



OPEN ACCESS

EDITED BY

Tyler Derr,
Vanderbilt University, United States

REVIEWED BY

Shaohua Tao,
Xuchang University, China
Sead Delalic,
University of Sarajevo, Bosnia and Herzegovina
Yuying Zhao,
Vanderbilt University, United States

*CORRESPONDENCE

Tomislav Duricic
✉ tduricic@tugraz.at
Dominik Kowald
✉ dkowald@tugraz.at
Elisabeth Lex
✉ elisabeth.lex@tugraz.at

RECEIVED 30 June 2023

ACCEPTED 29 November 2023

PUBLISHED 19 December 2023

CITATION

Duricic T, Kowald D, Lacic E and Lex E (2023)
Beyond-accuracy: a review on diversity,
serendipity, and fairness in recommender
systems based on graph neural networks.
Front. Big Data 6:1251072.
doi: 10.3389/fdata.2023.1251072

COPYRIGHT

© 2023 Duricic, Kowald, Lacic and Lex. This is
an open-access article distributed under the
terms of the [Creative Commons Attribution
License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or
reproduction in other forums is permitted,
provided the original author(s) and the
copyright owner(s) are credited and that the
original publication in this journal is cited, in
accordance with accepted academic practice.
No use, distribution or reproduction is
permitted which does not comply with these
terms.

Beyond-accuracy: a review on diversity, serendipity, and fairness in recommender systems based on graph neural networks

Tomislav Duricic^{1,2*}, Dominik Kowald^{1,2*}, Emanuel Lacic³ and Elisabeth Lex^{1*}

¹Institute of Interactive Systems and Data Science, Graz University of Technology, Graz, Austria, ²Know Center, Graz, Austria, ³Infobip, Zagreb, Croatia

By providing personalized suggestions to users, recommender systems have become essential to numerous online platforms. Collaborative filtering, particularly graph-based approaches using Graph Neural Networks (GNNs), have demonstrated great results in terms of recommendation accuracy. However, accuracy may not always be the most important criterion for evaluating recommender systems' performance, since beyond-accuracy aspects such as recommendation diversity, serendipity, and fairness can strongly influence user engagement and satisfaction. This review paper focuses on addressing these dimensions in GNN-based recommender systems, going beyond the conventional accuracy-centric perspective. We begin by reviewing recent developments in approaches that improve not only the accuracy-diversity trade-off but also promote serendipity, and fairness in GNN-based recommender systems. We discuss different stages of model development including data preprocessing, graph construction, embedding initialization, propagation layers, embedding fusion, score computation, and training methodologies. Furthermore, we present a look into the practical difficulties encountered in assuring diversity, serendipity, and fairness, while retaining high accuracy. Finally, we discuss potential future research directions for developing more robust GNN-based recommender systems that go beyond the unidimensional perspective of focusing solely on accuracy. This review aims to provide researchers and practitioners with an in-depth understanding of the multifaceted issues that arise when designing GNN-based recommender systems, setting our work apart by offering a comprehensive exploration of beyond-accuracy dimensions.

KEYWORDS

survey, recommender systems, graph neural networks, beyond-accuracy, diversity, serendipity, novelty, fairness

1 Introduction

With their ability to provide personalized suggestions, recommender systems have become an integral part of numerous online platforms by helping users find relevant products and content (Aggarwal et al., 2016). There are various methods employed to implement recommender systems, among which collaborative filtering (CF) has proven to be particularly effective due to its ability to leverage user-item interaction data to generate personalized recommendations (Koren et al., 2021). Recent advances in Graph Neural Networks (GNNs) have also had a significant impact on the field of recommender systems,

and especially on collaborative filtering. GNN-based CF approaches have demonstrated exceptional results in terms of recommendation accuracy, which has traditionally been the main criterion for evaluating the performance of recommender systems (Pu et al., 2012; He et al., 2020).

However, most studies have focused only on accuracy and have often neglected other equally or sometimes even more important aspects of recommender systems, such as diversity, serendipity, and fairness. The importance of these *beyond-accuracy* dimensions is increasingly being recognized, as studies have shown that these aspects can have a significant impact on user satisfaction (Abdollahpouri et al., 2019). For example, diverse and serendipitous recommendations can prevent the over-specialization of content and enhance user discovery. Novelty, a closely related concept to serendipity, introduces fresh and unexpected options to users, further enriching the discovery process. Fairness, on the other hand, ensures that the system does not discriminate against certain users or item providers, thereby promoting equitable user experiences (Gao et al., 2023).

This review paper further explores these dimensions in the context of GNN-based recommender systems, going beyond the traditional accuracy-centric viewpoint. We discuss recent advances in approaches that not only improve the accuracy-diversity trade-off, but also promote serendipity, novelty and fairness. Furthermore, we highlight the practical issues encountered in assuring these dimensions when constructing GNN-based CF approaches, while preserving high recommendation accuracy. This review is intended to provide researchers and practitioners with a comprehensive understanding of the multifaceted optimization issues that arise when designing GNN-based recommender systems, thereby contributing to the development of more robust and user-centric recommender systems.

2 Background

Graph neural networks (GNNs) have recently emerged as an effective way to learn from graph-structured data by capturing complex patterns and relationships (Hamilton, 2020). Through the propagation and transformation of feature information among interconnected nodes in a graph, GNNs can effectively capture the local and global structure of the given graphs. Consequently, they emerge as an ideal method especially suitable for dealing with tasks involving interconnected, relational data such as social network analysis, molecular chemistry, and recommender systems among others.

In recommender systems, integrating Graph Neural Networks (GNNs) with traditional collaborative filtering techniques has been shown beneficial. Representing users and items as nodes in a graph with interactions acting as edges allows GNNs to provide more accurate personalized recommendations by discovering and utilizing intricate connections that would otherwise remain undetected (Wang X. et al., 2019). In particular, higher-order connectivity together with transitive relationships play an essential role when trying to extract user preferences in certain scenarios.

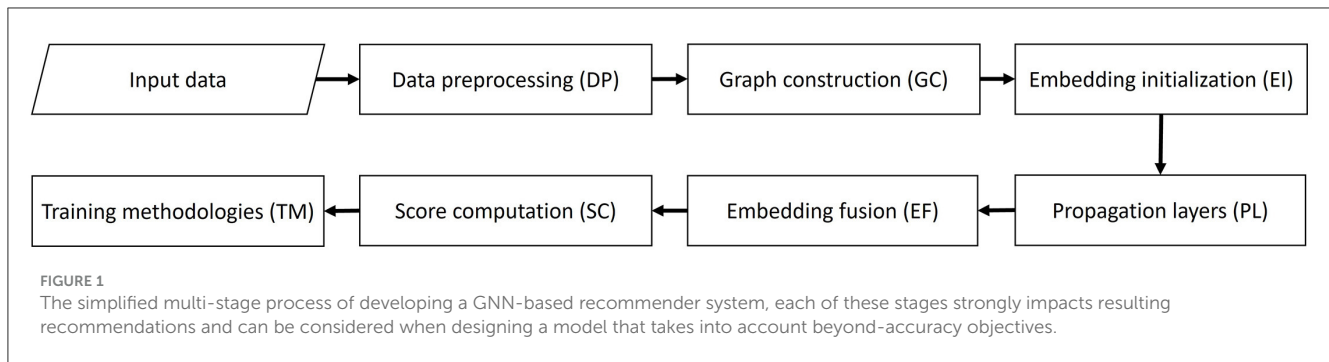
GNN-based recommender systems represent an evolving field with continuous advancements and innovations. Recent research has focused on multiple aspects of GNNs in recommender systems,

ranging from optimizing propagation layers to effectively managing large-scale graphs and integration of auxiliary information (Zhou et al., 2022). Aside from these aspects, an expanding interest lies in exploring beyond-accuracy objectives for recommender systems. Such objectives include diversity, explainability/interpretability, fairness, serendipity/novelty, privacy/security, and robustness which offer a more comprehensive evaluation of the system's performance (Wu S. et al., 2022; Gao et al., 2023). However, our work focuses primarily on three key aspects: diversity, serendipity, and fairness, since these aspects have a significant impact on user satisfaction, while also considering ethical concerns in the field of recommender systems. Ensuring diversity amongst recommendations minimizes over-specialization effects, benefiting users in product/content discovery and exploration (Kunaver and Požrl, 2017). Considering serendipity and novelty also helps to overcome the over-specialization problem by allowing the system to recommend novel and unexpected yet relevant items, thus improving user satisfaction (Kaminskas and Bridge, 2016). The aspect of fairness ensures that the system does not discriminate against certain users or item providers, thereby promoting equitable user experiences (Deldjoo et al., 2023).

Diversity, serendipity, novelty, and fairness in recommender systems are interconnected and often influence each other. For instance, increasing diversity can lead to more serendipitous and novel recommendations, since users are exposed to a wider range of unexpected and less-known items (Kotkov et al., 2020). Some studies occasionally use the terms “diversity” and “novelty” interchangeably, highlighting a common overlap in their conceptual usage (Sun et al., 2020; Dhawan et al., 2022). It's important to note that novelty and serendipity are closer related concepts, as they both compare the recommended items with a user's history, emphasizing the discovery of unexpected content that aligns with personal preferences. Furthermore, focusing on diversity and serendipity can also promote fairness, since it ensures a more equitable distribution of recommendations across items and prevents the system from consistently suggesting only popular items (Mansoury et al., 2020). However, it's important to note that these aspects need to be balanced with the system's accuracy and relevance to maintain user satisfaction. Considering beyond-accuracy dimensions contributes to supporting the development of GNN-based recommender systems that are not only robust and accurate but also user-centric and ethically considerate.

While GNNs have seen rapid advancements, their application in recommender systems has also been the subject of several surveys. Wu S. et al. (2022) and Gao et al. (2023) provide a broad overview of GNN methods in recommender systems, touching upon aspects of diversity and fairness. Dai et al. (2022) delves into fairness in graph neural networks in general, briefly discussing fairness in GNN-based recommender systems. Meanwhile, Fu et al. (2023) explores serendipity in deep learning recommender systems, with limited focus on GNN-based recommenders. Building on these insights, our review distinctively emphasizes the importance of diversity, serendipity, novelty, and fairness in GNN-based recommender systems, offering a deeper dive into these dimensions.

To conduct our review, we searched for literature on Google Scholar using keywords such as “diversity”, “serendipity”, “novelty”, “fairness”, “beyond-accuracy”, “graph neural networks”



or “recommender system”. We manually checked the resulting papers for their relevance and retrieved 20 publications overall from relevant journals and conferences in the field (see Table 1). While re-ranking and post-processing methods are often used when optimizing beyond-accuracy metrics in recommender systems (Gao et al., 2023), this paper specifically concentrates on advancements within GNN-based models, thus leaving these methods outside the discussion. Finally, it is important to highlight that diversity, serendipity, and fairness are extensively researched in recommender systems beyond GNNs. Broader literature across various architectures has provided insights into these challenges and their overarching solutions. While our paper primarily focuses on GNN-based recommender systems, we direct readers to consult these works for a comprehensive perspective (Kaminskas and Bridge, 2016; Castells et al., 2021; Li et al., 2022; Dong et al., 2023; Wang et al., 2023a; Zhao et al., 2023).

3 Model development

The construction of a GNN-based recommender system is a complex, multi-stage process that requires careful planning and execution at each step. These stages include data preprocessing (DP), graph construction (GC), embedding initialization (EI), propagation layers (PL), embedding fusion (EF), score computation (SC), and training methodologies (TM). In this section, we provide an overview of this multi-stage process as it is crucial for understanding the specific stages at which current research has concentrated efforts to address the beyond-accuracy aspects of diversity, serendipity, and fairness in GNN-based recommender systems, as shown in Figure 1.

3.1 Data preprocessing, graph construction, embedding initialization

The initial stage of developing a GNN-based collaborative filtering model is data preprocessing, where user-item interaction data and auxiliary information such as user/item features or social connections are collected and processed (Lacic et al., 2015a; Duricic et al., 2018, 2020; Fan et al., 2019b; Wang H. et al., 2019). Techniques like data imputation ensure that missing data is filled, providing a more complete dataset, while outlier detection helps in maintaining the data’s integrity. Feature normalization ensures

consistent data scales, enhancing model performance. Addressing the cold-start problem at this stage ensures that new users or items without sufficient interaction history can still receive meaningful recommendations (Lacic et al., 2015b; Liu et al., 2020).

The graph construction stage is crucial, as the graph’s structure directly influences the model’s efficacy. Choosing the type of graph determines the nature of relationships between nodes. Adjusting edge weights can prioritize certain interactions, while adding virtual nodes/edges can introduce auxiliary information to improve recommendation quality (Kim et al., 2022; Wang et al., 2023b).

In the embedding initialization stage, nodes are assigned low-dimensional vectors or embeddings. The choice of embedding size balances computational efficiency and representation power. Different initialization methods offer trade-offs between convergence speed and stability. Including diverse information in the embeddings can capture richer user-item relationships, enhancing recommendation quality Wang et al. (2021). This initialization can be represented as $H^{(0)} = [h_{\text{user}}^{(0)}; h_{\text{item}}^{(0)}]$, where $h_{\text{user}}^{(0)}$ and $h_{\text{item}}^{(0)}$ are the initial embeddings of the user and item nodes, respectively.

3.2 Propagation layers, embedding fusion, score computation, training methodologies

Propagation layers in GNNs aggregate and transform features of neighboring nodes to generate node embeddings, represented as $H^{(l+1)} = \sigma(D^{-1}AH^{(l)}W^{(l)})$, where $H^{(l)}$ is the matrix of node features at layer l , A is the adjacency matrix, D is the degree matrix, $W^{(l)}$ is the weight matrix at layer l , and σ is the activation function (Hamilton, 2020). There are numerous approaches built on this concept. For instance, He et al. (2020) adopt a simplified approach, emphasizing straightforward neighborhood aggregation to enhance the quality of node embeddings; whereas Fan et al. (2019b) integrate user-item interactions with user-user and item-item relations, capturing complex interactions through a comprehensive graph structure.

Afterward, these embeddings are combined during the embedding fusion stage, forming a latent user-item representation used for score computation by applying a weighted summation, concatenation, or a more complex method of combining user and item embeddings (Wang X. et al., 2019; He et al., 2020).

TABLE 1 This table summarizes key literature on GNN-based recommender systems, emphasizing beyond-accuracy metrics: diversity, serendipity, novelty, and fairness.

Beyond-accuracy goal	References and venue/ journal	Topic/ contribution	Model development stages utilized to tackle metric	Metric
Diversity	Isufi et al. (2021) Information Processing and Management	Neighbor-based mechanism	GC, PL, EF, TM	C, PD
	Ye et al. (2021) ACM RecSys conf.	Dynamic graph construction	EI, GC, TM	PD
	Yang L. et al. (2023) ACM WSDM conf.	Neighbor-based mechanisms	PL, EF, TM	C
	Zuo et al. (2023) MDPI Applied Sciences	Adversarial learning	GC, PL, TM	C, PD
	Ma et al. (2022) IEEE IJCNN conf.	Contrastive learning	EI, GC, PL, TM	C, PD
	Zheng et al. (2021) ACM Web Conf.	Neighbor-based mechanism, Adversarial learning	PL, TM	C, E, GC
	Xie et al. (2021) IEEE Trans. on Big Data	Heterogeneous GNNs	GC, PL, SC, TM	C, LTR, NOV
Serendipity/Novelty	Dhawan et al. (2022) Electronic Commerce Research and Applications	General GNN architecture enhancements	-	SRDP, NOV
	Sun et al. (2020) ACM SIGKDD conf.	General GNN architecture enhancements	GC, PL, EF, SC, TM	SRDP, NOV
	Zhao et al. (2022) ACM SIGIR conf.	Normalization techniques	PL	NOV
	Boo et al. (2023) ACM IUI conf.	Neighbor-based mechanisms	EI, EF, SC, TM	SRDP
Fairness	Xu et al. (2023) Information Sciences	Contrastive learning	GC, TM	ARP
	Li et al. (2019) ACM CIKM conf.	Multimodal feature learning	GC, PL, EF	LTR
	Liu et al. (2022a) Applied Soft Computing	Self-training mechanisms	PL, TM	GF
	Kim et al. (2022) ACM CIKM conf.	Neighbor-based mechanisms	PL, SC, TM	LTR
	Yang Y. et al. (2023) ACM Web Conf.	Contrastive learning	GC, PL, EF, TM	LTR
	Wu K. et al. (2022) ACM ASONAM conf.	Neighbor-based mechanisms	GC, PL, EF, TM	GF
	Gupta et al. (2019) ACM CIKM conf.	Long-tail recommendations	PL, SC, TM	ARP
	Liu and Zheng (2020) ACM RecSys conf.	Long-tail recommendations	DP, EF, SC, TM	C, LTR
	Liu et al. (2022b) Neural Computing and Applications	Neighbor-based mechanisms, Adversarial learning	GC, PL, TM	GF

Each entry specifies the paper’s publication venue/journal, a broad strategy categorization, and the model development stages the method utilizes or adapts to enhance the respective metric, including data preprocessing (DP), graph construction (GC), embedding initialization (EI), propagation layers (PL), embedding fusion (EF), score computation (SC), and training methodologies (TM). Additionally, the table highlights which concrete metrics were assessed: Coverage (C), Gini Coefficient (GC), Entropy (E), Pairwise dissimilarity (PD) for Diversity; Serendipity (SRDP); Novelty (NOV); Average Recommendation Popularity (ARP), Group Fairness (GF), and Long Tail Recommendation (LTR) for Fairness.

The score computation stage involves a scoring function to output a score for each user-item pair based on the fused embeddings. The scoring function can be as simple as a dot product between user and item embeddings, or it can be a more complex

function that takes into account additional factors ([Wang X. et al., 2019](#); [He et al., 2020](#)).

Finally, in the training methodologies stage, a suitable loss function is selected, and an optimization algorithm, typically a

variant of stochastic gradient descent, is used to update model parameters (Rendle et al., 2012; Fan et al., 2019a).

Understanding the unique strengths of each stage outlined in this section is essential, and a comparative evaluation can guide the selection of the most suitable approach for specific collaborative filtering scenarios, such as addressing the challenges associated with beyond-accuracy metrics. In Table 1, we provide a comprehensive overview of existing literature, aiding readers in navigating the diverse methodologies and findings discussed throughout this review.

4 Diversity in GNN-based recommender systems

4.1 Definition and importance of diversity

Diversity in recommender systems is a measure of the dissimilarity among the set of items recommended to a user. It prevents over-specialization and enhances user discovery, exposing users to a broader range of items and potentially increasing satisfaction and engagement with the system (Kunaver and Požrl, 2017; Duricic et al., 2021). Diversity can be intra-list, referring to variety within a single recommendation list, or inter-list, concerning variety across different users' lists (Kaminskas and Bridge, 2016). When items are categorized, diversity also entails ensuring a balanced representation of different categories in the recommendations.

Common metrics for measuring diversity include Item Coverage, calculated as the ratio of unique items recommended to the total items in the catalog. The Gini Coefficient reflects recommendation inequality and is given by:

$$\text{Gini Coefficient} = 1 - \sum_{i=1}^n P_i^2 \quad (1)$$

where P_i is the proportion of recommendations for item i . Entropy measures unpredictability or randomness in recommendations and is computed as:

$$\text{Entropy} = - \sum_{i=1}^n P_i \log P_i \quad (2)$$

with P_i as the probability of item i being recommended (Zheng et al., 2021). Another important metric, Pairwise Dissimilarity, quantifies the average dissimilarity between all pairs of items in a recommendation list (Chen et al., 2018). It is calculated using the formula:

$$\text{Pairwise Dissimilarity} = \frac{2}{N(N-1)} \sum_{i=1}^{N-1} \sum_{j=i+1}^N d(i,j) \quad (3)$$

where N is the number of items in the recommendation list, and $d(i,j)$ represents the measure of dissimilarity between item i and item j .

4.2 Review of recent developments in improving accuracy-diversity trade-off

Several approaches have emerged recently to tackle recommendation diversity using graph neural networks (GNNs). These methods can be broadly categorized based on the specific mechanisms or strategies they employ:

- **Neighbor-based mechanisms¹**: An approach introduced by Isufi et al. (2021) combines nearest neighbors (NN) and furthest neighbors (FN) with a joint convolutional framework. The *DGRec* method diversifies embedding generation through submodular neighbor selection, layer attention, and loss reweighting (Yang L. et al., 2023). Additionally, *DGCN* model leverages graph convolutional networks for capturing collaborative effects in the user-item bipartite graph, ensuring diverse recommendations through rebalanced neighbor discovery (Zheng et al., 2021).
- **Dynamic graph construction²**: *DDGraph* approach involves dynamically constructing a user-item graph to capture both user-item interactions and non-interactions, and then applying a novel candidate item selection operator to choose items from different sub-regions based on distance metrics (Ye et al., 2021).
- **Adversarial learning³**: To improve the accuracy-diversity trade-off in tag-aware systems, the *DTGCF* model utilizes personalized category-boosted negative sampling, adversarial learning for category-free embeddings, and specialized regularization techniques (Zuo et al., 2023). Furthermore, the above-mentioned *DGCN* model also employs adversarial learning to make item representations more category-independent.
- **Contrastive learning⁴**: The Contrastive Co-training (*CCT*) method by Ma et al. (2022) employs an iterative pipeline that augments recommendation and contrastive graph views with pseudo edges, leveraging diversified contrastive learning to address popularity and category biases in recommendations.
- **Heterogeneous graph neural networks⁵**: The *GraphDR* approach by Xie et al. (2021) utilizes a heterogeneous graph

1 Neighbor-based mechanisms aggregate and propagate information from neighboring nodes (users or items) to enhance the representation of a target node, capturing intricate relational patterns for improved recommendations (Wu S. et al., 2022).

2 Dynamic graph construction involves continuously updating and evolving the graph structure to incorporate new interactions and/or entities (Skarding et al., 2021).

3 Adversarial examples in recommender systems, as a form of data augmentation, bolster data diversity for improved generalization, counteract inherent biases, and ensure fair node representation in GNNs for fairer recommendations (Deldjoo et al., 2021).

4 Contrastive learning pushes similar item or user embeddings closer and dissimilar ones apart to enhance recommendation quality (Liu et al., 2021).

5 Heterogeneous graph neural networks process diverse types of nodes and edges, capturing complex relationships using a heterogeneous graph as input (Wu S. et al., 2022).

neural network, capturing diverse interactions and prioritizing diversity in the matching module.

Each of these methods offers a unique approach to the accuracy-diversity challenge. While all aim to improve the trade-off, their strategies vary, highlighting the multifaceted nature of the challenge at hand.

5 Serendipity in GNN-based recommender systems

5.1 Definition and importance of serendipity and novelty

Serendipity and novelty are key aspects of recommender systems, essential for enhancing user discovery and engagement. These concepts are closely related and often evaluated together, as they complement each other by simultaneously assessing the unexpectedness and unfamiliarity of recommendations (Sun et al., 2020; Dhawan et al., 2022). Serendipity, indicating the unexpected nature of recommendations, encourages users to explore beyond their usual preferences and stimulates curiosity. The Serendipity Score, is a commonly used metric to assess this quality (Silveira et al., 2019):

$$\text{Serendipity} = \frac{1}{|U|} \sum_{u \in U} \left(\frac{1}{|I_k(u)|} \sum_{i \in I_k(u)} \max(P_i(u) - P_i(U), 0) \cdot \text{rel}_i(u) \right) \quad (4)$$

where $|U|$ denotes the cardinality of the user set, $I_k(u)$ the set of top k recommendations for user u , and $\text{rel}_i(u)$ the relevance of item i to user u . The difference $P_i(u) - P_i(U)$ captures the preference deviation of user u for item i from the mean user preference.

Conversely, novelty is concerned with how the recommended items are new or unfamiliar to a user, as quantified by the Novelty Score (Zhou et al., 2010):

$$\text{Novelty} = \frac{1}{|U|} \sum_{u \in U} \left(\sum_{i \in I_u(k)} \frac{-\log_2 D(i)}{|I_u(k)|} \right) \quad (5)$$

Here, $D(i)$ signifies the popularity of item i , inversely related to novelty. This measure ensures that recommendations are not only serendipitous but also novel, thus preventing recommendation over-specialization, enhancing user exploration and engagement (Kaminskas and Bridge, 2016).

5.2 Review of recent developments in promoting serendipity and novelty

Recent advancements in GNN-based recommender systems have shown promising results in promoting serendipity and novelty, although notably fewer efforts have been directed toward balancing the accuracy-serendipity and accuracy-novelty trade-offs in comparison to the accuracy-diversity trade-off. In our exploration, we identified several studies addressing these efforts

and have categorized them based on the primary theme of their contribution:

- **Neighbor-based mechanisms:** Approach proposed by Boo et al. (2023) enhances session-based recommendations by incorporating serendipitous session embeddings, leveraging session data and user preferences to amplify global embedding effects, enabling users to control explore-exploit tradeoffs.
- **Normalization techniques⁶:** Zhao et al. (2022) proposed *r-AdjNorm*, a simple and effective GNN improvement that can improve the accuracy-novelty trade-off by controlling the normalization strength in the neighborhood aggregation process.
- **General GNN architecture enhancements⁷:** Similarly to the popular *LightGCN* approach by He et al. (2020), the *ImprovedGCN* model by Dhawan et al. (2022) adapts and simplifies the graph convolution process in GCNs for item recommendation, inadvertently boosting serendipity. On the other hand, the *BGCF* framework by Sun et al. (2020), designed for diverse and accurate recommendations, also boosts serendipity and novelty through its joint training approach. These GNN-based models, while focusing on accuracy, inadvertently elevate recommendation serendipity and/or novelty.

These studies collectively demonstrate the potential of GNNs in enhancing the serendipity and novelty of recommender systems, while also highlighting the need for further research to address existing challenges.

6 Fairness in GNN-based recommender systems

6.1 Definition and importance of fairness

Fairness in recommender systems ensures no bias toward certain users or items. It can be divided into user fairness, which avoids algorithmic bias among users or demographics, and item fairness, which ensures equal exposure for items, countering popularity bias (Leonhardt et al., 2018; Kowald et al., 2020; Lex et al., 2020; Abdollahpouri et al., 2021; Lacic et al., 2022). Fairness helps to mitigate bias, supports diversity, and boosts user satisfaction. In GNN-based systems, which can amplify bias, fairness is crucial for balanced recommendations and optimal performance (Ekstrand et al., 2018; Chizari et al., 2022; Chen et al., 2023; Gao et al., 2023).

Key metrics for evaluating fairness include Average Recommendation Popularity (ARP) and Group Fairness (GF) (Yin

⁶ Normalization techniques in GNN-based recommender systems stabilize and scale node features or edge weights, ensuring consistent and improved model convergence and recommendation quality (Gupta et al., 2019).

⁷ We refer to general GNN architecture enhancements in recommender systems as the advancements in architectures, aggregators, or training procedures that better capture graph structures for improved recommendation accuracy.

et al., 2012; Fu et al., 2020). ARP, as defined below, assesses the tendency toward recommending popular items:

$$\text{ARP} = \frac{1}{|U|} \sum_{u \in U} \frac{1}{|I_u|} \sum_{i \in I_u} D(i)$$

where $D(i)$ is the popularity of item i , typically defined by the number of interactions or ratings it has received across the user base. On the other hand, GF measures the fairness of recommendations across different user groups:

$$\text{GF} = \left| \frac{1}{|S_0|} \sum_{u \in S_0} \mathcal{T}(Q_u) - \frac{1}{|S_1|} \sum_{u \in S_1} \mathcal{T}(Q_u) \right|$$

Here, S_0 and S_1 represent different user groups, Q_u denotes the list of items recommended to user u , and $\mathcal{T}(Q_u)$ is a metric that scores the quality of recommendations for user u . Lower GF values signify a fairer distribution of recommendations between the groups.

Beyond these metrics, focusing on the assessment of long-tail item recommendations also plays a role in ensuring that the system's suggestions are not limited to well-known or popular items, thus fostering a more inclusive recommendation environment.

6.2 Review of recent developments in promoting fairness

In the evolving landscape of GNN-based recommender systems, the pursuit of user and item fairness has become a prominent topic. Recent advancements can be broadly categorized based on the thematic emphasis of their contributions:

- **Neighbor-based mechanisms:** The *Navip* method debiases the neighbor aggregation process in GNNs using “neighbor aggregation via inverse propensity”, focusing on user fairness (Kim et al., 2022). Additionally, the *UGRec* framework by Liu et al. (2022b) employs an information aggregation component and a multihop mechanism to aggregate information from users' higher-order neighbors, ensuring user fairness by considering male and female discrimination. The *SKIPHOP* approach focuses on user fairness by introducing an approach that captures both direct user-item interactions and latent knowledge graph interests, capturing both first-order and second-order proximity. Using fairness for regularization, it ensures balanced recommendations for users with similar profiles (Wu K. et al., 2022).
- **Multimodal feature learning⁸:** The method proposed by Li et al. (2019) fuses hashtag embeddings with multi-modal features, considering interactions among users, micro-videos, and hashtags.

⁸ Multimodal feature learning integrates diverse data sources, like text, images, and graphs, into unified embeddings to enrich recommendation context and accuracy (Zhou et al., 2023).

- **Adversarial learning:** The *UGRec* model additionally incorporates adversarial learning to eliminate gender-specific features while preserving common features.
- **Contrastive learning:** The *DCRec* model by Yang Y. et al. (2023) leverages debiased contrastive learning to counteract popularity bias and addressing the challenge of disentangling user conformity from genuine interest, focusing on user fairness. The *TAGCL* framework also capitalizes on the contrastive learning paradigm, ensuring item fairness by reducing biases in social tagging systems (Xu et al., 2023).
- **Long-tail recommendations⁹:** The *TailNet* architecture is designed to enhance long-tail recommendation performance. It classifies items into short-head and long-tail based on click frequency and integrates a unique preference mechanism to balance between recommending niche items for serendipity and maintaining overall accuracy (Liu and Zheng, 2020). The *NISER* method by Gupta et al. (2019) addresses the long-tail issue by focusing on popularity bias in session-based recommendation systems. It aims to ensure item fairness by normalizing item and session representations, thereby improving recommendations, especially for less popular items. Additionally, the above-mentioned approach by Li et al. (2019) also focuses on long-tail recommendations.
- **Self-training mechanisms¹⁰:** The *Self-Fair* approach by Liu et al. (2022a) employs a self-training mechanism using unlabeled data with the goal of improving user fairness in recommendations for users of different genders. By iteratively refining predictions as pseudo-labels and incorporating fairness constraints, the model balances accuracy and fairness without relying heavily on labeled data.

In the broader context of graph neural networks, researchers have also tackled fairness in non-recommender systems tasks, such as classification (Dai and Wang, 2021; Ma et al., 2021; Dong et al., 2022; Zhang et al., 2022). Their insights provide valuable lessons for future development of fair recommender systems.

7 Discussion and future directions

In this paper, we present a review of the literature on diversity, serendipity/novelty, and fairness in GNN-based recommender systems, with a focus on optimizing for beyond-accuracy metrics. Throughout our analysis, we have explored various aspects of model development and discussed recent advancements in addressing these dimensions.

To further advance the field and guide future research, we have formulated three key questions:

Q1: What are the practical challenges in optimizing GNN-based recommender systems for beyond-accuracy metrics?

GNNs are able to capture complex relationships within graph structures. However, this sophistication can lead to overfitting,

⁹ Long-tail recommendations focus on suggesting less popular or niche items (Kowald et al., 2020).

¹⁰ Self-training mechanisms leverage unlabeled data by iteratively predicting and refining labels, enhancing the model's performance with augmented training data. (Yu et al., 2023).

especially when prioritizing accuracy (Fu et al., 2023). Data sparsity and the need for auxiliary data, such as demographic information, challenge the optimization of high-quality node representations, introducing biases (Dhawan et al., 2022). An overemphasis on past preferences can limit novel discoveries (Dhawan et al., 2022), and while addressing popularity bias is essential, it might inadvertently inject noise, reducing accuracy (Liu and Zheng, 2020). Balancing diverse objectives, like fairness, accuracy, and diversity, is nuanced, especially when optimizing one can compromise another (Liu et al., 2022b). These challenges emphasize the need for focused research on effective modeling of GNN-based recommender systems focused on beyond-accuracy optimization.

Q2: *Which model development stages of GNN-based recommender systems have seen the most innovation for tackling beyond-accuracy optimization, and which stages have been underutilized?*

By conducting a thorough analysis of the reviewed papers (see Table 1), we have observed that the graph construction, propagation layer, and training methodologies have seen significant innovation in GNN-based recommender systems. This includes advanced graph construction methods, innovative graph convolution operations, and unique training methodologies. However, stages like embedding initialization, embedding fusion, and score computation are relatively underutilized. These stages could offer potential avenues for future research and could provide novel ways to balance accuracy, fairness, diversity, and serendipity in recommendations.

Q3: *What are potentially unexplored areas of beyond-accuracy optimization in GNN-based recommender systems?*

A less explored aspect in GNN-based recommender systems is personalized diversity, which modifies the diversity in recommendations to match individual user preferences. Users favoring more diversity get more diverse recommendations, whereas those liking less diversity get less diverse ones (Eskandarian et al., 2017). This concept of personalized diversity, currently under-researched in GNN-based systems, hints at an intriguing future research direction. It can also relate to personalized serendipity or novelty, tailoring unexpected or novel recommendations to user preferences. Thus, incorporating personalized diversity, serendipity, and novelty in GNN-based systems could enrich beyond-accuracy optimization.

Overall, this review aims to help researchers and practitioners gain a deeper understanding of the multifaceted issues and potential avenues for future research in optimizing GNN-based recommender systems beyond traditional accuracy-centric approaches. By addressing the practical challenges, identifying

underutilized model development stages, and highlighting unexplored areas of optimization, we hope to contribute to the development of more robust, diverse, serendipitous, and fair recommender systems that cater to the evolving needs and expectations of users.

Author contributions

TD: literature analysis, conceptualization, and writing. ELA: conceptualization and writing. ELe and DK: conceptualization, writing, and supervision. All authors contributed to the article and approved the submitted version.

Funding

This work was supported by the “DDIA” COMET Module within the COMET—Competence Centers for Excellent Technologies Programme, funded by the Austrian Federal Ministry for Transport, Innovation and Technology (bmvit), the Austrian Federal Ministry for Digital and Economic Affairs (bmdw), FFG, SFG, and partners from industry and academia. The COMET Programme is managed by FFG. This research received support by the TU Graz Open Access Publishing Fund. Additional credit is given to OpenAI for the generative AI models, GPT-4, and ChatGPT, used in this work for text summarization and sentence rephrasing. Verification of accuracy and originality was performed for all content generated by these tools.

Conflict of interest

ELA was employed by Infobip.

The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

References

- Abdollahpouri, H., Adomavicius, G., Burke, R., Guy, I., Jannach, D., Kamishima, T., et al. (2019). Beyond personalization: research directions in multistakeholder recommendation. *arXiv*. doi: 10.48550/arXiv.1905.01986
- Abdollahpouri, H., Mansoury, M., Burke, R., Mobasher, B., and Malthouse, E. (2021). “User-centered evaluation of popularity bias in recommender systems,” in *Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization*, 119–129.
- Aggarwal, C. C. (2016). *Recommender Systems, Volume 1*. Cham: Springer.
- Boo, S., Kim, S., and Lee, S. (2023). “Serendipity into session-based recommendation: Focusing on unexpectedness, relevance, and usefulness of recommendations,” in *Companion Proceedings of the 28th International Conference on Intelligent User Interfaces*, 83–86.
- Castells, P., Hurley, N., and Vargas, S. (2021). “Novelty and diversity in recommender systems,” in *Recommender Systems Handbook*. Cham: Springer, 603–646.

- Chen, J., Wu, J., Chen, J., Xin, X., Li, Y., and He, X. (2023). How graph convolutions amplify popularity bias for recommendation? *arXiv*. doi: 10.48550/arXiv.2305.14886
- Chen, L., Zhang, G., and Zhou, E. (2018). "Fast greedy map inference for determinantal point process to improve recommendation diversity," in *Advances in Neural Information Processing Systems*, 31.
- Chizari, N., Shoebani, N., and Moreno-García, M. N. (2022). A comparative analysis of bias amplification in graph neural network approaches for recommender systems. *Electronics* 11, 3301. doi: 10.3390/electronics11203301
- Dai, E., Zhao, T., Zhu, H., Xu, J., Guo, Z., Liu, H., et al. (2022). A comprehensive survey on trustworthy graph neural networks: privacy, robustness, fairness, and explainability. *arXiv*. doi: 10.48550/arXiv.2204.08570
- Dai, E., and Wang, S. (2021). "Say no to the discrimination: Learning fair graph neural networks with limited sensitive attribute information," in *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*, 680–688.
- Deldjoo, Y., Jannach, D., Bellogin, A., Difonzo, A., and Zanzonelli, D. (2023). Fairness in recommender systems: research landscape and future directions. *User Modeling and User-Adapted Interaction*, 1–50.
- Deldjoo, Y., Noia, T. D., and Merra, F. A. (2021). A survey on adversarial recommender systems: from attack/defense strategies to generative adversarial networks. *ACM Comp. Surv. (CSUR)* 54, 1–38. doi: 10.1145/3439729
- Dhawan, S., Singh, K., Rabaea, A., and Batra, A. (2022). Improvedgcn: an efficient and accurate recommendation system employing lightweight graph convolutional networks in social media. *Electron. Commer. Res. Appl.* 5, 5:101191. doi: 10.1016/j.elerap.2022.101191
- Dong, Y., Liu, N., Jalaian, B., and Li, J. (2022). "Edits: Modeling and mitigating data bias for graph neural networks," in *Proceedings of the ACM Web Conference 2022*, 1259–1269.
- Dong, Y., Ma, J., Wang, S., Chen, C., and Li, J. (2023). "Fairness in graph mining: a survey," in *IEEE Transactions on Knowledge and Data Engineering (IEEE)*.
- Duricic, T., Hussain, H., Lacic, E., Kowald, D., Helic, D., and Lex, E. (2020). "Empirical comparison of graph embeddings for trust-based collaborative filtering," in *Foundations of Intelligent Systems: 25th International Symposium, ISMIS 2020*. Graz, Austria: Springer, 181–191.
- Duricic, T., Kowald, D., Schedl, M., and Lex, E. (2021). "My friends also prefer diverse music: homophily and link prediction with user preferences for mainstream, novelty, and diversity in music," in *Proceedings of the 2021 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, 447–454.
- Duricic, T., Lacic, E., Kowald, D., and Lex, E. (2018). "Trust-based collaborative filtering: Tackling the cold start problem using regular equivalence," in *Proceedings of the 12th ACM Conference on Recommender Systems*, 446–450.
- Ekstrand, M. D., Tian, M., Kazi, M. R. I., Mehrpouyan, H., and Kluver, D. (2018). "Exploring author gender in book rating and recommendation," in *Proceedings of the 12th ACM conference on recommender systems*, 242–250.
- Eskandarian, F., Mobasher, B., and Burke, R. (2017). "A clustering approach for personalizing diversity in collaborative recommender systems," in *Proceedings of the 25th Conference on User Modeling, Adaptation and Personalization*, 280–284.
- Fan, W., Derr, T., Ma, Y., Wang, J., Tang, J., and Li, Q. (2019a). "Deep adversarial social recommendation," in *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, 1351–1357.
- Fan, W., Ma, Y., Li, Q., He, Y., Zhao, E., Tang, J., et al. (2019b). "Graph neural networks for social recommendation," in *The World Wide Web Conference*, 417–426.
- Fu, Z., Niu, X., and Maher, M. L. (2023). "Deep learning models for serendipity recommendations: A survey and new perspectives," in *ACM Computing Surveys*. New York: Association for Computing Machinery.
- Fu, Z., Xian, Y., Gao, R., Zhao, J., Huang, Q., Ge, Y., et al. (2020). "Fairness-aware explainable recommendation over knowledge graphs," in *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 69–78.
- Gao, C., Zheng, Y., Li, N., Li, Y., Qin, Y., Piao, J., et al. (2023). A survey of graph neural networks for recommender systems: challenges, methods, and directions. *ACM Trans. Recommend. Syst.* 1, 1–51. doi: 10.1145/3568022
- Gupta, P., Garg, D., Malhotra, P., Vig, L., and Shroff, G. M. (2019). Niser: normalized item and session representations with graph neural networks. *arXiv*. doi: 10.48550/arXiv.1909.04276
- Hamilton, W. L. (2020). Graph representation learning. *Synth. Lect. Artif. Intell. Mach. Learn.* 14, 1–159. doi: 10.1007/978-3-031-01588-5
- He, X., Deng, K., Wang, X., Li, Y., Zhang, Y., and Wang, M. (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.
- Isufi, E., Pocchiari, M., and Hanjalic, A. (2021). Accuracy-diversity trade-off in recommender systems via graph convolutions. *Inform. Proc. Manage.* 58, 102459. doi: 10.1016/j.ipm.2020.102459
- Kaminskas, M., and Bridge, D. (2016). Diversity, serendipity, novelty, and coverage: a survey and empirical analysis of beyond-accuracy objectives in recommender systems. *ACM Trans. Interact. Intellig. Syst. (TiS)* 7, 1–42. doi: 10.1145/2926720
- Kim, M., Oh, J., Do, J., and Lee, S. (2022). "Debiasing neighbor aggregation for graph neural network in recommender systems," in *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, 4128–4132.
- Koren, Y., Rendle, S., and Bell, R. (2021). "Advances in collaborative filtering," in *Recommender Systems Handbook* (Springer), 91–142.
- Kotkov, D., Veijalainen, J., and Wang, S. (2020). How does serendipity affect diversity in recommender systems? a serendipity-oriented greedy algorithm. *Computing* 102, 393–411. doi: 10.1007/s00607-018-0687-5
- Kowald, D., Schedl, M., and Lex, E. (2020). "The unfairness of popularity bias in music recommendation: a reproducibility study," in *Advances in Information Retrieval: 42nd European Conference on IR Research, ECIR 2020* Lisbon, Portugal: Springer, 35–42.
- Kunaver, M., and Požrl, T. (2017). Diversity in recommender systems—a survey. *Knowl.-Based Syst.* 123, 154–162. doi: 10.1016/j.knsys.2017.02.009
- Lacic, E., Fadjljev, L., Weissenboeck, F., Lindstaedt, S., and Kowald, D. (2022). "What drives readership? an online study on user interface types and popularity bias mitigation in news article recommendations," in *Advances in Information Retrieval: 44th European Conference on IR Research, ECIR 2022*. Stavanger, Norway: Springer, 172–179.
- Lacic, E., Kowald, D., Eberhard, L., Trattner, C., Parra, D., and Marinho, L. B. (2015a). "Utilizing online social network and location-based data to recommend products and categories in online marketplaces," in *Mining, Modeling, and Recommending "Things" in Social Media: 4th International Workshops, MUSE 2013*. Paris, France: Springer, 96–115.
- Lacic, E., Kowald, D., Traub, M., Luzhnica, G., Simon, J. P., and Lex, E. (2015b). "Tackling cold-start users in recommender systems with indoor positioning systems," in *Proc. of ACM RecSys 15*, 2. New York: ACM.
- Leonhardt, J., Anand, A., and Khosla, M. (2018). "User fairness in recommender systems," in *Companion Proceedings of the The Web Conference 2018*, 101–102.
- Lex, E., Kowald, D., and Schedl, M. (2020). Modeling popularity and temporal drift of music genre preferences. *Trans. Int. Soc. Music. Inf. Retr.* 3, 17–30. doi: 10.5334/tismir.39
- Li, M., Gan, T., Liu, M., Cheng, Z., Yin, J., and Nie, L. (2019). "Long-tail hashtag recommendation for micro-videos with graph convolutional network," in *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, 509–518.
- Li, Y., Chen, H., Xu, S., Ge, Y., Tan, J., Liu, S., et al. (2022). Fairness in recommendation: a survey. *arXiv*. doi: 10.48550/arXiv.2205.13619
- Liu, H., Lin, H., Fan, W., Ren, Y., Xu, B., Zhang, X., et al. (2022a). Self-supervised learning for fair recommender systems. *Appl. Soft Comput.* 125, 109126. doi: 10.1016/j.asoc.2022.109126
- Liu, H., Wang, Y., Lin, H., Xu, B., and Zhao, N. (2022b). Mitigating sensitive data exposure with adversarial learning for fairness recommendation systems. *Neural Comp. Appl.* 34, 18097–18111. doi: 10.1007/s00521-022-07373-4
- Liu, S., Ounis, I., Macdonald, C., and Meng, Z. (2020). "A heterogeneous graph neural model for cold-start recommendation," in *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2029–2032.
- Liu, S., and Zheng, Y. (2020). "Long-tail session-based recommendation," in *Proceedings of the 14th ACM Conference on Recommender Systems*. New York: ACM, 509–514.
- Liu, Z., Ma, Y., Ouyang, Y., and Xiong, Z. (2021). Contrastive learning for recommender system. *arXiv*. doi: 10.48550/arXiv.2101.01317
- Ma, J., Deng, J., and Mei, Q. (2021). Subgroup generalization and fairness of graph neural networks. *Adv. Neural Inf. Process. Syst.* 34, 1048–1061.
- Ma, X., Hu, Q., Gao, Z., and AbdelHady, M. (2022). "Contrastive co-training for diversified recommendation," in *2022 International Joint Conference on Neural Networks (IJCNN)*. Padua: IEEE, 1–9.
- Mansoury, M., Abdollahpouri, H., Pechenizkiy, M., Mobasher, B., and Burke, R. (2020). "Fairmatch: A graph-based approach for improving aggregate diversity in recommender systems," in *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, New York: ACM, 154–162.
- Pu, P., Chen, L., and Hu, R. (2012). Evaluating recommender systems from the user's perspective: survey of the state of the art. *User Model. User-adapt. Interact.* 22, 317–355. doi: 10.1007/s11257-011-9115-7
- Rendle, S., Freudenthaler, C., Gantner, Z., and Schmidt-Thieme, L. (2012). "BPR: Bayesian personalized ranking from implicit feedback," in *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, 452–461.
- Silveira, T., Zhang, M., Lin, X., Liu, Y., and Ma, S. (2019). How good your recommender system is? A survey on evaluations in recommendation. *Int. J. Mach. Learn. Cybern.* 10, 813–831. doi: 10.1007/s13042-017-0762-9

- Skarding, J., Gabrys, B., and Musial, K. (2021). Foundations and modeling of dynamic networks using dynamic graph neural networks: a survey. *IEEE Access* 9, 79143–79168. doi: 10.1109/ACCESS.2021.3082932
- Sun, J., Guo, W., Zhang, D., Zhang, Y., Regol, F., Hu, Y., et al. (2020). “A framework for recommending accurate and diverse items using bayesian graph convolutional neural networks,” in *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2030–2039. doi: 10.1145/3394486.3403254
- Wang, C., Liang, Y., Liu, Z., Zhang, T., and Philip, S. Y. (2021). “Pre-training graph neural network for cross domain recommendation” in *2021 IEEE Third International Conference on Cognitive Machine Intelligence (CogMI)* (IEEE), 140–145.
- Wang, H., Zhang, F., Zhang, M., Leskovec, J., Zhao, M., Li, W., et al. (2019). “Knowledge-aware graph neural networks with label smoothness regularization for recommender systems” in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 968–977.
- Wang, X., He, X., Wang, M., Feng, F., and Chua, T.-S. (2019). “Neural graph collaborative filtering,” in *Proceedings of the 42nd international ACM SIGIR Conference on Research and development in Information Retrieval*, 165–174. doi: 10.1145/3331184.3331267
- Wang, Y., Ma, W., Zhang, M., Liu, Y., and Ma, S. (2023a). A survey on the fairness of recommender systems. *ACM Trans. Inform. Syst.* 41, 1–43. doi: 10.1145/3547333
- Wang, Y., Zhao, Y., Zhang, Y., and Derr, T. (2023b). Collaboration-aware graph convolutional network for recommender systems. *Proc. ACM Web Conf.* 2023, 91–101. doi: 10.1145/3543507.3583229
- Wu, K., Erickson, J., Wang, W. H., and Ning, Y. (2022). “Equipping recommender systems with individual fairness via second-order proximity embedding,” in *2022 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*. Istanbul: IEEE, 171–175.
- Wu, S., Sun, F., Zhang, W., Xie, X., and Cui, B. (2022). Graph neural networks in recommender systems: a survey. *ACM Comp. Surv.* 55, 1–37. doi: 10.1145/3535101
- Xie, R., Liu, Q., Liu, S., Zhang, Z., Cui, P., Zhang, B., et al. (2021). Improving accuracy and diversity in matching of recommendation with diversified preference network. *IEEE Trans. Big Data* 8, 955–967. doi: 10.1109/TBDDATA.2021.3103263
- Xu, C., Zhang, Y., Chen, H., Dong, L., and Wang, W. (2023). A fairness-aware graph contrastive learning recommender framework for social tagging systems. *Inf. Sci.* 640, 119064. doi: 10.1016/j.ins.2023.119064
- Yang, L., Wang, S., Tao, Y., Sun, J., Liu, X., Yu, P. S., et al. (2023). “Dgrec: Graph neural network for recommendation with diversified embedding generation,” in *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*, 661–669.
- Yang, Y., Huang, C., Xia, L., Huang, C., Luo, D., and Lin, K. (2023). “Debiased contrastive learning for sequential recommendation,” in *Proceedings of the ACM Web Conference 2023*. New York: ACM, 1063–1073.
- Ye, R., Hou, Y., Lei, T., Zhang, Y., Zhang, Q., Guo, J., et al. (2021). “Dynamic graph construction for improving diversity of recommendation,” in *Proceedings of the 15th ACM Conference on Recommender Systems*, 651–655.
- Yin, H., Cui, B., Li, J., Yao, J., and Chen, C. (2012). “Challenging the long tail recommendation,” in *Proceedings of the VLDB Endowment* (VLDB Endowment), 896–907.
- Yu, J., Yin, H., Xia, X., Chen, T., Li, J., and Huang, Z. (2023). “Self-supervised learning for recommender systems: a survey,” in *IEEE Transactions on Knowledge and Data Engineering* (IEEE).
- Zhao, Y., Wang, Y., Liu, Y., Cheng, X., Aggarwal, C., and Derr, T. (2023). Fairness and diversity in recommender systems: a survey. *arXiv preprint arXiv:2307.04644*. doi: 10.48550/arXiv.2307.04644
- Zhang, H., Wu, B., Yuan, X., Pan, S., Tong, H., and Pei, J. (2022). Trustworthy graph neural networks: Aspects, methods and trends. *arXiv*. doi: 10.48550/arXiv.2205.07424
- Zhao, M., Wu, L., Liang, Y., Chen, L., Zhang, J., Deng, Q., et al. (2022). “Investigating accuracy-novelty performance for graph-based collaborative filtering,” in *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 50–59.
- Zheng, Y., Gao, C., Chen, L., Jin, D., and Li, Y. (2021). “Dgcn: Diversified recommendation with graph convolutional networks,” in *Proceedings of the Web Conference 2021*, 401–412. doi: 10.1145/3442381.3449835
- Zhou, H., Zhou, X., Zeng, Z., Zhang, L., and Shen, Z. (2023). A comprehensive survey on multimodal recommender systems: taxonomy, evaluation, and future directions. *arXiv*.
- Zhou, T., Kuscsik, Z., Liu, J.-G., Medo, M., Wakeling, J. R., and Zhang, Y.-C. (2010). Solving the apparent diversity-accuracy dilemma of recommender systems. *Proc. Nat. Acad. Sci.* 107, 4511–4515. doi: 10.1073/pnas.1000488107
- Zhou, Y., Zheng, H., Huang, X., Hao, S., Li, D., and Zhao, J. (2022). Graph neural networks: Taxonomy, advances, and trends. *ACM Trans. Intel. Syst. Technol. (TIST)* 13, 1–54. doi: 10.1145/3495161
- Zuo, Y., Liu, S., and Zhou, Y. (2023). Dtgcf: Diversified tag-aware recommendation with graph collaborative filtering. *Appl. Sci.* 13, 2945. doi: 10.3390/app13052945