

GEMRec: A Graph-based Emotion-aware Music Recommendation Approach

Dongjing Wang^{1,2}, Shuiguang Deng¹ (✉), and Guandong Xu²

¹ College of Computer Science and Technology, Zhejiang University, China

² Advanced Analytics Institute, University of Technology Sydney, Australia
{tokyo1,dengsg}@zju.edu.cn, Guandong.Xu@uts.edu.au

Abstract. Music recommendation has gained substantial attention in recent times. As one of the most important context features, user emotion has great potential to improve recommendations, but this has not yet been sufficiently explored due to the difficulty of emotion acquisition and incorporation. This paper proposes a graph-based emotion-aware music recommendation approach (GEMRec) by simultaneously taking a user’s music listening history and emotion into consideration. The proposed approach models the relations between user, music, and emotion as a three-element tuple (user, music, emotion), upon which an Emotion Aware Graph (EAG) is built, and then a relevance propagation algorithm based on random walk is devised to rank the relevance of music items for recommendation. Evaluation experiments are conducted based on a real dataset collected from a Chinese microblog service in comparison to baselines. The results show that the emotional context from a user’s microblogs contributes to improving the performance of music recommendation in terms of hitrate, precision, recall, and F1 score.

Keywords: Music recommendation · Emotion analysis · Random walk · Emotion aware

1 Introduction

Nowadays, the lower cost of hardware and advances in technology have led to the growth in digital music, and there is an enormous amount of music available on the Internet. For example, currently Apple Music offers over 30 million songs (<http://support.apple.com/en-us/HT204951>). However, the vast amount of music has made it more difficult for users to find the music they really enjoy, which is known as the Paradox of Choice. Therefore, music recommendation has become an interesting topic both in research and application. Similar to recommender systems applied in various domains [21, 22], music recommendation has greatly benefited from the algorithmic advances of the recommender systems community, e.g., collaborative filtering, content-based and hybrid approaches, which have been predominantly adopted by simply treating music recommendation as a classic item, e.g., a book or movie, to solve the recommendation problem via user’s long-term music preferences. Music is not a neutral item but

a carrier of thoughts and emotions, so people will have different music tastes and preferences in different contexts [10]. For example, a user who likes both light music and rock & roll music may prefer the former when at rest. Therefore, incorporating contextual information like geo-location, emotional state, time, presence of other people, past and future events can help the recommender system to better understand and satisfy the users' real-time requirements [19].

Music is a specific type of content which can carry emotion features, hence our work focuses on utilizing emotion as a supplementary feature to make music recommendations. In particular, our work is inspired by the following three observations: (1) awareness: listening to music is a kind of emotion-sensitive behavior and direct associations exist between the users' mood and their favorite music [13]. For example, people who are experiencing a sad mood generally tend to listen to a different style of music compared to when they are experiencing a happy mood; (2) dynamics: people's emotions change constantly and dynamically so satisfactory music recommendation must reflect the real-time emotion status of the user; (3) indirect derivation: the emotional status of the user is a secondary context [1], since it cannot be measured directly, but needs to be derived from other types of contextual information. Thus, these observations make emotion-aware music recommendation a real challenge as to how to extract the users' emotion context and integrate this into recommendations, which constitute our major contributions. To the best of our knowledge, there are only a few existing studies which address music recommendation from this angle.

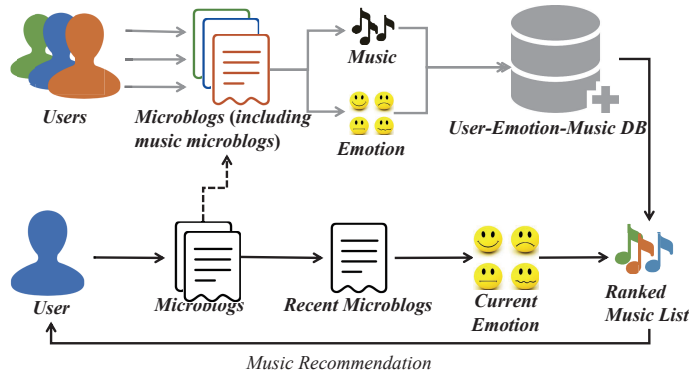


Fig. 1. The concepts of the proposed approach

In this paper, we focus on extracting and incorporating users' emotions into music recommendation. The key objective is to satisfy the users' music needs under certain emotional contexts with minimum user effort in relation to providing feedback. There are three key challenges to be tackled: (1) how to acquire users' current emotional contexts; (2) how to embed emotion in user-music relations; (3) how to infer personalized music preferences given a certain emotional context.

t. To address these challenges, we propose a graph-based emotion-aware music recommendation method (GEMRec), which can 1) extract users' music listening history along with the corresponding emotion from their microblog texts, 2) model users' music records along with the corresponding emotional contexts by an Emotion Aware Graph (EAG), and 3) adopt the relevance propagation algorithm based on random walk in order to reveal the explicit and hidden associations between users and music items under certain emotion circumstances, then rank and recommend music items. The framework of the proposed approach is shown in Fig. 1. Firstly, we extract the users' music listening history with the accompanying emotion information when they are listening to music from their microblogs and store this in a database in the form of three-element tuple associations (user, music, emotion). Then, when a user posts new microblogs, the proposed system will determine the user's current emotional status and recommend appropriate music items by referring to the three-tuple association database. To validate the performance of the proposed system, we conduct comprehensive experiments with a real-world dataset collected from a Chinese Microblog website named Sina Weibo (<http://www.weibo.com>). In addition, we also investigate how emotion can influence recommendations in detail.

To summarize, the main contributions of this work are three-fold: (1) we propose an Emotion Aware Graph which integrates users, their music listening history with their various emotion statuses in a unified manner; (2) we devise a graph-based relevance propagation algorithm to capture the overall user preferences of music items given a specific emotion status; (3) we conduct extensive experiments to evaluate the proposed method on a real-world dataset in comparison to the state-of-the-art techniques.

The rest of the paper is organized as follows. In the next section, we discuss the existing work on context-based music recommendation. In Section 3, we introduce the proposed recommendation method in detail. Then, we provide an evaluation of the proposed approach in Section 4. Finally, we conclude the paper with a summary and directions for future work in Section 5

2 Related Work

In this section, we discuss existing works on context-based music recommendation, including environment-related context approach and user-related context approach.

2.1 Environment-related Context Approach

This type of research is based on the fact that the environment has an influence on the state-of-mind or emotional state of users, and therefore indirectly influences users' music preferences [13]. For instance, people usually enjoy reflective and complex music during fall/winter, while they may prefer more energetic music during spring/summer [14]. Consequently, music recommendation approaches with environment-related parameters will perform better. The

environment-related context includes time [14, 19], location [9], and hybrid contexts [11]. Kaminskas and Ricci [9] explored the possibilities of adapting music to the place of interest (POIs) that the user is visiting. In detail, the author used emotional tags attached by users to both music and POIs and considered a set of similarity metrics for the tagged resources to establish a match between music tracks and POIs. Knees and Schedl [11] gave an overview of approaches to music recommendation which do not rely on the audio signal, but rather take into consideration various aspects of the context in which a music entity occurs. However, most of the prior approaches require researchers to label music with appropriate tags or map context with music, which restrict their application in the real world. Another limitation is that the prior methods have only been evaluated based upon synthetic or small-scale datasets.

2.2 User-related Context Approach

Compared with the environment-related context, the user-related context has a closer relationship with the users, and can influence the users' music preferences directly. The user-related contexts include activity [20], demographical information [12], and emotion state [3, 7, 16]. Wang et al. [20] described a novel probabilistic model for music recommendation that combines automated activity classification with automated music content analysis, with support for a rich set of activities and music content features. Chang et al. [3] proposed a correlation-coefficient-based approach to find emotional music sequences which may evoke a specific emotion in subjects. Then, they built a personal emotion-cognitive music database for an individual subject, and finally music from a personal emotion-cognitive music database is recommended to the user. Han et al. [7] proposed a context-aware music recommendation system in which music is recommended according to the user's current emotion state and the music's influence on changes in the user's emotion. Deng et al. [4] presented another contextual music recommendation approach, which can infer users' emotion from their microblogs, and then recommend music pieces appropriate for users' emotion. However, there is not much work which attempts to utilize emotional information, and the performance of existing emotion-aware approaches still need to be improved greatly. In addition, most of the existing work requires the users to input contextual information, which restrains their application in practice.

3 Proposed Approach

Listening to music is a type of emotion-sensitive behavior and direct associations exist between a users' mood and their favorite music [13]. However, the emotion status of the user is a secondary context [1] which cannot be measured directly. In addition, people's emotions change constantly and dynamically. Therefore, the main challenges are how to acquire the users' real-time contexts of emotion and incorporate them into music recommendation. As a social network service, microblogs have become increasingly popular in people's daily life. Users can

share many kinds of information, such as news, knowledge, resources, opinions, status, and feelings, anywhere and anytime. Therefore, users' interests and contextual information can be inferred from their microblogs. In the meantime, users may share the music items they are listening to, which implies the underlying correlation between the posted microblogs and the shared music items. The idea behind our approach is that the association in users' microblogs, more specifically the music items, and the corresponding emotions in the microblogs, can reflect the users' emotion-aware music preferences. As shown in Fig. 2, the proposed graph-based emotion-aware music recommendation method (GEMRec) consists of three main components: 1) *association extraction*, 2) *association modelling*, and 3) *music recommendation*. The *association extraction* component aims to extract the users' music listening history and the corresponding emotional contexts from their microblogs, and then obtains the associations between the user, the music items and the emotions, which can be represented as a three-element tuple (user, music, emotion). Then, in the *association modelling* component, all associations are modelled by an EAG, where the user, music, and emotion are represented as nodes and their relations are expressed as edges. Finally, in the *music recommendation* component, when a target user posts new microblogs or activates the recommendation manually, the music recommendation component performs recommendations on the Emotion Aware Graph using a random-walk-based relevance propagation algorithm and recommends appropriate music items according to the user's real-time emotion extracted from their recent microblogs.

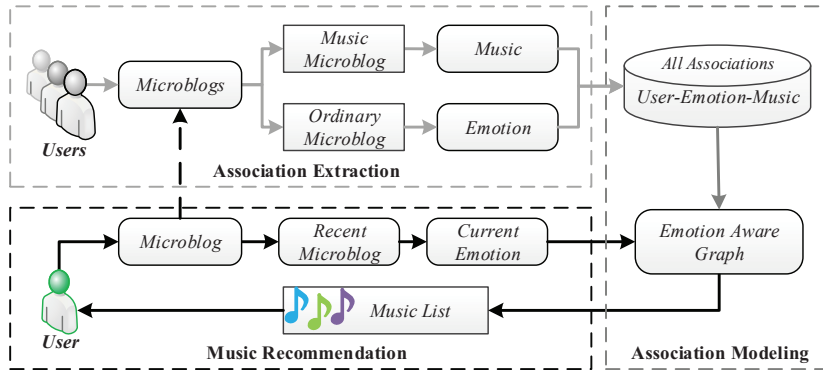


Fig. 2. Overview of GEMRec

3.1 Association Extraction

All the users' microblogs can be roughly categorized into two types, i.e., ordinary microblogs and music-sharing microblogs. As shown in Table 1, ordinary microblogs (the first two rows of Table 1) mainly give generic information

on the users’ opinions, feelings, status, and activities at a specific timestamp, while music-sharing microblogs contain information on music items, including the singer, music title, URL, comments and so on. We assume an inherent correlation exists between the music that the user listens to and the microblogs that they have recently posted. In other words, music listening behaviors are context-aware (emotion dependent). As for the example shown in Table 1, given that the polarity of emotion in microblogs posted before the music-sharing microblogs is positive in nature, we can speculate that this user (521156765) enjoys the music Home in a positive emotional context.

Table 1. Examples of Microblogs

UserID	Content	Time	Music
5211	[smile][smile]so excited about the good news!	2014-06-22 13:25:39	
5211	What a good day	2014-06-22 13:29:10	
5211	Share music Edward Sharpe & TheHome: music appears in modern family http://t.cn/zWNSPMS	2014-06-22 13:30:01	Home

In this *association extraction* component, the pieces of music are given explicitly, so the most important task is extracting the users’ emotion from their microblogs, which consists of three steps as follows.

Firstly, we construct a fine-grained Chinese emotion lexicon from three resources: 1) the original emotion lexicon from DUTIR (<http://ir.dlut.edu.cn>), and the emotion classification systems based on Ekman’s classical psychological research [5], which consists of seven kinds of emotion, namely, happy, like, surprise, angry, sad, fear, and hate; 2) the synonym lexicon from HITCIR (http://ir.hit.edu.cn/demo/1tp/Sharing_Plan.htm) to enhance the original emotion lexicon; 3) emoticons collected from the microblog website to extend the original emotion lexicon. In total, there are 48224 items in our emotion lexicon.

Secondly, we adopt a Chinese segmentation tool named ICTCLAS (<http://www.ictclas.org>) to segment microblog texts into words. Based on the constructed lexicon of emotion, we count the frequency of emotion words occurring in a text for each type of emotion, and the number of emotion words appearing in the text determines the emotion vector of the text. Then, each microblog can be represented as an emotion vector. In general, when we try to discover the user’s emotion association with the music, we take into consideration only the microblogs before the music listening timestamp for emotion mining. In doing so, we obtain a list of the user’s time-dependent microblogs by setting a certain time window around the timestamp when the music microblog was posted and compute the sum of all the emotion vectors of the microblogs within the time window as the emotion corresponding to the piece of music. Specifically, an appropriate time window is important because a narrow time window will result in too few microblogs from which to extract emotions, and a wide time window will cause emotional noise. In this paper, the time window is set as five hours.

Finally, we obtain all the emotion vectors at various music sharing timestamps, which forms the association data of emotion and music items for all users, represented by a three-element tuple (user, music, emotion).

3.2 Association Modeling

In this section, we illustrate how to model the association between the user, music, and the corresponding emotion on the graph by using different node types and treating edges with different weights.

As shown in Section 3.1, our association data are in the form of (user, music, emotion) triples which are usually modeled by a tri-partite graph or a tensor. However, both the tri-partite graph and tensor treat emotional information as a universal dimension shared by all users, while we argue that, in a recommender system, the dimension of emotions is a private effect and should not be shared and compared across all users arbitrarily. Specifically, we build an EAG to model our association data based on the following intuitions:

1. **Intuition 1:** Users have specific preferences to music, and tend to listen to their favorite music pieces [11].
2. **Intuition 2:** Music listening is a kind of emotion-sensitive behavior and direct associations exist between the users' emotional mood and their favorite music. Users' music preferences vary with differing emotions, and a user may listen to different music pieces when they experience different emotion [3].
3. **Intuition 3:** If two users have similar temporal preferences (in two emotions), they tend to listen to the same music pieces (when experiencing these two emotions) [7, 16].

As a bipartite graph, EAG is defined as:

$$\text{EAG} = G(U, EM, I, E, w) \quad (1)$$

where U denotes the set of user nodes, EM is the set of emotion nodes, I is the set of music item nodes, and $w : E \rightarrow R$ denotes a non-negative weight function for edges in E .

Nodes. In EAG, user node v_u connects to all items interacted by user u , denoted as $N(u)$, which represent u 's long-term preferences under all emotional contexts; emotion node $v_{u,em}$ only connects to items that user u viewed at emotion em , denoted as $N(u, em)$, representing u 's specific preference at emotion em . Therefore, if we start walking from user node v_u , we will go through $N(u)$ and then reach unknown items similar to items in $N(u)$ (**Intuition 1**). Similarly, we will reach unknown items similar to items in $N(u, em)$ if we start from the emotion nodes (**Intuition 2**). In short, the user node represents the user's long-term preference and the emotion node represents the preference under a specific emotional context.

Edges. The weight of the edges in EAG is defined as:

$$w(v, v') = \begin{cases} 1 & v \in U, v' \in I \text{ or } v \in I, v' \in U \\ \varphi_{i,em} & v \in I, v' \in EM \text{ or } v \in EM, v' \in I \end{cases} \quad (2)$$

This definition means, given an edge $e(v, v')$ between user nodes and item nodes, its weight will be 1. As for the edges between item nodes and emotion nodes, its weight is the weight of the emotion when the user is accessing specific item i , which is $\varphi_{i,em}$. The edges from the user nodes to the music item nodes bridge users who share similar music preferences. Furthermore, the edges from the emotion nodes to the music item nodes bridge users who share similar music preferences in certain emotions. All these edges enable us to exploit other users' (emotional) music preferences for recommendation (**Intuition 3**).

Table 2 shows an example of association data which only contains two types of emotion for brevity. This example contains 3 user nodes, 6 item nodes, and 6 emotion nodes. It shows user A interacted with items a, b, c , user B interacted with items c, d , and user C interacted with items d, e, f . Furthermore, item a has an emotion vector $(1.0, 0.0)$, which indicates that A interacted with this item under the context E_1 . Fig. 3 shows the corresponding EAG. In EAG, each user has their own emotion nodes, because the dimension of emotion is a private effect, which should not be shared and compared among all users, arbitrarily.

Table 2. Association Example

User	Music	Emotion Vector (E_1, E_2)
A	a	$(1.0, 0.0)$
A	b	$(0.5, 0.5)$
A	c	$(0.0, 1.0)$
B	c	$(0.0, 1.0)$
B	d	$(1.0, 0.0)$
C	d	$(0.0, 1.0)$
C	e	$(0.0, 1.0)$
C	f	$(1.0, 0.0)$

Then, we defined a weighted adjacent matrix \mathbf{M} to represent the bi-partite graph $G(U, EM, I, E, w)$, and each element $m_{vv'}$ of matrix \mathbf{M} is the weight of edge $e(v, v')$, whose value is $w(v, v')$. After this, we can compute the transition possibility matrix \mathbf{P} by normalizing matrix \mathbf{M} as follows:

$$p_{vv'} = \begin{cases} \frac{m_{vv'}}{\sum_{v' \in \text{outnodes}[v]} m_{vv'}} & \text{if } \text{outnodes}[v] \neq \emptyset \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

Next, we explain how to perform music recommendation on EAG represented by the transition possibility matrix \mathbf{P} .

3.3 Music Recommendation on EAG

As shown in Section 3.1, our association data are modelled with the proposed EAG. We then devise a relevance propagation algorithm based on random walk

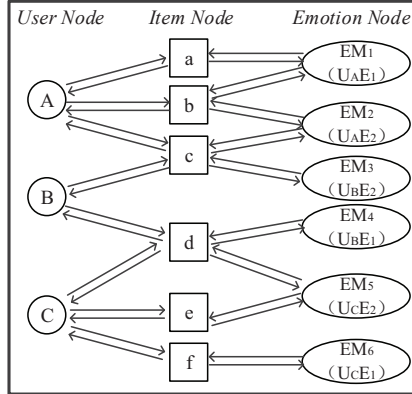


Fig. 3. An example of EAG

to capture the overall user preference on music items given a specific emotion status and to recommend appropriate music pieces. Specifically, the relevance propagation algorithm for ranking entities on EAG is a variation of the Personalized PageRank algorithm [8], where a node will get a higher rank score if it is connected to more important nodes with less outgoing links. By applying the propagation algorithm to EAG, we can find the target user’s potential preferences in certain emotions based on all emotion-aware music listening data. In general, the PageRank score can be calculated as follows:

$$\mathbf{r} = \alpha \mathbf{P}^T \mathbf{r} + (1 - \alpha) (1/n) \mathbf{e} \quad (4)$$

where r_i is the rank score of node v_i , n is the number of nodes, $\mathbf{e} = (1, 1, \dots, 1)^T$ represents the initial value for all nodes, and α is a damping factor constant that is normally set to 0.85. The algorithm is based on a random-walk with the restart process, which uniformly chooses a random node for restarting with possibility $(1 - \alpha)$ during the random walk process. To calculate the PageRank score of a node, we replace $(1/n) \mathbf{e}$ with a personalized initial vector \mathbf{r}_0 . This personalized initial vector normally represents a user’s interest. Thus, the corresponding PageRank score vector can be calculated as:

$$\mathbf{r} = c \mathbf{P}^T \mathbf{r} + (1 - c) \mathbf{q} \quad (5)$$

Random walk with the restart process gives each node a higher ranking score when the node is more closely related to the nodes in the query vector, considering the importance of nodes at the same time. As a result, we can use the value of r_i as the rank value for each entity, including item nodes. Finally, the item nodes are sorted according to their ranking score \mathbf{r} and the top-N items are recommended to the target user under certain emotional contexts.

There are several advantages of adapting the propagation algorithm to rank items in our approach. Firstly, we can utilize its propagation and attenuation

properties [6]. The propagation property means the relatedness of the nodes propagates following the links (edges), and the attenuation property is that the propagation strength decreases as the propagation goes further from the initial node. Secondly, through the graph and with the random-walk process, we can find a rank score of target entities by considering the influence of various hybrid factors, which enables our approach to incorporate many kinds of contextual information in different forms, including emotional context. Thirdly, the propagation algorithm can be pre-computed and it scales for large size datasets. In our work, the formula of the rank score can be transformed to $\mathbf{r} = (1 - c) \left(1 - c\mathbf{P}^T\right)^{-1} \mathbf{q}$, which shows that if we can pre-compute the matrix $(1 - c) \left(1 - c\mathbf{P}^T\right)^{-1}$ independently, then we can simply multiply the initial query vector and the matrix $(1 - c) \left(1 - c\mathbf{P}^T\right)^{-1}$ to compute the rank score for each node according to the query (the target user and their current emotion status) with acceptable time complexity $O(m * n)$. In our method, the query is given as a sparse vector with only several nonzero elements, which represent the target user and their emotional context. Thus, the matrix multiplication can be computed quickly and the recommendation can be performed online. In addition, although pre-computation requires matrix inversion which has very high time complexity, which is $O(n^3)$, [18] shows that we can estimate the inverse matrix with acceptable cost by using the dimension reduction approach.

4 Experimental Evaluation

In this section, we first investigate how the users' emotion affects the performance of music recommendation by adopting different emotion modeling methods. Then, we compare the proposed method against baseline methods. All experiments were performed on a PC with Intel Core i3-2120 running at 3.30 GHz, along with 8 GByte of memory and 64-bit Windows 8 operating system.

4.1 Dataset, Experiment Design, and Evaluation Metrics

We collect microblog data from Sina Weibo (<http://open.weibo.com/wiki/API>). The detailed statistics of the final dataset are shown in Table 3, and we split the dataset into a training and testing set according to the methodology of the 5-fold cross-validation.

Table 3. Complete Statistics of Dataset

#(Users)	#(Music)	#(Listening)	#(Listening per user)	#(Listening per music)
1,008	2,552	92,283	91.6	36.2

Besides, to evaluate the performance of our approaches and the baselines, we adopt four metrics [2] as follows:

HitRate. Hitrate is the fraction of hits, and a hit means the recommendation list contains at least one music pieces that the user actually listened to. The definition is given as

$$HitRate = \#(hits)/\#(R) \quad (6)$$

where:

- $\#(hits)$ is the number of hits.
- $\#(R)$ is the number of recommendations.

Precision, Recall, and F1 Score. Precision is the fraction of recommended music items in which the target user is actually interested. Recall is the fraction of the target user’s interested music that is recommended. F1 score is the harmonic mean of precision and recall. The definitions are given as follows:

$$Precision = \sum_{u \in U} (|R(u) \cap T(u)| / |R(u)|) \quad (7)$$

$$Recall = \sum_{u \in U} (|R(u) \cap T(u)| / |T(u)|) \quad (8)$$

$$F1 = 2 \times Precision \times Recall / (Precision + Recall) \quad (9)$$

where:

- U is the user set in the test data.
- $R(u)$ is the recommended music list.
- $T(u)$ is the set of all music items u actually listened to in the test data.

4.2 Effect of Emotion Modeling Methods

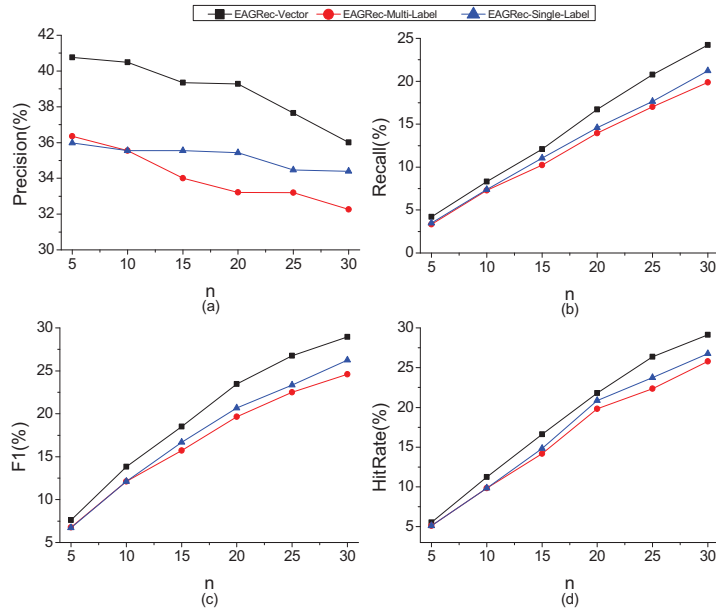
In order to investigate how users’ emotion can affect the performance of music recommendation, we evaluate three recommendation methods, namely EAGRec-Vector, EAGRec-Single-Label and EAGRec-Multi-Label, which adopt vector, single-label and multi-label, respectively, as their emotion modeling methods. As shown in Table 4, the vector can model all the emotions extracted from the users’ microblogs along with the emotion weight, the multi-label can model all emotions without a corresponding weight, and the single-label only incorporates the emotion with the highest weight value. In other words, the vector can represent all details of the original extracted emotion, while the multi-label can express all kinds of emotion appearing with the same weight and the single-label only considers the main emotion with the highest value.

As shown in Fig. 4, we can conclude that (1) EAGRec-Vector has the best performance in all four metrics. This is because the vector can represent all of the emotional information extracted from the users’ microblogs and incorporates all kinds of emotion along with their weights into recommendation. In

Table 4. Emotion Modeling Methods Example

Modeling Methods	Emotion Value ($E_1, E_2, E_3, E_4, E_5, E_6, E_7$)
Extracted Emotion	(0.7, 0.3, 0.0, 0.0, 0.0, 0.0, 0.0)
Vector	(0.7, 0.3, 0.0, 0.0, 0.0, 0.0, 0.0)
Multi-Label	(0.5, 0.5, 0.0, 0.0, 0.0, 0.0, 0.0)
Single-Label	(1.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0)

other words, all emotions play important roles in emotion-aware music recommendation; (2) EAGRec-Single-Label outperforms EAGRec-Multi-Label and the relative improvement in terms of precision with n set as 15 is 4.6%. In our view, the reason is that primary emotions and secondary emotions make a different contribution to the recommendation and should not be treated equally. In addition, primary emotions are more important than secondary emotions in users' emotion-aware music prediction and recommendation, and secondary emotions may have a negative impact if not utilized properly.

**Fig. 4.** Comparison of Different Methods

4.3 Comparison with Baselines

Our second experiment is designed to compare the proposed approach with different baseline recommendation methods, namely Bayesian Personalized Ranking

(BPR) [15] and Alternating Least Square for Personalized Ranking (RankALS) [17] which do not incorporate contextual information, together with User Collaborative Filtering with Emotion (UCFE) [4] which considers users' emotional context. We choose the best performing EAGRec-Vector as the representative for clear comparison.

As shown in Fig. 5, EAGRec-Vector outperforms baseline methods. Taking precision as an example, when compared with BPR, RankALS, and UCFE with recommendation number n set as 15, the relative performance improvement by EAGRec-Vector is around 62.6%, 70.3%, and 41.8%, respectively. The improvement over the baseline methods shows that: (1) emotional context plays an important role in the users' music preference, and the users' music listening behaviors and preferences are relevant to their emotional context; (2) extracting users' emotion from their microblogs is feasible and useful; (3) our proposed EAG is more effective for incorporating emotional data into the graph than UCFE and the random-walk-based recommendation approach is also effective for emotion aware music recommendation.

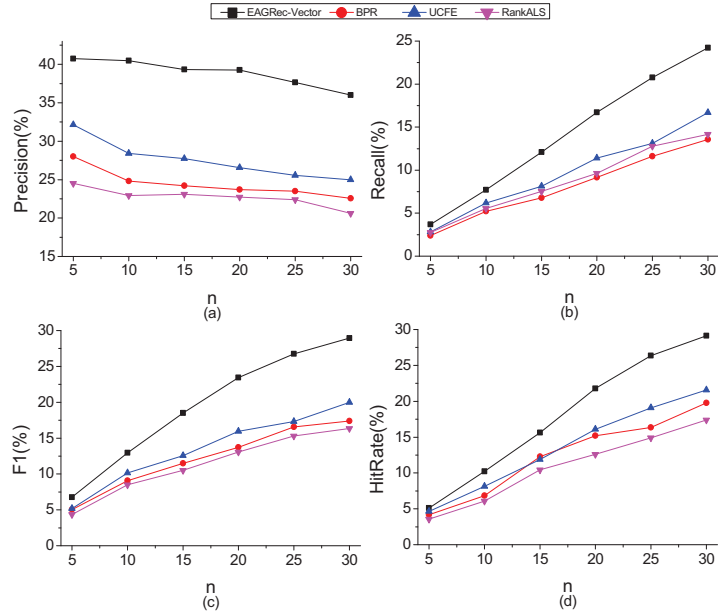


Fig. 5. Comparison of Different Methods

In conclusion, the emotional context information from the user's microblogs influences the users' music preferences. In addition, our proposed graph-based emotion-aware music recommendation methods can effectively incorporate users' emotional context and improve the performance of music recommendation.

5 Conclusion

This paper presents a graph-based approach for emotion-aware music recommendation, which can recommend music items appropriate for the users' current emotional context to satisfy their real-time requirements. The proposed approach models the relations between user, music, and emotion as a three-element tuple (user, music, emotion), upon which an EAG was built, and then a relevance propagation algorithm based on random walk was devised to rank the relevance of music items for recommendation. The experimental evaluation shows that the proposed method outperforms the traditional approach which does not consider the user's emotion. In addition, we study in detail how users' emotional context extracted from their microblogs can affect the performance of music recommendation. The results show that all parts of the extracted emotions are useful for music recommendation, and incorporating emotions along with their weight can improve the recommendation performance to the maximum extent possible.

Based on our current work, there are two possible future directions. First, we will attempt to connect a microblog service (such as Twitter, Sina Weibo) with music service websites (such as Xiami Music, Last.fm) to extract more tuples (user, music, emotion), and provide better recommendation results, especially when the user does not have many music microblogs. Secondly, we will explore if the users' satisfaction can be increased when users listen to recommended music using online experiments.

Acknowledgements This research work is supported in part by the National Key Technology Research and Development Program of China under Grant 2014BAD10B02, the Fundamental Research Funds for the Central Universities 2016FZA5012, and Australian Research Council (ARC) Linkage Project under No. LP140100937.

References

1. Abowd, G.D., Dey, A.K., Brown, P.J., Davies, N., Smith, M., Steggle, P.: Towards a better understanding of context and context-awareness. In: International Symposium on Handheld and Ubiquitous Computing. pp. 304–307. Springer (1999)
2. Baeza-Yates, R., Ribeiro-Neto, B., et al.: Modern information retrieval, vol. 463. ACM press New York (1999)
3. Chang, C.Y., Lo, C.Y., Wang, C.J., Chung, P.C.: A music recommendation system with consideration of personal emotion. In: International Computer Symposium (ICS). pp. 18–23. IEEE (2010)
4. Deng, S., Wang, D., Li, X., Xu, G.: Exploring user emotion in microblogs for music recommendation. *Expert Systems with Applications* 42(23), 9284–9293 (2015)
5. Ekman, P.E., Davidson, R.J.: The nature of emotion: Fundamental questions. Oxford University Press (1994)
6. Gori, M., Pucci, A.: Itemrank: a random-walk based scoring algorithm for recommender engines. In: Proceedings of the 20th International Joint Conference on Artificial Intelligence. pp. 2766–2771. Morgan Kaufmann Publishers Inc. (2007)

7. Han, B.j., Rho, S., Jun, S., Hwang, E.: Music emotion classification and context-based music recommendation. *Multimedia Tools and Applications* 47(3), 433–460 (2010)
8. Haveliwala, T.H.: Topic-sensitive pagerank: A context-sensitive ranking algorithm for web search. *IEEE Transactions on Knowledge and Data Engineering* 15(4), 784–796 (2003)
9. Kaminskas, M., Ricci, F.: Location-adapted music recommendation using tags. In: *International Conference on User Modeling, Adaptation, and Personalization*. pp. 183–194. Springer (2011)
10. Kaminskas, M., Ricci, F.: Contextual music information retrieval and recommendation: State of the art and challenges. *Computer Science Review* 6(2), 89–119 (2012)
11. Knees, P., Schedl, M.: A survey of music similarity and recommendation from music context data. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)* 10(1), 2 (2013)
12. North, A., Hargreaves, D.: *The social and applied psychology of music*. OUP Oxford (2008)
13. North, A.C., Hargreaves, D.J.: Situational influences on reported musical preference. *Psychomusicology: A Journal of Research in Music Cognition* 15(1-2), 30 (1996)
14. Pettijohn II, T.F., Williams, G.M., Carter, T.C.: Music for the seasons: seasonal music preferences in college students. *Current Psychology* 29(4), 328–345 (2010)
15. Rendle, S., Freudenthaler, C., Gantner, Z., Schmidt-Thieme, L.: Bpr: Bayesian personalized ranking from implicit feedback. In: *Proceedings of the twenty-fifth Conference on Uncertainty in Artificial Intelligence*. pp. 452–461. AUAI Press (2009)
16. Shan, M.K., Kuo, F.F., Chiang, M.F., Lee, S.Y.: Emotion-based music recommendation by affinity discovery from film music. *Expert Systems with Applications* 36(4), 7666–7674 (2009)
17. Takács, G., Tikk, D.: Alternating least squares for personalized ranking. In: *Proceedings of the sixth ACM Conference on Recommender Systems*. pp. 83–90. ACM (2012)
18. Tong, H., Faloutsos, C., Pan, J.y.: Fast random walk with restart and its applications. In: *Proceedings of the sixth International Conference on Data Mining*. pp. 613–622. IEEE (2006)
19. Wang, K., Zhang, R., Liu, X., Guo, X., Sun, H., Huai, J.: Time-aware travel attraction recommendation. In: *International Conference on Web Information Systems Engineering*. pp. 175–188. Springer (2013)
20. Wang, X., Rosenblum, D., Wang, Y.: Context-aware mobile music recommendation for daily activities. In: *Proceedings of the 20th ACM international Conference on Multimedia*. pp. 99–108. ACM (2012)
21. Wu, H., Shao, J., Yin, H., Shen, H.T., Zhou, X.: Geographical constraint and temporal similarity modeling for point-of-interest recommendation. In: *International Conference on Web Information Systems Engineering*. pp. 426–441. Springer (2015)
22. Yuan, B., Xu, B., Chung, T., Shuai, K., Liu, Y.: Mobile phone recommendation based on phone interest. In: *International Conference on Web Information Systems Engineering*. pp. 308–323. Springer (2014)