence

The GNN-based recommendation model firstly defines the initialised embedding features of the nodes: $\mathbf{h}_u^{(0)} = \text{Embed}(v_u)$ and $\mathbf{h}_i^{(0)} = \text{Emb}(v_i)$, where the $\text{Emb}(\cdot)$ is the embedding function. Then, it aggregates the neighbor information of users and items and updates the representation. After the L -th layer of the network, it can be represented as:

$$\mathbf{h}_*^{(L+1)} = \text{Update}^{(L)}(\mathbf{h}_*^{(L)}, \text{Aggregator}^{(L)}(\{\mathbf{h}_{v_i}^{(L)}, \forall v_i \in \mathcal{N}(v_*)\})) \quad (1)$$

where $\text{Update}^{(L)}$ and $\text{Aggregator}^{(L)}$ represent the function of aggregation and update at the L -th layer, respectively. $\mathcal{N}(v_*)$ indicates the neighbor nodes of v_* . We consider it as the final representation of users and items to predict interactions or ratings:

$$r_{u,i} = \text{score}(\mathbf{h}_u^{(L+1)}, \mathbf{h}_i^{(L+1)}) \quad (2)$$

To facilitate comparison of results, the results are fed into the normalization layer: $z_u = [\sigma(r_{u,1}), \sigma(r_{u,2}), \dots, \sigma(r_{u,M})]$, where $\sigma(\cdot)$ is the normalization function. For the user u , $\hat{y}_{u,i} = \arg \max_i z_{u,i}$ indicates the prediction and $\hat{p}_{u,i} = \max_i z_{u,i}$ indicates the corresponding confidence. We believe that perfect confidence should satisfy the following equation:

$$\mathbb{P}(\hat{y}_{u,i} = y_{u,i} | \hat{p}_{u,i} = p) = p, \forall p \in [0, 1] \quad (3)$$

For the top-k task, we use the mean value instead. Next, we visualize the relationship between confidence and accuracy using two representative GNN-based recommendations (lightGCN and KGCL) as an example, as shown in figure 2. Specifically, we divide the confidence interval $[0, 1]$ into ten equal bins with reference to [16], group the results into the corresponding bins based on the confidence, and then compute the average accuracy of each bin. For a perfect model, each bin's accuracy is exactly equal to the confidence, i.e., all blue bars match the brown diagonal. It is clear that the model outputs a confidence level that is too high and that the confidence level does not match the accuracy (yellow slashes).




Figure 2: Reliability diagrams for LightGCN (left) and KGCL (right). The blue bar is the output of model, the yellow bar is the output in the perfect case. Any deviation from a perfectly diagonal (i.e., yellow slashes) represents the miscalibration.

2.3 Conf-GNNRec

In this section, we propose the rating calibration method to reset the ratings. The function should 1) adapt to different user interaction habits. 2) be non-linear, and 3) maintain recommendation performance. Based on the above objectives, we design a segmented processing strategy to achieve rating calibration. A low rating usually indicates that the user has little interest in the item, and the original values can be retained directly. High ratings represent strong user preferences but may also overstate the weights of items. Therefore, we use logarithmic transformation to nonlinearly compress the portion that exceeds the mean value while retaining the relative ranking information. We reset the ratings r'_u as follows:

$$r'_u = [rc(r_{u,1}), rc(r_{u,2}), \dots, rc(r_{u,M})] \quad (4)$$

$$rc(r_{u,i}) = r_{u,i} \cdot \mathbb{I}(r_{u,i} \leq \bar{r}_u) + (\tau \log(r_{u,i} - \bar{r}_u) + \bar{r}_u) \cdot \mathbb{I}(r_{u,i} > \bar{r}_u) \quad (5)$$

where $rc()$ is responsible for calibrating overconfidence. $\mathbb{I}()$ is the indicator function that outputs 1 when the condition is true, and \bar{z}_u is the mean of all elements in z_u . τ is the temperature factor controlling the smoothness of ratings. So we can get a new probability distribution: $z'_u = [\sigma(r'_{u,1}), \sigma(r'_{u,2}), \dots, \sigma(r'_{u,M})]$. Accordingly, we reset the confidence of the prediction results: $\hat{p}'_{u,i} = \max_i z'_{u,i}$.

In addition, we separately penalize the items that are wrongly predicted but with high confidence, which not only effectively mitigates the over-confidence but also improves the accuracy of recommendation to some extent. It can be flexibly applied to all GNN-based recommender systems. The proposed regularization loss is as follows:

$$\mathcal{L}_{rcREC} = \sum_{u=1}^N \sum_{i=1}^M (1 - p_{u,i}) \hat{p}_{u,i} \log(1 + \exp(\hat{p}_{u,i})) \quad (6)$$

where $\hat{p}_{u,i} \log(1 + \exp(\hat{p}_{u,i}))$ increases the penalty for high confidence negative samples and increases nonlinearly with confidence. It also mitigates the effect of low confidence errors so that the model will make predictions about items more cautiously.

3 Experiment

In this section, we present the datasets, baseline methods, and evaluation protocols. Subsequently, we apply Conf-GNNRec on

different baseline models and compare the performance before and after, the related code¹ is publicly available. Finally, we visualize the effect of the improvement in the confidence.

3.1 Settings

3.1.1 Datasets. We perform experiments on three public datasets collected from different real-life platforms: Gowalla for location recommendation, Yelp2018 for restaurant recommendation, and Amazon-Book for product recommendation. Table 1 presents the statistical information of our experimented datasets.

Table 1: Statistics of Datasets

Datasets	User	Item	Interaction	Density
Gowalla	29,558	40,981	1,027,370	0.00084
Yelp2018	31,668	38,048	1,561,406	0.00130
Amazon-Book	52,643	91,599	2,984,108	0.00062

3.1.2 Baseline. We apply Conf-GNNRec to the following baselines and compare them:

- LightGCN[6]. This is a GCN-based recommendation method that simplifies the convolution operations during message passing among users and items.
- KGAT[17]. This model designs an attentive message passing scheme over the knowledge-aware collaborative graph for embedding fusion. Setting weights for different neighbor nodes based on the attention mechanism.
- MVIN[18]. It is a multi-view item embedding method based on GNN. Information from both the user and the entity side is considered to learn feature embeddings of items.
- KGCL[15]. This is a general framework that introduces knowledge graphs to guide recommender systems in denoising information as it is aggregated.

3.1.3 Evaluation Protocols. For all of the above methods, we use open source code and uniform datasets. We randomly split the dataset in the ratio of 7:1:2 to build the training, validation and test sets. Moreover, it is ensured that each user has at least one item in the test set. We followed the authors' advice on the method for the hyperparameters used in the baseline model and reported the results under the optimal hyperparameter settings. In particular, KGCL additionally introduces a knowledge graph, resulting in a recommended dataset inconsistent with Table 1, and we conduct the experiment according to the authors' settings. For the performance evaluation, we use two representative metrics: *Precision@N* and *Accuracy@N* are used to evaluate the accuracy of top-N recommended items. Average evaluation results across all users in the test set are reported with $N = 20$ by default.


3.2 Overall Performance

3.2.1 Recommendation. We compare the performance changes of the four recommendation models before and after the application of Conf-GNNRec and show the magnitude of improvement brought by Conf-GNNRec. Table 2 shows the experimental results.

¹<https://github.com/FairyMeng/Conf-GNNRec>

Table 2: Overall performance on Recommendation

METHOD	Gowalla		Yelp2018		Amazon-Book	
	PRE.	ACC.	PRE.	ACC.	PRE.	ACC.
LightGCN	5.267	15.327	2.716	4.313	3.295	7.305
+Conf-GNNRec	5.386	15.597	2.864	4.496	3.356	7.418
Improve (%)	2.259	1.762	5.449	4.243	1.851	1.547
KGAT	5.452	15.838	2.743	4.374	3.273	7.393
+Conf-GNNRec	5.577	16.225	2.882	4.512	3.346	7.547
Improve (%)	2.293	2.443	5.067	3.155	2.230	2.083
MVIN	5.527	15.935	2.883	4.419	3.401	7.428
+Conf-GNNRec	5.652	16.448	3.011	4.616	3.489	7.604
Improve (%)	2.262	3.219	4.093	4.458	2.587	2.369
KGCL	5.524	16.362	2.758	4.832	3.618	7.829
+Conf-GNNRec	5.644	16.871	2.942	5.093	3.735	8.139
Improve (%)	2.172	3.111	6.672	5.401	3.234	3.960

**Figure 3: Reliability diagrams for LightGCN (left) and KGCL (right) after applying Conf-GNNRec.**

It can be clearly observed that the application of Conf-GNNRec improves all models and datasets. In particular, the improvement of Conf-GNNRec is more significant in complex models where more information is introduced (e.g., KGCL) and datasets with high density (e.g., Yelp2018). The above illustrates the effectiveness of our model design in reducing false predictions generated by overconfidence.

3.2.2 Confidence. In this section, we visualize the relationship between the prediction confidence and the accuracy in lightGCN and KGCL on Gowalla, as shown in figure 3. All settings are consistent with those in 3.1.2. By comparing figure 2 and figure 3, it can be clearly seen that the model significantly reduces the error between confidence and accuracy. The KGCL model exhibits smoother relational consistency. In addition, as the confidence level increases, the accuracy of the recommender system shows a significant increase, which is in accordance with our intuitive understanding: the prediction with higher confidence tends to correspond to a more reliable recommendation result. However, we also recognize that achieving fully accurate confidence predictions in recommender systems is nearly impossible. This is mainly attributed to the high degree of personalization, uncertainty in user preferences, and the inevitable inclusion of noisy interactions (e.g., user incidental clicking behavior). Therefore, the goal of a trusted recommender system should focus on reducing the error between the prediction confidence and

the actual accuracy, thus improving the credibility of the model and the user experience.

4 Conclusion

In this paper, we explore the relationship between confidence and accuracy of GNN-based recommender systems. By analysing existing approaches, we find overconfidence in the models. To address this issue, we propose a new prediction confidence measure and design a loss function that significantly mitigates overconfidence while improving recommendation performance. We compare the performance changes of representative GNN-based recommendation methods before and after applying Conf-GNNRec on a public dataset, and the experimental results demonstrate the effectiveness of Conf-GNNRec.

It is worth noting that the research on prediction confidence for GNN-based recommendation is in its infancy. We plan to use Bayesian to improve the trustworthiness of recommendation results further. Meanwhile, we will endeavor to extend to other recommendation frameworks (e.g., large language models) for exploring a wider range of application scenarios and potential value.

Acknowledgments

This research was supported by the National Natural Science Foundation of China under Grants 62133012, 62472340, 62425605, 62303366. Key Research and Development Program of Shaanxi under Grant 2024CY2-GJHX-15.

References

- [1] Meng Yan, Haibin Huang, Ying Liu, Juan Zhao, Xiyue Gao, Cai Xu, Ziyu Guan, and Wei Zhao. 2024. Truths: trustworthy sequential recommender systems via user-generated multimodal content. In *International Conference on Database Systems for Advanced Applications*. Springer, 180–195.
- [2] Yunfan Gao, Tao Sheng, Youlin Xiang, Yun Xiong, Haofen Wang, and Jiawei Zhang. 2023. Chat-rec: Towards interactive and explainable llms-augmented recommender system. *arXiv preprint arXiv:2303.14524* (2023).
- [3] T Jemima Jebaseeli et al. 2024. Youtube Video Recommendation using user-based Collaborative Filtering and Graph Neural Networks Approaches to Improve users Personalized Experience. *Grenze International Journal of Engineering & Technology (GIJET)* 10 (2024).
- [4] Mahesh Thyluru Ramakrishna, Vinoth Kumar Venkatesan, Rajat Bhardwaj, Surbhi Bhatia, Mohammad Khalid Imam Rahmani, Saima Anwar Lashari, and Aliaa M Alabdali. 2023. HCoF: Hybrid Collaborative Filtering Using Social and Semantic Suggestions for Friend Recommendation. *Electronics* 12, 6 (2023), 1365.
- [5] Weigang Lu, Ziyu Guan, Wei Zhao, Yaming Yang, Yibing Zhan, Yiheng Lu, and Dapeng Tao. 2024. AGMixup: Adaptive Graph Mixup for Semi-supervised Node Classification. *arXiv preprint arXiv:2412.08144* (2024).
- [6] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*. 639–648.
- [7] Gang-Feng Ma, Xu-Hua Yang, Liang-Yu Gao, and Ling-Hang Lian. 2024. Adaptive denoising graph contrastive learning with memory graph attention for recommendation. *Neurocomputing* 610 (2024), 128595.
- [8] Bin Wu, Xiangnan He, Le Wu, Xue Zhang, and Yangdong Ye. 2023. Graph-augmented co-attention model for socio-sequential recommendation. *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 53, 7 (2023), 4039–4051.
- [9] Qi Wang, Hao Cui, Jiapeng Zhang, Yan Du, and Xiaojun Lu. 2024. Self-supervised global graph neural networks with enhance-attention for session-based recommendation. *Applied Soft Computing* 150 (2024), 111026.
- [10] Yan Ge, Jun Ma, Li Zhang, Xiang Li, and Haiping Lu. 2023. Trustworthiness-aware knowledge graph representation for recommendation. *Knowledge-Based Systems* 278 (2023), 110865.
- [11] Jiakai Tang, Zuxu Sun, Wei Yao, and Xu Chen. 2024. Robust Graph Recommendation via Noise-Aware Adversarial Perturbation. In *International Conference on Database Systems for Advanced Applications*. Springer, 134–150.
- [12] Guipeng Xu, Chen Lin, Wanxian Guan, Jinping Gou, Xubin Li, Hongbo Deng, Jian Xu, and Bo Zheng. 2023. E-commerce Search via Content Collaborative

- Graph Neural Network. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 2885–2897.
- [13] Weigang Lu, Yibing Zhan, Binbin Lin, Ziyu Guan, Liu Liu, Baosheng Yu, Wei Zhao, Yaming Yang, and Dacheng Tao. 2024. SkipNode: On alleviating performance degradation for deep graph convolutional networks. *IEEE Transactions on Knowledge and Data Engineering* (2024).
- [14] Weigang Lu, Ziyu Guan, Wei Zhao, Yaming Yang, and Long Jin. 2024. Nodemixup: Tackling under-reaching for graph neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 14175–14183.
- [15] Xiaohui Zhang, Huifang Ma, Fanyi Yang, Zhixin Li, and Liang Chang. 2023. KGCL: A Knowledge-enhanced Graph Contrastive learning framework for session-based recommendation. *Engineering Applications of Artificial Intelligence* 124 (2023), 106512.
- [16] Xiao Wang, Hongrui Liu, Chuan Shi, and Cheng Yang. 2021. Be confident! towards trustworthy graph neural networks via confidence calibration. *Advances in Neural Information Processing Systems* 34 (2021), 23768–23779.
- [17] Xiang Wang, Xiangnan He, Yixin Cao, Meng Liu, and Tat-Seng Chua. 2019. Kgat: Knowledge graph attention network for recommendation. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*. 950–958.
- [18] Chang-You Tai, Meng-Ru Wu, Yun-Wei Chu, Shao-Yu Chu, and Lun-Wei Ku. 2020. Mvin: Learning multiview items for recommendation. In *Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval*. 99–108.