A Survey of Music Recommendation Systems and Future Perspectives

Yading Song, Simon Dixon, and Marcus Pearce *

Centre for Digital Music
Queen Mary University of London
{yading.song, simon.dixon, marcus.pearce}@eecs.qmul.ac.uk

Abstract. Along with the rapid expansion of digital music formats, managing and searching for songs has become significant. Though music information retrieval (MIR) techniques have been made successfully in last ten years, the development of music recommender systems is still at a very early stage. Therefore, this paper surveys a general framework and state-of-art approaches in recommending music. Two popular algorithms: collaborative filtering (CF) and content-based model (CBM), have been found to perform well. Due to the relatively poor experience in finding songs in long tail and the powerful emotional meanings in music, two user-centric approaches: context-based model and emotion-based model, have been paid increasing attention. In this paper, three key components in music recommender - user modelling, item profiling, and match algorithms are discussed. Six recommendation models and four potential issues towards user experience, are explained. However, subjective music recommendation system has not been fully investigated. To this end, we propose a motivation-based model using the empirical studies of human behaviour, sports education, music psychology.

Keywords: Music recommendation; music information retrieval; collaborative filtering; content-based model; emotion-based model; motivation-based model; music psychology

1 Introduction

With the explosion of network in the past decades, internet has become the major source of retrieving multimedia information such as video, books, and music etc. People has considered that music is an important aspect of their lives and they listen to music, an activity they engaged in frequently. Previous research has also indicated that participants listened to music more often than any of the other activities [57] (i.e. watching television, reading books, and watching movies). Music, as a powerful communication and self-expression approach, therefore, has appealed a wealth of research.

* Yading is supported by the China Scholarship Council. We would like to thank Geraint A. Wiggins for his advices.

9th International Symposium on Computer Music Modelling and Retrieval (CMMR 2012)
19-22 June 2012, Queen Mary University of London

However, the problem now is to organise and manage the million of music titles produced by society [51]. MIR techniques have been developed to solve problems such as genre classification [42, 75], artist identification [46], and instrument recognition [49]. Since 2005, an annual evaluation event called Music Information Retrieval Evaluation eXchange (MIREX[1]) is held to facilitate the development of MIR algorithms.

Additionally, music recommender is to help users filter and discover songs according to their tastes. A good music recommender system should be able to automatically detect preferences and generate playlists accordingly. Meanwhile, the development of recommender systems provides a great opportunity for industry to aggregate the users who are interested in music. More importantly, it raises challenges for us to better understand and model users' preferences in music [76].

Currently, based on users' listening behaviour and historical ratings, collaborative filtering algorithm has been found to perform well [9]. Combined with the use of content-based model, the user can get a list of similar songs by low-level acoustic features such as rhythm, pitch or high-level features like genre, instrument etc [7].

Some music discovery websites such as Last.fm[2], Allmusic[3], Pandora[4] and Shazam[5] have successfully used these two approaches into reality. At the meantime, these websites provide an unique platform to retrieve rich and useful information for user studies.

Music is subjective and universal. It not only can convey emotion, but also can it modulate a listener's mood [23]. The tastes in music are varied from person to person, therefore, the previous approaches cannot always meet the users' needs. An emotion-based model and a context-based model have been proposed [18, 34]. The former one recommends music based on mood which allows the user to locate their expected perceived emotion on a 2D valence-arousal interface [22]. The latter one collects other contextual information such as comments, music review, or social tags to generate the playlist. Though hybrid music recommender systems would outperform the conventional models, the development is still at very early stage [88]. Due to recent studies in psychology, signal processing, machine learning and musicology, there is much room for future extension.

This paper, therefore, surveys a general music recommender framework from user profiling, item modelling, and item-user profile matching to a series of state-of-art approaches. Section 2 gives a brief introduction of components in music recommendation systems and in section 3, the state-of-art recommendation techniques are explained. To the end of this paper, we conclude and propose a new model based on users' motivation.

[1] http://www.music-ir.org/mirex/wiki/MIREX_HOME
[2] http://www.last.fm/
[3] http://www.allmusic.com/
[4] http://www.pandora.com
[5] http://www.shazam.com/

2 Components in Music Recommender System

Generally, a music recommender system consists of three key components - users, items and user-item matching algorithms. User profiling (see section 2.1) addresses the variation in users' profile. This step aims at differentiating their music tastes using basic information. Item profiling (see section 2.2) on the contrary, describes three different types of metadata - editorial, cultural and acoustic, which are used in different recommendation approaches. In section 2.3, we explain the query in music recommender systems, and the matching algorithms are presented in section 3.

2.1 User Modelling

A successful music recommender needs to meet users' various requirements. However, obtaining user information is expensive in terms of financial costs and human labor [74]. For user-oriented design, lots of efforts on user studies need to be investigated.

User modelling, as the one of the key elements, it models the difference in profile. For example, the difference in geographic region or age, their music preferences might be different. Interestingly, other factors such as gender, life styles, and interests could also determine their choices of music.

Recent research has revealed that intelligence, personality and the users' preference in music are linked [57]. According to Rentfrow and Gosling [26, 58] who had investigated the relationship between music preference and Big-Five Inventory (BFI: openness, conscientiousness, extraversion, agreeableness, and neuroticism), their studies showed a highly extraverted person would tend to choose the music which is energetic, while a greater preference for rhythmic and energetic music was associated with greater extraversion and agreeableness. User modelling, therefore, is essential in prediction of their music taste. It has been divided into two parts: user profile modelling and user experience modelling.

First Step - User Profile Modelling Celma [14] suggested that the user profile can be categorised into three domains: *demographic*, *geographic*, and *psychographic* (shown in Table 1). Based on the steadiness, psychological data has been further divided into stable attributes which are essential in making a long term prediction and fluid attributes which can change on an hour to hour basis [24].

Data type	**Example**
Demographic	Age, marital status, gender etc.
Geographic	Location, city, country etc.
Psychographic	*Stable*: interests, lifestyle, personality etc. *Fluid*: mood, attitude, opinions etc.

Table 1. User profile classification

Second Step - User Listening Experience Modelling Depending on the level of music expertise, their expectations in music are varied accordingly. Jennings [32] analysed the different types of listeners whose age range from 16-45 and categorised the listeners into four groups: *savent*, *enthusiasts*, *casuals*, *indifferents* (see Table 2).

Type	Percentage	Features
Savants	7	Everything in life seems to be tied up with music. Their musical knowledge is very extensive.
Enthusiasts	21	Music is a key part of life but is also balanced by other interests.
Casuals	32	Music plays a welcome role, but other things are far more important.
Indifferents	40	They would not lose much sleep if music ceased to exist, they are a predominant type of listeners of the whole population.

Table 2. Use listening experience categorisation

This information gives us a good example that their expertise needs to be considered when designing user-oriented recommendation systems. For instance, based on their expectation, we need to consider the amount of music to be discovered and filtered in long tail which represents interesting and unknown music but hidden in the tail of the popularity curve [3]. Other user information including access pattern, listening behaviour are also useful for user modelling and dynamic optimisation [50]. Exploring user information can be either done through the initial survey or observing their behaviour of music in long tail.

2.2 Item Profiling

The second component of recommender systems is music item. It defines a various of information that used in MIR. In 2005, Pachet [53] classified the music metadata into three categories: *editorial metadata* (EM), *cultural metadata* (CM), and *acoustic metadata* (AM).

- **Editorial metadata**: Metadata obtained by a single expert or group of experts. This is obtained literally by the editor, and also it can be seen as the information provided by them. E.g. the cover name, composer, title, or genre etc.
- **Cultural metadata**: Metadata obtained from the analysis of corpora of textual information, usually from the Internet or other public sources. This information results from an analysis of emerging patterns, categories or associations from a source of documents. E.g. Similarity between music items.
- **Acoustic metadata**: Metadata obtained from an analysis of the audio signal. This should be without any reference to a textual or prescribed information. E.g. Beat, tempo, pitch, instrument, mood etc.

Editorial metadata are mostly used in metadata information retrieval (see section 3.1), and cultural metadata have been largely used in context-based information retrieval (see section 3.5). However, most music recommendation systems are using acoustic metadata for discovering music which is named as content-based information retrieval (see section 3.3).

2.3 Query Type

Assuming that the users have already known the information about the music, the quickest way to search for music is via key editorial information such as title, the name of the singer and lyrics etc. However, it is not always the case of knowing them. In the past ten years, an advanced and more flexible music information retrieval system called "query by humming/singing system (QBSH)" was developed [25]. It allows the user to find the songs either by humming or singing.

Nevertheless, it still requires lots of human efforts. In recommender systems, a more appropriate way is to use listening histories or seed music as the query to detect their music preferences.

3 State-of-art Approaches in Music Recommendation

An ideal music recommender system should be able to automatically recommend personalised music to human listeners [36, 52]. Different from books or movies, the length of a piece of music is much shorter, and the times that listening their favourite songs are normally more than once.

The existing recommender systems such as *Amazon*, *Ebay* have gained a great success. It can recommend complementary goods, the buyer can compare the products (new-item/old-item) and negotiate with the sellers [69]. However, music recommender is not only giving products with reasonable price, but suggesting them personalised music.

So far, many music discovery websites such as *Last.fm*, *Allmusic*, *Pandora*, *Audiobaba*[6], *Mog*[7], *Musicovery* [8], *Spotify*[9], *Apple "Genius"* have aggregated millions of users, and the development is explosive [10, 11]. In this section, we present the most popular approaches, metadata information retrieval (see section 3.1), collaborative filtering (see section 3.2), content-based information retrieval (see section 3.3), emotion-based model (see section 3.4), context-based information retrieval (see section 3.5) and hybrid models (see section 3.6). At the end of each approach, their limitations are described.

3.1 Metadata Information Retrieval (Demographic-based Model)

As the most fundamental method, it is the easiest way to search for music. Metadata information retrieval uses textual metadata (editorial information)

[6] http://audiobaba.com/
[7] http://www.mog.com/
[8] http://www.musicovery.com/
[9] http://www.spotify.com/

supplied by the creators, such as the title of the song, artist name, and lyrics to find the target songs [20].

Limitation Though it is fast and accurate, the drawbacks are obvious. First of all, the user has to know about the editorial information for a particular music item. Secondly, it is also time consuming to maintain the increasing metadata. Moreover, the recommendation results is relatively poor, since it can only recommend music based on editorial metadata and none of the users' information has been considered.

3.2 Collaborative Filtering

To recommend items via the choice of other similar users, collaborative filtering technique has been proposed [28]. As one of the most successful approaches in recommendation systems, it assumes that if user X and Y rate n items similarly or have similar behaviour, they will rate or act on other items similarly [59].

Instead of calculating the similarity between items, a set of 'nearest neighbour' users for each user whose past ratings have the strongest correlation are found. Therefore, scores for the unseen items are predicted based on a combination of the scores known from the nearest neighbours [65]. Collaborative filtering is further divided into three subcategories: *memory-based*, *model-based*, and *hybrid* collaborative filtering [63, 68].

Memory-based Collaborative Filtering Memory-based collaborative filtering is to predict the item based on the entire collections of previous ratings. Every user is grouped with people with similar interests, so that a new item is produced by finding the nearest neighbour using a massive number of explicit user votes [9].

Model-based Collaborative Filtering In contrast to memory-based CF, model-based CF uses machine learning and data mining algorithms which allow the system to train and model the users' preferences. It represents the user preference by a set of rating scores and constructs a special prediction model [2]. Based on the known model, the system makes prediction for test and real-world data.

Hybrid Collaborative Filtering Hybrid CF model is to make prediction by combining different CF models. It has been proved that hybrid CF model outperforms any individual method [83].

Limitations Because of the subjectivity in music, the assumption that users with similar behaviours may have same tastes has not been widely studied. Though collaborative filtering recommender works well, the key problems such as cold start, popularity bias are unavoidable [27].

- **Popularity bias** Generally, popular music can get more ratings. The music in long tail, however, can rarely get any. As a result, collaborative filtering mainly recommend the popular music to the listeners. Though giving popular items are reliable, it is still risky, since the user rarely get pleasantly surprised.
- **Cold start** It is also known as data sparsity problems. At an early stage, few ratings is provided. Due to the lack of these ratings, prediction results are poor.
- **Human effort** A perfect recommender system should not involve too much human efforts, since the users are not always willing to rate. The ratings may also grow towards those who do rate, but it may not be representative. Because of this absence of even distributed ratings, it can either give us false negative or false positive results.

3.3 Content/Audio/Signal-based Music Information Retrieval

Different from collaborative filtering technique, content-based approach makes prediction by analysing the song tracks [2, 41]. It is rooted in information retrieval and information filtering [13] that recommends a song which is similar to those the user has listened to in the past rather than what the user have rated 'like' [4, 43]. Lots of research have been paid attention on extracting and comparing the acoustic features in finding perceptual similar tracks [8, 45]. The most representative ones so far are timbre, rhythm [7, 10, 11].

Based on the extracted features, the distance between songs are measured [43]. Three typical similarity measurements are listed below [44].

- **K-means clustering with Earth-Mover's Distance**: It computes a general distance between Gaussian Mixture Models (GMM) by combining individual distance between gaussian components [62].
- **Expectation-Maximization with Monte Carlo Sampling**: This measurement makes use of vectors sampled directly from the GMMs of the two songs to be compared; the sampling is performed computationally via random number generation [51].
- **Average Feature Vectors with Euclidean Distance**: It calculates low-order statistics such as mean and variance over segments [16].

Query by Humming (QBSH) Humming and singing are the natural way to express the songs [31]. In the early 1990s, based on content-based model, query by humming system was proposed [25, 80]. Early query by humming systems were using melodic contour which had been seen as the most discriminative features in songs.

It follows three steps: construction of the songs database, transcription of the users' melodic information query and pattern matching algorithms which are used to get the closest results from collections [1]. In the past few years, except melody, a better performance has also been achieved by embedding with lyrics and enhancing the main voice [19, 77].

Limitations To some extent, content-based model solves the problems in collaborative filtering. For instance, by measuring the similarity of acoustic features between songs, the system can recommend music using distance measurements. Therefore, no human rating is needed. However, similarity-based method has not been fully investigated in terms of listeners' preference. None of the research proved that similar behaviour leads to the choice of same music.

Since content-based model largely depends on acoustic features, the number of selected features needs to be further considered. Moreover, other user information and non-acoustic information should be included for future modification and augmentation.

3.4 Emotion-based Model

Music as a self-expression tool, it always performs with affection. Rich in content and expressivity [86], the conventional approaches for music information retrieval are no longer sufficient. Music emotion has appealed lots of research and it has become the main trend for music discovery and recommendation [34]. A commercial web service called '*Musicovery*' uses the fundamental emotion model (2D valence-arousal) found by psychologists. It allows users to locate their expected perceived emotion in a 2D space: *valence* (how positive or negative) and *arousal* (how exciting or calming).

Similar to content-based model, the emotion perception is associated with different patterns of acoustic cues [6, 35, 48, 61]. Different perceptual features such as energy, rhythm, temporal, spectral, and harmony have been widely used in emotion recognition [84].

Limitations

- **Data collection** In order to accurately model the system, a great amount of dataset are needed. However, finding the reliable ground truth is expensive and requires a lot of human efforts [67]. Instead of human annotation [73], social tags [36, 74], annotation games like *MajorMiner* [47] and *TagATune* [37], lyrics or music review are being used for data collection.
- **Ambiguity and granularity** Emotion itself is hard to define and describe. The same affective feeling experienced by different people may give different emotion expression (i.e. cheerful, happy) and there is no perfect relationship between affective terms with emotions [64, 85]. Some research were based on basic taxonomy (sad, happy, angry etc.), but it cannot describe the richness of our human perception. MIREX evaluation has categorised emotion into 5 mood clusters [29] (see Table 3). Russell [60] found a circumflex model which affective concepts fall in a circle in the following order: pleasure (0°), excitement (45°), arousal (90°), distress (135°), displeasure (180°), depression (225°), sleepiness (270°), and relaxation (315°). It can represent the structure of affective experience and now it has been become the most noted 2D valence-arousal emotion model. The problem of classifying emotion, therefore has been solved, since each point on the plane represents an affective term.

Cluster 1	Passionate, rousing, confident, boisterous, rowdy
Cluster 2	Rollicking, cheerful, fun, sweet, amiable/good natured
Cluster 3	Literate, poignant, wistful, bittersweet, autumnal, brooding
Cluster 4	Humorous, silly, campy, quirky, whimsical, witty, wry
Cluster 5	Aggressive, fiery, tense/anxious, intense, volatile, visceral

Table 3. MIREX five mood categories

3.5 Context-based Information Retrieval

Rather than using acoustic features in content-based model and ratings in collaborative filtering, context-based information retrieval model uses the public opinion to discover and recommend music [18]. Along with the development of social networks such as *Facebook*[10], *Youtube*[11], and *Twitter*[12], these websites provide us rich human knowledge such as comments, music review, tags and friendship networks [36].

Context-based information retrieval, therefore, uses web/document mining techniques to filter out important information to support problems like artist similarity, genre classification, emotion detection [82], semantic space [39, 40] etc. Some researchers have suggested that the use of social information has outperformed content-based model [70, 81].

However, the same problems as collaborative filtering, the popular music can always get more public opinions than those in long tail [21]. Eventually, rich music gets richer feedback, it again results in a popularity bias problem.

3.6 Hybrid Model Information Retrieval

Hybrid model aims at combining two or more models to increase the overall performance. Burke [9] pointed out several methods to build a hybrid model such as *weighted*, *switching*, *mixed*, *feature combination*, and *cascade*. There is no doubt that a proper hybrid model would outperform a single approach, since it can incorporate the advantages of both methods while inheriting the disadvantages of neither [65, 87, 88].

3.7 Other Issues

We have discussed above the essential problems in music recommender systems, the other issues such as dynamic evolvement, playlist generation, user interface design and evaluation need to be further considered. Though it doesn't affect the recommendation performance, it certainly influence the user listening experience.

[10] https://www.facebook.com/
[11] http://www.youtube.com/
[12] https://twitter.com/

Dynamic Evolvement As the users aggregate in the recommender systems, it needs to be able to adapt to new data such as user listening histories and listening behaviour to further personalised their music taste [30]. This procedure is called evolvement. It addresses the problem that when the new user comes and new items into the system, it can dynamically and automatically evolve itself [65].

Playlist Generation Another issue is the sequence of the playlist [38]. Most of the recommender systems are not flexible, because the playlist is ordered by the similarity distance between seed songs. Though the most similar songs are given in order, the theme and mood can be dramatically changed in between. This may result in the dissatisfaction and discontinuation of the songs.

Research indicates that a playlist should have a main theme (mood, event, activity) evolve with time [17]. Rather than randomly shuffling, human skipping behaviour can be considered for dynamic playlist generation [15, 54]. For example, assuming that the users dislike the song when they skipped it, the system therefore, removes the songs which are similar to the song which they skipped [55, 56, 78].

User Interface Design A bad design of user interface cannot affect the accuracy of the system, it does influence the ratings and listening experience. A clear design always gives the user a better understanding of the system. Moreover, an overall control of the system and less human efforts required for operation should be considered during designing.

Evaluation There is no common objective evaluation in music recommendation systems [72]. Most of the evaluation techniques are based on subjective system testing which let users to rank the systems given the playlist generated by different approaches [5, 79]. However, it is very expensive in terms of financial costs and human labor. Another important factor is that the evaluation in different regions (i.e. different background, age, language) might give different results. Hence, a proper evaluation criteria is essential and highly recommended.

4 Conclusion and Future Work

In this paper, we explain a basic metadata-based model and two popular music recommender approaches: collaborative filtering and content-based model. Though they have achieved great success, their drawbacks such as popularity bias and human efforts are obvious. Moreover, the use of hybrid model would outperform a single model since it incorporates the advantages of both methods. Its complexity is not fully studied yet.

Due to the subjective nature in music and the issues existing in the previous methods, two human-centred approaches are proposed. By considering affective and social information, emotion-based model and context-based model largely improved the quality of recommendation. However, this research is still at an early stage.

As we can see from the development of music recommenders over the past years, the given results tend to be more personalised and subjective. Only considering the music itself and human ratings are no longer sufficient. A great amount of work in recent years have been done in music perception, psychology, neuroscience and sport which study the relationship between music and the impact of human behaviour. David Huron also mentioned music has sex and drug-like qualities. Undoubtably, music always has been an important component of our life, and now we have greater access to it.

Researches in psychology pointed out that music not only improves mood, increases activation, visual and auditory imagery, but also recalls of associated films or music videos and relieves stress [33]. Moreover, the empirical experiments in sport mentioned that the main benefits for listening to the music which include work output extension, performance enhancement, and dissociation from unpleasant feelings etc [71]. For example, athletes prefer uptempo, conventional, intense, rebellious, energetic, and rhythmic music rather than reflective and complex music [66]. An important fact found by psychologists is that users' preference in music is linked to their personality. Also worth mentioning that fast, upbeat music produces a stimulative effect whereas slow, while soft music produces a sedative effects [12]. All of these highlight that music recommender is not only a tool for relaxing, but also acts as an effective tool to meet our needs under different contexts. To our knowledge, there is few research based on these empirical results.

Designing a personalised music recommender is complicated, and it is challenging to thoroughly understand the users' needs and meet their requirements. As discussed above, the future research direction will be mainly focused on user-centric music recommender systems. A survey among athletes showed practitioners in sport and exercise environments tend to select music in a rather arbitrary manner without full consideration of its motivational characteristics. Therefore, future music recommender should be able to lead the users reasonably choose music. To the end, we are hoping that through this study we can build the bridge among isolated research in all the other disciplines.

Acknowledgments. Yading is supported by the China Scholarship Council. We would like to thank Geraint A. Wiggins, Steven Hargreaves and Emmanouil Benetos for their advices.

References

1. Ricardo A. Baeza-Yates and Chris H. Perleberg. Fast and Practical Approximate String Matching. In *Combinatorial Pattern Matching, Third Annual Symposium*, pages 185–192, 1992.
2. G. Adomavicius and A. Tuzhilin. Toward the Next Generation of Recommender Systems: A Survey of the State-of-the-art and Possible Extensions. *IEEE Transactions on Knowledge and Data Engineering*, 17(6):734–749, June 2005.
3. C Anderson. *The Long Tail. Why the Future of Business is selling less of more.* Hyperion Verlag, 2006.

4. Jean-julien Aucouturier and Francois Pachet. Music Similarity Measures: What is the Use. In *Proceedings of the ISMIR*, pages 157–163, 2002.
5. Luke Barrington, Reid Oda, and G. Lanckriet. Smarter Than Genius? Human Evaluation of Music Recommender Systems. In *10th International Society for Music Information Retrieval Conference*, number ISMIR, pages 357–362, 2009.
6. Kerstin Bischoff, Claudiu S Firan, Raluca Paiu, Wolfgang Nejdl, L S De, Cyril Laurier, and Mohamed Sordo. Music Mood and Theme Classification - A Hybrid Approach. In *10th International Society for Music Information Retrieval Conference*, number Ismir, pages 657–662, 2009.
7. Dmitry Bogdanov and Perfecto Herrera. How Much Metadata Do We Need in Music Recommnendation? A Subjective Evaluation Using Preference Sets. In *12th International Society for Music Information Retrieval Conference*, number ISMIR 2011, pages 97–102, 2011.
8. Dmitry Bogdanov, J. Serra, Nicolas Wack, Perfecto Herrera, and Xavier Serra. Unifying Low-level and High-level Music Similarity Measures. *IEEE Transactions on Multimedia*, 13(99):1–1, 2011.
9. Robin Burke. Hybrid Recommender Systems: Survey and Experiments. *User Modeling and User-Adapted Interaction*, 12(4):331–370, 2002.
10. Pedro Cano, Markus Koppenberger, and Nicolas Wack. An Industrial-strength Content-based Music Recommendation System. In *Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval - SIGIR '05*, page 673, New York, New York, USA, 2005. ACM Press.
11. Pedro Cano, Markus Koppenberger, and Nicolas Wack. Content-based Music Audio Recommendation. In *Proceedings of the 13th annual ACM international conference on Multimedia*, number ACM, pages 211–212, 2005.
12. FRD Carpentier. Effects of Music on Physological Arousal, Exploration into Tempo and Genre. *Media Psychology*, 10(3):339–363, 2007.
13. M.A. Casey, Remco Veltkamp, Masataka Goto, Marc Leman, Christophe Rhodes, and Malcolm Slaney. Content-based Music Information Retrieval: Current Directions and Future Challenges. *Proceedings of the IEEE*, 96(4):668–696, 2008.
14. O. Celma Herrada. Music Recommendation and Discovery in the Long Tail. *PhD Thesis*, 2009.
15. Zeina Chedrawy and S. Abidi. A Web Recommender System for Recommending, Predicting and Personalizing Music Playlists. In *Web Information Systems Engineering-WISE 2009*, pages 335–342. Springer, 2009.
16. Parag Chordia, Mark Godfrey, and Alex Rae. Extending Content-Based Recommendation: The Case of Indian Classical Music. In *ISMIR 2008: proceedings of the 9th International Conference of Music Information Retrieval*, pages 571–576, 2008.
17. Sally Jo Cunningham, David Bainbridge, and Annette Falconer. More of an Art than a Science : Supporting the Creation of Playlists and Mixes. In *Seventh International Conference on Music Information Retrieval*, Victoria, Canada, 2006.
18. M. Datcu, H. Daschiel, A. Pelizzari, M. Quartulli, A. Galoppo, A. Colapicchioni, M. Pastori, K. Seidel, P.G. Marchetti, and S. D'Elia. Music Recommendation Using Content and Context Information Mining. *Geoscience and Remote Sensing, IEEE Transactions on*, 41(12):2923–2936, 2003.
19. C. de la Bandera, A.M. Barbancho, L.J. Tardón, Simone Sammartino, and Isabel Barbancho. Humming Method for Content-Based Music Information Retrieval. *ISMIR 2011*, (Ismir):49–54, 2011.
20. J. Stephen Downie. Music Information Retrieval. *Annual Review of Information Science and Technology*, 37(1):295–340, January 2005.

21. Douglas Eck, Paul Lamere, T. Bertin-Mahieux, and Stephen Green. Automatic Generation of Social Tags for Music Recommendation. *Advances in neural information processing systems*, 20:385–392, 2007.
22. T. Eerola and J. K. Vuoskoski. A Comparison of the Discrete and Dimensional Models of Emotion in Music. *Psychology of Music*, 39(1):18–49, August 2010.
23. Yazhong Feng and Y Zhuang. Popular Music Retrieval by Detecting Mood. In *International Society for Music Information Retrieval 2003*, volume 2, pages 375–376, 2003.
24. Alan Page Fiske, Shinobu Kitayama, Hazel Rose Markus, and R E Nisbett. *The Cultural Matrix of Social Psychology.* 1998.
25. Asif Ghias, Jonathan Logan, David Chamberlin, and Brian C. Smith. Query by Humming. *Proceedings of the third ACM international conference on Multimedia - MULTIMEDIA '95*, pages 231–236, 1995.
26. S Gosling. A Very Brief Measure of the Big-Five Personality Domains. *Journal of Research in Personality*, 37(6):504–528, December 2003.
27. Jonathan L. Herlocker, Joseph a. Konstan, Loren G. Terveen, and John T. Riedl. Evaluating Collaborative Filtering Recommender Systems. *ACM Transactions on Information Systems*, 22(1):5–53, January 2004.
28. Will Hill, Larry Stead, Mark Rosenstein, George Furnas, and South Street. Recommending and Evaluating Choices in a Vitual Community of Use. *Mosaic A Journal For The Interdisciplinary Study Of Literature*, pages 5–12, 1995.
29. X Hu and J. Stephen Downie. Exploring Mood Metadata: Relationships with Genre, Artist and Usage Metadata. In *8th International Conference on Music Information Retrieval*, 2007.
30. Yajie Hu and Mitsunori Ogihara. Nextone Player: A Music Recommendation System Based on User Behavior. In *12th International Society for Music Information Retrieval Conferenceh*, number Ismir, pages 103–108, 2011.
31. Jyh-Shing Roger Jang and Hong-Ru Lee. A General Framework of Progressive Filtering and Its Application to Query by Singing/Humming. *IEEE Transactions on Audio, Speech, and Language Processing*, 16(2):350–358, February 2008.
32. D Jennings. *Net, Blogs and Rock 'n' Rolls: How Digital Discovery Works and What It Means for Consumers.* 2007.
33. Patrik N Juslin and Daniel Västfjäll. Emotional Responses to Music: the Need to Consider Underlying Mechanisms. *The Behavioral and brain sciences*, 31(5):559–621, October 2008.
34. Y.E. Kim, E.M. Schmidt, Raymond Migneco, B.G. Morton, Patrick Richardson, Jeffrey Scott, J.A. Speck, and Douglas Turnbull. Music Emotion Recognition: A State of the Art Review. In *Proc. of the 11th Intl. Society for Music Information Retrieval (ISMIR) Conf*, number Ismir, pages 255–266, 2010.
35. Fang-Fei Kuo, Meng-Fen Chiang, Man-Kwan Shan, and Suh-Yin Lee. Emotion-based Music Recommendation by Association Discovery from Film Music. In *Proceedings of the 13th annual ACM international conference on Multimedia - MULTIMEDIA '05*, page 507, New York, New York, USA, 2005. ACM Press.
36. Paul Lamere. Social Tagging and Music Information Retrieval. *Journal of New Music Research*, 37(2):101–114, June 2008.
37. Edith L. M. Law, Luis Von Ahn, Roger B. Dannenberg, and Mike Crawford. Tagatune: A Game for Music and Sound Annotation. In *8th International Conference on Music Information Retrieval*, 2007.
38. Jin Ha Lee, Bobby Bare, and Gary Meek. How Similar is too Similar? Exploring Users' Perception of Similarity in Playlist Evaluation. In *International Conference on Music Information Retrieval 2011*, number ISMIR, pages 109–114, 2011.

39. Mark Levy. A Semantic Space for Music Derived from Social Tags. *Austrian Compuer Society*, 1:12, 2007.
40. Mark Levy and Mark Sandler. Music Information Retrieval Using Social Tags and Audio. *IEEE Transactions on Multimedia*, 11(3):383–395, 2009.
41. Qing Li, Byeong Man Kim, Dong Hai Guan, and Duk Oh. A Music Recommender Based on Audio Features. In *Proceedings of the 27th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 532–533, Sheffield, United Kingdom, 2004. ACM.
42. Bass Lines, Emiru Tsunoo, George Tzanetakis, and Nobutaka Ono. Beyond Timbral Statistics : Improving Music Classification Using Percussive. *IEEE Transactions on Audio, Speech and Language Processing*, 19(4):1003–1014, 2011.
43. Beth Logan. Music Recommendation from Song Sets. In *International Conference on Music Information Retrieval 2004*, number October, pages 10–14, Barcelona, Spain, 2004.
44. Terence Magno and Carl Sable. A Comparison of Signal of Signal-Based Music Recommendation to Genre Labels, Collaborative Filtering, Musicological Analysis, Human Recommendation and Random Baseline. In *ISMIR 2008: proceedings of the 9th International Conference of Music Information Retrieval*, pages 161–166, 2008.
45. Chun-man Mak, Tan Lee, Suman Senapati, Yu-ting Yeung, and Wang-kong Lam. Similarity Measures for Chinese Pop Music Based on Low-level Audio Signal Attributes. In *11th International Society for Music Information Retrieval Conference*, number ISMIR, pages 513–518, 2010.
46. M Mandel. Song-level Features and Support Vector Machines for Music Classification. In *Proc. International Conference on Music*, 2005.
47. MI Mandel. A Web-based Game for Collecting Music Metadata. In *In 8th International Conference on Music Information Retrieval (ISMIR)*, 2008.
48. M Mann, TJ Cox, and FF Li. Music Mood Classification of Television Theme Tunes. In *12th International Society for Music Information Retrieval Conference*, number Ismir, pages 735–740, 2011.
49. Janet Marques and Pedro J Moreno. A Study of Musical Instrument Classification Using Gaussian Mixture Models and Support Vector Machines, 1999.
50. M. Ogihara, Bo Shao, Dingding Wang, and Tao Li. Music Recommendation Based on Acoustic Features and User Access Patterns. *IEEE Transactions on Audio, Speech, and Language Processing*, 17(8):1602–1611, November 2009.
51. F. Pachet and J.J. Aucouturier. Improving Timbre Similarity: How High is the Sky? *Journal of negative results in speech and audio sciences*, 1(1):1–13, 2004.
52. François Pachet and Daniel Cazaly. A Taxonomy of Musical Genres. In *Content-Based Multimedia Information Retrieval Access Conference (RIAO)*, number April, 2000.
53. Francois Pachet. Knowledge Management and Musical Metadata. In *Encyclopedia of Knowledge Management*. 2005.
54. Elias Pampalk, Tim Pohle, and Gerhard Widmer. Dynamic Playlist Generation Based on Skipping Behavior. In *Proc. of the 6th ISMIR Conference*, volume 2, pages 634–637, 2005.
55. Steffen Pauws, Berry Eggen, and Miles Davis. PATS : Realization and User Evaluation of an Automatic Playlist Generator PATS : Realization and User Evaluation of an Automatic Playlist Generator. In *3rd International Conference on Music Information Retrieval*, 2002.
56. CBE Plaza. Uncovering Affinity of Artist to Multiple Genres from Social Behavior Data. In *ISMIR 2008: proceedings of the 9th*, pages 275–280, 2008.

57. Peter J. Rentfrow and Samuel D. Gosling. The Do Re Mi's of Everyday Life: The structure and personality correlates of music preferences. *Journal of Personality and Social Psychology*, 84(6):1236–1256, 2003.
58. Peter J Rentfrow and Samuel D Gosling. Message in a Ballad: the Role of Music Preferences in Interpersonal Perception. *Psychological science*, 17(3):236–42, March 2006.
59. Paul Resnick, Hal R Varian, and Guest Editors. Recommender Systems. *Communications of the ACM*, 40(3):56–58, 1997.
60. J.A. Russell. A Circumplex Model of Affect. *Journal of personality and social psychology*, 39(6):1161–1178, 1980.
61. Pasi Saari, Tuomas Eerola, and Olivier Lartillot. Generalizability and Simplicity as Criteria in Feature Selection: Application to Mood Classification in Music. *Audio, Speech, and Language Processing, IEEE Transactions on*, 19(99):1–1, 2011.
62. A Salomon. A Content-based Music Similarity Function. In *Cambridge Research Labs-Tech Report*, number June, 2001.
63. Badrul Sarwar, George Karypis, and Joseph Konstan. Item-based Collaborative Filtering Recommendation Algorithms. *Proceedings of the 10th*, pages 285–295, 2001.
64. Bo Shao, Tao Li, and M. Ogihara. Quantify Music Artist Similarity Based on Style and Mood. In *Proceeding of the 10th ACM workshop on Web Information and Data Management*, pages 119–124. ACM, 2008.
65. Yoav Shoham and Marko Balabannovic. Content-Based, Collaborative Recommendation. *Communications of the ACM*, 40(3):66–72, 1997.
66. Stuart D Simpson and Costas I Karageorghis. The Effects of Synchronous Music on 400-m Sprint Performance. *Journal of sports sciences*, 24(10):1095–102, October 2006.
67. Janto Skowronek and M McKinney. Ground Truth for Automatic Music Mood Classification. In *Proc. ISMIR*, pages 4–5, 2006.
68. Xiaoyuan Su and Taghi M. Khoshgoftaar. A Survey of Collaborative Filtering Techniques. *Advances in Artificial Intelligence*, 2009(Section 3):1–19, 2009.
69. Neel Sundaresan. Recommender Systems at the Long Tail. In *of the fifth ACM conference on Recommender systems*, number RecSys 2011, pages 1–5, 2011.
70. Panagiotis Symeonidis, Maria Ruxanda, Alexandros Nanopoulos, and Yannis Manolopoulos. Ternary Semantic Analysis of Social Tags for Personalized Music Recommendation. In *Proc. 9th ISMIR Conf*, pages 219–224. Citeseer, 2008.
71. P.C. Terry and C.I. Karageorghis. Psychophysical Effects of Music in Sport and Exercise: An Update on Theory, Research and Application. In *Proceedings of the 2006 Joint Conference of the Australian Psychological Society and the New Zealand Psychological Society: Psychology Bridging the Tasman: Science, Culture and Practice*, pages 415–419. Australian Psychological Society, 2006.
72. Nava Tintarev and Judith Masthoff. Effective Explanations of Eecommendations: User-centered Design. In *Proceedings of the 2007 ACM conference on Recommender systems*, pages 153–156. ACM, 2007.
73. KTG Tsoumakas and George Kalliris. Multi-Label Classification of Music into Emotions. In *ISMIR 2008: proceedings of the 9th International Conference of Music Information Retrieval*, pages 325–330, 2008.
74. Douglas Turnbull, Luke Barrington, and Gert Lanckriet. Five Approaches to Collecting Tags for Music. In *ISMIR 2008: proceedings of the 9th International Conference of Music Information Retrieval*, pages 225–230, 2008.

75. George Tzanetakis, Student Member, and Perry Cook. Musical Genre Classification of Audio Signals. *IEEE Transactions on Speech and Audio Processing*, 10(5):293–302, 2002.
76. Alexandra Uitdenbogerd and van Schyndel Ron. A Review of Factors Affecting Music Recommender. In *3rd International Conference on Music Information Retrieval (2002)*, 2002.
77. Erdem Unal, Elaine Chew, Panayiotis G. Georgiou, and Shrikanth S. Narayanan. Challenging Uncertainty in Query by Humming Systems: A Fingerprinting Approach. *IEEE Transactions on Audio, Speech, and Language Processing*, 16(2):359–371, February 2008.
78. Rob van Gulik and Fabio Vignoli. Visual Playlist Generation on the Artist Map. In *5th International Conference on Music Infomation Retrieval*, number ISMIR2005, pages 520–523, 2005.
79. Fabio Vignoli. A Music Retrieval System Based on User-driven Similarity and its Evaluation. In *International Conference on Music Information Retrieval 2005*, 2005.
80. C.C. Wang, J.S.R. Jang, and Wennen Wang. An Improved Query by Singing/Humming System Using Melody and Lyrics Information. In *11th International Society for Music Information Retrieval Conference*, number Ismir, pages 45–50, 2010.
81. Dingding Wang, Tao Li, and Mitsunori Ogihara. Tags Better Than Audio Features? The Effect of Joint use of Tags and Audio Content Features for Artistic Style Clutering. In *International Conference on Music Information Retrieval 2010*, number ISMIR, pages 57–62, 2010.
82. Ju-chiang Wang, Hung-shin Lee, Hsin-min Wang, and Shyh-kang Jeng. Learning the Similarity of Audio Msuic in Bag-of-Frames Representation from Tagged Music Data. In *International Conference on Music Information Retrieval 2011*, number ISMIR, pages 85–90, 2011.
83. Jun Wang, A.P. De Vries, and M.J.T. Reinders. Unifying User-based and Item-based Collaborative Filtering Approaches by Similarity Fusion Categories. In *Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 501–508. ACM, 2006.
84. Xing Wang, Xiaoou Chen, Deshun Yang, and Yuqian Wu. Music Emotion Classification of Chinese Songs Based on Lyrics using TF*IDF and Rhyme. In *12th International Society for Music Information Retrieval Conference*, number Ismir, pages 765–770, 2011.
85. Dan Yang and W.S. Lee. Disambiguating Music Emotion Using Software Agents. In *Proceedings of the 5th International Conference on Music Information Retrieval (ISMIR04)*, pages 52–58, 2004.
86. Yi-Hsuan Yang. *Music Emotion Recognition*. Tayler and Francis Group, 2011.
87. Kazuyoshi Yoshii, Masataka Goto, Kazunori Komatani, Tetsuya Ogata, and Hiroshi G Okuno. Hybrid Collaborative and Content-based Music Recommendation Using Probabilistic Model with Latent User Preferences. In *Proceedings of the 7th International Conference on Music Information Retrieval*, pages 296–301, 2006.
88. Kazuyoshi Yoshii, Masataka Goto, Kazunori Komatani, Tetsuya Ogata, and Hiroshi G Okuno. Improving Efficiency and Scalability of Model-Based Music Recommender System Based on Incremental Training. In *ISMIR 2007: proceedings of the 8th International Conference of Music Information Retrieval*, number ISMIR, 2007.