



Current challenges and visions in music recommender systems research

Markus Schedl¹ · Hamed Zamani² · Ching-Wei Chen³ ·
Yashar Deldjoo⁴ · Mehdi Elahi⁵

Received: 28 September 2017 / Revised: 17 March 2018 / Accepted: 27 March 2018 / Published online: 5 April 2018
© The Author(s) 2018

Abstract

Music recommender systems (MRSs) have experienced a boom in recent years, thanks to the emergence and success of online streaming services, which nowadays make available almost all music in the world at the user's fingertip. While today's MRSs considerably help users to find interesting music in these huge catalogs, MRS research is still facing substantial challenges. In particular when it comes to build, incorporate, and evaluate recommendation strategies that integrate information beyond simple user–item interactions or content-based descriptors, but dig deep into the very essence of listener needs, preferences, and intentions, MRS research becomes a big endeavor and related publications quite sparse. The purpose of this trends and survey article is twofold. We first identify and shed light on what we believe are the most pressing challenges MRS research is facing, from both academic and industry perspectives. We review the state of the art toward solving these challenges and discuss its limitations. Second, we detail possible future directions and visions we contemplate for the further evolution of the field. The article should therefore serve two purposes: giving the interested reader an overview of current challenges in MRS research and providing guidance for young researchers by identifying interesting, yet under-researched, directions in the field.

Keywords Music recommender systems · Challenges · Automatic playlist continuation · User-centric computing

This research was supported in part by the Center for Intelligent Information Retrieval. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the sponsors.

✉ Markus Schedl
markus.schedl@jku.at
Hamed Zamani
zamani@cs.umass.edu
Ching-Wei Chen
cw@spotify.com
Yashar Deldjoo
yashar.deldjoo@polimi.it
Mehdi Elahi
meelahi@unibz.it

¹ Department of Computational Perception,
Johannes Kepler University Linz, Linz, Austria

² Center for Intelligent Information Retrieval,
University of Massachusetts Amherst, Amherst, USA

³ Spotify USA Inc., New York, USA

1 Introduction

Research in music recommender systems (MRSs) has recently experienced a substantial gain in interest both in academia and in industry [162]. Thanks to music streaming services like Spotify, Pandora, or Apple Music, music aficionados are nowadays given access to tens of millions music pieces. By filtering this abundance of music items, thereby limiting choice overload [20], MRSs are often very successful to suggest songs that fit their users' preferences. However, such systems are still far from being perfect and frequently produce unsatisfactory recommendations. This is partly because of the fact that users' tastes and musical needs are highly dependent on a multitude of factors, which are not considered in sufficient depth in current MRS approaches, which are typically centered on the core concept of user–item interactions, or sometimes content-based item descriptors. In contrast, we

⁴ Department of Computer Science, Politecnico di Milano,
Milan, Italy

⁵ Free University of Bozen-Bolzano, Bolzano, Italy

argue that satisfying the users' musical entertainment needs requires taking into account intrinsic, extrinsic, and contextual aspects of the listeners [2], as well as more decent interaction information. For instance, personality and emotional state of the listeners (intrinsic) [71,147] as well as their activity (extrinsic) [75,184] are known to influence musical tastes and needs. So are users' contextual factors including weather conditions, social surrounding, or places of interest [2,100]. Also the composition and annotation of a music playlist or a listening session reveal information about which songs go well together or are suited for a certain occasion [126,194]. Therefore, researchers and designers of MRS should reconsider their users in a holistic way in order to build systems tailored to the specificities of each user.

Against this background, in this trends and survey article, we elaborate on what we believe to be among the most pressing current challenges in MRS research, by discussing the respective state of the art and its restrictions (Sect. 2). Not being able to touch all challenges exhaustively, we focus on *cold start*, *automatic playlist continuation*, and *evaluation* of MRS. While these problems are to some extent prevalent in other recommendation domains too, certain characteristics of music pose particular challenges in these contexts. Among them are the short duration of items (compared to movies), the high emotional connotation of music, and the acceptance of users for duplicate recommendations. In the second part, we present our visions for future directions in MRS research (Sect. 3). More precisely, we elaborate on the topics of *psychologically inspired music recommendation* (considering human personality and emotion), *situation-aware music recommendation*, and *culture-aware music recommendation*. We conclude this article with a summary and identification of possible starting points for the interested researcher to face the discussed challenges (Sect. 4).

The composition of the authors allows to take academic as well as industrial perspectives, which are both reflected in this article. Furthermore, we would like to highlight that particularly the ideas presented as *Challenge 2: Automatic playlist continuation* in Sect. 2 play an important role in the task definition, organization, and execution of the ACM Recommender Systems Challenge 2018¹ which focuses on this use case. This article may therefore also serve as an entry point for potential participants in this challenge.

2 Grand challenges

In the following, we identify and detail a selection of the grand challenges, which we believe the research field of music recommender systems is currently facing, i.e., overcoming the cold start problem, automatic playlist contin-

uation, and properly evaluating music recommender systems. We review the state of the art of the respective tasks and its current limitations.

2.1 Particularities of music recommendation

Before we start digging deeper into these challenges, we would first like to highlight the major aspects that make music recommendation a particular endeavor and distinguishes it from recommending other items, such as movies, books, or products. These aspects have been adapted and extended from a tutorial on music recommender systems [161], co-presented by one of the authors at the ACM Recommender Systems 2017 conference.²

Duration of items In traditional movie recommendation, the items of interest have a typical duration of 90 min or more. In book recommendation, the consumption time is commonly even much longer. In contrast, the duration of music items usually ranges between 3 and 5 min (except maybe for classical music). Because of this, music items may be considered more disposable.

Magnitude of items The size of common commercial music catalogs is in the range of tens of millions music pieces, while movie streaming services have to deal with much smaller catalog sizes, typically thousands up to tens of thousands of movies and series.³ Scalability is therefore a much more important issue in music recommendation than in movie recommendation.

Sequential consumption Unlike movies, music pieces are most frequently consumed sequentially, more than one at a time, i.e., in a listening session or playlist. This yields a number of challenges for a MRS, which relate to identifying the right arrangement of items in a recommendation list.

Recommendation of previously recommended items Recommending the same music piece again, at a later point in time, may be appreciated by the user of a MRS, in contrast to a movie or product recommender, where repeated recommendations are usually not preferred.

Consumption behavior Music is often consumed passively, in the background. While this is not a problem per se, it can affect preference elicitation. In particular when using implicit feedback to infer listener preferences, the fact that

² http://www.cp.jku.at/tutorials/mrs_recsys_2017.

³ Spotify reports about 30 million songs in 2017 (<https://press.spotify.com/at/about>); Amazon's advanced search for books reports 10 million hardcover and 30 million paperback books in 2017 (<https://www.amazon.com/Advanced-Search-Books/b?node=241582011>), whereas Netflix, in contrast, offers about 5,500 movies and TV series as of 2016 (<http://time.com/4272360/the-number-of-movies-on-netflix-is-dropping-fast>).

¹ <http://www.recsyschallenge.com/2018>.

a listener is not paying attention to the music (therefore, e.g., not skipping a song) might be wrongly interpreted as a positive signal.

Listening intent and purpose Music serves various purposes for people and hence shapes their intent to listen to it. This should be taken into account when building a MRS. In extensive literature and empirical studies, Schäfer et al. [155] distilled three fundamental intents of music listening out of 129 distinct music uses and functions: self-awareness, social relatedness, and arousal and mood regulation. *Self-awareness* is considered as a very private relationship with music listening. The self-awareness dimension “helps people think about who they are, who they would like to be, and how to cut their own path” [154]. *Social relatedness* [153] describes the use of music to feel close to friends and to express identity and values to others. *Mood regulation* is concerned with managing emotions, which is a critical issue when it comes to the well-being of humans [77,110,176]. In fact, several studies found that mood and emotion regulation is the most important purpose why people listen to music [18,96,122,155], for which reason we discuss the particular role emotions play when listening to music separately below.

Emotions Music is known to evoke very strong emotions.⁴ This is a mutual relationship, though, since also the emotions of users affect musical preferences [17,77,144]. Due to this strong relationship between music and emotions, the problem of automatically describing music in terms of emotion words is an active research area, commonly referred to as music emotion recognition (MER), e.g., [14,103,187]. Even though MER can be used to tag music by emotion terms, how to integrate this information into MRS is a highly complicated task, for three reasons. First, MER approaches commonly neglect the distinction between *intended emotion* (i.e., the emotion the composer, songwriter, or performer had in mind when creating or performing the piece), *perceived emotion* (i.e., the emotion recognized while listening), and *induced emotion* that is felt by the listener. Second, the preference for a certain kind of emotionally laden music piece depends on whether the user wants to enhance or to modulate her mood. Third, emotional changes often occur within the same music piece, whereas tags are commonly extracted for the whole piece. Matching music and listeners in terms of emotions therefore requires to model the listener’s musical preference as a time-dependent function of their emotional experiences, also considering the intended purpose (mood enhancement or regulation). This

⁴ Please note that the terms “emotion” and “mood” have different meanings in psychology, whereas they are commonly used as synonyms in music information retrieval (MIR) and recommender systems research. In psychology, in contrast, “emotion” refers to a short-time reaction to a particular stimulus, whereas “mood” refers to a longer-lasting state without relation to a specific stimulus.

is a highly challenging task and usually neglected in current MRS, for which reason we discuss emotion-aware MRS as one of the main future directions in MRS research, cf. Sect. 3.1.

Listening context Situational or contextual aspects [15,48] have a strong influence on music preference, consumption, and interaction behavior. For instance, a listener will likely create a different playlist when preparing for a romantic dinner than when warming-up with friends to go out on a Friday night [75]. The most frequently considered types of context include location (e.g., listening at workplace, when commuting, or relaxing at home) [100] and time (typically categorized into, for example, morning, afternoon, and evening) [31]. Context may, in addition, also relate to the listener’s activity [184], weather [140], or the use of different listening devices, e.g., earplugs on a smartphone vs. hi-fi stereo at home [75], to name a few. Since music listening is also a highly social activity, investigating the social context of the listeners is crucial to understand their listening preferences and behavior [45,134]. The importance of considering such contextual factors in MRS research is acknowledged by discussing situation-aware MRS as a trending research direction, cf. Sect. 3.2.

2.2 Challenge 1: Cold start problem

Problem definition One of the major problems of recommender systems in general [64,151], and music recommender systems in particular [99,119] is the *cold start* problem, i.e., when a new user registers to the system or a new item is added to the catalog and the system does not have sufficient data associated with these items/users. In such a case, the system cannot properly recommend existing items to a new user (*new user* problem) or recommend a new item to the existing users (*new item* problem) [3,62,99,164].

Another subproblem of cold start is the *sparsity* problem which refers to the fact that the number of given ratings is much lower than the number of possible ratings, which is particularly likely when the number of users and items is large. The inverse of the ratio between given and possible ratings is called sparsity. High sparsity translates into low rating *coverage*, since most users tend to rate only a tiny fraction of items. The effect is that recommendations often become unreliable [99]. Typical values of sparsity are quite close to 100% in most real-world recommender systems. In the music domain, this is a particularly substantial problem. Dror et al. [51], for instance, analyzed the Yahoo! Music dataset, which as of time of writing represents the largest music recommendation dataset. They report a sparsity of 99.96%. For

comparison, the Netflix dataset of movies has a sparsity of “only” 98.82%.⁵

State of the art

A number of approaches have already been proposed to tackle the cold start problem in the music recommendation domain, foremost content-based approaches, hybridization, cross-domain recommendation, and active learning.

Content-based recommendation (CB) algorithms do not require ratings of users other than the target user. Therefore, as long as some pieces of information about the user’s own preferences are available, such techniques can be used in cold start scenarios. Furthermore, in the most severe case, when a new item is added to the catalog, content-based methods enable recommendations, because they can extract features from the new item and use them to make recommendations. It is noteworthy that while collaborative filtering (CF) systems have cold start problems both for new users and new items, content-based systems have only cold start problems for new users [5].

As for the new item problem, a standard approach is to extract a number of features that define the acoustic properties of the audio signal and use content-based learning of the user interest (user profile learning) in order to effect recommendations. Feature extraction is typically done automatically, but can also be effected manually by musical experts, as in the case of Pandora’s Music Genome Project.⁶ Pandora uses up to 450 specific descriptors per song, such as “aggressive female vocalist,” “prominent backup vocals,” “abstract lyrics,” or “use of unusual harmonies.”⁷ Regardless of whether the feature extraction process is performed automatically or manually, this approach is advantageous not only to address the new item problem but also because an accurate feature representation can be highly predicative of users’ tastes and interests which can be leveraged in the subsequent information filtering stage [5]. An advantage of music to video is that features in music are limited to a single audio channel, compared to audio and visual channels for videos adding a level complexity to the content analysis of videos explored individually or multimodal in different research works [46,47,59,128].

Automatic feature extraction from audio signals can be done in two main manners: (1) by extracting a feature vector from each item individually, independent of other items, or (2) by considering the cross-relation between items in the training dataset. The difference is that in (1) the same process

is performed in the training and testing phases of the system, and the extracted feature vectors can be used off-the-shelf in the subsequent processing stage; for example, they can be used to compute similarities between items in a one-to-one fashion at testing time. In contrast, in (2) first a model is built from all features extracted in the training phase, whose main role is to map the features into a new (acoustic) space in which the similarities between items are better represented and exploited. An example of approach (1) is the block-level feature framework [167,168], which creates a feature vector of about 10,000 dimensions, independently for each song in the given music collection. This vector describes aspects such as spectral patterns, recurring beats, and correlations between frequency bands. An example of strategy (2) is to create a low-dimensional i-vector representation from the Mel-frequency cepstral coefficients (MFCCs), which model musical timbre to some extent [56]. To this end, a universal background model is created from the MFCC vectors of the whole music collection, using a Gaussian mixture model (GMM). Performing factor analysis on a representation of the GMM eventually yields i-vectors.

In scenarios where some form of semantic labels, e.g., genres or musical instruments, are available, it is possible to build models that learn the intermediate mapping between low-level audio features and semantic representations using machine learning techniques, and subsequently use the learned models for prediction. A good point of reference for such *semantic-inferred* approaches can be found in [19,36].

An alternative technique to tackle the new item problem is *hybridization*. A review of different hybrid and ensemble recommender systems can be found in [6,26]. In [50], the authors propose a music recommender system which combines an acoustic CB and an item-based CF recommender. For the content-based component, it computes acoustic features including spectral properties, timbre, rhythm, and pitch. The content-based component then assists the collaborative filtering recommender in tackling the cold start problem since the features of the former are automatically derived via audio content analysis.

The solution proposed in [189] is a hybrid recommender system that combines CF and acoustic CB strategies also by feature hybridization. However, in this work the feature-level hybridization is not performed in the original feature domain. Instead, a set of latent variables referred to as *conceptual genre* are introduced, whose role is to provide a common shared feature space for the two recommenders and enable hybridization. The weights associated with the latent variables reflect the musical taste of the target user and are learned during the training stage.

In [169], the authors propose a hybrid recommender system incorporating item–item CF and acoustic CB based on similarity metric learning. The proposed metric learning is

⁵ Note that Dror et al.’s analysis was conducted in 2011. Even though the general character (rating matrices for music items being sparser than those of movie items) remained the same, the actual numbers for today’s catalogs are likely slightly different.

⁶ <http://www.pandora.com/about/mgp>.

⁷ <http://enacademic.com/dic.nsf/enwiki/3224302>.

an optimization model that aims to learn the weights associated with the audio content features (when combined in a linear fashion) so that a degree of consistency between CF-based similarity and the acoustic CB similarity measure is established. The optimization problem can be solved using quadratic programming techniques.

Another solution to cold start is *cross-domain recommendation* techniques, which aim at improving recommendations in one domain (here music) by making use of information about the user preferences in an auxiliary domain [28,67]. Hence, the knowledge of the preferences of the user is transferred from an auxiliary domain to the music domain, resulting in a more complete and accurate user model. Similarly, it is also possible to integrate additional pieces of information about the (new) users, which are not directly related to music, such as their personality, in order to improve the estimation of the user's music preferences. Several studies conducted on user personality characteristics support the conjecture that it may be useful to exploit this information in music recommender systems [69,73,86,130,147]. For a more detailed literature review of cross-domain recommendation, we refer to [29,68,102].

In addition to the aforementioned approaches, *active learning* has shown promising results in dealing with the cold start problem in single domain [60,146] or cross-domain recommendation scenario [136,192]. Active learning addresses this problem at its origin by identifying and eliciting (high quality) data that can represent the preferences of users better than by what they provide themselves. Such a system therefore interactively demands specific user feedback to maximize the improvement of system performance.

Limitations The state-of-the-art approaches elaborated on above are restricted by certain limitations. When using *content-based filtering*, for instance, almost all existing approaches rely on a number of predefined audio features that have been used over and over again, including spectral features, MFCCs, and a great number of derivatives [106]. However, doing so assumes that (all) these features are predictive of the user's music taste, while in practice it has been shown that the acoustic properties that are important for the perception of music are highly subjective [132]. Furthermore, listeners' different tastes and levels of interest in different pieces of music influence perception of item similarity [158]. This subjectiveness demands for CB recommenders that incorporate personalization in their mathematical model. For example, in [65] the authors propose a hybrid (CB+CF) recommender model, namely regression-based latent factor models (RLFM). In [4], the authors propose a user-specific feature-based similarity model (UFSM), which defines a similarity function for each user, leading to a high degree of personalization. Although not designed specifically for the

music domain, the authors of [4] provide an interesting literature review of similar user-specific models.

While *hybridization* can therefore alleviate the cold start problem to a certain extent, as seen in the examples above, respective approaches are often complex, computationally expensive, and lack transparency [27]. In particular, results of hybrids employing latent factor models are typically hard to understand for humans.

A major problem with *cross-domain recommender systems* is their need for data that connects two or more target domains, e.g., books, movies, and music [29]. In order for such approaches to work properly, items, users, or both therefore need to overlap to a certain degree [40]. In the absence of such overlap, relationships between the domains must be established otherwise, e.g., by inferring semantic relationships between items in different domains or assuming similar rating patterns of users in the involved domains. However, whether respective approaches are capable of transferring knowledge between domains is disputed [39]. A related issue in cross-domain recommendation is that there is a lack of established datasets with clear definitions of domains and recommendation scenarios [102]. Because of this, the majority of existing work on cross-domain RS uses some type of conventional recommendation dataset transformation to suit it for their need.

Finally, also *active learning* techniques suffer from a number of issues. First of all, the typical active learning techniques propose to a user to rate the items that the system has predicted to be interesting for them, i.e., the items with highest predicted ratings. This indeed is a default strategy in recommender systems for eliciting ratings since users tend to rate what has been recommended to them. Even when users browse the item catalog, they are more likely to rate items which they like or are interested in, rather than those items that they dislike or are indifferent to. Indeed, it has been shown that doing so creates a strong bias in the collected rating data as the database gets populated disproportionately with high ratings. This in turn may substantially influence the prediction algorithm and decrease the recommendation accuracy [63].

Moreover, not all the active learning strategies are necessarily personalized. The users differ very much in the amount of information they have about the items, their preferences, and the way they make decisions. Hence, it is clearly inefficient to request all the users to rate the same set of items, because many users may have a very limited knowledge, ignore many items, and will therefore not provide ratings for these items. Properly designed active learning techniques should take this into account and propose different items to different users to rate. This can be highly beneficial and increase the chance of acquiring ratings of higher quality [57].

Moreover, the traditional interaction model designed for active learning in recommender systems can support building the initial profile of a user mainly in the sign-up process. This is done by generating a user profile by requesting the user to rate a set of selected items [30]. On the other hand, the users must be able to also update their profile by providing more ratings anytime they are willing to. This requires the system to adopt a conversational interaction model [30], e.g., by exploiting novel interactive design elements in the user interface [38], such as explanations that can describe the benefits of providing more ratings and motivating the user to do so.

Finally, it is important to note that in an up-and-running recommender system, the ratings are given by users not only when requested by the system (active learning) but also when a user voluntarily explores the item catalog and rates some familiar items (natural acquisition of ratings) [30,61,63,127,146]. While this could have a huge impact on the performance of the system, it has been mostly ignored by the majority of the research works in the field of active learning for recommender systems. Indeed, almost all research works have been based on a rather non-realistic assumption that the only source for collecting new ratings is through the system requests. Therefore, it is crucial to take into account a more realistic scenario when studying the active learning techniques in recommender systems, which can better picture how the system evolves over time when ratings are provided by users [143,146].

2.3 Challenge 2: Automatic playlist continuation

Problem definition In its most generic definition, a playlist is simply a sequence of tracks intended to be listened to together. The task of automatic playlist generation (APG) then refers to the automated creation of these sequences of tracks. In this context, the ordering of songs in a playlist to generate is often highlighted as a characteristic of APG, which is a highly complex endeavor. Some authors have therefore proposed approaches based on Markov chains to model the transitions between songs in playlists, e.g., [32, 125]. While these approaches have been shown to outperform approaches agnostic of the song order in terms of log-likelihood, recent research has found little evidence that the exact order of songs actually matters to users [177], while the ensemble of songs in a playlist [181] and direct song-to-song transitions [92] do matter.

Considered a variation of APG, the task of *automatic playlist continuation* (APC) consists of adding one or more tracks to a playlist in a way that fits the same target characteristics of the original playlist. This has benefits in both the listening and creation of playlists: users can enjoy listening to continuous sessions beyond the end of a finite-length playlist, while also finding it easier to create longer, more com-

puting playlists without needing to have extensive musical familiarity.

A large part of the APC task is to accurately infer the intended purpose of a given playlist. This is challenging not only because of the broad range of these intended purposes (when they even exist), but also because of the diversity in the underlying features or characteristics that might be needed to infer those purposes.

Related to Challenge 1, an extreme cold start scenario for this task is where a playlist is created with some metadata (e.g., the title of a playlist), but no song has been added to the playlist. This problem can be cast as an *ad hoc information retrieval task*, where the task is to rank songs in response to a user-provided metadata query.

The APC task can also potentially benefit from user profiling, e.g., making use of previous playlists and the long-term listening history of the user. We call this *personalized playlist continuation*.

According to a study carried out in 2016 by the Music Business Association⁸ as part of their Music Biz Consumer Insights program,⁹ playlists accounted for 31% of music listening time among listeners in the USA, more than albums (22%), but less than single tracks (46%). Other studies, conducted by MIDiA,¹⁰ show that 55% of streaming music service subscribers create music playlists, with some streaming services such as Spotify currently hosting over 2 billion playlists.¹¹ In a 2017 study conducted by Nielsen,¹² it was found that 58% of users in the USA create their own playlists, 32% share them with others. Studies like these suggest a growing importance of playlists as a mode of music consumption, and as such, the study of APG and APC has never been more relevant.

State of the art APG has been studied ever since digital multimedia transmission made huge catalogs of music available to users. Bonnin and Jannach provide a comprehensive survey of this field in [21]. In it, the authors frame the APG task as the creation of a sequence of tracks that fulfill some “target characteristics” of a playlist, given some “background knowledge” of the characteristics of the catalog of tracks from which the playlist tracks are drawn. Existing APG systems tackle both of these problems in many different ways.

⁸ <https://musicbiz.org/news/playlists-overtake-albums-listenership-says-loop-study>.

⁹ <https://musicbiz.org/resources/tools/music-biz-consumer-insights/consumer-insights-portal>.

¹⁰ <https://www.midiaresearch.com/blog/announcing-midias-state-of-the-streaming-nation-2-report>.

¹¹ <https://press.spotify.com/us/about>.

¹² <http://www.nielsen.com/us/en/insights/reports/2017/music-360-2017-highlights.html>.

In early approaches [9,10,135] the *target characteristics* of the playlist are specified as multiple explicit constraints, which include musical attributes or metadata such as artist, tempo, and style. In others, the target characteristics are a single seed track [121] or a start and an end track [9,32,74]. Other approaches create a circular playlist that comprises all tracks in a given music collection, in such a way that consecutive songs are as similar as possible [105,142]. In other works, playlists are created based on the context of the listener, either as single source [157] or in combination with content-based similarity [35,149].

A common approach to build the *background knowledge* of the music catalog for playlist generation is using machine learning techniques to extract that knowledge from manually curated playlists. The assumption here is that curators of these playlists are encoding rich latent information about which tracks go together to create a satisfying listening experience for an intended purpose. Some proposed APG and APC systems are trained on playlists from sources such as online radio stations [32,123], online playlist websites [126,181], and music streaming services [141]. In the study by Pichl et al. [141], the names of playlists on Spotify were analyzed to create contextual clusters, which were then used to improve recommendations.

An approach to specifically address song ordering within playlists is the use of generative models that are trained on hand-curated playlists. McFee and Lanckriet [125] represent songs by metadata, familiarity, and audio content features, adopting ideas from statistical natural language processing. They train various Markov chains to model transitions between songs. Similarly, Chen et al. [32] propose a logistic Markov embedding to model song transitions. This is similar to matrix decomposition methods and results in an embedding of songs in Euclidean space. In contrast to McFee and Lanckriet's model, Chen et al.'s model does not use any audio features.

Limitations While some work on automated playlist continuation highlights the special characteristics of playlists, i.e., their *sequential order*, it is not well understood to which extent and in which cases taking into account the order of tracks in playlists helps create better models for recommendation. For instance, in [181] Vall et al. recently demonstrated on two datasets of hand-curated playlists that the song order seems to be negligible for accurate playlist continuation when a lot of popular songs are present. On the other hand, the authors argue that order does matter when creating playlists with tracks from the long tail. Another study by McFee and Lanckriet [126] also suggests that transition effects play an important role in modeling playlist continuity. This is in line with a study presented by Kamehkhosh et al. in [92], in which users identified song order as being the second but last impor-

tant criterion for playlist quality.¹³ In another recent user study [177] conducted by Tintarev et al. the authors found that many participants did not care about the order of tracks in recommended playlists, sometimes they did not even notice that there is a particular order. However, this study was restricted to 20 participants who used the Discover Weekly service of Spotify.¹⁴

Another challenge for APC is evaluation: in other words, how to assess the quality of a playlist. Evaluation in general is discussed in more detail in the next section, but there are specific questions around evaluation of playlists that should be pointed out here. As Bonnin and Jannach [21] put it, the ultimate criterion for this is *user satisfaction*, but that is not easy to measure. In [125], McFee and Lanckriet categorize the main approaches to APG evaluation as human evaluation, semantic cohesion, and sequence prediction. Human evaluation comes closest to measuring user satisfaction directly, but suffers from problems of scale and reproducibility. Semantic cohesion as a quality metric is easily measurable and reproducible, but assumes that users prefer playlists where tracks are similar along a particular semantic dimension, which may not always be true, see, for instance, the studies carried out by Slaney and White [172] and by Lee [115]. Sequence prediction casts APC as an information retrieval task, but in the domain of music, an inaccurate prediction needs not be a bad recommendation, and this again leads to a potential disconnect between this metric and the ultimate criterion of user satisfaction.

Investigating which factors are potentially important for a positive user perception of a playlist, Lee conducted a qualitative user study [115], investigating playlists that had been automatically created based on content-based *similarity*. They made several interesting observations. A concern frequently raised by participants was that of consecutive songs being too similar, and a general lack of *variety*. However, different people had different interpretations of variety, e.g., variety in genres or styles vs. different artists in the playlist. Similarly, different criteria were mentioned when listeners judged the coherence of songs in a playlist, including lyrical content, tempo, and mood. When creating playlists, participants mentioned that similar lyrics, a common theme (e.g., music to listen to in the train), story (e.g., music for the Independence Day), or era (e.g., rock music from the 1980s) are important and that tracks not complying negatively effect the flow of the playlist. These aspects can be extended by responses of participants in a study conducted by Cunningham et al. [42],

¹³ The ranking of criteria (from most to least important) was: homogeneity, artist diversity, transition, popularity, lyrics, order, and freshness.

¹⁴ <https://www.spotify.com/discoverweekly>.

who further identified the following categories of playlists: same artist, genre, style, or orchestration, playlists for a certain event or activity (e.g., party or holiday), romance (e.g., love songs or breakup songs), playlists intended to send a message to their recipient (e.g., protest songs), and challenges or puzzles (e.g., cover songs liked more than the original or songs whose title contains a question mark).

Lee also found that *personal preferences* play a major role. In fact, already a single song that is very much liked or hated by a listener can have a strong influence on how they judge the entire playlist [115]. This seems particularly true if it is a highly disliked song [44]. Furthermore, a good *mix of familiar and unknown songs* was often mentioned as an important requirement for a good playlist. Supporting the discovery of interesting new songs, still contextualized by familiar ones, increases the likelihood of realizing a *serendipitous encounter* in a playlist [160,193]. Finally, participants also reported that their familiarity with a playlist's genre or theme influenced their judgment of its quality. In general, listeners were more picky about playlists whose tracks they were familiar with or they liked a lot.

Supported by the studies summarized above, we argue that the question of what makes a great playlist is highly subjective and further depends on the intent of the creator or listener. Important criteria when creating or judging a playlist include track similarity/coherence, variety/diversity, but also the user's personal preferences and familiarity with the tracks, as well as the intention of the playlist creator. Unfortunately, current automatic approaches to playlist continuation are agnostic of the underlying psychological and sociological factors that influence the decision of which songs users choose to include in a playlist. Since knowing about such factors is vital to understand the intent of the playlist creator, we believe that algorithmic methods for APC need to holistically learn such aspects from manually created playlists and integrate respective intent models. However, we are aware that in today's era where billions of playlists are shared by users of online streaming services,¹⁵ a large-scale analysis of psychological and sociological background factors is impossible. Nevertheless, in the absence of explicit information about user intent, a possible starting point to create intent models might be the metadata associated with user-generated playlists, such as title or description. To foster this kind of research, the playlists provided in the dataset for the ACM Recommender Systems Challenge 2018 include playlist titles.¹⁶

¹⁵ <https://press.spotify.com/us/about>.

¹⁶ <https://recsys-challenge.spotify.com>.

2.4 Challenge 3: Evaluating music recommender systems

Problem definition Having its roots in machine learning (cf. rating prediction) and information retrieval (cf. “retrieving” items based on implicit “queries” given by user preferences), the field of recommender systems originally adopted evaluation metrics from these neighboring fields. In fact, accuracy and related quantitative measures, such as precision, recall, or error measures (between predicted and true ratings), are still the most commonly employed criteria to judge the recommendation quality of a recommender system [11,78]. In addition, novel measures that are tailored to the recommendation problem have emerged in recent years. These so-called beyond-accuracy measures [98] address the particularities of recommender systems and gauge, for instance, the utility, novelty, or serendipity of an item. However, a major problem with these kinds of measures is that they integrate factors that are hard to describe mathematically, for instance, the aspect of surprise in case of serendipity measures. For this reason, there sometimes exist a variety of different definitions to quantify the same beyond-accuracy aspect.

State of the art In the following, we discuss performance measures which are most frequently reported when evaluating recommender systems. An overview of these is given in Table 1. They can be roughly categorized into accuracy-related measures, such as prediction error (e.g., MAE and RMSE) or standard IR measures (e.g., precision and recall), and beyond-accuracy measures, such as diversity, novelty, and serendipity. Furthermore, while some of the metrics quantify the ability of recommender systems to find good items, e.g., precision and recall, others consider the ranking of items and therefore assess the system's ability to position good recommendations at the top of the recommendation list, e.g., MAP, NDCG, or MPR.

Mean absolute error (MAE) is one of the most common metrics for evaluating the prediction power of recommender algorithms. It computes the average absolute deviation between the predicted ratings and the actual ratings provided by users [81]. Indeed, MAE indicates how close the rating predictions generated by an MRS are to the real user ratings. MAE is computed as follows:

$$MAE = \frac{1}{|T|} \sum_{r_{u,i} \in T} |r_{u,i} - \hat{r}_{u,i}| \quad (1)$$

where $r_{u,i}$ and $\hat{r}_{u,i}$, respectively, denote the actual and the predicted ratings of item i for user u . MAE sums over the absolute prediction errors for all ratings in a test set T .

Table 1 Evaluation measures commonly used for recommender systems

Measure	Abbreviation	Type	Ranking-aware
Mean absolute error	MAE	Error/accuracy	No
Root-mean-square error	RMSE	Error/accuracy	No
Precision at top K recommendations	P@K	Accuracy	No
Recall at top K recommendations	R@K	Accuracy	No
Mean average precision at top K recommendations	MAP@K	Accuracy	Yes
Normalized discounted cumulative gain	NDCG	Accuracy	Yes
Half-life utility	HLU	Accuracy	Yes
Mean percentile rank	MPR	Accuracy	Yes
Spread	–	Beyond	No
Coverage	–	Beyond	No
Novelty	–	Beyond	No
Serendipity	–	Beyond	No
Diversity	–	Beyond	No

Root-mean-square error (RMSE) is another similar metric that is computed as:

$$RMSE = \sqrt{\frac{1}{|T|} \sum_{r_{u,i} \in T} (r_{u,i} - \hat{r}_{u,i})^2}. \tag{2}$$

It is an extension to MAE in that the error term is squared, which penalizes larger differences between predicted and true ratings more than smaller ones. This is motivated by the assumption that, for instance, a rating prediction of 1 when the true rating is 4 is much more severe than a prediction of 3 for the same item.

Precision at top K recommendations (P@K) is a common metric that measures the accuracy of the system in commanding relevant items. In order to compute P@K, for each user, the top K recommended items whose ratings also appear in the test set T are considered. This metric was originally designed for binary relevance judgments. Therefore, in case of availability of relevance information at different levels, such as a five-point Likert scale, the labels should be binarized, e.g., considering the ratings greater than or equal to 4 (out of 5) as relevant. For each user u, P_u@K is computed as follows:

$$P_u@K = \frac{|L_u \cap \hat{L}_u|}{|\hat{L}_u|} \tag{3}$$

where L_u is the set of relevant items for user u in the test set T and \hat{L}_u denotes the recommended set containing the K items in T with the highest predicted ratings for the user

u. The overall P@K is then computed by averaging P_u@K values for all users in the test set.

Mean average precision at top K recommendations (MAP@K) is a rank-based metric that computes the overall precision of the system at different lengths of recommendation lists. MAP is computed as the arithmetic mean of the average precision over the entire set of users in the test set. Average precision for the top K recommendations (AP@K) is defined as follows:

$$AP@K = \frac{1}{N} \sum_{i=1}^K P@i \cdot rel(i) \tag{4}$$

where rel(i) is an indicator signaling if the ith recommended item is relevant, i.e., rel(i) = 1, or not, i.e., rel(i) = 0; N is the total number of relevant items. Note that MAP implicitly incorporates recall, because it also considers the relevant items not in the recommendation list.¹⁷

Recall at top K recommendations (R@K) is presented here for the sake of completeness, even though it is not a crucial measure from a consumer’s perspective. Indeed, the listener is typically not interested in being recommended all or a large number of relevant items, rather in having good recommendations at the top of the recommendation list. For a user u, R_u@K is defined as:

¹⁷ We should note that in the recommender systems community, another variation of average precision is gaining popularity recently, formally defined by: $AP@K = \frac{1}{\min(K,N)} \sum_{i=1}^K P@i \cdot rel(k)$ in which N is the total number of relevant items and K is the size of recommendation list. The motivation behind the minimization term is to prevent the AP scores to be unfairly suppressed when the number of recommendations is too low to capture all the relevant items. This variation of MAP was popularized by Kaggle competitions [97] about recommender systems and has been used in several other research works, consider for example [8,124].

$$R_u@K = \frac{|L_u \cap \hat{L}_u|}{|L_u|} \quad (5)$$

where L_u is the set of relevant items of user u in the test set T and \hat{L}_u denotes the recommended set containing the K items in T with the highest predicted ratings for the user u . The overall $R@K$ is calculated by averaging $R_u@K$ values for all the users in the test set.

Normalized discounted cumulative gain (NDCG) is a measure for the ranking quality of the recommendations. This metric has originally been proposed to evaluate the effectiveness of information retrieval systems [93]. It is nowadays also frequently used for evaluating music recommender systems [120, 139, 185]. Assuming that the recommendations for user u are sorted according to the predicted rating values in descending order. DCG_u is defined as follows:

$$DCG_u = \sum_{i=1}^N \frac{r_{u,i}}{\log_2(i+1)} \quad (6)$$

where $r_{u,i}$ is the true rating (as found in test set T) for the item ranked at position i for user u , and N is the length of the recommendation list. Since the rating distribution depends on the users' behavior, the DCG values for different users are not directly comparable. Therefore, the cumulative gain for each user should be normalized. This is done by computing the ideal DCG for user u , denoted as $IDCG_u$, which is the DCG_u value for the best possible ranking, obtained by ordering the items by true ratings in descending order. Normalized discounted cumulative gain for user u is then calculated as:

$$NDCG_u = \frac{DCG_u}{IDCG_u}. \quad (7)$$

Finally, the overall normalized discounted cumulative gain $NDCG$ is computed by averaging $NDCG_u$ over the entire set of users.

In the following, we present common quantitative evaluation metrics, which have been particularly designed or adopted to assess recommender systems performance, even though some of them have their origin in information retrieval and machine learning. The first two (HLU and MRR) still belong to the category of accuracy-related measures, while the subsequent ones capture beyond-accuracy aspects

Half-life utility (HLU) measures the utility of a recommendation list for a user with the assumption that the likelihood of viewing/choosing a recommended item by the user exponentially decays with the item's position in the ranking [24, 137]. Formally written, HLU for user u is defined as:

$$HLU_u = \sum_{i=1}^N \max(r_{u,i} - d, 0) \frac{1}{2^{(rank_{u,i}-1)/(h-1)}} \quad (8)$$

where $r_{u,i}$ and $rank_{u,i}$ denote the rating and the rank of item i for user u , respectively, in the recommendation list of length N ; d represents a default rating (e.g., average rating); and h is the half-time, calculated as the rank of a music item in the list, such that the user can eventually listen to it with a 50% chance. HLU_u can be further normalized by the maximum utility (similar to NDCG), and the final HLU is the average over the half-time utilities obtained for all users in the test set. A larger HLU may correspond to a superior recommendation performance.

Mean percentile rank (MPR) estimates the users' satisfaction with items in the recommendation list and is computed as the average of the percentile rank for each test item within the ranked list of recommended items for each user [89]. The percentile rank of an item is the percentage of items whose position in the recommendation list is equal to or lower than the position of the item itself. Formally, the percentile rank PR_u for user u is defined as:

$$PR_u = \frac{\sum_{r_{u,i} \in T} r_{u,i} \cdot rank_{u,i}}{\sum_{r_{u,i} \in T} r_{u,i}} \quad (9)$$

where $r_{u,i}$ is the true rating (as found in test set T) for item i rated by user u and $rank_{u,i}$ is the percentile rank of item i within the ordered list of recommendations for user u . MPR is then the arithmetic mean of the individual PR_u values over all users. A randomly ordered recommendation list has an expected MPR value of 50%. A smaller MPR value is therefore assumed to correspond to a superior recommendation performance.

Spread is a metric of how well the recommender algorithm can spread its attention across a larger set of items [104]. In more detail, spread is the entropy of the distribution of the items recommended to the users in the test set. It is formally defined as:

$$spread = - \sum_{i \in I} P(i) \log P(i) \quad (10)$$

where I represents the entirety of items in the dataset and $P(i) = count(i) / \sum_{i' \in I} count(i')$, such that $count(i)$ denotes the total number of times that a given item i showed up in the recommendation lists. It may be infeasible to expect an algorithm to achieve the perfect spread (i.e., recommending each item an equal number of times) without avoiding

irrelevant recommendations or unfulfillable rating requests. Accordingly, moderate spread values are usually preferable.

Coverage of a recommender system is defined as the proportion of items over which the system is capable of generating recommendations [81]:

$$coverage = \frac{|\hat{T}|}{|T|} \tag{11}$$

where $|T|$ is the size of the test set and $|\hat{T}|$ is the number of ratings in T for which the system can predict a value. This is particularly important in cold start situations, when recommender systems are not able to accurately predict the ratings of new users or new items and hence obtain low coverage. Recommender systems with lower coverage are therefore limited in the number of items they can recommend. A simple remedy to improve low coverage is to implement some default recommendation strategy for an unknown user–item entry. For example, we can consider the average rating of users for an item as an estimate of its rating. This may come at the price of accuracy, and therefore, the trade-off between coverage and accuracy needs to be considered in the evaluation process [7].

Novelty measures the ability of a recommender system to recommend new items that the user did not know about before [1]. A recommendation list may be accurate, but if it contains a lot of items that are not novel to a user, it is not necessarily a useful list [193].

While novelty should be defined on an individual user level, considering the actual freshness of the recommended items, it is common to use the self-information of the recommended items relative to their global popularity:

$$novelty = \frac{1}{|U|} \sum_{u \in U} \sum_{i \in L_u} \frac{-\log_2 pop_i}{N} \tag{12}$$

where pop_i is the popularity of item i measured as percentage of users who rated i , L_u is the recommendation list of the top N recommendations for user u [193,195]. The above definition assumes that the likelihood of the user selecting a previously unknown item is proportional to its global popularity and is used as an approximation of novelty. In order to obtain more accurate information about novelty or freshness, explicit user feedback is needed, in particular since the user might have listened to an item through other channels before.

It is often assumed that the users prefer recommendation lists with more novel items. However, if the presented items are too novel, then the user is unlikely to have any knowledge

of them, nor to be able to understand or rate them. Therefore, moderate values indicate better performances [104].

Serendipity aims at evaluating MRS based on the *relevant and surprising* recommendations. While the need for serendipity is commonly agreed upon [82], the question of how to measure the degree of serendipity for a recommendation list is controversial. This particularly holds for the question of whether the factor of surprise implies that items must be novel to the user [98]. On a general level, serendipity of a recommendation list L_u provided to a user u can be defined as:

$$serendipity(L_u) = \frac{|L_u^{unexp} \cap L_u^{useful}|}{|L_u|} \tag{13}$$

where L_u^{unexp} and L_u^{useful} denote subsets of L that contain, respectively, recommendations unexpected to and useful for the user. The usefulness of an item is commonly assessed by explicitly asking users or taking user ratings as proxy [98]. The unexpectedness of an item is typically quantified by some measure of distance from expected items, i.e., items that are similar to the items already rated by the user. In the context of MRS, Zhang et al. [193] propose an “unserendipity” measure that is defined as the average similarity between the items in the user’s listening history and the new recommendations. Similarity between two items in this case is calculated by an adapted cosine measure that integrates co-liking information, i.e., number of users who like both items. It is assumed that lower values correspond to more surprising recommendations, since lower values indicate that recommendations deviate from the user’s traditional behavior [193].

Diversity is another beyond-accuracy measure as already discussed in the limitations part of Challenge 1. It gauges the extent to which recommended items are different from each other, where difference can relate to various aspects, e.g., musical style, artist, lyrics, or instrumentation, just to name a few. Similar to serendipity, diversity can be defined in several ways. One of the most common is to compute pairwise distance between all items in the recommendation set, either averaged [196] or summed [173]. In the former case, the diversity of a recommendation list L is calculated as follows:

$$diversity(L) = \frac{\sum_{i \in L} \sum_{j \in L, i \neq j} dist_{i,j}}{|L| \cdot (|L| - 1)} \tag{14}$$

where $dist_{i,j}$ is the some distance function defined between items i and j . Common choices are inverse cosine similarity [150], inverse Pearson correlation [183], or Hamming distance [101].

When it comes to the task of evaluating playlist recommendation, where the goal is to assess the capability of the

recommender in providing proper transitions between subsequent songs, the conventional error or accuracy metrics may not be able to capture this property. There is hence a need for *sequence-aware evaluation* measures. For example, consider the scenario where a user who likes both classical and rock music is recommended a rock music right after she has listened to a classic piece. Even though both music styles are in agreement with her taste, the transition between songs plays an important role toward user satisfaction. In such a situation, given a currently played song and in the presence of several equally likely good options to be played next, a RS may be inclined to rank songs based on their popularity. Hence, other metrics such as *average log-likelihood* have been proposed to better model the transitions [33,34]. In this regard, when the goal is to suggest a sequence of items, alternative *multi-metric* evaluation approaches are required to take into consideration multiple quality factors. Such evaluation metrics can consider the ranking order of the recommendations or the internal coherence or diversity of the recommended list as a whole. In many scenarios, adoption of such quality metrics can lead to a trade-off with accuracy which should be balanced by the RS algorithm [145].

Limitations As of today, the vast majority of evaluation approaches in recommender systems research focus on quantitative measures, either accuracy-like or beyond-accuracy, which are often computed in offline studies.

Doing so has the advantage of facilitating the reproducibility of evaluation results. However, limiting the evaluation to quantitative measures means to forgo another important factor, which is user experience. In other words, in the absence of user-centric evaluations, it is difficult to extend the claims to the more important objective of the recommender system under evaluation, i.e., giving users a pleasant and useful personalized experience [107].

Despite acknowledging the need for more user-centric evaluation strategies [158], the factor human, user, or, in the case of MRS, listener is still way too often neglected or not properly addressed. For instance, while there exist quantitative objective measures for serendipity and diversity, as discussed above, *perceived* serendipity and diversity can be highly different from the measured ones [182] as they are subjective user-specific concepts. This illustrates that even beyond-accuracy measures cannot fully capture the real *user satisfaction* with a recommender system. On the other hand, approaches that address user experience (UX) can be investigated to evaluate recommender systems. For example, a MRS can be evaluated based on *user engagement*, which provides a restricted explanation of UX that concentrates on judgment of product quality during interaction [79,118,133]. User satisfaction, user engagement, and more generally user experience are commonly assessed through user studies [13,116,117].

Addressing both objective and subjective evaluation criteria, Knijnenburg et al. [108] propose a holistic framework for user-centric evaluation of recommender systems. Figure 1 provides an overview of the components. The objective system aspects (OSAs) are considered unbiased factors of the RS, including aspects of the user interface, computing time of the algorithm, or number of items shown to the user. They are typically easy to specify or compute. The OSAs influence the subjective system aspects (SSAs), which are caused by momentary, primary evaluative feelings while interacting with the system [80]. This results in a different perception of the system by different users. SSAs are therefore highly individual aspects and typically assessed by user questionnaires. Examples of SSA include general appeal of the system, usability, and perceived recommendation diversity or novelty. The aspect of experience (EXP) describes the user's attitude toward the system and is commonly also investigated by questionnaires. It addresses the user's perception of the interaction with the system. The experience is highly influenced by the other components, which means changing any of the other components likely results in a change of EXP aspects. Experience can be broken down into the evaluation of the system, the decision process, and the final decisions made, i.e., the outcome. The interaction (INT) aspects describe the observable behavior of the user, time spent viewing an item, as well as clicking or purchasing behavior. In a music context, examples further include liking a song or adding it to a playlist. Therefore, interactions aspects belong to the objective measures and are usually determined via logging by the system. Finally, Knijnenburg et al.'s framework mentions personal characteristics (PC) and situational characteristics (SC), which influence the user experience. PC include aspects that do not exist without the user, such as user demographics, knowledge, or perceived control, while SC include aspects of the interaction context, such as when and where the system is used, or situation-specific trust or privacy concerns. Knijnenburg et al. [108] also propose a questionnaire to assess the factors defined in their framework, for instance, perceived recommendation quality, perceived system effectiveness, perceived recommendation variety, choice satisfaction, intention to provide feedback, general trust in technology, and system-specific privacy concern.

While this framework is a generic one, tailoring it to MRS would allow for user-centric evaluation thereof. In particular, the aspects of personal and situational characteristics should be adapted to the particularities of music listeners and listening situations, respectively, cf. Sect. 2.1. To this end, researchers in MRS should consider the aspects relevant to the perception and preference of music, and their implications on MRS, which have been identified in several studies, e.g., [43,113,114,158,159]. In addition to the general ones mentioned by Knijnenburg et al., of great importance in the

music domain seem to be psychological factors, including affect and personality, social influence, musical training and experience, and physiological condition.

We believe that carefully and holistically evaluating MRS by means of accuracy and beyond-accuracy, objective and subjective measures, in offline and online experiments, would lead to a better understanding of the listeners' needs and requirements vis-à-vis MRS, and eventually a considerable improvement of current MRS.

3 Future directions and visions

While the challenges identified in the previous section are already researched on intensely, in the following, we provide a more forward-looking analysis and discuss some MRS-related trending topics, which we assume influential for the next generation of MRS. All of them have in common that their aim is to create more personalized recommendations. More precisely, we first outline how psychological constructs such as personality and emotion could be integrated into MRS. Subsequently, we address situation-aware MRS and argue for the need of multifaceted user models that describe contextual and situational preferences. To round off, we discuss the influence of users' cultural background on recommendation preferences, which needs to be considered when building culture-aware MRS.

3.1 Psychologically inspired music recommendation

Personality and emotion are important psychological constructs. While personality characteristics of humans are a predictable and stable measure that shapes human behaviors, emotions are short-term affective responses to a particular stimulus [179]. Both have been shown to influence music tastes [71,154,159] and user requirements for MRS [69,73]. However, in the context of (music) recommender systems, personality and emotion do not play a major role yet. Given the strong evidence that both influence listening preferences [147,159] and the recent emergence of approaches to accurately predict them from user-generated data [111,170], we believe that psychologically inspired MRS is an upcoming area.

3.1.1 Personality

In psychology research, personality is often defined as a "consistent behavior pattern and interpersonal processes originating within the individual" [25]. This definition accounts for the individual differences in people's emotional, interpersonal, experiential, attitudinal, and motivational styles [95]. Several prior works have studied the relation of decision making and personality factors. In [147], as an example, it

has been shown that personality can influence the human decision-making process as well as the tastes and interests. Due to this direct relation, people with similar personality factors are very likely to share similar interests and tastes.

Earlier studies conducted on the user personality characteristics support the potential benefits that personality information could have in recommender systems [22,23,58,85,87,178,180]. As a known example, psychological studies [147] have shown that extravert people are likely to prefer the upbeat and conventional music. Accordingly, a personality-based MRS could use this information to better predict which songs are more likely than others to please extravert people [86]. Another example of potential usage is to exploit personality information in order to compute similarity among users and hence identify the like-minded users [178]. This similarity information could then be integrated into a neighborhood-based collaborative filtering approach.

In order to use personality information in a recommender system, the system first has to elicit this information from the users, which can be done either explicitly or implicitly. In the former case, the system can ask the user to complete a personality questionnaire using one of the personality evaluation inventories, e.g., the ten-item personality inventory [76] or the big five inventory [94]. In the latter case, the system can learn the personality by tracking and observing users' behavioral patterns, for instance, liking behavior on Facebook [111] or applying filters to images posted on Instagram [170]. Not too surprisingly, it has shown that systems that explicitly elicit personality characteristics achieve superior recommendation outcomes, e.g., in terms of user satisfaction, ease of use, and prediction accuracy [52]. On the downside, however, many users are not willing to fill in long questionnaires before being able to use the RS. A way to alleviate this problem is to ask users only the most informative questions of a personality instrument [163]. Which questions are most informative, though, first needs to be determined based on existing user data and is dependent on the recommendation domain. Other studies showed that users are to some extent willing to provide further information in return for a better quality of recommendations [175].

Personality information can be used in various ways, particularly, to generate recommendations when traditional rating or consumption data is missing. Otherwise, the personality traits can be seen as an additional feature that extends the user profile, that can be used mainly to identify similar users in neighborhood-based recommender systems or directly fed into extended matrix factorization models [67].

3.1.2 Emotion

The emotional state of the MRS user has a strong impact on his or her short-time musical preferences [99]. Vice versa,

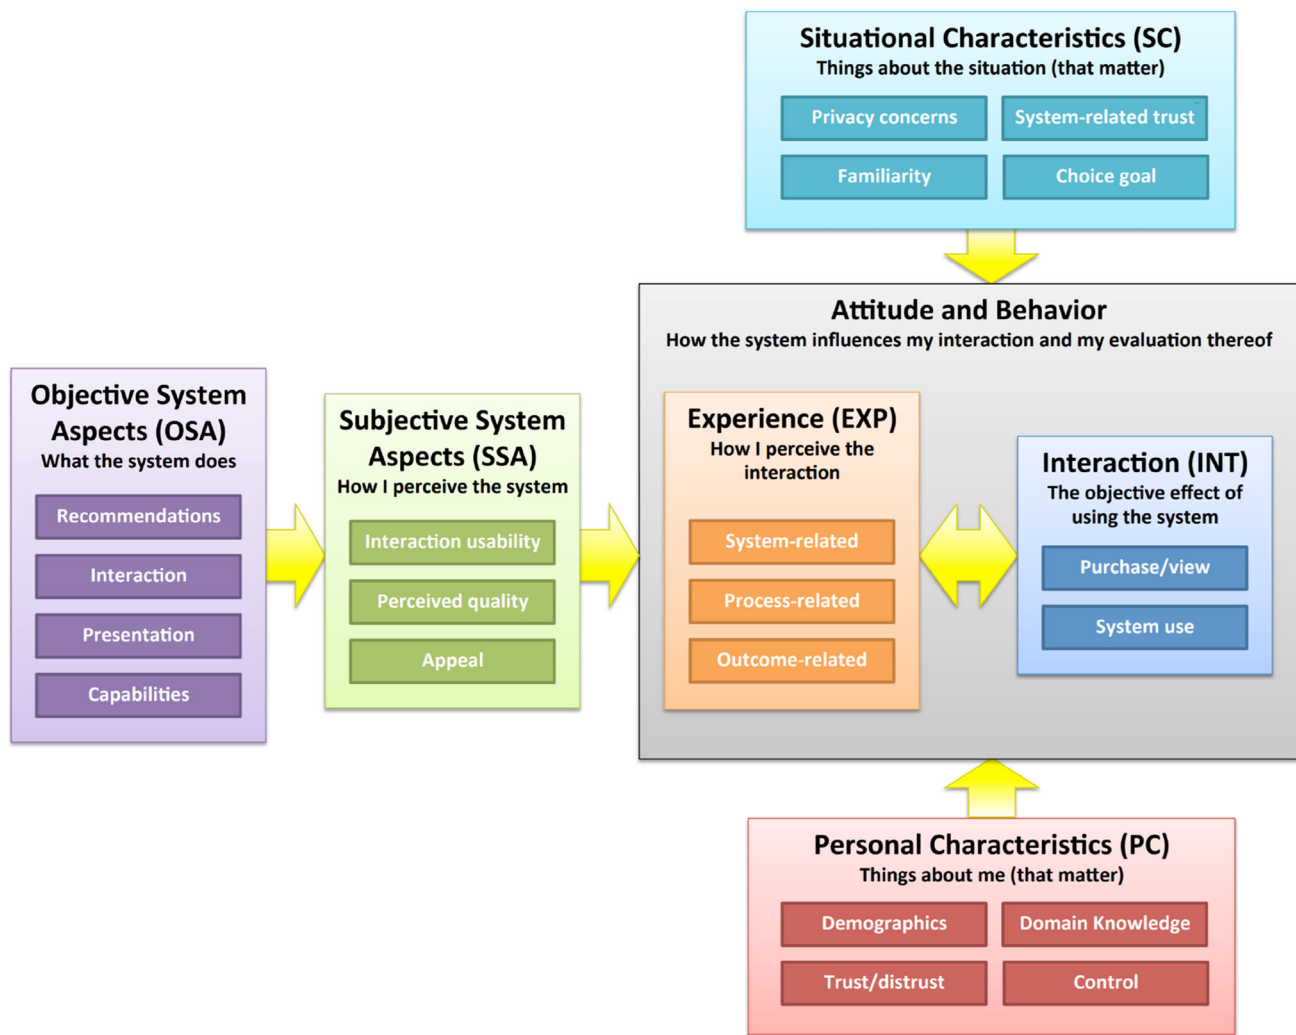


Fig. 1 Evaluation framework of the user experience for recommender systems, according to [108]

music has a strong influence on our emotional state. It therefore does not come as a surprise that emotion regulation was identified as one of the main reasons why people listen to music [122,155]. As an example, people may listen to completely different musical genres or styles when they are sad in comparison with when they are happy. Indeed, prior research on music psychology discovered that people may choose the type of music which moderates their emotional condition [109]. More recent findings show that music can be mainly chosen so as to augment the emotional situation perceived by the listener [131]. In order to build emotion-aware MRS, it is therefore necessary to (i) infer the emotional state the listener is in, (ii) infer emotional concepts from the music itself, and (iii) understand how these two interrelate. These three tasks are detailed below.

Eliciting the emotional state of the listener Similar to personality traits, the emotional state of a user can be elicited explicitly or implicitly. In the former case, the user is typically

presented one of the various *categorical models* (emotions are described by distinct emotion words such as happiness, sadness, anger, or fear) [84,191] or *dimensional models* (emotions are described by scores with respect to two or three dimensions, e.g., valence and arousal) [152]. For a more detailed elaboration on emotion models in the context of music, we refer to [159,186]. The implicit acquisition of emotional states can be effected, for instance, by analyzing user-generated text [49], speech [66], or facial expressions in video [55].

Emotion tagging in music The music piece itself can be regarded as an emotion-laden content and in turn can be described by emotion words. The task of automatically assigning such emotion words to a music piece is an active research area, often referred to as music emotion recognition (MER), e.g., [14,91,103,187,188,191]. How to integrate such emotion terms created by MER tools into a MRS is, however, not an easy task, for several reasons. First, early

MER approaches usually neglected the distinction between *intended emotion*, *perceived emotion*, and *induced or felt emotion*, cf. Sect. 2.1. Current MER approaches focus on perceived or induced emotions. However, musical content still contains various characteristics that affect the emotional state of the listener, such as lyrics, rhythm, and harmony, and the way how they affect the emotional state is highly subjective. This so even though research has detected a few general rules, for instance, a musical piece that is in major key is typically perceived brighter and happier than those in minor key, or a piece in rapid tempo is perceived more exciting or more tense than slow tempo ones [112].

Connecting listener emotions and music emotion tags Current emotion-based MRSs typically consider emotional scores as contextual factors that characterize the situation the user is experiencing. Hence, the recommender systems exploit emotions in order to pre-filter the preferences of users or post-filter the generated recommendations. Unfortunately, this neglects the psychological background, in particular on the subjective and complex interrelationships between expressed, perceived, and induced emotions [159], which is of special importance in the music domain as music is known to evoke stronger emotions than, for instance, products [161]. It has also been shown that personality influences in which emotional state which kind of emotionally laden music is preferred by listeners [71]. Therefore, even if automated MER approaches would be able to accurately predict the perceived or induced emotion of a given music piece, in the absence of deep psychological listener profiles, matching emotion annotations of items and listeners may not yield satisfying recommendations. This is so because how people judge music and which kind of music they prefer depends to a large extent on their current psychological and cognitive states. We hence believe that the field of MRS should embrace psychological theories, elicit the respective user-specific traits, and integrate them into recommender systems, in order to build decent emotion-aware MRS.

3.2 Situation-aware music recommendation

Most of the existing music recommender systems make recommendations solely based on a set of user-specific and item-specific signals. However, in real-world scenarios, many other signals are available. These additional signals can be further used to improve the recommendation performance. A large subset of these additional signals includes *situational signals*. In more detail, the music preference of a user depends on the situation at the moment of recommendation.¹⁸ *Location* is an example of situational signals;

for instance, the music preference of a user would differ in libraries and in gyms [35]. Therefore, considering location as a situation-specific signal could lead to substantial improvements in the recommendation performance. *Time of the day* is another situational signal that could be used for recommendation; for instance, the music a user would like to listen to in mornings differs from those in nights [41]. One situational signal of particular importance in the music domain is *social context* since music tastes and consumption behaviors are deeply rooted in the users' social identities and mutually affect each other [45,134]. For instance, it is very likely that a user would prefer different music when being alone than when meeting friends. Such social factors should therefore be considered when building situation-aware MRS. Other situational signals that are sometimes exploited include the user's current *activity* [184], the *weather* [140], the user's *mood* [129], and the *day of the week* [83]. Regarding time, there is also another factor to consider, which is that most music that was considered trendy years ago is now considered old. This implies that ratings for the same song or artist might strongly differ, not only between users, but in general as a function of time. To incorporate such aspects in MRS, it would be crucial to record a timestamp for all ratings.

It is worth noting that situational features have been proven to be strong signals in improving retrieval performance in search engines [16,190]. Therefore, we believe that researching and building situation-aware music recommender systems should be one central topic in MRS research.

While several situation-aware MRSs already exist, e.g., [12,35,90,100,157,184], they commonly exploit only one or very few such situational signals, or are restricted to a certain usage context, e.g., music consumption in a car or in a tourist scenario. Those systems that try to take a more comprehensive view and consider a variety of different signals, on the other hand, suffer from a low number of data instances or users, rendering it very hard to build accurate context models [75]. What is still missing, in our opinion, are (commercial) systems that integrate a variety of situational signals on a very large scale in order to truly understand the listeners needs and intents in any given situation and recommend music accordingly. While we are aware that data availability and privacy concerns counteract the realization of such systems on a large commercial scale, we believe that MRS will eventually integrate decent multifaceted user models inferred from contextual and situational factors.

3.3 Culture-aware music recommendation

While most humans share an inclination to listen to music, independent on their location or cultural background, the way music is performed, perceived, and interpreted evolves in a culture-specific manner. However, research in MRS seems to be agnostic of this fact. In music information retrieval

¹⁸ Please note that music taste is a relatively stable characteristic, while music preferences vary depending on the context and listening intent.

(MIR) research, on the other hand, cultural aspects have been studied to some extent in recent years, after preceding (and still ongoing) criticisms of the predominance of Western music in this community. Arguably the most comprehensive culture-specific research in this domain has been conducted as part of the CompMusic project,¹⁹ in which five non-Western music traditions have been analyzed in detail in order to advance automatic description of music by emphasizing cultural specificity. The analyzed music traditions included Indian Hindustani and Carnatic [53], Turkish Makam [54], Arab-Andalusian [174], and Beijing Opera [148]. However, the project's focus was on music creation, content analysis, and ethnomusicological aspects rather than on the music consumption side [37, 165, 166]. Recently, analyzing content-based audio features describing rhythm, timbre, harmony, and melody for a corpus of a larger variety of world and folk music with given country information, Panteli et al. found distinct acoustic patterns of the music created in individual countries [138]. They also identified geographical and cultural proximities that are reflected in music features, looking at outliers and misclassifications in a classification experiments using country as target class. For instance, Vietnamese music was often confused with Chinese and Japanese, South African with Botswana.

In contrast to this—meanwhile quite extensive—work on culture-specific analysis of music traditions, little effort has been made to analyze cultural differences and patterns of music consumption behavior, which is, as we believe, a crucial step to build culture-aware MRS. The few studies investigating such cultural differences include [88], in which Hu and Lee found differences in perception of moods between American and Chinese listeners. By analyzing the music listening behavior of users from 49 countries, Ferwerda et al. found relationships between music listening diversity and Hofstede's cultural dimensions [70, 72]. Skowron et al. used the same dimensions to predict genre preferences of listeners with different cultural backgrounds [171]. Schedl analyzed a large corpus of listening histories created by Last.fm users in 47 countries and identified distinct preference patterns [156]. Further analyses revealed countries closest to what can be considered the global mainstream (e.g., the Netherlands, UK, and Belgium) and countries farthest from it (e.g., China, Iran, and Slovakia). However, all of these works define culture in terms of country borders, which often makes sense, but is sometimes also problematic, for instance, in countries with large minorities of inhabitants with different cultures.

In our opinion, when building MRS, the analysis of cultural patterns of music consumption behavior, subsequent creation of respective cultural listener models, and their integration into recommender systems are vital steps to improve personalization and serendipity of recommendations. Culture

should be defined on various levels though, not only country borders. Other examples include having a joint historical background, speaking the same language, sharing the same beliefs or religion, and differences between urban vs. rural cultures. Another aspect that relates to culture is a temporal one since certain cultural trends, e.g., what defines the “youth culture,” are highly dynamic in a temporal and geographical sense. We believe that MRS which are aware of such cross-cultural differences and similarities in music perception and taste, and are able to recommend music a listener in the same or another culture may like, would substantially benefit both users and providers of MRS.

4 Conclusions

In this trends and survey paper, we identified several grand challenges the research field of music recommender systems (MRS) is facing. These are, among others, in the focus of current research in the area of MRS. We discussed (1) the *cold start problem* of items and users, with its particularities in the music domain, (2) the challenge of *automatic playlist continuation*, which is gaining importance due to the recently emerged user request of being recommended musical experiences rather than single tracks [161], and (3) the challenge of holistically *evaluating* music recommender systems, in particular, capturing aspects beyond accuracy.

In addition to the grand challenges, which are currently highly researched, we also presented a visionary outlook of what we believe to be the most interesting future research directions in MRS. In particular, we discussed (1) *psychologically inspired MRS*, which consider in the recommendation process factors such as listeners' emotion and personality, (2) *situation-aware MRS*, which holistically model contextual and environmental aspects of the music consumption process, infer listener needs and intents, and eventually integrate these models at large scale in the recommendation process, and (3) *culture-aware MRS*, which exploit the fact that music taste highly depends on the cultural background of the listener, where culture can be defined in manifold ways, including historical, political, linguistic, or religious similarities.

We hope that this article helped pinpointing major challenges, highlighting recent trends, and identifying interesting research questions in the area of music recommender systems. Believing that research addressing the discussed challenges and trends will pave the way for the next generation of music recommender systems, we are looking forward to exciting, innovative approaches and systems that improve user satisfaction and experience, rather than just accuracy measures.

¹⁹ <http://compmusic.upf.edu>.

Acknowledgements Open access funding provided by Johannes Kepler University Linz. We would like to thank all researchers in the fields of recommender systems, information retrieval, music research, and multimedia, with whom we had the pleasure to discuss and collaborate in recent years, and whom in turn influenced and helped shaping this article. Special thanks go to Peter Knees and Fabien Gouyon for the fruitful discussions while preparing the ACM Recommender Systems 2017 tutorial on music recommender systems. In addition, we would like to thank the reviewers of our manuscript, who provided useful and constructive comments to improve the original draft and turn it into what it is now. We would also like to thank Eelco Wiechert for providing additional pointers to relevant literature. Furthermore, the many personal discussions with actual users of MRS unveiled important shortcomings of current approaches and in turn were considered in this article.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

1. Adamopoulos P, Tuzhilin A (2015) On unexpectedness in recommender systems: or how to better expect the unexpected. *ACM Trans Intell Syst Technol* 5(4):54
2. Adomavicius G, Mobasher B, Ricci F, Tuzhilin A (2011) Context-aware recommender systems. *AI Mag* 32:67–80
3. Adomavicius G, Tuzhilin A (2005) Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions. *IEEE Trans Knowl Data Eng* 17(6):734–749. <https://doi.org/10.1109/TKDE.2005.99>
4. Agarwal D, Chen BC (2009) Regression-based latent factor models. In: Proceedings of the 15th ACM SIGKDD international conference on knowledge discovery and data mining. ACM, pp 19–28
5. Aggarwal CC (2016) Content-based recommender systems. In: *Recommender systems*. Springer, pp 139–166
6. Aggarwal CC (2016) Ensemble-based and hybrid recommender systems. In: *Recommender systems*. Springer, pp 199–224
7. Aggarwal CC (2016) Evaluating recommender systems. In: *Recommender systems*. Springer, pp 225–254
8. Aioli F (2013) Efficient top-n recommendation for very large scale binary rated datasets. In: Proceedings of the 7th ACM conference on recommender systems. ACM, pp. 273–280
9. Alghoniemy M, Tewfik A (2001) A network flow model for playlist generation. In: Proceedings of the IEEE international conference on multimedia and expo (ICME), Tokyo, Japan
10. Alghoniemy M, Tewfik AH (2000) User-defined music sequence retrieval. In: Proceedings of the eighth ACM international conference on multimedia, pp 356–358. ACM
11. Baeza-Yates R, Ribeiro-Neto B (2011) *Modern information retrieval—the concepts and technology behind search*, 2nd edn. Addison-Wesley, Pearson
12. Baltrunas L, Kaminskas M, Ludwig B, Moling O, Ricci F, Lüke KH, Schwaiger R (2011) InCarMusic: Context-Aware Music Recommendations in a Car. In: International conference on electronic commerce and web technologies (EC-Web), Toulouse, France
13. Barrington L, Oda R, Lanckriet GRG. Smarter than genius? Human evaluation of music recommender systems. In: Proceedings of the 10th international society for music information retrieval conference, ISMIR 2009, Kobe International Conference Center, Kobe, Japan, 26–30 October 2009, pp 357–362
14. Barthet M, Fazekas G, Sandler M (2012) Multidisciplinary perspectives on music emotion recognition: Implications for content and context-based models. In: Proceedings of international symposium on computer music modelling and retrieval, pp 492–507
15. Bauer C, Novotny A (2017) A consolidated view of context for intelligent systems. *J Ambient Intell Smart Environ* 9(4):377–393. <https://doi.org/10.3233/ais-170445>
16. Bennett PN, Radlinski F, White RW, Yilmaz E (2011) Inferring and using location metadata to personalize web search. In: Proceedings of the 34th international ACM SIGIR conference on research and development in information retrieval, SIGIR’11. ACM, New York, NY, USA, pp 135–144. <https://doi.org/10.1145/2009916.2009938>
17. Bodner E, Iancu I, Gilboa A, Sarel A, Mazor A, Amir D (2007) Finding words for emotions: the reactions of patients with major depressive disorder towards various musical excerpts. *Arts Psychother* 34(2):142–150
18. Boer D, Fischer R (2010) Towards a holistic model of functions of music listening across cultures: a culturally decentred qualitative approach. *Psychol Music* 40(2):179–200
19. Bogdanov D, Haro M, Fuhrmann F, Xambó A, Gómez E, Herrera P (2013) Semantic audio content-based music recommendation and visualization based on user preference examples. *Inf Process Manag* 49(1):13–33
20. Bollen D, Knijnenburg BP, Willemsen MC, Graus M (2010) Understanding choice overload in recommender systems. In: Proceedings of the 4th ACM conference on recommender systems, Barcelona, Spain
21. Bonnin G, Jannach D (2015) Automated generation of music playlists: survey and experiments. *ACM Comput Surv* 47(2):26
22. Braunhofer M, Elahi M, Ricci F (2014) Techniques for cold-starting context-aware mobile recommender systems for tourism. *Intelli Artif* 8(2):129–143. <https://doi.org/10.3233/IA-140069>
23. Braunhofer M, Elahi M, Ricci F (2015) User personality and the new user problem in a context-aware point of interest recommender system. In: Tussyadiah I, Inversini A (eds) *Information and communication technologies in tourism 2015*. Springer, Cham, pp 537–549
24. Breese JS, Heckerman D, Kadie C (1998) Empirical analysis of predictive algorithms for collaborative filtering. In: Proceedings of the 14th conference on uncertainty in artificial intelligence. Morgan Kaufmann Publishers Inc., pp 43–52
25. Burger JM (2010) *Personality*. Wadsworth Publishing, Belmont
26. Burke R (2002) Hybrid recommender systems: survey and experiments. *User Model User-Adap Interact* 12(4):331–370
27. Burke R (2007) Hybrid web recommender systems. Springer Berlin Heidelberg, Berlin, pp 377–408. https://doi.org/10.1007/978-3-540-72079-9_12
28. Cantador I, Cremonesi P (2014) Tutorial on cross-domain recommender systems. In: Proceedings of the 8th ACM conference on recommender systems, RecSys’14. ACM, New York, NY, USA, pp 401–402. <https://doi.org/10.1145/2645710.2645777>
29. Cantador I, Fernández-Tobías I, Berkovsky S, Cremonesi P (2015) Cross-domain recommender systems. Springer, Boston, pp 919–959. https://doi.org/10.1007/978-1-4899-7637-6_27
30. Carenini G, Smith J, Poole D (2003) Towards more conversational and collaborative recommender systems. In: Proceedings of the 8th international conference on intelligent user interfaces, IUI’03. ACM, New York, NY, USA, pp. 12–18. <https://doi.org/10.1145/604045.604052>
31. Cebrián T, Planagumà M, Villegas P, Amatriain X (2010) Music recommendations with temporal context awareness. In: Proceedings of the 4th ACM conference on recommender systems (RecSys), Barcelona, Spain
32. Chen S, Moore JL, Turnbull D, Joachims T (2012) Playlist prediction via metric embedding. In: Proceedings of the 18th ACM

- SIGKDD international conference on knowledge discovery and data mining, KDD'12. ACM, New York, NY, USA, pp 714–722. <https://doi.org/10.1145/2339530.2339643>
33. Chen S, Moore JL, Turnbull D, Joachims T (2012) Playlist prediction via metric embedding. In: Proceedings of the 18th ACM SIGKDD international conference on knowledge discovery and data mining. ACM, pp 714–722
 34. Chen S, Xu J, Joachims T (2013) Multi-space probabilistic sequence modeling. In: Proceedings of the 19th ACM SIGKDD international conference on knowledge discovery and data mining. ACM, pp 865–873
 35. Cheng Z, Shen J (2014) Just-for-me: an adaptive personalization system for location-aware social music recommendation. In: Proceedings of the 4th ACM international conference on multimedia retrieval (ICMR), Glasgow, UK
 36. Cheng Z, Shen J (2016) On effective location-aware music recommendation. *ACM Trans Inf Syst* 34(2):13
 37. Cornelis O, Six J, Holzapfel A, Leman M (2013) Evaluation and recommendation of pulse and tempo annotation in ethnic music. *J New Music Res* 42(2):131–149. <https://doi.org/10.1080/09298215.2013.812123>
 38. Cremonesi P, Elahi M, Garzotto F (2017) User interface patterns in recommendation-empowered content intensive multimedia applications. *Multimed Tools Appl* 76(4):5275–5309. <https://doi.org/10.1007/s11042-016-3946-5>
 39. Cremonesi P, Quadrana M (2014) Cross-domain recommendations without overlapping data: Myth or reality? In: Proceedings of the 8th ACM conference on recommender systems, RecSys'14. ACM, New York, NY, USA, pp. 297–300. <https://doi.org/10.1145/2645710.2645769>
 40. Cremonesi P, Tripodi A, Turrin R (2011) Cross-domain recommender systems. In: IEEE 11th international conference on data mining workshops, pp 496–503. <https://doi.org/10.1109/ICDMW.2011.57>
 41. Cunningham S, Caulder S, Grout V (2008) Saturday night or fever? Context-aware music playlists. In: Proceedings of the 3rd international audio mostly conference: sound in motion, Piteå, Sweden
 42. Cunningham SJ, Bainbridge D, Falconer A (2006) 'More of an art than a science': supporting the creation of playlists and mixes. In: Proceedings of the 7th international conference on music information retrieval (ISMIR), Victoria, BC, Canada
 43. Cunningham SJ, Bainbridge D, McKay D (2007) Finding new music: a diary study of everyday encounters with novel songs. In: Proceedings of the 8th international conference on music information retrieval, Vienna, Austria, pp 83–88
 44. Cunningham SJ, Downie JS, Bainbridge D (2005) "The Pain, The Pain": modelling music information behavior and the songs we hate. In: Proceedings of the 6th international conference on music information retrieval (ISMIR 2005), London, UK, pp 474–477
 45. Cunningham SJ, Nichols DM (2009) Exploring social music behaviour: an investigation of music selection at parties. In: Proceedings of the 10th international society for music information retrieval conference (ISMIR 2009), Kobe, Japan
 46. Deldjoo Y, Cremonesi P, Schedl M, Quadrana M (2017) The effect of different video summarization models on the quality of video recommendation based on low-level visual features. In: Proceedings of the 15th international workshop on content-based multimedia indexing. ACM, p. 20
 47. Deldjoo Y, Elahi M, Cremonesi P, Garzotto F, Piazzolla P, Quadrana M (2016) Content-based video recommendation system based on stylistic visual features. *J Data Semant*. <https://doi.org/10.1007/s13740-016-0060-9>
 48. Dey AK (2001) Understanding and using context. *Pers Ubiquitous Comput* 5(1):4–7. <https://doi.org/10.1007/s007790170019>
 49. Dey L, Asad MU, Afroz N, Nath RPD (2014) Emotion extraction from real time chat messenger. In: 2014 International conference on informatics, electronics vision (ICIEV), pp 1–5. <https://doi.org/10.1109/ICIEV.2014.6850785>
 50. Donaldson J (2007) A hybrid social-acoustic recommendation system for popular music. In: Proceedings of the ACM conference on recommender systems (RecSys), Minneapolis, MN, USA
 51. Dror G, Koenigstein N, Koren Y, Weimer M (2011) The yahoo! music dataset and kdd-cup'11. In: Proceedings of the 2011 international conference on KDD Cup 2011, vol 18, pp 3–18. JMLR.org
 52. Dunn G, Wiersema J, Ham J, Aroyo L (2009) Evaluating interface variants on personality acquisition for recommender systems. In: Proceedings of the 17th international conference on user modeling, adaptation, and Personalization: formerly UM and AH, UMAP'09. Springer, Berlin, Heidelberg, pp 259–270
 53. Dutta S, Murthy HA (2014) Discovering typical motifs of a raga from one-liners of songs in carnatic music. In: Proceedings of the 15th international society for music information retrieval conference (ISMIR), Taipei, Taiwan, pp 397–402
 54. Dzhambazov G, Srinivasamurthy A, Şentürk S, Serra X (2016) On the use of note onsets for improved lyrics-to-audio alignment in turkish makam music. In: 17th International society for music information retrieval conference (ISMIR 2016), New York, USA
 55. Ebrahimi Kahou S, Michalski V, Konda K, Memisevic R, Pal C (2015) Recurrent neural networks for emotion recognition in video. In: Proceedings of the 2015 ACM on international conference on multimodal interaction, ICM'15. ACM, New York, NY, USA, pp 467–474. <https://doi.org/10.1145/2818346.2830596>
 56. Eghbal-zadeh H, Lehner B, Schedl M, Widmer G (2015) I-Vectors for timbre-based music similarity and music artist classification. In: Proceedings of the 16th international society for music information retrieval conference (ISMIR), Malaga, Spain
 57. Elahi M (2011) Adaptive active learning in recommender systems. *User Model Adapt Pers* 414–417
 58. Elahi M, Braunhofer M, Ricci F, Tkalcic M (2013) Personality-based active learning for collaborative filtering recommender systems. In: AI* IA 2013: advances in artificial intelligence. Springer, pp 360–371. https://doi.org/10.1007/978-3-319-03524-6_31
 59. Elahi M, Deldjoo Y, Bakhshandegan Moghaddam F, Cella L, Cereda S, Cremonesi P (2017) Exploring the semantic gap for movie recommendations. In: Proceedings of the eleventh ACM conference on recommender systems. ACM, pp 326–330
 60. Elahi M, Reppsys V, Ricci F (2011) Rating elicitation strategies for collaborative filtering. In: Huemer C, Setzer T (eds) EC-Web, Lecture Notes in Business Information Processing, vol 85. Springer, pp 160–171. https://doi.org/10.1007/978-3-642-23014-1_14
 61. Elahi M, Ricci F, Rubens N (2012) Adapting to natural rating acquisition with combined active learning strategies. In: ISMIS'12: Proceedings of the 20th international conference on foundations of intelligent systems. Springer, Berlin, Heidelberg, pp 254–263
 62. Elahi M, Ricci F, Rubens N (2014) Active learning in collaborative filtering recommender systems. In: Hepp M, Hoffner Y (eds) E-commerce and web technologies, Lecture Notes in Business Information Processing, vol 188. Springer, pp 113–124. https://doi.org/10.1007/978-3-319-10491-1_12
 63. Elahi M, Ricci F, Rubens N (2014) Active learning strategies for rating elicitation in collaborative filtering: a system-wide perspective. *ACM Trans Intell Syst Technol* 5(1):13:1–13:33. <https://doi.org/10.1145/2542182.2542195>
 64. Elahi M, Ricci F, Rubens N (2016) A survey of active learning in collaborative filtering recommender systems. *Comput Sci Rev* 20:29–50

65. Elbadrawy A, Karypis G (2015) User-specific feature-based similarity models for top-n recommendation of new items. *ACM Trans Intell Syst Technol* 6(3):33
66. Erdal M, Kächele M, Schwenker F (2016) Emotion recognition in speech with deep learning architectures. Springer, Cham, pp 298–311. https://doi.org/10.1007/978-3-319-46182-3_25
67. Fernandez Tobias I, Braunhofer M, Elahi M, Ricci F, Ivan C (2016) Alleviating the new user problem in collaborative filtering by exploiting personality information. *User Model User-Adap Interact (Personality in Personalized Systems)*. <https://doi.org/10.1007/s11257-016-9172-z>
68. Fernández-Tobías I, Cantador I, Kaminskas M, Ricci F (2012) Cross-domain recommender systems: a survey of the state of the art. In: Spanish conference on information retrieval, p 24
69. Ferwerda B, Graus M, Vall A, Tkalčič M, Schedl M (2016) The influence of users' personality traits on satisfaction and attractiveness of diversified recommendation lists. In: Proceedings of the 4th workshop on emotions and personality in personalized services (EMPIRE 2016), Boston, USA
70. Ferwerda B, Schedl M (2016) Investigating the relationship between diversity in music consumption behavior and cultural dimensions: a cross-country analysis. In: Workshop on surprise, opposition, and obstruction in adaptive and personalized systems
71. Ferwerda B, Schedl M, Tkalčič M (2015) Personality & emotional states: understanding users' music listening needs. In: Extended proceedings of the 23rd international conference on user modeling, adaptation and personalization (UMAP), Dublin, Ireland
72. Ferwerda B, Vall A, Tkalčič M, Schedl M (2016) Exploring music diversity needs across countries. In: Proceedings of the UMAP
73. Ferwerda B, Yang E, Schedl M, Tkalčič M (2015) Personality traits predict music taxonomy preferences. In: ACM CHI'15 extended abstracts on human factors in computing systems, Seoul, Republic of Korea
74. Flexer A, Schnitzer D, Gasser M, Widmer G (2008) Playlist generation using start and end songs. In: Proceedings of the 9th international conference on music information retrieval (ISMIR), Philadelphia, PA, USA
75. Gillhofer M, Schedl M (2015) Iron maiden while jogging, debussy for dinner? An analysis of music listening behavior in context. In: Proceedings of the 21st international conference on multimedia modeling (MMM), Sydney, Australia
76. Gosling SD, Rentfrow PJ, Swann WB Jr (2003) A very brief measure of the big-five personality domains. *J Res Personal* 37(6):504–528
77. Gross J (2007) Emotion regulation: conceptual and empirical foundations. In: Gross J (ed) *Handbook of emotion regulation*, 2nd edn. The Guilford Press, New York, pp 1–19
78. Gunawardana A, Shani G (2015) Evaluating recommender systems. In: Ricci F, Rokach L, Shapira B, Kantor PB (eds) *Recommender systems handbook*, chap. 8, 2nd edn. Springer, Heidelberg, pp 256–308
79. Hart J, Sutcliffe AG, di Angeli A (2012) Evaluating user engagement theory. In: CHI conference on human factors in computing systems. Paper presented in workshop 'Theories behind UX Research and How They Are Used in Practice' 6 May 2012
80. Hassenzahl M (2005) The thing and I: understanding the relationship between user and product. Springer, Dordrecht, pp 31–42. https://doi.org/10.1007/1-4020-2967-5_4
81. Herlocker JL, Konstan JA, Terveen LG, Riedl JT (2004) Evaluating collaborative filtering recommender systems. *ACM Trans Inf Syst* 22(1):5–53. <https://doi.org/10.1145/963770.963772>
82. Herlocker JL, Konstan JA, Terveen LG, Riedl JT (2004) Evaluating collaborative filtering recommender systems. *ACM Trans Inf Syst* 22(1):5–53
83. Herrera P, Resa Z, Sordo M (2010) Rocking around the clock eight days a week: an exploration of temporal patterns of music listening. In: Proceedings of the ACM conference on recommender systems: workshop on music recommendation and discovery (WOMRAD 2010), pp 7–10
84. Hevner K (1935) Expression in music: a discussion of experimental studies and theories. *Psychol Rev* 42:186–204
85. Hu R, Pu P (2009) A comparative user study on rating vs. personality quiz based preference elicitation methods. In: Proceedings of the 14th international conference on intelligent user interfaces, IUI'09. ACM, New York, NY, USA, pp 367–372. <https://doi.org/10.1145/1502650.1502702>
86. Hu R, Pu P (2010) A study on user perception of personality-based recommender systems. In: Bra PD, Kobsa A, Chin DN (eds) *UMAP, Lecture Notes in Computer Science*, vol 6075. Springer, pp 291–302
87. Hu R, Pu P (2011) Enhancing collaborative filtering systems with personality information. In: Proceedings of the fifth ACM conference on recommender systems, RecSys'11. ACM, New York, NY, USA, pp 197–204. <https://doi.org/10.1145/2043932.2043969>
88. Hu X, Lee JH (2012) A cross-cultural study of music mood perception between American and Chinese listeners. In: Proceedings of the ISMIR
89. Hu Y, Koren Y, Volinsky C (2008) Collaborative filtering for implicit feedback datasets. In: Proceedings of the 8th IEEE international conference on data mining. IEEE, pp. 263–272
90. Hu Y, Oghihara M (2011) NextOne player: a music recommendation system based on user behavior. In: Proceedings of the 12th international society for music information retrieval conference (ISMIR 2011), Miami, FL, USA
91. Huq A, Bello J, Rowe R (2010) Automated music emotion recognition: a systematic evaluation. *J New Music Res* 39(3):227–244
92. Iman Kamehkhosh Dietmar Jannach GB (2018) How automated recommendations affect the playlist creation behavior of users. In: Joint proceedings of the 23rd ACM conference on intelligent user interfaces (ACM IUI 2018) workshops: intelligent music interfaces for listening and creation (MILC), Tokyo, Japan
93. Järvelin K, Kekäläinen J (2002) Cumulated gain-based evaluation of ir techniques. *ACM Trans Inf Syst* 20(4):422–446. <https://doi.org/10.1145/582415.582418>
94. John O, Srivastava S (1999) The big five trait taxonomy: history, measurement, and theoretical perspectives. In: Pervin LA, John OP (eds) *Handbook of personality: theory and research*, 510, 2nd edn. Guilford Press, New York, pp 102–138
95. John OP, Srivastava S (1999) The big five trait taxonomy: history, measurement, and theoretical perspectives. In: *Handbook of personality: theory and research*, vol 2, pp. 102–138
96. Juslin PN, Sloboda J (2011) *Handbook of music and emotion: theory, research, applications*. OUP, Oxford
97. Kaggle Official Homepage. <https://www.kaggle.com>. Accessed 11 March 2018
98. Kaminskas M, Bridge D (2016) Diversity, serendipity, novelty, and coverage: a survey and empirical analysis of beyond-accuracy objectives in recommender systems. *ACM Trans Interact Intell Syst* 7(1):2:1–2:42. <https://doi.org/10.1145/2926720>
99. Kaminskas M, Ricci F (2012) Contextual music information retrieval and recommendation: state of the art and challenges. *Comput Sci Rev* 6(2):89–119
100. Kaminskas M, Ricci F, Schedl M (2013) Location-aware music recommendation using auto-tagging and hybrid matching. In: Proceedings of the 7th ACM conference on recommender systems (RecSys), Hong Kong, China
101. Kelly JP, Bridge D (2006) Enhancing the diversity of conversational collaborative recommendations: a comparison. *Artif Intell Rev* 25(1):79–95. <https://doi.org/10.1007/s10462-007-9023-8>
102. Khan MM, Ibrahim R, Ghani I (2017) Cross domain recommender systems: a systematic literature review. *ACM Comput Surv* 50(3):36

103. Kim YE, Schmidt EM, Migneco R, Morton BG, Richardson P, Scott J, Speck J, Turnbull D (2010) Music emotion recognition: a state of the art review. In: Proceedings of the international society for music information retrieval conference
104. Kluver D, Konstan JA (2014) Evaluating recommender behavior for new users. In: Proceedings of the 8th ACM conference on recommender systems. ACM, pp 121–128. <https://doi.org/10.1145/2645710.2645742>
105. Knees P, Pohle T, Schedl M, Widmer G (2006) Combining audio-based similarity with web-based data to accelerate automatic music playlist generation. In: Proceedings of the 8th ACM SIGMM international workshop on multimedia information retrieval (MIR), Santa Barbara, CA, USA
106. Knees P, Schedl M (2016) Music similarity and retrieval: an introduction to audio- and web-based strategies. The information retrieval series. Springer Berlin Heidelberg. <https://books.google.it/books?id=MdRhjwEACAJ>
107. Knijnenburg BP, Willemsen MC (2015) Evaluating recommender systems with user experiments. In: Recommender systems handbook. Springer, pp 309–352
108. Knijnenburg BP, Willemsen MC, Gantner Z, Soncu H, Newell C (2012) Explaining the user experience of recommender systems. *User Model User-Adapt Interact* 22(4–5):441–504
109. Konecni VJ (1982) Social interaction and musical preference. In: *The psychology of music*, pp 497–516
110. Koole SL (2009) The psychology of emotion regulation: an integrative review. *Cogn Emot* 23:4–41
111. Kosinski M, Stillwell D, Graepel T (2013) Private traits and attributes are predictable from digital records of human behavior. *Proc Natl Acad Sci* 110(15):5802–5805
112. Kuo FF, Chiang MF, Shan MK, Lee SY (2005) Emotion-based music recommendation by association discovery from film music. In: Proceedings of the 13th annual ACM international conference on multimedia. ACM, pp 507–510
113. Laplante A (2014) Improving music recommender systems: What we can learn from research on music tastes? In: 15th International society for music information retrieval conference, Taipei, Taiwan
114. Laplante A, Downie JS (2006) Everyday life music information-seeking behaviour of young adults. In: Proceedings of the 7th international conference on music information retrieval, Victoria (BC), Canada
115. Lee JH (2011) How similar is too similar? Exploring users' perceptions of similarity in playlist evaluation. In: Proceedings of the 12th international society for music information retrieval conference (ISMIR 2011), Miami, FL, USA
116. Lee JH, Cho H, Kim YS (2016) Users' music information needs and behaviors: design implications for music information retrieval systems. *J Assoc Inf Sci Technol* 67(6):1301–1330
117. Lee JH, Wishkoski R, Aase L, Meas P, Hubbles C (2017) Understanding users of cloud music services: selection factors, management and access behavior, and perceptions. *J Assoc Inf Sci Technol* 68(5):1186–1200
118. Lehmann J, Lalmas M, Yom-Tov E, Dupret G (2012) Models of user engagement. In: Proceedings of the 20th international conference on user modeling, adaptation, and personalization, UMAP'12. Springer, Berlin, Heidelberg, pp 164–175. https://doi.org/10.1007/978-3-642-31454-4_14
119. Li Q, Myaeng SH, Guan DH, Kim BM (2005) A probabilistic model for music recommendation considering audio features. In: Asia information retrieval symposium. Springer, pp 72–83
120. Liu NN, Yang Q (2008) Eigenrank: a ranking-oriented approach to collaborative filtering. In: SIGIR'08: proceedings of the 31st annual international ACM SIGIR conference on research and development in information retrieval. ACM, New York, NY, USA, pp 83–90. <https://doi.org/10.1145/1390334.1390351>
121. Logan B (2002) Content-based playlist generation: exploratory experiments. In: Proceedings of the 3rd international symposium on music information retrieval (ISMIR), Paris, France
122. Lonsdale AJ, North AC (2011) Why do we listen to music? A uses and gratifications analysis. *Br J Psychol* 102(1):108–134
123. Maillat F, Eck D, Desjardins G, Lamere P et al (2009) Steerable playlist generation by learning song similarity from radio station playlists. In: ISMIR, pp 345–350
124. McFee B, Bertin-Mahieux T, Ellis DP, Lanckriet GR (2012) The million song dataset challenge. In: Proceedings of the 21st international conference on world wide web. ACM, pp 909–916
125. McFee B, Lanckriet G (2011) The natural language of playlists. In: Proceedings of the 12th international society for music information retrieval conference (ISMIR 2011), Miami, FL, USA
126. McFee B, Lanckriet G (2012) Hypergraph models of playlist dialects. In: Proceedings of the 13th international society for music information retrieval conference (ISMIR), Porto, Portugal
127. McNee SM, Lam SK, Konstan JA, Riedl J (2003) Interfaces for eliciting new user preferences in recommender systems. In: Proceedings of the 9th international conference on user modeling, UM'03. Springer, Berlin, Heidelberg, pp. 178–187. <http://dl.acm.org/citation.cfm?id=1759957.1759988>
128. Mei T, Yang B, Hua XS, Li S (2011) Contextual video recommendation by multimodal relevance and user feedback. *ACM Trans Inf Syst* 29(2):10
129. North A, Hargreaves D (1996) Situational influences on reported musical preference. *Psychomusicol Music Mind Brain* 15(1–2):30–45
130. North A, Hargreaves D (2008) *The social and applied psychology of music*. Oxford University Press, Oxford
131. North AC, Hargreaves DJ (1996) Situational influences on reported musical preference. *Psychomusicology A J Res Music Cogn* 15(1–2):30
132. Novello A, McKinney MF, Kohlrausch A (2006) Perceptual Evaluation of Music Similarity. In: Proceedings of the 7th international conference on music information retrieval (ISMIR), Victoria, BC, Canada
133. O'Brien HL, Toms EG (2010) The development and evaluation of a survey to measure user engagement. *J Am Soc Inf Sci Technol* 61(1):50–69. <https://doi.org/10.1002/asi.v61:1>
134. O'Hara K, Brown B (eds) (2006) *Consuming music together: social and collaborative aspects of music consumption technologies, computer supported cooperative work*, vol 35. Springer, Dordrecht
135. Pachet F, Roy P, Cazaly D (1999) A combinatorial approach to content-based music selection. In: IEEE international conference on multimedia computing and systems, 1999, vol 1. IEEE, pp 457–462
136. Pagano R, Quadra M, Elahi M, Cremonesi P (2017) Toward active learning in cross-domain recommender systems. [CoRR. arXiv:1701.02021](https://arxiv.org/abs/1701.02021)
137. Pan R, Zhou Y, Cao B, Liu NN, Lukose R, Scholz M, Yang Q (2008) One-class collaborative filtering. In: Proceedings of the 8th IEEE international conference on data mining. IEEE, pp 502–511
138. Panteli M, Benetos E, Dixon S (2016) Learning a feature space for similarity in world music. In: Proceedings of the 17th international society for music information retrieval conference (ISMIR 2016), New York, NY, USA
139. Park ST, Chu W (2009) Pairwise preference regression for cold-start recommendation. In: Proceedings of the third ACM conference on recommender systems, RecSys'09. ACM, New York, NY, USA, pp 21–28. <https://doi.org/10.1145/1639714.1639720>
140. Pettijohn T, Williams G, Carter T (2010) Music for the seasons: seasonal music preferences in college students. *Curr Psychol* 29(4):328–345

141. Pichl M, Zangerle E, Specht G (2015) Towards a context-aware music recommendation approach: what is hidden in the playlist name? In: 2015 IEEE international conference on data mining workshop (ICDMW). IEEE, pp 1360–1365
142. Pohle T, Knees P, Schedl M, Pampalk E, Widmer G (2007) “Reinventing the Wheel”: a novel approach to music player interfaces. *IEEE Trans Multimed* 9:567–575
143. Pu P, Chen L, Hu R (2012) Evaluating recommender systems from the user’s perspective: survey of the art. *User Model User-Adapt Interact* 22(4–5):317–355. <https://doi.org/10.1007/s11257-011-9115-7>
144. Punkanen M, Eerola T, Erkkilä J (2011) Biased emotional recognition in depression: perception of emotions in music by depressed patients. *J Affect Disord* 130(1–2):118–126
145. Quadrana M, Cremonesi P, Jannach D (2018) Sequence-aware recommender systems. *arXiv preprint arXiv:1802.08452*
146. Rashid AM, Karypis G, Riedl J (2008) Learning preferences of new users in recommender systems: an information theoretic approach. *SIGKDD Explor Newsl* 10:90–100. <https://doi.org/10.1145/1540276.1540302>
147. Rentfrow PJ, Gosling SD (2003) The do re mi’s of everyday life: the structure and personality correlates of music preferences. *J Personal Soc Psychol* 84(6):1236–1256
148. Repetto RC, Serra X (2014) Creating a corpus of Jingju (Beijing opera) music and possibilities for melodic analysis. In: 15th International society for music information retrieval conference, Taipei, Taiwan, pp 313–318
149. Reynolds G, Barry D, Burke T, Coyle E (2007) Towards a personal automatic music playlist generation algorithm: the need for contextual information. In: Proceedings of the 2nd international audio mostly conference: interaction with sound, Ilmenau, Germany, pp 84–89
150. Ribeiro MT, Lacerda A, Veloso A, Ziviani N (2012) Pareto-efficient hybridization for multi-objective recommender systems. In: Proceedings of the Sixth ACM conference on recommender systems, RecSys’12. ACM, New York, NY, USA, pp 19–26. <https://doi.org/10.1145/2365952.2365962>
151. Rubens N, Elahi M, Sugiyama M, Kaplan D (2015) Active learning in recommender systems. In: Recommender systems handbook—chapter 24: recommending active learning. Springer US, pp 809–846
152. Russell JA (1980) A circumplex model of affect. *J Personal Soc Psychol* 39(6):1161–1178
153. Schäfer T, Auerswald F, Bajorat IK, Ergemlidze N, Frille K, Gehrigk J, Gusakova A, Kaiser B, Pätzold RA, Sanahuja A, Sari S, Schramm A, Walter C, Wilker T (2016) The effect of social feedback on music preference. *Musicae Sci* 20(2):263–268. <https://doi.org/10.1177/1029864915622054>
154. Schäfer T, Mehlhorn C (2017) Can personality traits predict musical style preferences? A meta-analysis. *Personal Individ Differ* 116:265–273. <https://doi.org/10.1016/j.paid.2017.04.061>
155. Schäfer T, Sedlmeier P, Stdtler C, Huron D (2013) The psychological functions of music listening. *Front Psychol* 4(511):1–34
156. Schedl M (2017) Investigating country-specific music preferences and music recommendation algorithms with the LFM-1b dataset. *Int J Multimed Inf Retr* 6(1):71–84. <https://doi.org/10.1007/s13735-017-0118-y>
157. Schedl M, Breitschopf G, Ionescu B (2014) Mobile music genius: reggae at the beach, metal on a Friday night? In: Proceedings of the 4th ACM international conference on multimedia retrieval (ICMR), Glasgow, UK
158. Schedl M, Flexer A, Urbano J (2013) The neglected user in music information retrieval research. *J Intell Inf Syst* 41:523–539
159. Schedl M, Gómez E, Trent ES, Tkalčič M, Eghbal-Zadeh H, Martorell A (2017) On the Interrelation between listener characteristics and the perception of emotions in classical orchestra music. *IEEE Trans Affect Comput*. <https://doi.org/10.1109/TAFFC.2017.2663421>
160. Schedl M, Hauger D, Schnitzer D (2012) A model for serendipitous music retrieval. In: Proceedings of the 2nd workshop on context-awareness in retrieval and recommendation (CaRR), Lisbon, Portugal
161. Schedl M, Knees P, Gouyon F (2017) New paths in music recommender systems research. In: Proceedings of the 11th ACM conference on recommender systems (RecSys 2017), Como, Italy
162. Schedl M, Knees P, McFee B, Bogdanov D, Kaminskas M (2015) Music recommender systems. In: Ricci F, Rokach L, Shapira B, Kantor PB (eds) Recommender systems handbook, chap. 13, 2nd edn. Springer, Berlin, pp 453–492
163. Schedl M, Melenhorst M, Liem CC, Martorell A, Mayor O, Tkalčič M (2016) A personality-based adaptive system for visualizing classical music performances. In: Proceedings of the 7th ACM multimedia systems conference (MMSys), Klagenfurt, Austria
164. Schein AI, Popescul A, Ungar LH, Pennock DM (2002) Methods and metrics for cold-start recommendations. In: SIGIR’02: Proceedings of the 25th annual international ACM SIGIR conference on research and development in information retrieval. ACM, New York, NY, USA, pp 253–260. <https://doi.org/10.1145/564376.564421>
165. Serra X (2014) Computational approaches to the art music traditions of India and Turkey. *J New Music Res* 43(1):1–2. <https://doi.org/10.1080/09298215.2014.894083>
166. Serra X (2014) Creating research corpora for the computational study of music: the case of the compmusic project. In: AES 53rd international conference on semantic audio. AES, AES, London, UK, pp 1–9
167. Seyerlehner K, Schedl M, Pohle T, Knees P (2010) Using block-level features for genre classification, tag classification and music similarity estimation. In: Extended abstract to the music information retrieval evaluation eXchange (MIREX 2010)/11th international society for music information retrieval conference (ISMIR 2010), Utrecht, the Netherlands
168. Seyerlehner K, Widmer G, Schedl M, Knees P (2010) Automatic music tag classification based on block-level features. In: Proceedings of the 7th sound and music computing conference (SMC), Barcelona, Spain
169. Shao B, Wang D, Li T, Ogihara M (2009) Music recommendation based on acoustic features and user access patterns. *IEEE Trans Audio Speech Lang Process* 17(8):1602–1611
170. Skowron M, Ferwerda B, Tkalčič M, Schedl M (2016) Fusing social media cues: personality prediction from Twitter and Instagram. In: Proceedings of the 25th international world wide web conference (WWW), Montreal, Canada
171. Skowron M, Lemmerich F, Ferwerda B, Schedl M (2017) Predicting genre preferences from cultural and socio-economic factors for music retrieval. In: Proceedings of the ECIR
172. Slaney M, White W (2006) Measuring playlist diversity for recommendation systems. In: Proceedings of the 1st ACM workshop on Audio and music computing multimedia. ACM, pp 77–82
173. Smyth B, McClave P (2001) Similarity vs. diversity. In: Proceedings of the 4th international conference on case-based reasoning: case-based reasoning research and development, ICCBR’01. Springer, London, UK, pp 347–361. <http://dl.acm.org/citation.cfm?id=646268.758890>
174. Sordo M, Chaachoo A, Serra X (2014) Creating corpora for computational research in arab-andalusian music. In: 1st International workshop on digital libraries for musicology, London, UK, pp. 1–3. <https://doi.org/10.1145/2660168.2660182>
175. Swearingen K, Sinha R (2001) Beyond algorithms: an hci perspective on recommender systems. In: ACM SIGIR 2001 workshop on recommender systems, vol 13, pp 1–11

176. Tamir M (2011) The maturing field of emotion regulation. *Emot Rev* 3:3–7
177. Tintarev N, Lofi C, Liem CC (2017) Sequences of diverse song recommendations: an exploratory study in a commercial system. In: Proceedings of the 25th conference on user modeling, adaptation and personalization, UMAP'17. ACM, New York, NY, USA, pp 391–392. <https://doi.org/10.1145/3079628.3079633>
178. Tkalcic M, Kosir A, Tasic J (2013) The Idos-peraff-1 corpus of facial-expression video clips with affective, personality and user-interaction metadata. *J Multimodal User Interfaces* 7(1–2):143–155. <https://doi.org/10.1007/s12193-012-0107-7>
179. Tkalčić M, Quercia D, Graf S (2016) Preface to the special issue on personality in personalized systems. *User Model User-Adapt Interact* 26(2):103–107. <https://doi.org/10.1007/s11257-016-9175-9>
180. Uitdenbogerd A, Schyndel R (2002) A review of factors affecting music recommender success. In: 3rd International conference on music information retrieval, ISMIR 2002. IRCAM-Centre Pompidou, pp 204–208
181. Vall A, Quadrana M, Schedl M, Widmer G, Cremonesi P (2017) The importance of song context in music playlists. In: Proceedings of the poster track of the 11th ACM conference on recommender systems (RecSys), Como, Italy
182. Vargas S, Baltrunas L, Karatzoglou A, Castells P (2014) Coverage, redundancy and size-awareness in genre diversity for recommender systems. In: Proceedings of the 8th ACM conference on recommender systems, RecSys'14. ACM, New York, NY, USA, pp 209–216. <https://doi.org/10.1145/2645710.2645743>
183. Vargas S, Castells P (2011) Rank and relevance in novelty and diversity metrics for recommender systems. In: Proceedings of the 5th ACM conference on recommender systems (RecSys), Chicago, IL, USA
184. Wang X, Rosenblum D, Wang Y (2012) Context-aware mobile music recommendation for daily activities. In: Proceedings of the 20th ACM international conference on multimedia. ACM, Nara, Japan, pp 99–108
185. Weimer M, Karatzoglou A, Smola A (2008) Adaptive collaborative filtering. In: RecSys'08: proceedings of the 2008 ACM conference on recommender systems. ACM, New York, NY, USA, pp. 275–282. <https://doi.org/10.1145/1454008.1454050>
186. Yang YH, Chen HH (2011) Music emotion recognition. CRC Press, Boca Raton
187. Yang YH, Chen HH (2012) Machine recognition of music emotion: a review. *ACM Trans Intell Syst Technol* 3(4):40
188. Yang YH, Chen HH (2013) Machine recognition of music emotion: a review. *Trans Intell Syst Technol* 3(3):40:1–40:30
189. Yoshii K, Goto M, Komatani K, Ogata T, Okuno HG (2006) Hybrid collaborative and content-based music recommendation using probabilistic model with latent user preferences. In: ISMIR, vol 6, p 7th
190. Zamani H, Bendersky M, Wang X, Zhang M (2017) Situational context for ranking in personal search. In: Proceedings of the 26th international conference on world wide web, WWW'17. International world wide web conferences steering committee, Republic and Canton of Geneva, Switzerland, pp 1531–1540. <https://doi.org/10.1145/3038912.3052648>
191. Zentner M, Grandjean D, Scherer KR (2008) Emotions evoked by the sound of music: characterization, classification, and measurement. *Emotion* 8(4):494
192. Zhang Z, Jin X, Li L, Ding G, Yang Q (2016) Multi-domain active learning for recommendation. In: AAAI, pp 2358–2364
193. Zhang YC, O Seaghdha D, Quercia D, Jambor T (2012) Auralist: introducing serendipity into music recommendation. In: Proceedings of the 5th ACM international conference on web search and data mining (WSDM), Seattle, WA, USA
194. Zheleva E, Guiver J, Mendes Rodrigues E, Milić-Frayling N (2010) Statistical models of music-listening sessions in social media. In: Proceedings of the 19th international conference on world wide web (WWW), Raleigh, NC, USA, pp 1019–1028
195. Zhou T, Kuscsik Z, Liu JG, Medo M, Wakeling JR, Zhang YC (2010) Solving the apparent diversity-accuracy dilemma of recommender systems. *Proc Natl Acad Sci* 107(10):4511–4515
196. Ziegler CN, McNee SM, Konstan JA, Lausen G (2005) Improving recommendation lists through topic diversification. In: Proceedings of the 14th international conference on the world wide web. ACM, pp 22–32