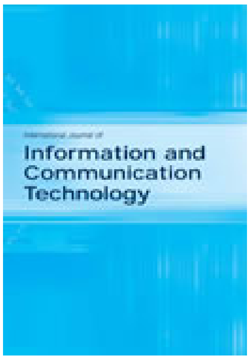

International Journal of Information and Communication Technology

ISSN online: 1741-8070 - ISSN print: 1466-6642
https://www.inderscience.com/ijict

An intelligent recommendation system for music therapy resources based on a knowledge graph

Ya Ni, Yingjin Shan

DOI: 10.1504/IJICT.2025.10074332

Article History:

Received:	26 July 2025
Last revised:	16 September 2025
Accepted:	18 September 2025
Published online:	13 November 2025

An intelligent recommendation system for music therapy resources based on a knowledge graph

Ya Ni*

School of Humanities and Law,
China University of Petroleum (East China),
Qingdao, 266580, China
Email: our471556@sina.com
*Corresponding author

Yingjin Shan

Academic Affairs Office,
China University of Petroleum (East China),
Qingdao, 266580, China
Email: syjupc@163.com

Abstract: To address the limitation of insufficient flexibility in current music therapy recommendation systems for capturing critical inter-entity relationships, this paper first uses the pre-trained language model BERT to enrich representation vectors of entities and relationships, and integrates the user-music therapy resource interaction graph into the knowledge graph, extracting collaborative information. The graph attention mechanism is introduced, allowing the weight of each neighbour node to be dynamically adjusted according to its relationship with the target node. Finally, the score that the user gives to the music therapy resource is predicted by calculating the dot product between the user representation and the music therapy resource representation, and the top N music therapy resources are recommended. Experimental results show that the AUC of the proposed model is improved by 4.85–21% compared to the baseline model, 21%, which can accurately recommend music therapy resources that match the user's preferences.

Keywords: music healing resource; intelligent recommender system; BERT model; knowledge graph; graph attention network.

Reference to this paper should be made as follows: Ni, Y. and Shan, Y. (2025) 'An intelligent recommendation system for music therapy resources based on a knowledge graph', *Int. J. Information and Communication Technology*, Vol. 26, No. 40, pp.18–33.

Biographical notes: Ya Ni received her Master's degree from Central Conservatory of Music, Beijing, China in 2014. She is currently working at the School of Humanities and Law in China University of Petroleum (East China). Her research interests include music education and violin performance.

Yingjin Shan received his Master's degree from Université Paris VI, Paris, France in 2013. He is currently working at Academic Affairs Office in China University of Petroleum (East China). His research interests include teaching and learning methods in higher education.

1 Introduction

Music healing, as an interdisciplinary intervention integrating art and psychology, has gradually become an important supplement to the mental health service system by virtue of its universality, non-traumatisation and emotional resonance characteristics (Langdon et al., 2018). However, the current market for music therapy resources has seen a surge in growth, and users often face selection dilemmas due to information overload. Conventional recommendation systems primarily depend on collaborative filtering or content-based approaches, but both have issues with cold start and sparse matrices (Assuncao et al., 2022). As a powerful tool for knowledge representation and management, knowledge graphs can clearly present entities, concepts, and their relationships in a visual way, providing an effective means for understanding and processing complex knowledge (Sakurai et al., 2022). In this knowledge network, each musical work is no longer an isolated entity, but is closely connected with related musical elements, therapeutic properties, and other relevant knowledge, thus providing a rich information basis for in-depth understanding and exploration of the internal patterns of music therapy resources (Bellini et al., 2024). Therefore, it is of great research significance to efficiently apply knowledge graphs to intelligent recommendation systems for music therapy resources.

Content-based recommendation approaches and collaborative filtering-based approaches constitute the main categories of conventional recommendation systems (Velankar and Kulkarni, 2022). Kathavate (2021) recommended music resources based on the content of music resources, combined with resource characteristics such as popularity and freshness. Xu (2020) integrated matrix factorisation with decision trees; this method can effectively rate music scores during the cold start of the recommendation system and recommend music that meets user needs. Li et al. (2022) suggested a learning framework based on stacked autoencoders to simultaneously infer the hidden representations of users and musical items, and used the learned representations to reduce the error in the training data. Mao et al. (2022) combined marginalised denoising autoencoders with matrix factorisation and generalised them. Jing et al. (2024) proposed a Bayesian binary neural network, which not only considers the auxiliary information of users and music resources, but also considers the weight ratio within them, enabling more robust predictions.

The above recommendation algorithms rely on historical interaction information between users and music. However, when a new user joins an application, the user has no interaction information with the items, which can lead to the cold start issue of the system. To cope with this, recent studies have introduced deep learning approaches to enhance recommender system performance. Schedl (2019) suggested a music resource recommendation algorithm in light of deep neural networks, which inputs user information and music information in the form of embeddings into the deep neural network, alleviating the data sparsity and cold start problem to some extent. Hou (2024) established a music recommendation system based on convolutional neural networks (CNN), performed sentiment analysis on the singers selected by users, and recommended a music list that meets the comprehensive sentiment of users, achieving satisfactory recommendation results. The hybrid attention mechanism independent recurrent neural network music therapy resource recommendation algorithm (AIRNN) proposed by Han et al. (2024) preprocessed audio data using scattering transform and introduced an attention mechanism, significantly improving the accuracy of recommendations. The hierarchical attention-based recommendation model (HARM), introduced by Lu (2022), leverages multi-layer attention mechanisms to effectively capture users' nuanced preferences for multi-domain song features. Additionally, it discerns the varying influence of different historical behaviours on user preference formation.

Despite their advantages, deep learning recommendation models remain dependent on substantial user-item interaction data for effective training. When the recommendation system starts or new users or items are added, due to the lack of historical behaviour data, the cold start problem occurs, making it impossible to accurately recommend items to users. Introducing knowledge graphs into the recommendation system can provide rich semantic information, semantically link various information, and enable the recommendation approach to more accurately comprehend user demands and inclinations, thereby improving the recommendation effect. Cui et al. (2023) put forward a deep knowledge-aware music resource recommendation model, using the knowledge graph as side information, and used CNN to fuse the embedding representations of words in music and related entities in the knowledge graph. They designed an attention module to consider changes in user interests, and dynamically adjusted weights according to user interaction history. Liu et al. (2024) employed knowledge graph convolutional network (KGCN) architecture to enhance music recommendation tasks, and by leveraging KGCN to model high-order structural relationships and semantic item representations, the approach effectively captures users' latent long-term preferences, they improved the effectiveness and interpretability of recommendations. Bai and Zhang (2025) proposed an end-to-end graph neural network recommendation model driven by knowledge graphs. By stacking multiple layers of neural networks to obtain high-order neighbour information of nodes, the model generates final recommendations by fusing user and music embedding representations through an attention-based scoring mechanism, improving the accuracy of recommendation results.

As analysed from the current research status, the knowledge graph, as a structured knowledge representation method, can model relationships between different entities and enhance the recommendation system's input data with higher-dimensional feature representations. Naturally integrating the feature representations learned from knowledge graph into music recommendation can better capture users' interests and behaviour patterns, thus improving the accuracy of music recommendation. This paper aims to build a smart recommendation system for music therapy resources based on the knowledge graph, aiming to cope with the issue of insufficient flexibility in learning important relationships between entities in current research, which leads to low recommendation accuracy. Within the entity-relation embedding learning component, the BERT model is used to encode sentences to study the representation vectors of entities and relations. The entity and relation representation vectors are transformed into fixed-dimensional entity and relation representation vectors by a multi-layer perceptron (MLP). The user-music therapy resource interaction graph is integrated into the knowledge graph to obtain a collaborative knowledge graph. A computationally efficient graph convolution operator extracts collaborative filtering patterns. Moreover, a graph attention mechanism is introduced during the knowledge propagation process to dynamically assign weights to neighbour nodes based on the relationships between nodes. In this way, the important relationships between entities can be learned more flexibly, thereby further improving the accuracy of recommendations. Finally, the user's preference score for the music therapy resource is computed via the inner product of their latent representation vector with the resource's embedding. The top N music therapy resources based on the score are recommended. Experimental outcome indicates that the AUC of the suggested model is 0.9894, which is at least 4.85% higher than that of the baseline model. It enables the delivery of accurate music therapy resource recommendations to users and provides technical support to facilitate the efficient exploitation of music therapy resources.

2 Relevant technologies

2.1 Knowledge graph

A knowledge graph represents a structured semantic network where heterogeneous entities, their attributes, and inter-entity relations are formally modelled as a directed, labelled graph (Zhang et al., 2024). Within graph-based representations, a knowledge graph formally models real-world subjects and objects as entities, while their interactions constitute edges, and inter-entity relationships are modelled as directed edges with semantic labels, preserving the relational structure of the source domain. In this sense, a knowledge graph also serves as a symbolic representation of real-world relationships, establishing a disambiguated ontological projection of real-world semantic networks. It illustrates real-world connections through simple nodes and edges. A knowledge graph constitutes a typed heterogeneous graph where nodes represent entities and edges model their semantic relationships, represented by $G = \{V, E\}$, $V = \{P\}$ denotes the collection of entities within the knowledge graph, while E stands for edge set.

Within knowledge graph-powered recommendation frameworks, item attributes can be regarded as a node that is a single jump in the knowledge graph, and it has a direct connection with the entity. A minimal knowledge graph can be constructed from item attributes alone. Knowledge graphs are typically encoded as resource description framework (RDF) triples, adopting either (subject, predicate, object) or (entity, property, value) schemas for structured knowledge representation. An entity represents a unique, distinguishable thing that exists independently. Triples are formalised as $G = \{(h, r, t)|h, t \in \varepsilon, r \in R\}$, where ε is the set of entities, R is the set of relations, h is the head entity, t is the tail entity, and r is the relationship between them. A knowledge graph can not only add more semantic information to the recommendation algorithm but also compensating for existing methods' interpretability shortcomings to generate more explainable recommendations (Wang et al., 2024).

2.2 Graph neural network

Graph neural networks (GNN) are deep learning models used to process graph data structures. Unlike traditional neural network models, GNNs can capture both local and global information of nodes, thus improving the understanding and processing ability of graph structures. This capability has led to widespread applications of GNNs in the domain of artificial intelligence. The basic model of GNN chiefly makes up of a propagation unit and an aggregation unit. Early GNNs had some limitations and were difficult to adapt to heterogeneous and dynamically changing graph structures (Veličković, 2023). For example, due to the limitations of traditional convolutional filters on non-Euclidean structured data, they could not be well applied to graph structured data. To overcome this problem, researchers developed GCN. GCN represents graph data as a graph adjacency matrix, using the adjacency matrix to extract relationships between nodes, and uses convolution operations to extract node features (Bhatti et al., 2023).

Although GCN updates node representations by aggregating first-order neighbour information, it can effectively capture local structural features in the graph (Yue et al., 2022). However, GCN employs isotropic aggregation where all neighbour nodes contribute equally, neglecting their relative importance to the target node. The graph attention network (GAT) introduces an attention mechanism, which can adaptively assign different weights to neighbours of each node (Vrahatis et al., 2024). By calculating attention coefficients between nodes, the model can automatically learn which neighbour nodes are more important to the current node, thereby more flexibly aggregating neighbour information. In recommendation systems, GAT can dynamically recommend items that are most relevant to users' interests based on users' historical behaviour and item attributes, enhancing recommendation precision and customisation.

3 Knowledge representation of music healing resources based on pre-trained language models

The music therapy resource knowledge graph (MHRKG) framework is designed to unify heterogeneous music therapy data through structured semantic representation, enabling multidimensional knowledge integration. Its core goal is to achieve semantic associations between user needs, music resources, and therapeutic theories, thereby laying the foundation for personalised recommendations.

According to the analysis in the basic knowledge section, the key to constructing MHRKG lies in the representation of entities and relationships. In MHRKG, entities include users, music resources, therapeutic methods, emotion tags, etc., and relationships include the relationships between users and emotion tags, users and music resources, music resources and emotion tags, and music resources and therapeutic methods. The construction of existing knowledge graphs mostly relies on manual or semi-automatic methods, leading to incompleteness of entities and relationships. Although many knowledge graph completion models based on representation learning have been proposed, they are still affected by data sparsity, and no model can well learn all relationship patterns. To this end, this paper uses pre-trained language models to represent entities and relationships in MHRKG. In the entity and relationship representation vector learning module, the BERT model (Aum and Choe, 2021) is used to enrich the representation vectors of entities and relationships using text descriptions, alleviating the problem of data sparsity.

MHRKG contains a large amount of text descriptions about entities. Entity-relationship text embeddings provide supplementary semantic features that enhance input representation dimensionality, addressing sparsity-induced performance degradation. First, in the entity and relationship representation vector learning module, the form of entity sentences is 'entity name: entity description'. The sentence form of relationships has two types: 'relationship name' and 'reversed: relationship name', where the word 'reversed' indicates a reversed relationship. After obtaining the sentences of entities and relationships separately, the BertTokenizer tokeniser of the BERT model is used to tokenise the sentences.

In the entity and relationship representation vector learning module, the BERT model is used to encode sentences to obtain latent feature vectors for entities and their interactions. E_1, E_2, E_n are the inputs of BERT, which are the combined representation through additive aggregation of token, segment, and position vectors. T_1, T_2, T_n are the final output vectors after multiple layers of transformer. To infuse textual semantics into entity and relation embeddings, we concatenate their vector representations with BERT-encoded features before fusion. The calculation of the representation vector of the entity input into BERT is as follows:

$$e_{emb} = Embedding(e) \tag{1}$$

$$E_e = g\left(w_{mlp,e}, e_{emb}\right) \tag{2}$$

where *Embedding* is the Pytorch *Embedding* model, *embedding* model converts the numerical features of entities into vectors, e_{emb} is the entity embedding vector of *embedding* model output with d_{emb} dimensions, $w_{mlp,e}$ is the trainable parameters of the MLP, and MLP maps the entity embedding vector of d_{emb} dimensions into the entity representation vector R of d_{bert} dimensions.

$$r_{emb} = Embedding(r) \tag{3}$$

$$E_r = g\left(w_{mlp,r}, r_{emb}\right) \tag{4}$$

where r_{emb} is the d_{emb} dimensional relation embedding vector output from the *embedding* model, and $w_{mlp,r}$ denotes the trainable parameters of the MLP. MLP maps the relation embedding vector of d_{emb} dimensions into the relation representation vector E_r of d_{bert} dimensions, E_e is the entity representation vector input into BERT, and E_r is the relation representation vector input into BERT.

MLP converts the entity feature vector T_e and the relation feature vector T_r into the entity feature vector e and the relation feature vector r of d_{emb} dimensions. The closed-form solution for joint entity-relation representations derives as bellow, where $w_{mlp,e}$ and $w_{mlp,r}$ are trainable parameters.

$$e = g\left(w_{mlp,e}, T_e\right) \tag{5}$$

$$r = g\left(w_{mlp,r}, T_r\right) \tag{6}$$

4 Construction of intelligent recommendation system for music healing resources based on knowledge graph

4.1 Collaborative knowledge graph construction for music healing resources

Existing music healing resource recommendation models fully learn latent behaviour-preference associations, significantly enhances recommendation precision. However, when aggregating neighbourhood information, there is a problem of insufficiently learning the important relationships between entities. Introducing the user-music healing resource interaction graph into MHRKG can also better obtain collaborative information. This approach more effectively models user-item interactions, consequently enhancing recommendation precision.

Figure 1 The model of music healing resource recommendation system based on knowledge graph construction (see online version for colours)

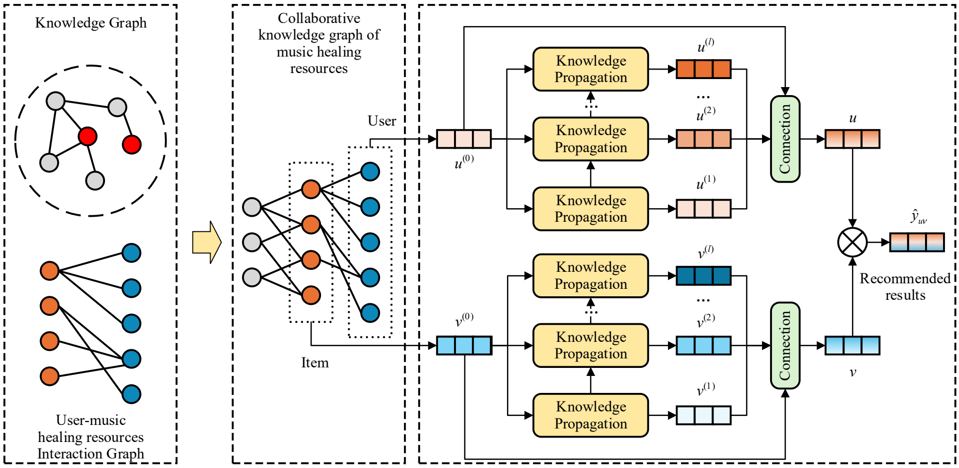

After representing entities and relations in MHRKG, this article builds an intelligent recommendation system for music healing resources based on the knowledge graph. The suggested model is indicated in Figure 1. First, this article constructs a collaborative knowledge graph by integrating the domain knowledge graph with user-resource interaction graphs through relational fusion. Subsequently, knowledge propagation is performed layer by layer, obtaining high-order representations of users and music healing resources, and finally, the predicted recommendation results are obtained.

Assuming in the music healing resource recommendation scenario, there is a set of M users $U = \{u_1, u_2, \ldots, u_M\}$ and a set of N music healing resources $V = \{v_1, v_2, , v_N\}$. According to the user's implicit feedback, such as click, browsing, and listening behaviours, an user-music healing resource interaction matrix Y can be obtained. When $y_{uv} = 1$ indicates that user u likes music healing resource v, conversely, $y_{uv} = 0$ indicates that user u does not like music healing resource v.

In addition, the knowledge graph serves as supplementary information for music therapy resources is defined as $G = \{(h, r, t)|h, t \in E, r \in R\}$, in which the triple (h, r, t) denotes a relational triple where entity h is linked to entity t via relation r. E and R represent the sets of entities and relations, respectively. For the recommendation task, utilising the interaction matrix Y (user-music healing resource) and knowledge graph G, the core task involves estimating the likelihood of user u exhibiting interest in resource v. The final purpose is to study a forecasting function $\hat{y}_{uv} = F(u, v|\theta, Y, G)$, in which $\hat{y}_{uv}$ is the forecasting score of user u for music healing resource v, and θ stands foe the model parameters of the function.

Collaborative MHRKG combines the user-music healing resource interaction graph with the knowledge graph of music healing resources. It merges the two graphs into a unified relational graph, thus achieving more accurate and effective recommendations. In the collaborative knowledge graph, user, music healing resource, and other entity nodes are connected through relationships. The benefit of this is to better utilise higher-order connectivity. In addition to direct relationships between users and music healing resources, more distant relationships, such as those connected through common interests or similar attributes, should also be considered.

4.2 Knowledge dissemination based on graph attention network

First, to more effectively extract collaborative information from user historical interaction records while keeping the computational complexity of the model at a reasonable level, this paper adopts a lightweight graph convolution method (Wang et al., 2023) to obtain enhanced representations for both users and music healing resource nodes, denoted as u' and v' respectively, as shown in equations (7) and (8). Here, $N(u)$ and $N(v)$ stand for the neighbourhoods of u and v, respectively.

$$u' = \sum_{v \in N(u)} \frac{1}{\sqrt{|N(u)||N(v)|}} v^{(l-1)} \tag{7}$$

$$v' = \sum_{u \in N(v)} \frac{1}{\sqrt{|N(u)||N(v)|}} u^{(l-1)} \tag{8}$$

Through the above calculation steps, collaborative information in the user's historical records can be obtained. Next, high-order information in the collaborative MHRKG is learned. For the head node h, there exists a knowledge triple connection relationship (h, r, t) between it and its neighbouring nodes. Here, the head node h includes user nodes, music healing resource nodes, and other entity nodes. The main work of knowledge propagation is to aggregate h and its neighbouring nodes to enrich the entity node information. To achieve this goal, the neighbourhood information of the node is modelled as a reception domain. This chapter adopts the idea of GAT, and during the propagation process, learns the attention weights of each neighbour to determine the influence of each neighbour node on the head node. Subsequently, the neighbour characteristics are transformed through attention-weighted aggregation to generate context-aware representations. Finally, the weighted neighbour feature representations are aggregated to obtain the new representation of the head node h.

In this way, the importance of high-order connections can be revealed, thereby better enriching the information of the head node. The sampled neighbourhood representation is $N(h)$, which contains the characteristic information of the head node h and its closed nodes. By aggregating $N(h)$, a new representation of the head node h can be obtained, thus achieving knowledge propagation. For a head node h, the attention weights of the neighbourhood nodes are first calculated. Here, the tanh operation is used as the activation operation, as shown in equation (9), where W_r is the relation learning matrix.

$$\pi(h, r, t) = (W_r t)^T \tanh(W_r h + r) \tag{9}$$

After calculating the attention weights of the neighbourhood nodes, these weights need to be normalised to ensure their sum is 1. This can make the influence of every closed node on the head node more accurate and reasonable. Specifically, the softmax function can be used to normalise the attention weights, as shown in equation (10).

$$\pi(h, r, t) = \frac{\exp(\pi(h, r, t))}{\sum_{(h,r,t) \in N(h)} \exp(\pi(h, r, t))} \tag{10}$$

Based on the normalised neighbourhood node attention weights, the reception domain data of the head entity $S(h)$ can be obtained. The reception domain information is a weighted sum of the characteristic representations of the closed nodes, as shown in equation (11).

$$S(h) = \sum_{(h,r,t)\in N(h)} \pi(h, r, t)t \tag{11}$$

After obtaining $S(h)$, they need to be aggregated into a single vector for subsequent recommendation tasks. Specifically, a linear transformation can be used to aggregate the head node representation h and $S(h)$. The activation function selected here is the LeakyReLU function, as shown in equation (12), where W is a trainable weight matrix.

$$f = \text{LeakyReLU}(W(h + S(h))) \tag{12}$$

For the goal of capturing higher-order relational patterns, the architecture employs stacked propagation layers that recursively aggregate multi-hop neighbourhood information. The entity representation at the l^{th} hop (i.e., the l^{th} propagation) is recursively derived from the entity representation at the $(l-1)^{\text{th}}$ hop, as specifically calculated in equation (13). Here, $h^{(l-1)}$ stands for the entity representation at the $(l-1)^{\text{th}}$ hop, and $S(h)^{(l-1)}$ represents the receptive domain message of the head entity at the $(l-1)^{\text{th}}$ hop. By stacking multiple propagation layers, higher-order connection information can be gradually explored, thus enhancing the head entity's feature representation through contextual aggregation.

$$h^{(l)} = f\left(h^{(l-1)}, S(h)^{(l-1)}\right) \tag{13}$$

After the above steps, an enhanced collaborative MHRKG can be obtained after knowledge propagation. In the collaborative MHRKG, there are enhanced vector representations of user u and music healing resource v. That is, the l^{th} propagation layer generates user embedding $u^{(l)}$. A and therapeutic music embedding $v^{(l)}$ through graph convolution. By concatenating all enhanced vector representations of the user and the music healing resource, the final user representation u and music healing resource representation v can be obtained, which contain rich feature information of the user and the music healing resource. The representations of u and v are as follows:

$$u = u^{(0)} \| \cdots \| u^{(l)} \tag{14}$$

$$v = v^{(0)} \| \cdots \| v^{(l)} \tag{15}$$

4.3 Intelligent recommendation prediction for music healing resources

Upon obtaining the final user representation vector u and music therapy resource embedding v, the score that the user gives to the music healing resource can be predicted by calculating their inner product. The prediction function is computed by taking the inner product between the aggregated representation vector of the user and that of the music healing resource, as specifically shown in equation (16), where $\hat{y}_{uv}$ represents the interest preference of the model for user u towards the music healing resource v.

$$\hat{y}_{uv} = u^T v \tag{16}$$

Figure 2 Schematic diagram of TransR (see online version for colours)

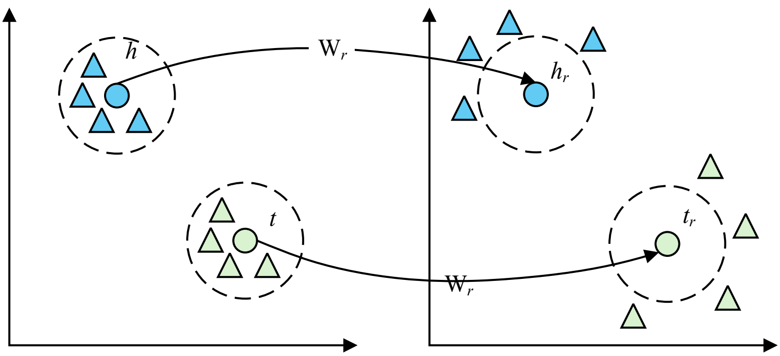

The model's total loss is the sum of two parts: the knowledge graph embedding loss and the recommendation prediction loss. Regarding the knowledge graph embedding loss operation, the TransR algorithm (Hou et al., 2024) is adopted. Compared with the TransE algorithm, the TransR algorithm can take into account that relations and vectors exist in different spaces. It maps entities to the space of relations through a mapping matrix, as shown in Figure 2. The knowledge graph embedding learning function $g(h, r, t)$ is shown in equation (17), where W_r is the relation mapping learning matrix, $h_r = hW_r$, $t_r = tW_r$.

$$g(h, r, t) = \|h_r + r - t_r\|_2^2 \tag{17}$$

Therefore, the knowledge graph embedding loss is shown in equation (18), where σ is the sigmoid function, and $T = \{(h, r, t, t')|(h, r, t) \in G, (h, r, t') \notin G\}$ is a triple constructed through randomly substituting a single entity within a valid knowledge graph triple. By minimising the knowledge graph embedding loss, optimisation techniques can be applied to improve the model's efficacy on knowledge graph-related tasks, thereby improving the model's generalisation ability.

$$L_{KG} = \sum_{(h,r,t,t')\in T} -\ln \sigma\left(g\left(h, r, t'\right) - g(h, r, t)\right) \tag{18}$$

On the other hand, the recommendation loss is shown in equation (19), where J represents the loss value calculated using cross-entropy. For a given user u, assuming that the user has interaction behaviour with some music healing resources, T^u represents the number of negative instances selected when performing negative sampling for the user, where $\{v: y_{uv} = 1\}$ represents the set of all music healing resources with which the user has interacted, and the operation of counting the number of elements in the set. In negative sampling, the selected negative samples are chosen according to a probability distribution, which is denoted as P, and it is assumed that P adopts a uniform distribution.

$$L_{REC} = \sum_{u\in U}\left(\sum_{v:y_{uv}=1} J\left(y_{uv}, \hat{y}_{uv}\right) - \sum_{i=1}^{T^u} E_{v_i \sim P(v_i)} J\left(y_{uv_i}, \hat{y}_{uv_i}\right)\right) + \lambda\|F\|_2^2 \tag{19}$$

In summary, the final loss function is expressed as equation (20), which consists of two parts: the knowledge graph embedding loss L_{KG} and the recommendation prediction loss L_{REC}. The last term is the $L2$ regularisation term, which is adopted to mitigate the overfitting issue during model training.

$$L = L_{REC} + L_{KG} + \lambda \|F\|_2^2 \tag{20}$$

5 Experimental results and analyses

5.1 Datasets and preprocessing

The experiment is implemented using the PyTorch 1.11.0 framework and Python 3.8. The CPU model used for the experiment is Intel(R) Xeon(R) Platinum 8358P CPU @2.60 GHz, and the GPU model is NVIDIA RTX A5000 GPU, with CUDA version 11.3 and 42 G of memory. During model training, the embedding dimension of the vector is set to 128, the epoch is 100, and the model parameters are updated using the Adam optimiser, with the learning rate set to 0.001.

Table 1 Preprocessed dataset information

Dataset	*Number of music healing resources*	*Number of sessions*	*Number of knowledge graph entities*	*Number of knowledge graph relationships*	*Number of ternary groups*
LFM-1b	21,902	369,129	121,131	45	434,713

Figure 3 Subset of LFM-1b dataset

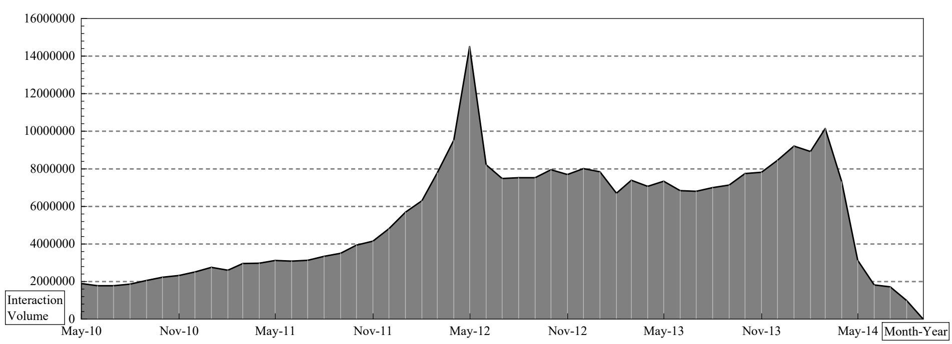

This paper uses music healing resources from the LFM-1b (Schedl, 2017) music recommendation dataset as the experimental dataset. This dataset covers listening behaviour of millions of users on millions of music healing resources. Due to the large volume of the LFM-1b dataset, we selected the subset for May 2023, which has the most intensive interactions as shown in Figure 3. The music playback data of users per hour is segmented into sessions, and the session length is truncated to between 3 and 20. For the knowledge graph data of the dataset, the dataset is used to match the experimental dataset and the Freebase triple knowledge base to obtain knowledge graph triple information, and entity relations that appear less than 1,000 times are removed. By this step, entity relations with low frequency can be filtered out, thereby improving the data quality and reliability of the knowledge graph. The information of the preprocessed dataset is shown in Table 1. 80% of the data is adopted as the training set, 10% as the validation set, and the remaining 10% as the test set.

5.2 Comparative experiments

To fully assess the effectiveness of the proposed model SMHRKG, this article selects MCCNN (Hou, 2024), MSC-AM (Han et al., 2024), DM-MAR (Lu, 2022), MKGCN (Cui et al., 2023), KG-MFKL (Liu et al., 2024), and MGCN-KG (Bai and Zhang, 2025) as comparison models. The model is trained for top-*K* recommendation tasks. After obtaining the forecasting outcome, the top-*K* music healing resources are sorted and generated into a recommendation list. Ultimately, the model's recommendation efficacy is rigorously validated through standard metrics including recall, F1-score and computational efficiency benchmarks, as shown in Figure 4. The performance of the MCCNN, MSC-AM, and DM-MAR models is relatively average, with recall and F1 values at a moderate level. This is because the MCCNN, MSC-AM, and DM-MAR models only use deep learning models to extract features of music healing resources, relying on user tags, but do not utilise the knowledge graph to mine interaction features among various entities. Therefore, the indicators of these models are not as good as the other four models. MKGCN combines the knowledge graph and CNN. Although it integrates semantic information from the knowledge graph when representing item features, it achieves good results in the recall rate of recommendation results. However, the recommendation accuracy is still relatively low compared to other models. The comprehensive indicators perform well in recommendation tasks when the amount of recommendations is less than 20, but the recommendation results become worse as the number of recommendations increases. Finally, the comprehensive indicators are at an average level. MGCN-KG performs well in recall rate and F1 indicators. It can be seen that the addition of structural information from the knowledge graph effectively enhances the recommendation performance. However, when the amount of recommendations exceeds a certain amount, the model's F1 indicator no longer increases, and the recommendation effect no longer has an advantage compared to other models. It can be seen that in-depth knowledge paths do not significantly improve the results, but instead increase the algorithm complexity and consume more computing time.

SMHRKG cleverly combines the semantic information of user behaviour and the knowledge graph, exploring users' interest preferences at different levels. Regardless of the number of recommendation tasks, the recommendation accuracy of this model is far higher than that of all other models. It not only compensates for the defect of model cold start but also brings interpretability to the recommendation results while increasing personalisation. The operation times of all models are shown in Figure 4(c). SMHRKG takes the least time, and the time spent by deep learning-based recommendation models increases significantly, indicating that the introduction of the knowledge graph can shorten the model operation time to some extent.

Moreover, to further estimate the performance of the SMHRKG model, this paper also selects quantitative indicators HR@K, NDCG@K, and AUC to compare and analyse the recommendation effects of different models, where *K* is the number of music healing resources, and here it is taken as 5, 10. The comparison outcomes are shown in Table 2. The HR@5 and HR@10 of SMHRKG are 0.8268 and 0.9037, respectively, which are at least 4.34% and 2.1% higher than the comparison models. The NDCG@5 and

NDCG@10 of SMHRKG are 0.5186 and 0.7638, respectively, which are at least 8.2% and 6.69% higher than the baseline model. The AUC of SMHRKG is 0.9894, which is 21.21%, 16.25%, 13.05%, 9.84%, 4.85%, 1.95%, SMHRKG uses GAT to dynamically assign weights to neighbouring nodes. The benefit of this is that it can more accurately capture the importance of each neighbouring node, thus efficiently aggregating neighbouring information to gain more accurate vector representations of users and items. Secondly, introducing the user-music healing resource interaction graph into the knowledge graph can better obtain collaborative information. This approach allows the model to fully utilise the interaction relationships between users and music healing resources, mining potential association patterns, thus providing a more powerful basis for recommendations, greatly improving the accuracy and personalisation of recommendations.

Figure 4 Performance comparison of music healing resource recommender systems (see online version for colours)

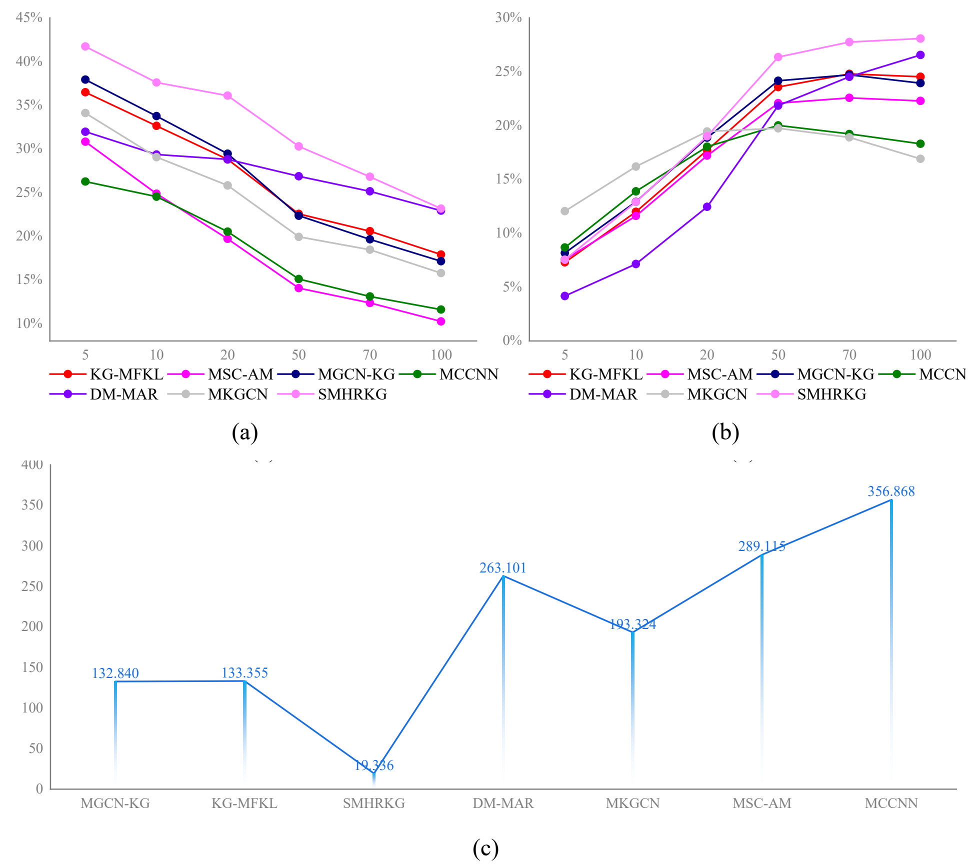

Table 2 Comparison of recommendation performance metrics for different models

Model	*HR@5*	*HR@10*	*NDCG@5*	*NDCG@10*	*AUC*
MCCNN	0.5824	0.7104	0.3169	0.5418	0.8163
MSC-AM	0.6492	0.7569	0.3507	0.5647	0.8511
DM-MAR	0.6825	0.7716	0.3741	0.6158	0.8752
MKGCN	0.7051	0.8274	0.4085	0.6349	0.9008
KG-MFKL	0.7618	0.8436	0.4552	0.6852	0.9436
MGCN-KG	0.7924	0.8851	0.4793	0.7159	0.9705
SMHRKG	0.8268	0.9037	0.5186	0.7638	0.9894

6 Conclusions

With the increasing attention to mental health issues, music healing has been widely applied as a non-pharmacological intervention method due to its ease of use and effectiveness. To address the problem that existing music healing recommendation systems are not flexible enough in learning important relationships between entities and have low recommendation accuracy, this paper builds an intelligent recommendation system for music healing resources in light of a knowledge graph. First, in the entity and relation representation vector learning module, the BERT model is used to enrich the representation vectors of entities and relations, alleviating the problem of data sparsity. Then, the interaction graph between users and music healing resources is integrated into the knowledge graph, constructing a collaborative knowledge graph and extracting collaborative information. When aggregating neighbourhood information, a graph attention mechanism is introduced, allowing the weight of each neighbour node to be dynamically adjusted according to its relationship with the target node, capturing more complex relationships between nodes. Finally, the inner product between the user representation and the music healing resource representation is calculated to forecast the user's rating for the music healing resource, and the top N music healing resources with the highest scores are recommended. Experimental results show that the HR@5 and HR@10 of the proposed model are at least 4.34% and 2.1% higher than those of the baseline model, respectively, which can achieve relatively accurate recommendations for music healing resources.

Owing to constraints imposed by research time and experimental resources, the music healing resource recommendation system in light of the knowledge graph in this article has great development potential in the future and can be explored and improved in the following aspects:

1 Our model architecture incorporates knowledge graph-derived features through cross-domain information fusion, thereby enhancing recommendation performance, but the auxiliary information adopted is still relatively single. In future work, other information, such as audio features of songs, will be considered to achieve multimodal music recommendation, and integrate various types of information to help improve the effect of the recommendation algorithm.

2 In real life, users' interactions and preferences are constantly changing. Recommendation systems based on static preference modelling cannot capture users' changing interests in a timely manner. To make the model more dynamic, some appropriate techniques and algorithms need to be used to capture changes in user behaviour and preferences and continuously update the model over time.

Declarations

This work is supported by the Shandong Province Undergraduate Teaching Reform Research Project (No. 23202502200) and the University-Industry Collaborative Education Program Initiated by the Ministry of Education of China (No. 231107179240829).

All authors declare that they have no conflicts of interest.

References

Assuncao, W.G., Piccolo, L.S. and Zaina, L.A. (2022) 'Considering emotions and contextual factors in music recommendation: a systematic literature review', *Multimedia Tools and Applications*, Vol. 81, No. 6, pp.8367–8407.

Aum, S. and Choe, S. (2021) 'srBERT: automatic article classification model for systematic review using BERT', *Systematic Reviews*, Vol. 10, No. 1, pp.28–35.

Bai, G. and Zhang, T. (2025) 'Dynamic node graph neural network for multimodal music recommendation', *Journal of Internet Technology*, Vol. 26, No. 3, pp.407–414.

Bellini, V., Di Sciascio, E., Donini, F.M., Pomo, C., Ragone, A. and Schiavone, A. (2024) 'A qualitative analysis of knowledge graphs in recommendation scenarios through semantics-aware autoencoders', *Journal of Intelligent Information Systems*, Vol. 62, No. 3, pp.787–807.

Bhatti, U.A., Tang, H., Wu, G., Marjan, S. and Hussain, A. (2023) 'Deep learning with graph convolutional networks: an overview and latest applications in computational intelligence', *International Journal of Intelligent Systems*, Vol. 3, No. 1, pp.83–91.

Cui, X., Qu, X., Li, D., Yang, Y., Li, Y. and Zhang, X. (2023) 'Mkgcn: multi-modal knowledge graph convolutional network for music recommender systems', *Electronics*, Vol. 12, No. 12, pp.26–38.

Han, X., Chen, F. and Ban, J. (2024) 'A GAI-based multi-scale convolution and attention mechanism model for music emotion recognition and recommendation from physiological data', *Applied Soft Computing*, Vol. 164, pp.12–19.

Hou, R. (2024) 'Music content personalized recommendation system based on a convolutional neural network', *Soft Computing*, Vol. 28, No. 2, pp.1785–1802.

Hou, W., Bai, B. and Cai, C. (2024) 'CR-TransR: a knowledge graph embedding model for cultural domain', *ACM Journal on Computing and Cultural Heritage*, Vol. 17, No. 1, pp.1–11.

Jing, E., Liu, Y., Chai, Y., Yu, S., Liu, L., Jiang, Y. and Wang, Y. (2024) 'Emotion-aware personalized music recommendation with a heterogeneity-aware deep Bayesian network', *ACM Transactions on Information Systems*, Vol. 43, No. 5, pp.1–43.

Kathavate, S. (2021) 'Music recommendation system using content and collaborative filtering methods', *International Journal of Engineering Research & Technology (IJERT)*, Vol. 10, No. 2, pp.167–171.

Langdon, G.S., Margolis, F. and Muenzenmaier, K. (2018) 'Weaving words and music: healing from trauma for people with serious mental illness', *Music and Medicine*, Vol. 10, No. 3, pp.157–161.

Li, L., Tao, D., Zheng, C. and Gao, R. (2022) 'Attentive auto-encoder for content-aware music recommendation', *CCF Transactions on Pervasive Computing and Interaction*, Vol. 4, No. 1, pp.76–87.

Liu, X., Yang, Z. and Cheng, J. (2024) 'Music recommendation algorithms based on knowledge graph and multi-task feature learning', *Scientific Reports*, Vol. 14, No. 1, p. 20-32.

Lu, W. (2022) 'Design of a music recommendation model on the basis of multilayer attention representation', *Scientific Programming*, Vol. 20, No. 4, pp.63–77.

Mao, Y., Zhong, G., Wang, H. and Huang, K. (2022) 'Music-CRN: an efficient content-based music classification and recommendation network', *Cognitive Computation*, Vol. 14, No. 6, pp.2306–2316.

Sakurai, K., Togo, R., Ogawa, T. and Haseyama, M. (2022) 'Deep reinforcement learning-based music recommendation with knowledge graph using acoustic features', *ITE Transactions on Media Technology and Applications*, Vol. 10, No. 1, pp.8–17.

Schedl, M. (2017) 'Investigating country-specific music preferences and music recommendation algorithms with the LFM-1b dataset', *International Journal of Multimedia Information Retrieval*, Vol. 6, No. 1, pp.71–84.

Schedl, M. (2019) 'Deep learning in music recommendation systems', *Frontiers in Applied Mathematics and Statistics*, Vol. 5, pp.35–44.

Velankar, M. and Kulkarni, P. (2022) 'Music recommendation systems: overview and challenges', *Advances in speech and music technology: Computational Aspects and Applications*, Vol. 10, pp.51–69.

Veličković, P. (2023) 'Everything is connected: graph neural networks', *Current Opinion in Structural Biology*, Vol. 79, pp.10–18.

Vrahatis, A.G., Lazaros, K. and Kotsiantis, S. (2024) 'Graph attention networks: a comprehensive review of methods and applications', *Future Internet*, Vol. 16, No. 9, pp.31–45.

Wang, R., Lou, J. and Jiang, Y. (2023) 'LightGCAN: a lightweight graph convolutional attention network for user preference modeling and personalized recommendation', *Expert Systems with Applications*, Vol. 23, pp.12–21.

Wang, W., Shen, X., Yi, B., Zhang, H., Liu, J. and Dai, C. (2024) 'Knowledge-aware fine-grained attention networks with refined knowledge graph embedding for personalized recommendation', *Expert Systems with Applications*, Vol. 249, pp.37–44.

Xu, D. (2020) 'Research on music culture personalized recommendation based on factor decomposition machine', *Personal and Ubiquitous Computing*, Vol. 24, No. 2, pp.247–257.

Yue, C., Wang, Y., Tang, X. and Chen, Q. (2022) 'DRGCNN: dynamic region graph convolutional neural network for point clouds', *Expert Systems with Applications*, Vol. 20, pp.11–17.

Zhang, J-C., Zain, A.M., Zhou, K-Q., Chen, X. and Zhang, R-M. (2024) 'A review of recommender systems based on knowledge graph embedding', *Expert Systems with Applications*, Vol. 250, pp.68–76.