Chapter

Peter Knees*, Markus Schedl, Bruce Ferwerda, and Audrey Laplante

User Awareness in Music Recommender Systems

Abstract: Music recommender systems are a widely adopted application of personalized systems and interfaces. By tracking the listening activity of their users and building preference profiles, a user can be given recommendations based on the preference profiles of all users (collaborative filtering), characteristics of the music listened to (content-based methods), meta-data and relational data (knowledge-based methods; sometimes also considered content-based methods) or a mixture of these with other features (hybrid methods). In this chapter, we focus on the listener's aspects of music recommender systems. We discuss different factors influencing relevance for recommendation on both the listener's and the music's side and categorize existing work. In more detail, we then review aspects of (i) listener background in terms of individual, i.e., personality traits and demographic characteristics, and cultural features, i.e., societal and environmental characteristics, (ii) listener context, in particular modeling dynamic properties and situational listening behavior, and (iii) listener intention, in particular by studying music information behavior, i.e., how people seek, find, and use music information. This is followed by a discussion of user-centric evaluation strategies for music recommender systems. We conclude the chapter with a reflection on current barriers, by pointing out current and longer-term limitations of existing approaches and outlining strategies for overcoming these.

Keywords: music recommender systems, personalization, user modeling, user context, user intent

***Corresponding author: Peter Knees,** Institute of Information Systems Engineering, Faculty of Informatics, TU Wien, Vienna, Austria
Markus Schedl, Institute of Computational Perception, Johannes Kepler University Linz, Linz, Austria
Bruce Ferwerda, Department of Computer Science and Informatics, School of Engineering, Jönköping University, Jönköping, Sweden
Audrey Laplante, École de bibliothéconomie et des sciences de l'information, Université de Montréal, Quebéc, Canada

1 Introduction

Music recommender systems are a widely adopted application of personalized systems and interfaces [118]. On a technical level, large-scale music recommender systems became feasible through online music distribution channels and collective platforms that track users' listening events.[1] By tracking the listening activity of their users and building preference profiles, a user can be given recommendations based on the pool of preference profiles of all users (collaborative filtering, e.g., [124, 15]), characteristics of the music listened to (content-based methods, e.g., [12, 129]), expert- and user-generated (relational) meta-data (knowledge-based methods, e.g., [134, 70, 103]; sometimes also considered content-based methods), or a mixture of these, potentially extended by other features (hybrid methods, e.g., [24, 94, 60]).

The main research area exploring these opportunities, i.e., music information retrieval (MIR), historically, has predominantly followed content-based approaches [71]. This can facilitate music recommendation starting from preferred examples and then following the query-by-example paradigm, which is central to information retrieval tasks. While aspects of user adaptivity and relevance feedback can be addressed, e.g., [104, 72], modeling of the listener was underrepresented in the majority of work [116].

For developing recommender systems, traditionally, static collections and recorded user interactions have served as offline ground truth. This permits researchers to optimize retrieval and recommendation system performance, e.g., by maximizing precision or minimizing error, cf. [16, 120]. More recently, with the establishment of dedicated online music streaming platforms such as Spotify[2] and Pandora,[3] more dynamic and user-oriented criteria, assessed by means of massive online A/B testing, have driven the industrial development, cf. [126, 1]. However, both offline and online approaches operate on the basis of a system-centric view and therefore neglect user- and usage-centric perspectives on the process of music listening. Such perspectives involve, e.g., factors of listening context such as activity or social setting, listening intent, or the listener's personality, background, and preferences. Incorporating this information can enhance the process of music recommendation in a variety of situations, from mitigating cold-start scenarios, i.e. when usage data of

1 Early examples from the pre-streaming era are peer-to-peer networks and platforms like Last.fm (http://last.fm).

2 http://www.spotify.com

3 http://www.pandora.com

new users is missing, to mood- and situation-tailored suggestions, to adaptive and personalized interfaces that support the listener in his or her music information-seeking activities.

In this chapter, we focus on aspects of the listener in music recommendation. In Sec. 2, we discuss different factors that influence the relevance of recommendations. This covers both aspects of the listener and of the musical items to recommend. Additionally, we briefly outline the development from search scenarios to approaches to personalization and user adaptation. Sec. 3 deals with aspects of *listener background* and discusses variables that influence differences in music preferences of listeners, divided into individual (i.e., personality traits and demographic characteristics) and cultural features (i.e., societal and environmental characteristics). In Sec. 4, we focus on the *listener context*, i.e., contextual and situational listening behavior. To this end, we elaborate on the modeling and elicitation of the listener's emotion, on the emotion assigned to music items, and on the relationship between these two. We further discuss methods that exploit various sensor data for user modeling. Sec. 5 then focuses on *listener intention*, in particular by studying music information behavior, i.e., how people seek, find, and use music information. This includes studies conducted in the information science field on how people discover new music artists or new music genres in everyday life, as well as studies that examine how people use and perceive music recommender systems.

To round this chapter off, we give an overview of user-centric evaluation strategies for music recommender systems in Sec. 6, before concluding with a discussion of current barriers in Sec. 7, where we point out current and longer-term limitations of existing approaches and outline strategies for overcoming these. Despite presenting existing technical academic work, we highlight findings from non-technical disciplines to call attention to currently missing facets of music recommender systems. These identified but yet not technically covered requirements should help the reader in identifying potential new research directions.

2 Relevant Factors in Music Recommendation

In the field of recommender systems research, the interaction of two factors is relevant for making recommendations: the user and the item, i.e., in our case a music entity, such as a track, an artist, or a playlist. In traditional recommender systems, based on previous interactions of users and items, future interactions are predicted, either by identifying similar users or items (memory-

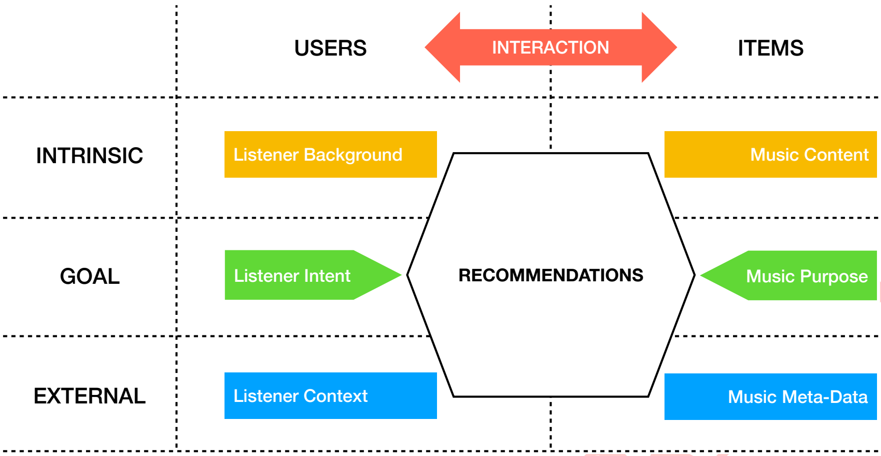

Fig. 1: User and item factors influencing recommendations. Recommender systems typically model and predict the interaction between users and items. Both users and items themselves can be described by different factors that can be categorized into *intrinsic properties*, *goal*, and *external factors*. While intrinsic factors refer to mostly stable features, goal, and external factors are more dynamic.

based collaborative filtering) or by learning latent representations of users and items by decomposing the matrix of interactions (model-based collaborative filtering). While model-based methods have the advantage of resulting in representations of users and items that permit effective prediction of future interactions, a major drawback of these latent representations is that they are hard to interpret and, while describing the data, typically cannot be connected to actual properties of neither users nor items. In an attempt to connect models to such properties (e.g., rating biases of users, popularity biases of products, or domain-specific properties like preference for the music genre of a track), more factors and degrees of freedom are included to fit the observed data (e.g., [75, 56, 26]). However, particularly in scenarios where no prior interaction data has been observed ("cold start"), such purely data-driven models show their weakness, making explicit modeling of user properties, usage context, etc. and their effects desirable.

While different aspects of music items are well covered by research in MIR, modeling different facets of the listener have found less entry into recommendation systems. For both user and item, we can identify different categories of these facets that impact recommendations, namely *intrinsic* properties, *goal*, and *external* aspects. Fig. 1 shows these six finer grained factors underlying the interaction between users and items.[4]

In terms of *item factors*, we can distinguish between:

Music Content referring to everything that is contained in or, more pragmatically, that can be extracted from the audio signal itself, such as aspects of rhythm, timbre, melody, harmony, structure, or even the mood of a piece.

Music Purpose referring to the intended usage of the music which can have a spiritual or political purpose (e.g., an anthem) or created for the purpose of playing in the background (e.g., muzak). This also relates to aspects of *associative coding* in Juslin's theory [63], which conveys connotations between the music and other "arbitrary" events or objects. The role of this facet for recommendation has remained largely unexplored, apart from special treatment of certain events, such as offering special playlists for holiday seasons like Christmas, containing tracks usually filtered out during the remaining time.

Music Meta-Data (and Cultural Context) referring to aspects of the music that cannot be inferred directly from the audio signal, such as meta-data like year of recording or country of origin, as well as different types of community meta-data: user-generated content such as reviews or tags; additional multi-modal contextual information such as album artwork, liner notes, or music videos; and diverse outcomes and impacts of marketing strategies. This also captures elements categorized as associative coding (see above).

Correspondingly, in terms of *user facets*, we can distinguish between:

Listener Background referring to the listener's personality traits, such as preference and taste, musical knowledge, training, and experience, as well as to demographics and cultural background. Generally, this comprises more static and stable characteristics of the user.

Listener Intent referring to the goal of the listener in consuming music. Potential goals span from evocation of certain emotions, i.e., emotional self-regulation, to the desire to demonstrate knowledge, musical sophistication, or musical taste in a social setting.

4 This extends the four categories of influences of music similarity perception by Schedl et al. [116] and further integrates aspects of the model of music perception and conveyed emotions by Juslin [63].

Listener Context referring to the current situation and setting of the user, including location, time, activity, social context, or mood. This generally describes more dynamic aspects of the listener.

In the following, we focus on the listener and review work dealing with these different dimensions. First, considering aspects of listener background, we give an overview of work exploring music preference, personality, and cultural characteristics. Second, we focus on contextual and more dynamic factors, namely modeling the listener's emotional state, as well as deriving a listener's context from sensor data of personal devices. Finally, we deal with the least explored area in terms of music recommender systems, i.e., the listener's intent, in the context of music information behavior.

3 Correlates of Music Preferences

Music plays an important part in our lives. We actively use music as a resource to support us in everyday life activities. Hence, music can have different functions (cf. Sec. 5). Merriam and Merriam [95] defined several functions of music amongst which are: emotion expression, aesthetic enjoyment, entertainment, communication, and symbolic representation. Given the different functions of music, the kind of music that is appropriate for a certain function may be a personal matter. What we like to listen to is shaped by our personal tastes and preferences as well as by our cultural preconceptions [100]. Although there is ample research done in traditional psychology on individual differences of music preferences, it is important to investigate to what extent these findings still hold in a technological mediated context as well as investigating new relationships that have become available through new interaction opportunities that technologies facilitate. In the sections below we discuss work on music preferences in a technological setting that deals with preference correlations with individual as well as cultural aspects with which we try to draw parallels with results from traditional psychology.

3.1 Individual aspects

The existence of individual differences in music preferences has been investigated quite extensively in traditional psychological research already (for an overview see [109]). However, with recent technological advances, current online music systems (e.g., online music streaming services) provide their users with an almost unlimited amount of content that is directly at their disposal. This abundance of available music may have deviating effects on our prior knowledge of how people listen to music. For example, users may be prone to try out different content more than they would do in the offline world and even their preference may change more often or becomes more versatile [42].

Prior psychological work argued that age may play an important role in identifying individual differences in music preferences due to varying influences that shape music preferences across the course of life. For example, an individual may develop their music taste through the influence of parents in the early ages, but get influenced by the taste of peers later on in life [110]. Recent work investigated the change of music preferences over age by analyzing the music listening histories of an online music streaming service [41]. By being able to trace back the music listening histories, a mapping was able to be made on the change of music preferences. Although the online music listening behaviors reflected more diversity and versatility, the general trends are in line with prior psychological work [110]. The results showed that over time music preferences became more stable, while in the younger age groups music preferences are more exploratory and scattered across genres. Recent work has shown that, conversely, demographic information of users can also be predicted from online listening behavior [76] as well as the musical sophistication of music listeners [30].

Aside of the identification of age differences, another way that is often used to segment online music listeners is based on their personality. Personality has shown to be a stable construct and is often used as a general model to relate behavior, preferences, and needs of people to [62]. A common way to segment people on personality traits is being done based on the five factor model (FFM). The FFM describes personality traits based on five general dimensions: openness to experience, conscientiousness, extraversion, agreeableness, and neuroticism (see Tab. 1).

General Dimensions	Primary Factors
Openness to Experience	Artistic, curious, imaginative, insightful, original, wide interest
Conscientiousness	Efficient, organized, planful, reliable, responsible, thorough
Extraversion	Active, assertive, energetic, enthusiastic, outgoing, talkative
Agreeableness	Appreciative, forgiving, generous, kind, sympathetic, trusting
Neuroticism	Anxious, self-pitying, tense, touchy, unstable, worrying

Tab. 1: Five-factor model, adopted from [62].

Several works have shown the relationship between personality traits and online music preferences, but mainly investigated the relationship of personality traits with ways of interacting within a system. Although these results do not allow for a comparison with results from traditional psychology, they do provide new insights in user interactions with music systems and how these interactions can be personalized. Ferwerda et al. [38] investigated how personality and emotional states influence the preference for certain kinds of music. Other studies have shown that personality is related to the way we browse (i.e., by genre, mood, or activity) for online music [44]. In their study they simulated an online music streaming service in which they observed how users navigated through the service to find music that they liked to listen to. Tkalčič et al. [128] investigated the relationship between personality traits and digital concert notes. They looked at whether personality influences the preferences of the amount of content presented. Others looked at the music diversity needs based on personality traits [31], and have proposed ways to incorporate personality traits into music recommender systems to improve music recommendations to users [33, 35].

3.2 Cultural aspects

Aside from individual aspects, preference differences can already occur on a cultural level. The environments that we are exposed to have a big influence in how our preferences are shaped [100, 115]. Especially with services being online and widespread, analyses of more global behaviors are possible. For example, artist preference differences have been found based on linguistic distance [91]. A known way to investigate cultures is by relying on Hofstede's cultural dimensions [54]. Although this model originates from 1968, it is still being actualized. Hofstede's cultural dimensions are based on data of 97 countries. The data showed patterns that resulted in the following six dimensions: power distance index, individualism, uncertainty avoidance index, masculinity, long-term orientation, and indulgence, as described in the following.

Power distance defines the extent to which power is distributed unequally by less powerful members of institutions (e.g., family). High power distance indicates that a hierarchy is clearly established and executed in society. Low power distance indicates that authority is questioned and power attempted to be distributed equally.

Individualism defines the degree of integration of people into societal groups. High individualism is defined by loose social ties—the main emphasis is on the “I” instead of the “we”—while this is the opposite for low individualistic cultures.

Masculinity defines a society’s preference for achievement, heroism, assertiveness and material rewards for success (countries scoring high in this dimension). Conversely, low masculinity represents a preference for cooperation, modesty, caring for the weak and quality of life.

Uncertainty avoidance defines a society’s tolerance for ambiguity. High scoring countries in this scale are more inclined to opt for stiff codes of behavior, guidelines, and laws, whereas more acceptance of different thoughts and/or ideas are more common for those scoring low in this dimension.

Long-term orientation is associated with the connection of the past with the current and future actions and/or challenges. Lower scoring countries tend to believe that traditions are honored and kept, and value steadfastness. High scoring countries believe more that adaptation and circumstantial, pragmatic problem-solving are necessary.

Indulgence defines in general the happiness of a country. Countries scoring high in this dimension are related to a society that allows relatively free gratification of basic and natural human desires related to enjoying life and having fun (e.g., be in control of their own life and emotions), whereas low scoring countries show more controlled gratification of needs and regulate it by means of strict social norms.

Studies that looked at Hofstede’s cultural dimensions found differences on several aspects, for example, diversity in music listening. Countries scoring high on the power distance tend to show less diversity in the artists and genres they listened to. The individualism dimension was found to negatively correlate with music diversity [43]. Extended analysis can be found in [34]. Others have shown that Hofstede’s cultural dimensions and socio-economic factors can be used to predict genre preferences of music listeners [92, 119, 122]. By applying a random forest algorithm, they were able to achieve an improvement of 16.4% in genre prediction over the baseline [122].

The identification of individual and cultural differences with regards to music contributes to new and deeper understanding of behaviors, preferences, and needs in online music environments. Aside of that, the findings also provide insights on how these differences can be exploited for personalizing experiences. For example, a persistent problem for personalized systems is implicit preference elicitation for new users. Relying on identified individual and cultural differences may contribute to mitigate these preference elicitation problems. For example, research has shown that personality can be predicted from behavioral trails on social media (e.g., Facebook [14, 35, 49], Twitter [48, 108], Instagram [36, 37, 39, 40, 83], and a combination of social media sources [123]). With the increased implementation of single sign-on (SSO) mechanisms [5] allow users to easily login and register to an application, but also let applications import user information from the connected application. Hence, these personalization prediction methods could play an important role in personalization strategies to mitigate preference elicitation for new users. When there are no external information sources available to extract user information from, personalization strategies based on cultural findings may be the second best option for personalization. Country information often already consists in a standard user profile and is therefore easy to acquire.

4 Contextual and Situational Music Listening Behavior

The situation or context a person is in when listening to music—or deciding what to listen to—is known to strongly affect the music preferences as well as consumption and interaction behavior [7, 22]. To give an example, a person is likely to listen to different music or create a different playlist when preparing for a romantic dinner than when preparing to go out on a Saturday night [47].

The most frequently considered types of context include *location* (e.g., listening at workplace, when commuting, or relaxing at home) [66] and *time* (typically categorized into, e.g., morning, afternoon, and evening) [13]. In addition, context may also relate to the listener's *activity* [133], *weather* [106], *listening device*, e.g., earplugs on a smartphone vs. hi-fi stereo at home [47], and various *social aspects* [20, 101], just to name a few.

Another type of context is *interactional context* with sequences, which is particularly important for the tasks of session-based recommendation and sequence-aware recommendation. In this case, context refers to the sequence of music pieces a listener decides to consume consecutively. In the music domain, such tasks are often referred to as automatic playlist generation or automatic playlist continuation [10, 17]. Sequence learning and natural language processing techniques applied to playlist names are typically used to infer contextual aspects.

5 Buttons that allow users to register or login with accounts of other applications. For example, social networking services: "Login with your Facebook account."

A particularly important situational characteristic is that of *emotion*, both from a user's perspective [113] and song annotations [137]. In the following, we therefore first introduce in Sec. 4.1 the most common approaches to model listeners' moods and emotions, emotions perceived while listening to music, and ways to affectively connect listeners and music pieces. In Sec. 4.2, we subsequently review methods that exploit various sensor data for user modeling in music recommender systems, e.g., from sensors built into smart devices. Such sensor data can either be used directly to learn contextual music preferences, or to infer higher-level context categories such as the target user's activity.

4.1 Emotion and mood: connecting listeners and music

The affective state of the listener has a strong impact on his or her short-time musical preferences [65]. Vice versa, music strongly influences our affective state. It therefore does not come as a surprise that affect regulation is regarded as one of the main reasons why people listen to music [93, 113]. As an example, people may listen to completely different musical genres or styles when they are sad in comparison to when they are happy. Indeed, prior research on music psychology discovered that people usually choose the type of music which moderates their affective condition [74]. More recent findings show that music is often chosen for the purpose of augmenting the emotional situation perceived by the listener [99].

Note that in psychology—often in contrast to recommender systems or MIR research, but also everyday use—it is common to distinguish between *mood* and *emotion* as two different affective constructs. The most important differences are that a mood is characterized as an experience of longer but less intense duration without a particular stimulus, whereas an emotion is a short experience with an identifiable stimulus event that causes it.

In order to build affect-aware music recommenders, it is necessary to (i) infer the emotional state or mood 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. In the context of (i), we also introduce the most important ways to describe emotions.

4.1.1 Modeling the listener's emotional state

The emotional state of a human can be elicited explicitly or implicitly. In the former case, the person is typically presented a questionnaire or user interface

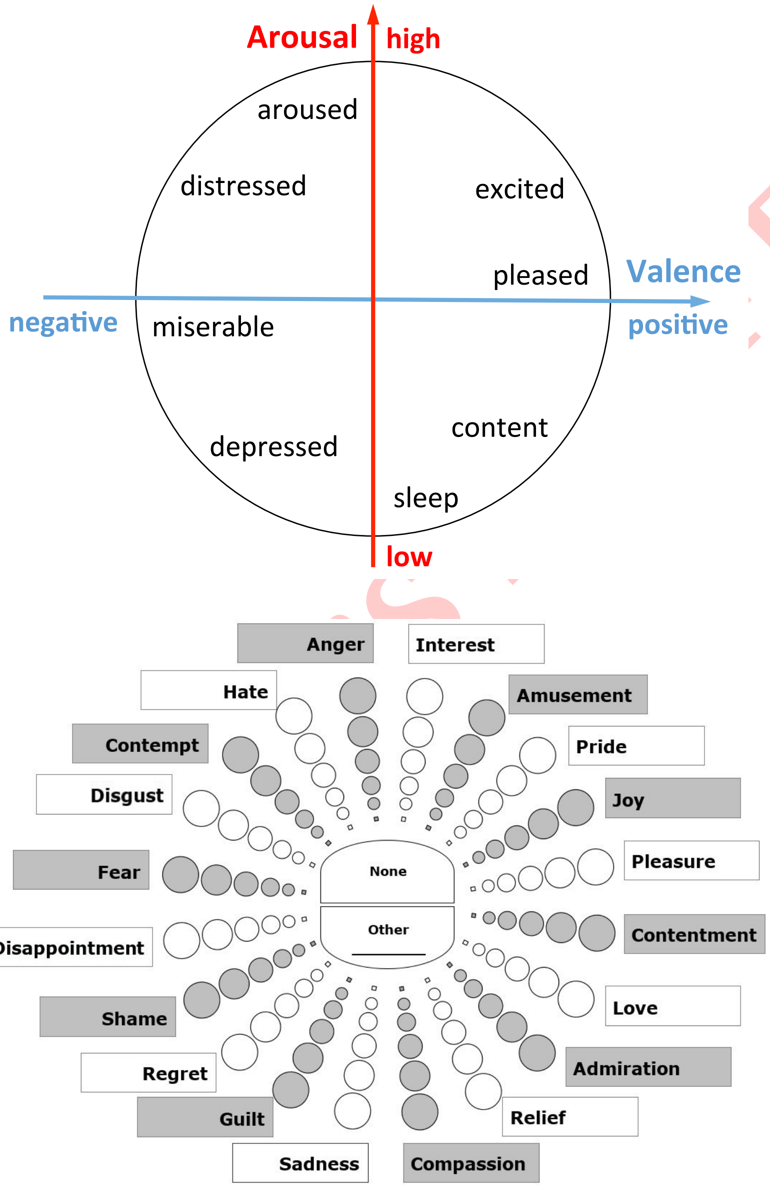

Fig. 2: Emotion models. Top: emotional terms expressed in the valence–arousal space, according to [111]. Bottom: Geneva emotion wheel, according to [121].

that maps the user's explicit input to an emotion representation according to one of the various *categorical models* or *dimensional models*. Categorical models describe emotions by distinct words such as happiness, sadness, anger, or fear [139, 53], while dimensional models described emotions by scores with respect to two or three dimensions. One of the most prominent dimensional models is Russel's two-dimensional circumplex model [111], which represents valence and arousal as orthogonal dimensions, cf. Fig. 2 (top). Into this model, categorical models can be integrated, e.g., by mapping emotion terms to certain positions within the continuous emotion space. The exact positions are commonly determined through empirical studies with humans. For a more detailed elaboration on emotion models in the context of music, we refer to [137, 117]. One prominent example is the Geneva emotion wheel,[6] depicted in Fig. 2 (bottom). It is a hybrid model that uses emotion terms as dimensions and describes the intensity of each of these emotions on a continuous scale.

Besides explicit emotion elicitation, the implicit acquisition of people's emotional states can be effected, for instance, by analyzing user-generated text [23], speech [29], or facial expressions in video [27] as well as a combination of audio and visual cues [67, 98].

4.1.2 Modeling the emotion perceived in music

Music can be regarded as an emotion-laden content and can hence also be described by emotion words, similar to listeners. The task of automatically assigning to a given music piece such emotion terms (in case of categorical emotion models) or intensities (in case of dimensional models) is an active research area, typically refereed to as music emotion recognition (MER), e.g. [138, 6, 68, 57, 139]. If categorical emotion models are adopted, the MER task is treated as a classification task, whereas it is considered a regression task in case of dimensional models. Even though a variety of datasets, feature modeling approaches, and machine learning methods have been created, devised, and applied, respectively, integrating the emotion terms or intensities predicted by MER tools into a music recommender system is, however, not an easy task due to several reasons.

In the beginning of MER research, the problem was approached from a pure machine learning perspective, taking audio features as input to predict labels, which in this case constituted emotion terms. These approaches were

6 http://www.affective-sciences.org/gew

agnostic of the actual meaning of the emotion terms as they failed to distinguish between *intended emotion*, *perceived emotion*, and *induced or felt emotion* [120]. Intended emotion refers to the emotion the composer, songwriter, or performer had in mind when creating or performing the music piece; perceived emotion refers to the emotion recognized by a person listening to a piece; induced emotion refers to the emotion felt by the listener. Current MER approaches commonly target perceived or induced emotions.

Second, the ways musical characteristics reflected by content descriptors (rhythm, tonality, lyrics, etc.) influence the emotional state of the listener remains highly subjective, even though some general rules have been identified [77]. For instance, a musical piece in major key is typically perceived brighter and happier than a piece in minor key. A fast piece is perceived more exciting or more tense than a slow one. However, the perception of music emotion also depends on other psychological constructs such as the listener's personality [117, 38], cf. Sec. 3.1.

4.1.3 Relating human emotions and music emotion annotations

As a result of the previous discussion on the relationship between listeners and emotions (Sec. 4.1.1) and between music pieces and emotions (Sec. 4.1.2), we assume to have information about user–emotion assignments and item–emotion assignments. Towards building emotion-aware music recommender systems, we next have to connect users and items through emotions, in the affectively intended way, which is a challenging endeavor. To this end, knowing about the user's intent is crucial.

Three fundamental intents or purposes of music listening have been identified in a decent study conducted by Schäfer et al. [113]: *self-awareness* (e.g., stimulating a reflection of people on their identity), *social relatedness* (e.g., feeling closeness to friends and expressing identity), and *arousal and mood regulation* (e.g., managing emotions). Several studies found that affect regulation is indeed the most important purpose why people listen to music [113, 93, 9]. Nevertheless, modeling music preferences as a function of the listener's mood, listening intent, and affective impact of listening to a certain emotionally laden music piece is still insufficiently understood.

This is the likely reason why full-fledged emotion-aware systems still do not exist, to the best of the authors' knowledge. Preliminary approaches integrate content and mood-based filtering [2] or implement cross-modal recommendation, such as matching the mood of an input video with that of music pieces and recommending matching pieces [112]. Other work infers the user's emotional state from sensor data (cf. Sec. 4.2) and matches it with explicit user-specific preference indications. For instance, Park et al. [105] gather information about temperature, humidity, noise, light level, weather, season, and time of day and subsequently use these features to predict whether the user is depressed, content, exuberant, or anxious. Based on explicit user preference feedback about which type of music he or she prefers in a given emotional state, the proposed system then adapts recommendations.

To conclude, without decent psychological listener profiles and a comprehensive understanding of the listener's affective state, listening intent, and affective impact of a song on the listener, emotion-aware recommender systems are unlikely to produce recommendations that truly satisfy the user. Gaining such insights, elaborating methods to create respective listener profiles, and subsequently devising approaches to integrate them into systems can therefore be considered open research challenges.

4.2 Sensor data for context modeling

Contextual modeling for music recommender systems can also be achieved by exploiting various sensor data, where we understand sensors in a broad sense, not only as physical or hardware devices, but also including virtual sensors like active apps or running background tasks on a personal device. Today's smart devices are packed with sensors, ranging from motion to proximity to light sensors. It has therefore become easier than ever to gather large amounts of sensor data and exploit them for various purposes, such as gait recognition [132], human activity classification [78], or personal health assistance [89].

Tab. 2 provides a categorization of some sensor data that can be gathered from smart devices [47]. Data attributes are either of categorical or numerical type. In addition to the frequently used temporal, spatial, and motion signals, the table lists hardware-specific attributes (device and phone data), environmental information about the surrounding of the user (ambient), connectivity information (network), information about used applications (task), and application-specific information (in our context, of a music player).

Such sensor data has been exploited to some extent to build user models that are subsequently integrated into context- or situation-aware music recommender systems. Most earlier approaches are characterized by taking only a single category of context sensors into account though, most often spatial and temporal features. To give a few examples, Lee and Lee [87] exploit weather conditions alongside listening histories. Cebrian et al. [13] use temporal features. Addressing the task of supporting sports and fitness training, a considerable

Category	Attributes
Time	day of week (N), hour of day (N)
Location	provider (C), latitude (C), longitude (C), accuracy (N), altitude (N)
Weather	temperature (N), wind direction (N), wind speed (N), precipitation (N), humidity (N), visibility (N), pressure (N), cloud cover (N), weather code (N)
Device	battery level (N), battery status (N), available internal/external storage (N), volume settings (N), audio output mode (C)
Phone	service state (C), roaming (C), signal strength (N), GSM indicator (N), network type (N)
Task	recently used tasks/apps (C), screen on/off (C), docking mode (C)
Network	*mobile network*: available (C), connected (C); *active network*: type (C), subtype (C), roaming (C); *Bluetooth*: available (C), enabled (C); *Wi-Fi*: enabled (C), available (C), connected (C), BSSID (C), SSID (C), IP (N), link speed (N), RSSI (N)
Ambient	light (N), proximity (N), temperature (N), pressure (N), noise (N)
Motion	acceleration force (N), rate of rotation (C), orientation of user (N), orientation of device (C)
Player	repeat mode (C), shuffle mode (C), *sound effects*: equalizer present (C), equalizer enabled (C), bass boost enabled (C), bass boost strength (N), virtualizer enabled (C), virtualizer strength (N), reverb enabled (C), reverb strength (N)

Tab. 2: Categories of common sensor data used in context modeling, adapted from [47]. Letters in parenthesis indicate whether the attribute is **C**ategorical or **N**umerical.

amount of work uses sensors to gauge steps per minute or heart rate to match the music played with the pace of the listener, or to stimulate a particular exercising behavior, for example, Biehl et al. [8], Elliott and Tomlinson [28], Dornbush et al. [25], Cunningham et al. [18], de Oliveira and Oliver [21], and Moens et al. [96]. Besides the use case of sports and exercising, there further exists research work that targets other specific task, e.g., music recommendation while working [136], driving a car [4], or for multiple activity classes such as running, eating, sleeping, studying, working, or shopping [133, 131]. An approach to identify music to accompany the daily activities of relaxing, studying, and workout is proposed in [135].

More recent research works integrate a larger variety of sensor data into user models. For instance, Okada et al. [102] present a mobile music recommender that exploits sensor data to predict the user's activity, environment, and location. Activity is inferred from the device's accelerometer and classified into idle, walking, bicycling, running, etc. The user's environment is predicted by recording audio on the device and matching it to a database of audio snippets, which are labeled as meeting, office, bus, etc. As for location, latitude and longitude GPS data is clustered into common user locations. Integrating activity, environment, location, and quantized temporal data with respect to time of day and week days, the proposed system learns rules such as "Every Sunday afternoon, the user goes jogging and listens to metal songs.", which are used among other information to effect recommendations.

Wang et al. [133] present a system, which records time of day, accelerometer data, and ambient noise. Using these features, the recommender predicts the user's activity, i.e., running, walking, sleeping, working, or shopping. Activity-aware recommendations are eventually effected by matching music pieces labeled with activity tags with the user's current activity.

Schedl et al. [114] propose a mobile music recommender called Mobile Music Genius, which acquires and monitors various context features during playback. It uses a decision tree classifier to learn to predict music preferences for given contexts. The user preference is modeled at various levels, i.e., genre, artist, track, and mood; the context is modeled as an approx. 100-dimensional feature vector, including the attributes listed in Tab. 2. While playing, the user context is continuously monitored and compared to the temporally preceding context vector. If context changes exceed a sensitivity parameter, a track that fits the new context is requested from the classifier and added as next track to the playlist.

Hong et al. [55] exploit day of the week, location, weather, and device vibration as features to build a bipartite graph, in which nodes either represent contexts or music pieces. Edges connect the two categories of nodes and the weight of an edge indicates how often a user has listened to a certain piece in a certain context. A random walk with restart algorithm is then used to identify music pieces that fit a given context. This algorithm takes into account the number, lengths, and edge weights of paths between the given context node and the music nodes.

In summary, early works that exploit sensor data for context modeling for music recommendation focused on single and typically simple sensors, such as time or weather, while more recent ones consider a variety of sensor data, extending the above by motion, environment, or location information, among others. Context models are then created either based on the entirety of the considered sensor data [114, 55], or based on inferred information such as the user's activity [102, 133, 131] or mood [105].

5 Music information behavior

To deepen our understanding the listeners' perception and uses of music recommender systems, we should also examine their music information behavior in everyday life. The term "information behavior" encompasses a wide range of activities including seeking, finding, using, selecting, avoiding, and stumbling upon information [11]. The information can come from formal sources (e.g., books, magazines, music recordings) or from informal ones (e.g. friends, family members). In other words, information behavior research does not limit its scope to users' interaction with information systems. It also looks more broadly at the information practices that surround the use (or non-use) of these systems.

In this section, we look at the music-related information behavior of people in everyday life. In the first subsection, we review the literature on how people discover music in their daily life and the place music recommender systems play in it; in the second, we focus more specifically on studies on users' perception of music recommender systems.

5.1 Discovering music in everyday life

Task-based experiments and transaction logs are useful for identifying usability problems in a system. However, the intent and experience of users can only be inferred from the traces of their interactions with the system. To know about the true users' needs, we need to take a step back to get a broader perspective on information practices in real life. To understand how music recommender systems can better support their users, we suggest looking at the studies on how people discover music in their daily life. For the most part, these studies have been conducted by researchers in information sciences and employ qualitative interviewing, observations, diaries, and surveys, with the objective of learning about the real-life behaviors of real-life users, oftentimes in real-life settings.

5.1.1 Importance of friends and families

Consistently and across age groups, studies show that friends, family, and acquaintances were and remain the main source of music discovery in everyday life [127, 79, 61, 69]. Even as music streaming services are pervasive, people prefer approaching friends and relatives to ask for music suggestions rather than seeking recommendations in online services. In a survey conducted by Lee and her colleagues [84], 82.8% reported turning to friends or family members when searching for music information. Qualitative studies reveal that people do not only turn to people out of convenience. They appreciate receiving recommendations specifically tailored to their tastes from a source they trust, that is a close friend, a relative, or an acquaintance they consider as being more knowledgeable in the music domain than they are and whose music tastes they value [79, 61, 80]. Additionally, along with the recommendations, the informant will often willingly provide information about the artist and the music, and convey her appreciation and enthusiasm along the way, thus turning the social interaction into a "musical experience in its own right" [69].

5.1.2 Prevalence of serendipitous encounters

Research on music information behavior also uncovered that people discover music primarily by chance. People rarely go to music streaming services with the specific objective of looking for the perfect gem. They stumble upon it during their daily routine (e.g., music heard on the radio, in a cafe, in a friend's car) or en route, while looking for something else. Indeed, studies with younger adults show that a majority of music discoveries (63.3% in [19]) are the result of passive information behavior or serendipitous encounters [19, 79]. Music has become a nearly constant soundtrack in many people's lives. Opportunities for encountering music are numerous. However, the strong engagement adolescents and young adults have with music might also explain the prevalence of these events. Research suggests that serendipitous encountering of information does not occur completely randomly, "but from circumstances brought about by unconscious motives which lead ultimately to the serendipitous event" [45]. In [79], it was found that many avid music listeners were also "super-encounterers," for they regularly engaged in activities likely to produce serendipity (e.g., wandering in a music festival) and were constantly monitoring their environment for interesting music.

5.1.3 The role of music recommender systems

Although all surveys converge to show the extreme popularity of music streaming services, the adoption of the discovery functionalities of these services is somewhat slower. In a survey conducted by Lee et al. [84], 64.6% of the participants reported using cloud music services to discover music. Liikkanen and Åman [90] get similar results with 65% of the respondents reporting using

YouTube[7] to discover new artists. Interestingly, Lee et al.'s [84] survey also reveal that a much smaller proportion of women (36.4%) use these functionalities compared with men (77.1%). Moreover, the interviews conducted in [61, 69] reveal what seems to be a prevalent pattern: people first get introduced to new music in their daily life (e.g., in the media, through friends), then they use music recommender systems to expand their exploration to other songs and/or artists.

5.2 User studies of music recommender systems

Considering the rapid and widespread adoption of music streaming services, several recent user studies on these systems have been published. In this section, we put our main focus on the studies on users' experience with and perception of music recommendations provided by these systems. This studies consist mainly in qualitative research conducted by social scientists and user experience studies on music streaming services conducted by researchers in information sciences.

5.2.1 Users' perception of music recommendations

Since the major players in the music streaming industry have comparable (and very extensive) catalogs, the branding of these services now lies in the services they offer, including personalized recommendations and curated playlists [97]. As mentioned above, many users make use of the discovery functions of music streaming services (cf. Sec. 5.1.3). Therefore, it seems logical that the respondents of a survey considered "Exposure to new things/serendipity" as the most important quality for a music service [84]. But how do users perceive these recommendations? Studies do not provide a clear answer to that question. Mainly because the perception seems to vary from user to user. Participants in a study comparing YouTube and Spotify had an overall positive perception of the music discovery functionalities of both services [90]. Results from a large-scale study by Avdeeff [3] highlight one interesting advantage of system recommendations from the users' perspective: for younger people who often perceive music genres as being confusing, the suggestions YouTube provides represent a useful alternative for discovery. Another perceived advantage lies in

7 https://www.youtube.com

the fact that system recommendations reduce the load for the users by assisting them in digital curation tasks. Indeed, among the heavy music service users interviewed by Hagen [50], many considered that the radio function music services offer was a useful way to expand a playlist, which can sometimes lead to further exploration. Other users, however, were not as enthusiastic about the radio function, which had played songs they disliked or recommendations they did not understand. But the main criticisms targeted the lack of novelty and true exploration, which prompted one of Johansson's participants to say that she felt "stuck in [her] own circles" [61]. Likewise, Kjus found that users "lose interest in the large databases after a period of initial fascination," a symptom he attributes to the inability of recommender systems to lead users to long tail items [69].

In the same line, several studies reveal a general lack of trust towards music recommender systems. Lee and Price [85] found that most users want more transparency: they want the systems to explain the recommendations. For some users, when Spotify partnered with artists and labels, distrust came with the realization that some services infuse their own commercial interests into their algorithms [69, 97]. Privacy issues also contribute to mistrust among some users [85, 61]. In [61], Johansson reports on the strong negative reaction of some participants regarding Spotify sharing its users' activities on Facebook[8] as a default setting following a partnership deal with the social networking site.

5.2.2 Users' engagement with music recommender systems

Many studies have focused, voluntarily or not, on avid music listeners. Recent large-scale studies and smaller qualitative studies with more diversified samples have uncovered a wider array of engagement practices with music streaming services. In [46], the researchers used mixed methods to study music services users. Their analysis resulted in seven personas with various levels of music expertise and involvement with music systems, including the "Guided Listener" who "wants to engage with a music streaming service minimally." Indeed, user studies demonstrate that many users wish to be able to listen to music continuously, without interacting much with the system. Johansson [61] reports on young users feeling "lazy" or "less active," who do not want to make the effort of creating playlists or browsing to find music to listen to. Hagen [50] also notices that users' engagements vary considerably.

8 https://www.facebook.com

This behavior may reflect the abundance (or overabundance) of choices in music services. Having access to such large music collections can seem exhilarating at first. But it can be intimidating to some. Users have expressed feeling "stressed" or "overwhelmed" by the number of items to chose from, or have used the term "drowning" to refer to how they felt [61, 69], a problem known in psychology as choice overload. Strategies for dealing with choice overload include withdrawing/escaping or surrendering the selection to someone else [59]. In this context, for lesser-engaged users, recommender systems become their gateway to music. Which is why it seems important to take the needs of these users into consideration in the design of music recommender systems.

6 User-centric Evaluation of Music Recommender Systems

Most evaluation approaches in current music recommender systems research and recommender systems research in general focus on quantitative measures, both geared towards retrieval accuracy (e.g., precision, recall, F-measure, or NDCG) and qualities beyond pure accuracy that cover further factors of perceived recommendation quality (e.g., spread, coverage, diversity, or novelty) [120]. While such beyond-accuracy measures aim to quantify and objectively assess parameters desired by users, they can not fully capture actual user satisfaction with a recommender system. For instance, while operational measures have been defined to quantify serendipity and diversity, they can only capture certain criteria. Serendipity and diversity as *perceived* by the user, however, can differ substantially from these measures, since they are highly subjective concepts [130]. Thus, despite the advantages of facilitating automation of evaluation and reproducibility of results, limiting recommender systems evaluation to such quantitative measures means to forgo essential factors related to user experience (UX).

Hence, in order to improve user awareness in music recommender systems, it is essential to incorporate evaluation strategies that consider factors of UX. To overcome the tendency to measure aspects such as user satisfaction or user engagement [51, 88] individually, evaluation frameworks for recommender systems aim at providing a more holistic view. One such framework is the ResQue model by Pu et al. [107], which proposes to evaluate the perceived qualities of recommender systems according to 15 constructs pertaining to four evaluation layers. The first layer deals with perceived system qualities and aims at evaluating aspects of recommendation quality, namely *accuracy*, *novelty*, and *diversity*, *interaction adequacy*, *interface adequacy*, *information sufficiency*, and *explicability*. The second layer deals with beliefs, evaluating *perceived ease of use*, *control*, *transparency*, and *perceived usefulness*. Layers three and four deal with attitudes (*overall satisfaction* and *confidence and trust*) and behavioral intentions (*use intentions* and *purchase intentions*), respectively. These constructs are evaluated using questionnaires consisting of up to 32 questions to be rated on a Likert scale.

Another evaluation framework is presented by Knijnenburg et al. [73]. In contrast to the model by Pu et al., which focuses on outcome experience of recommender systems, Knijnenburg et al. aim at providing insight into the relationships of six constructs that impact user experience: objective system aspects (OSA), subjective system aspects (SSA), subjective user experiences (EXP), objective interaction (INT), and personal and situational characteristics (PC and SC). This includes considerations of users' expertise of the domain or concerns for privacy, which are not reflected in the model by Pu et al. In particular, Knijnenburg et al. explicitly link INT, i.e., observable user behavior such as clicks, to OSA, i.e., unbiased factors such as response time or user interface parameters, through several subjective constructs, i.e., SSA (momentary, primary evaluative feelings evoked during interacting with the system [52]) and EXP (the user's attitude towards the system), and argue that EXP is caused by OSA (through SSA) and PC (e.g., user demographics, knowledge, or perceived control) and/or SC (e.g., interaction context or situation-specific trust or privacy concerns). To assess subjective factors, such as perceived qualities of recommendations and system effectiveness, variety, choice satisfaction, or intention to provide feedback, Knijnenburg et al. also propose a questionnaire.

Both frameworks can be applied to music recommender systems, however, neither is specifically designed for modeling the processes particular to these systems. Therefore, some of the assumptions do not hold for the requirements of today's music recommender systems. For instance, the ResQue model by Pu et al. has a strong focus towards commercial usage intentions, which, in music recommender systems, are by far more complex than the mere goal of using the system or intending to purchase (cf. Sec. 5). In particular with flat-rate streaming subscriptions being the dominating business model, selecting items does not have the same significance as a purchase. On a higher level, the commercial usage intention could be translated into or interpreted as continued subscription or monthly renewal of the service.

The model by Knijnenburg et al. can be easier adapted, as it was built around multimedia recommender systems, therefore offering more flexibility to incorporate the factors discussed in Sec. 2 and exemplarily detailed throughout this chapter. For instance, the INT aspects in the model can be adapted to refer to typical observable behavior of the user like favoring a song or adding it to a playlist, while PC aspects should reflect psychological factors, including affect and personality (as described in sections 4.1 and 3, respectively), social influence, musical training and experience, and physiological condition, to mention a few more examples. SC, on the other hand, should be adapted to the particularities of music listening context and situations. To get a better understanding of the various evaluation needs in the specific scenario of music recommendation and tailor new strategies, researchers in music recommender systems should therefore increasingly resort to the findings of user-centric evaluations, in-situ interviews, and ethnographic studies, cf. [5, 19, 32, 58, 64, 81, 82, 84, 86, 125].

7 Conclusions

In this chapter, we have identified different factors relevant for making music recommendations, namely intrinsic aspects, external factors, and goals, of both users and items. Focusing on the listener's aspects, we have highlighted exemplary technical approaches to model these factors and exploit knowledge about them for recommendation scenarios.

Regarding intrinsic aspects, i.e., the listener's background, we have reviewed work investigating the connection between a user's demographic information, personality traits, or cultural background on one hand, and musical preference, interest in diversity, and browsing preference on the other. While correlations of different user characteristics with music consumption behavior could be uncovered, for a lack of comprehensive data, it has not yet been explored whether these findings are consistent over different experiments. Hence it needs to be investigated if these or other interactions emerge when evaluating several indicators simultaneously in a larger study and with the same subjects. Exploiting more diverse data sources, such as from social media and other online activities, should give a more holistic picture of the user. In practice, the resulting challenge is to connect and match user activities across different platforms. While single sign-on options have facilitated the tracing of individuals across several services for the syndicated platforms, impartial academic research does not have comparable means. It remains, however, at least ethically questionable to which extent such profiling of users is justifiable and necessary on the premise of providing improved user experience, specifically music listening experiences.

Similar considerations can be taken on matters of modeling external factors, i.e., the listener's context. We have reviewed work dealing with estimating the current emotional state of listeners and connecting this to the emotions estimated to be conveyed by a piece, as well as work dealing with estimation of situational context of the listener based on sensor data. Both aspects can be highly dynamic and obtaining a ground truth and a general basis for repeatable experiments is challenging, if not elusive. Explicit assessments of mood and context requires introspection and reflection by the user and might be considered intrusive. On the other hand, taking a purely data-driven approach and exploring a wealth of accessible logging data to uncover latent patterns without proper means of validation might give rise to false assumptions and models and further raise issues concerning privacy.

To gain a better understanding of listeners' intents, we have reviewed work dealing with music information behavior to examine how people discover music in their daily life and how users perceive system recommendations in music streaming services. The key role friends and family play in discovering music suggest that people highly value the trustworthiness of the source of a music recommendation. Indeed, one of the most important criticisms leveled at music recommender systems is their lack of transparency and a breach of trust that comes from the role some services have taken in the promotion of specific artists. This impression contributes to the perception that music recommender systems are not the independent discovery tools they pretend to be. In terms of system design, this means that music recommender systems should pay close attention to building and maintaining the trust of their users, for instance by providing explanations to users as to why items are being recommended to them and by clearly identifying promotional recommendations. Furthermore, music information behavior studies revealed the prevalence of passive information behavior and serendipitous encounters in discovering music in daily life. In the same line, the review of studies on users' perception of and experience with music streaming services showed that users have various expectations regarding how much they want to engage with music recommender systems. Although certain users, especially the highly devoted music fans, are willing to spend time actively engaging with a system to keep a high level of control, others feel submerged by the millions of tracks available in a music service and prefer giving out a larger part of the control to the system in order to interact only minimally with it. This means that user-centered music recommender systems should let the users decide how much control they want to surrender to the system in order to cater to all users.

To conclude, we believe that a deepened understanding of the different factors of both user and music and their interplay is the key to improved music recommendation services and listening experiences. User awareness is therefore an essential aspect to adapt and balance systems between exploitation and exploration settings and not only identify the “right music at the right time,” but also help in discovering new artists and styles, deepening knowledge, refining tastes, broadening horizons—and generally be a catalyst to enable people to enjoy listening to music.

Acknowledgment: We thank Fabien Gouyon for the discussions leading to the updated factor model. Peter Knees acknowledges support by the Austrian Research Promotion Agency (FFG) under the BRIDGE 1 project *SmarterJam* (858514).

References

[1] X. Amatriain and J. Basilico. Recommender systems in industry: A netflix case study. In F. Ricci, L. Rokach, and B. Shapira, editors, *Recommender Systems Handbook*, pages 385–419. Springer US, Boston, MA, 2015.

[2] I. Andjelkovic, D. Parra, and J. O'Donovan. Moodplay: Interactive mood-based music discovery and recommendation. In *Proceedings of the 2016 Conference on User Modeling Adaptation and Personalization*, UMAP '16, pages 275–279, New York, NY, USA, 2016. ACM.

[3] M. Avdeeff. Technological engagement and musical eclecticism: An examination of contemporary listening practices. 2012.

[4] L. Baltrunas, M. Kaminskas, B. Ludwig, O. Moling, F. Ricci, K.-H. Lüke, and R. Schwaiger. InCarMusic: Context-Aware Music Recommendations in a Car. In *International Conference on Electronic Commerce and Web Technologies (EC-Web)*, Toulouse, France, 2011.

[5] L. Barrington, R. Oda, and G. Lanckriet. Smarter than genius? human evaluation of music recommender systems. In *Proceedings of the 10th International Society for Music Information Retrieval Conference*, ISMIR'09, Nara, Japan, 2009.

[6] M. Barthet, G. Fazekas, and M. Sandler. Multidisciplinary perspectives on music emotion recognition: Implications for content and context-based models. In *Proceedings of International Symposium on Computer Music Modelling and Retrieval*, pages 492–507, 2012.

[7] C. Bauer and A. Novotny. A consolidated view of context for intelligent systems. *Journal of Ambient Intelligence and Smart Environments*, 9(4):377–393, 2017.

[8] J. T. Biehl, P. D. Adamczyk, and B. P. Bailey. DJogger: A Mobile Dynamic Music Device. In *Proceedings of the 24th Annual ACM SIGCHI Conference on Human Factors in Computing Systems Extended Abstracts (CHI EA)*, Montréal, QC, Canada, 2006.

[9] D. Boer and R. Fischer. Towards a holistic model of functions of music listening across cultures: A culturally decentred qualitative approach. *Psychology of Music*, 40(2):179–200, 2010.

[10] G. Bonnin and D. Jannach. Automated generation of music playlists: Survey and experiments. *ACM Computing Surveys (CSUR)*, 47(2):26, 2015.

[11] D. O. Case. *Looking for information : a survey of research on information seeking needs, and behavior.* Emerald Group Publishing, Bingley, UK, 3rd edition, 2012.

[12] M. A. Casey, R. Veltkamp, M. Goto, M. Leman, C. Rhodes, and M. Slaney. Content-based music information retrieval: Current directions and future challenges. *Proceedings of the IEEE*, 96(4):668–696, 2008.

[13] T. Cebrián, M. Planagumà, P. Villegas, and X. Amatriain. Music Recommendations with Temporal Context Awareness. In *Proceedings of the 4th ACM Conference on Recommender Systems (RecSys)*, Barcelona, Spain, 2010.

[14] F. Celli, E. Bruni, and B. Lepri. Automatic personality and interaction style recognition from facebook profile pictures. In *Proceedings of the 22nd ACM international conference on Multimedia*, pages 1101–1104. ACM, 2014.

[15] Ò. Celma. *Music Recommendation and Discovery – The Long Tail, Long Fail, and Long Play in the Digital Music Space.* Springer, Berlin, Heidelberg, Germany, 2010.

[16] Ò. Celma and P. Herrera. A New Approach to Evaluating Novel Recommendations. In *Proceedings of the 2nd ACM Conference on Recommender Systems (RecSys)*, Lausanne, Switzerland, 2008.

[17] C.-W. Chen, P. Lamere, M. Schedl, and H. Zamani. RecSys Challenge 2018: Automatic Music Playlist Continuation. In *Proceedings of the 12th ACM Conference on Recommender Systems*, RecSys '18, pages 527–528, New York, NY, USA, 2018. ACM.

[18] S. Cunningham, S. Caulder, and V. Grout. Saturday Night or Fever? Context-Aware Music Playlists. In *Proceedings of the 3rd International Audio Mostly Conference: Sound in Motion*, Piteå, Sweden, 2008.

[19] S. J. Cunningham, D. Bainbridge, and D. Mckay. Finding new music: A diary study of everyday encounters with novel songs. In *Proceedings of the 8th International Conference on Music Information Retrieval*, pages 83–88, Vienna, Austria, September 23-27 2007.

[20] S. J. Cunningham and D. M. Nichols. 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, October 2009.

[21] R. de Oliveira and N. Oliver. TripleBeat: Enhancing Exercise Performance with Persuasion. In *Proceedings of the 10th International Conference on Human Computer Interaction with Mobile Devices and Services (Mobile CHI)*, Amsterdam, the Netherlands, 2008.

[22] A. K. Dey. Understanding and using context. *Personal and Ubiquitous Computing*, 5(1):4–7, 2001.

[23] L. Dey, M. U. Asad, N. Afroz, and R. P. D. Nath. Emotion extraction from real time chat messenger. In *2014 International Conference on Informatics, Electronics Vision (ICIEV)*, pages 1–5, May 2014.

[24] J. Donaldson. A hybrid social-acoustic recommendation system for popular music. In *Proceedings of the ACM Conference on Recommender Systems (RecSys)*, Minneapolis, MN, USA, 2007.

[25] S. Dornbush, J. English, T. Oates, Z. Segall, and A. Joshi. XPod: A Human Activity Aware Learning Mobile Music Player. In *20th International Joint Conference on Artificial Intelligence (IJCAI): Proceedings of the 2nd Workshop on Artificial Intelligence Techniques for Ambient Intelligence*, Hyderabad, India, 2007.

[26] G. Dror, N. Koenigstein, and Y. Koren. Yahoo! Music Recommendations: Modeling Music Ratings with Temporal Dynamics and Item Taxonomy. In *Proceedings of the 5th ACM Conference on Recommender Systems (RecSys)*, Chicago, IL, USA, 2011.

[27] S. Ebrahimi Kahou, V. Michalski, K. Konda, R. Memisevic, and C. Pal. Recurrent neural networks for emotion recognition in video. In *Proceedings of the 2015 ACM on International Conference on Multimodal Interaction*, ICMI '15, pages 467–474, New York, NY, USA, 2015. ACM.

[28] G. T. Elliott and B. Tomlinson. PersonalSoundtrack: Context-aware playlists that adapt to user pace. In *Proceedings of the 24th Annual ACM SIGCHI Conference on Human Factors in Computing Systems Extended Abstracts (CHI EA)*, Montréal, QC, Canada, 2006.

[29] M. Erdal, M. Kächele, and F. Schwenker. Emotion recognition in speech with deep learning architectures. In F. Schwenker, H. M. Abbas, N. El Gayar, and E. Trentin, editors, *Proceedings of Artificial Neural Networks in Pattern Recognition: 7th IAPR TC3 Workshop*, pages 298–311. Springer International Publishing, 2016.

[30] B. Ferwerda and M. Graus. Predicting musical sophistication from music listening behaviors: A preliminary study. *arXiv preprint arXiv:1808.07314*, 2018.

[31] B. Ferwerda, M. Graus, A. Vall, M. Tkalčič, and M. Schedl. The influence

of users' personality traits on satisfaction and attractiveness of diversified recommendation lists. In *4 th Workshop on Emotions and Personality in Personalized Systems (EMPIRE) 2016*, page 43, 2016.

[32] B. Ferwerda, M. P. Graus, A. Vall, M. Tkalcic, and M. Schedl. How item discovery enabled by diversity leads to increased recommendation list attractiveness. In *Proceedings of the Symposium on Applied Computing*, pages 1693–1696. ACM, 2017.

[33] B. Ferwerda and M. Schedl. Enhancing music recommender systems with personality information and emotional states: A proposal. In *UMAP Workshops*, 2014.

[34] B. Ferwerda and M. Schedl. Investigating the relationship between diversity in music consumption behavior and cultural dimensions: A cross-country analysis. In *UMAP (Extended Proceedings)*, 2016.

[35] B. Ferwerda and M. Schedl. Personality-based user modeling for music recommender systems. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 254–257. Springer, 2016.

[36] B. Ferwerda, M. Schedl, and M. Tkalcic. Predicting personality traits with instagram pictures. In *Proceedings of the 3rd Workshop on Emotions and Personality in Personalized Systems 2015*, pages 7–10. ACM, 2015.

[37] B. Ferwerda, M. Schedl, and M. Tkalcic. Using instagram picture features to predict users' personality. In *International Conference on Multimedia Modeling*, pages 850–861. Springer, 2016.

[38] B. Ferwerda, M. Schedl, and M. Tkalčič. 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, June–July 2015.

[39] B. Ferwerda and M. Tkalcic. Predicting users' personality from instagram pictures: Using visual and/or content features? In *The 26th Conference on User Modeling, Adaptation and Personalization, Singapore (2018)*, 2018.

[40] B. Ferwerda and M. Tkalcic. You are what you post: What the content of instagram pictures tells about users' personality. In *The 23rd International on Intelligent User Interfaces*, 2018.

[41] B. Ferwerda, M. Tkalčič, and M. Schedl. Personality traits and music genre preferences: How music taste varies over age groups. In *Proceedings of the 1st Workshop on Temporal Reasoning in Recommender Systems (RecTemp) at the 11th ACM Conference on Recommender Systems, Como, August 31, 2017.*, 2017.

[42] B. Ferwerda, M. Tkalčič, and M. Schedl. Personality traits and music

genres: What do people prefer to listen to? In *Proceedings of the 25th Conference on User Modeling, Adaptation and Personalization*, pages 285–288. ACM, 2017.

[43] B. Ferwerda, A. Vall, M. Tkalčič, and M. Schedl. Exploring music diversity needs across countries. In *Proceedings of the 2016 Conference on User Modeling Adaptation and Personalization*, pages 287–288. ACM, 2016.

[44] B. Ferwerda, E. Yang, M. Schedl, and M. Tkalčič. Personality traits predict music taxonomy preferences. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*, pages 2241–2246. ACM, 2015.

[45] A. Foster. Serendipity and information seeking: an empirical study. *Journal of Documentation*, 59(3):321–340, 2003.

[46] J. Fuller, L. Hubener, Y.-S. Kim, and J. H. Lee. Elucidating user behavior in music services through persona and gender. In *Proceedings of the 17th International Society for Music Information Retrieval Conference*, pages 626–632, 2016.

[47] M. Gillhofer and M. Schedl. 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, January 2015.

[48] J. Golbeck, C. Robles, M. Edmondson, and K. Turner. Predicting personality from twitter. In *Privacy, Security, Risk and Trust (PASSAT) and 2011 IEEE Third Inernational Conference on Social Computing (SocialCom), 2011 IEEE Third International Conference on*, pages 149–156. IEEE, 2011.

[49] J. Golbeck, C. Robles, and K. Turner. Predicting personality with social media. In *CHI'11 extended abstracts on human factors in computing systems*, pages 253–262. ACM, 2011.

[50] A. N. Hagen. The playlist experience: Personal playlists in music streaming services. *Popular Music and Society*, 38(5):625–645, 2015.

[51] J. Hart, A. G. Sutcliffe, and A. di Angeli. Evaluating user engagement theory. In *CHI Conference on Human Factors in Computing Systems*, May 2012. Paper presented in Workshop 'Theories behind UX Research and How They Are Used in Practice' 6 May 2012.

[52] M. Hassenzahl. The thing and i: Understanding the relationship between user and product. In M. A. Blythe, K. Overbeeke, A. F. Monk, and P. C. Wright, editors, *Funology: From Usability to Enjoyment*, pages 31–42. Springer Netherlands, Dordrecht, 2005.

[53] K. Hevner. Expression in Music: A Discussion of Experimental Studies and Theories. *Psychological Review*, 42, March 1935.

[54] G. Hofstede, G. J. Hofstede, and M. Minkov. *Cultures and Organizations: Software of the Mind*. McGraw-Hill, New York, NY, United States, 3rd edition, 2010.

[55] J. Hong, W.-S. Hwang, J.-H. Kim, and S.-W. Kim. Context-aware music recommendation in mobile smart devices. In *Proceedings of the 29th Annual ACM Symposium on Applied Computing*, SAC '14, pages 1463–1468, New York, NY, USA, 2014. ACM.

[56] Y. Hu, Y. Koren, and C. Volinsky. Collaborative Filtering for Implicit Feedback Datasets. In *Proceedings of the 8th IEEE International Conference on Data Mining (ICDM)*, Pisa, Italy, 2008.

[57] A. Huq, J. Bello, and R. Rowe. Automated Music Emotion Recognition: A Systematic Evaluation. *Journal of New Music Research*, 39(3):227–244, November 2010.

[58] G. B. Iman Kamehkhosh, Dietmar Jannach. 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, March 2018.

[59] S. S. Iyengar and M. R. Lepper. When choice is demotivating: Can one desire too much of a good thing? *Journal of Personality and Social Psychology*, 79(6):995–1006, 2000.

[60] D. Jannach, L. Lerche, and I. Kamehkhosh. Beyond "hitting the hits": Generating coherent music playlist continuations with the right tracks. In *Proceedings of the 9th ACM Conference on Recommender Systems*, RecSys '15, pages 187–194, New York, NY, USA, 2015. ACM.

[61] S. Johansson. *Music as part of connectivity culture*. Taylor & Francis Group, Milton, UNITED KINGDOM, 2018.

[62] O. P. John, E. M. Donahue, and R. L. Kentle. The big five inventory—versions 4a and 54, 1991.

[63] P. N. Juslin. What Does Music Express? Basic Emotions and Beyond. *Frontiers in Psychology*, 4(596), 2013.

[64] I. Kamehkhosh and D. Jannach. User perception of next-track music recommendations. In *Proceedings of the 25th Conference on User Modeling, Adaptation and Personalization*, UMAP '17, pages 113–121, New York, NY, USA, 2017. ACM.

[65] M. Kaminskas and F. Ricci. Contextual music information retrieval and recommendation: State of the art and challenges. *Computer Science Review*, 6(2):89–119, 2012.

[66] M. Kaminskas, F. Ricci, and M. Schedl. Location-aware Music Recommendation Using Auto-Tagging and Hybrid Matching. In *Proceedings of the 7th ACM Conference on Recommender Systems (RecSys)*, Hong Kong, China, October 2013.

[67] H. Kaya, F. Görpinar, and A. A. Salah. Video-based emotion recognition in the wild using deep transfer learning and score fusion. *Image and Vision Computing*, 65:66–75, 2017. Multimodal Sentiment Analysis and Mining in the Wild Image and Vision Computing.

[68] Y. E. Kim, E. M. Schmidt, R. Migneco, B. G. Morton, P. Richardson, J. Scott, J. Speck, and D. Turnbull. Music emotion recognition: A state of the art review. In *Proceedings of the International Society for Music Information Retrieval Conference*, 2010.

[69] Y. Kjus. Musical exploration via streaming services: The norwegian experience. *Popular Communication*, 14(3):127–136, 2016.

[70] P. Knees and M. Schedl. A survey of music similarity and recommendation from music context data. *ACM Transactions on Multimedia Computing, Communications, and Applications*, 10(1):2:1–2:21, Dec. 2013.

[71] P. Knees and M. Schedl. *Music Similarity and Retrieval - An Introduction to Audio- and Web-based Strategies*, volume 36 of *The Information Retrieval Series*. Springer, 2016.

[72] P. Knees and G. Widmer. Searching for music using natural language queries and relevance feedback. In N. Boujemaa, M. Detyniecki, and A. Nürnberger, editors, *Adaptive Multimedia Retrieval: Retrieval, User, and Semantics*, pages 109–121, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg.

[73] B. P. Knijnenburg, M. C. Willemsen, Z. Gantner, H. Soncu, and C. Newell. Explaining the user experience of recommender systems. *User Modeling and User-Adapted Interaction*, 22(4-5):441–504, 2012.

[74] V. J. Konecni. Social interaction and musical preference. *The psychology of music*, pages 497–516, 1982.

[75] Y. Koren. Factorization Meets the Neighborhood: A Multifaceted Collaborative Filtering Model. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, Las Vegas, NV, USA, 2008.

[76] T. Krismayer, M. Schedl, P. Knees, and R. Rabiser. Predicting user demographics from music listening information. *Multimedia Tools and Applications*, 2018.

[77] F.-F. Kuo, M.-F. Chiang, M.-K. Shan, and S.-Y. Lee. Emotion-based music recommendation by association discovery from film music. In *Pro-*

ceedings of the 13th annual ACM international conference on Multimedia, pages 507–510. ACM, 2005.

[78] J. R. Kwapisz, G. M. Weiss, and S. A. Moore. Activity recognition using cell phone accelerometers. *SIGKDD Explor. Newsl.*, 12(2):74–82, Mar. 2011.

[79] A. Laplante. *Everyday life music information-seeking behaviour of young adults: An exploratory study.* Doctoral dissertation, 2008.

[80] A. Laplante. Who influence the music tastes of adolescents? a study on interpersonal influence in social networks. In *Proceedings of the 2nd International ACM Workshop on Music information retrieval with user-centered and multimodal strategies (MIRUM)*, pages 37–42, 2012.

[81] A. Laplante. Improving music recommender systems: What we can learn from research on music tastes? In *15th International Society for Music Information Retrieval Conference*, Taipei, Taiwan, October 2014.

[82] A. Laplante and J. S. Downie. Everyday life music information-seeking behaviour of young adults. In *Proceedings of the 7th International Conference on Music Information Retrieval*, Victoria (BC), Canada, October 8-12 2006.

[83] A. Lay and B. Ferwerda. Predicting users' personality based on their 'liked'images on instagram. In *The 23rd International on Intelligent User Interfaces*, 2018.

[84] J. H. Lee, H. Cho, and Y.-S. Kim. Users' music information needs and behaviors: Design implications for music information retrieval systems. *Journal of the Association for Information Science and Technology*, 67(6):1301–1330, 2016.

[85] J. H. Lee and R. Price. Understanding users of commercial music services through personas: Design implications. In *Proceedings of the 16th International Society for Music Information Retrieval Conference (ISMIR 2015)*, pages 476–482, 2015.

[86] J. H. Lee, R. Wishkoski, L. Aase, P. Meas, and C. Hubbles. Understanding users of cloud music services: Selection factors, management and access behavior, and perceptions. *Journal of the Association for Information Science and Technology*, 68(5):1186–1200, 2017.

[87] J. S. Lee and J. C. Lee. Context Awareness by Case-Based Reasoning in a Music Recommendation System. In H. Ichikawa, W.-D. Cho, I. Satoh, and H. Youn, editors, *Ubiquitous Computing Systems*, volume 4836 of *LNCS*. Springer, 2007.

[88] J. Lehmann, M. Lalmas, E. Yom-Tov, and G. Dupret. Models of user engagement. In *Proceedings of the 20th International Conference on User*

Modeling, Adaptation, and Personalization, UMAP'12, pages 164–175, Berlin, Heidelberg, 2012. Springer-Verlag.

[89] H. Li and M. Trocan. Deep learning of smartphone sensor data for personal health assistance. *Microelectronics Journal*, 2018.

[90] L. A. Liikkanen and P. Åman. Shuffling services: Current trends in interacting with digital music. *Interacting with Computers*, 28(3):352–371, 2016.

[91] M. Liu, X. Hu, and M. Schedl. Artist preferences and cultural, socio-economic distances across countries: A big data perspective. In *The 18th International Society for Music Information Retrieval Conference, Suzhou, China, October 23-27, 2017*, 2017.

[92] M. Liu, X. Hu, and M. Schedl. The relation of culture, socio-economics, and friendship to music preferences: A large-scale, cross-country study. *PLOS ONE*, 13(12):e0208186, 2018.

[93] A. J. Lonsdale and A. C. North. Why do we listen to music? A uses and gratifications analysis. *British Journal of Psychology*, 102(1):108–134, February 2011.

[94] B. McFee and G. Lanckriet. Hypergraph Models of Playlist Dialects. In *Proceedings of the 13th International Society for Music Information Retrieval Conference (ISMIR)*, Porto, Portugal, 2012.

[95] A. P. Merriam and V. Merriam. Uses and functions. In *The Anthropology of Music*, chapter 11. Northwestern University Press, 1964.

[96] B. Moens, L. van Noorden, and M. Leman. D-Jogger: Syncing Music with Walking. In *Proceedings of the 7th Sound and Music Computing Conference (SMC)*, Barcelona, Spain, 2010.

[97] J. W. Morris and D. Powers. Control, curation and musical experience in streaming music services. *Creative Industries Journal*, 8(2):106–122, 2015.

[98] F. Noroozi, M. Marjanovic, A. Njegus, S. Escalera, and G. Anbarjafari. Audio-visual emotion recognition in video clips. *IEEE Transactions on Affective Computing*, pages 1–1, 2017.

[99] A. C. North and D. J. Hargreaves. Situational influences on reported musical preference. *Psychomusicology: A Journal of Research in Music Cognition*, 15(1-2):30, 1996.

[100] A. C. North, D. J. Hargreaves, and J. J. Hargreaves. Uses of music in everyday life. *Music Perception: An Interdisciplinary Journal*, 22(1):41–77, 2004.

[101] K. O'Hara and B. Brown, editors. *Consuming Music Together: Social and Collaborative Aspects of Music Consumption Technologies*, volume 35 of *Computer Supported Cooperative Work*. Springer Netherlands, 2006.

[102] K. Okada, B. F. Karlsson, L. Sardinha, and T. Noleto. Contextplayer: Learning contextual music preferences for situational recommendations. In *SIGGRAPH Asia 2013 Symposium on Mobile Graphics and Interactive Applications*, SA '13, pages 6:1–6:7, New York, NY, USA, 2013. ACM.

[103] S. Oramas, V. C. Ostuni, T. D. Noia, X. Serra, and E. D. Sciascio. Sound and music recommendation with knowledge graphs. *ACM Transactions on Intelligent Systems and Technology*, 8(2):21:1–21:21, Oct. 2016.

[104] E. Pampalk, T. Pohle, and G. Widmer. Dynamic Playlist Generation Based on Skipping Behavior. In *Proceedings of the 6th International Conference on Music Information Retrieval (ISMIR)*, London, UK, 2005.

[105] H.-S. Park, J.-O. Yoo, and S.-B. Cho. A context-aware music recommendation system using fuzzy Bayesian networks with utility theory. In *Proceedings of the 3rd International Conference on Fuzzy Systems and Knowledge Discovery (FSKD)*, Xi'an, China, 2006.

[106] T. Pettijohn, G. Williams, and T. Carter. Music for the seasons: Seasonal music preferences in college students. *Current Psychology*, pages 1–18, 2010.

[107] P. Pu, L. Chen, and R. Hu. A user-centric evaluation framework for recommender systems. In *Proceedings of the Fifth ACM Conference on Recommender Systems*, RecSys '11, pages 157–164, New York, NY, USA, 2011. ACM.

[108] D. Quercia, M. Kosinski, D. Stillwell, and J. Crowcroft. Our twitter profiles, our selves: Predicting personality with twitter. In *Privacy, Security, Risk and Trust (PASSAT) and 2011 IEEE Third Inernational Conference on Social Computing (SocialCom), 2011 IEEE Third International Conference on*, pages 180–185. IEEE, 2011.

[109] P. J. Rentfrow. The role of music in everyday life: Current directions in the social psychology of music. *Social and personality psychology compass*, 6(5):402–416, 2012.

[110] P. J. Rentfrow and S. D. Gosling. The do re mi's of everyday life: The structure and personality correlates of music preferences. *Journal of personality and social psychology*, 84(6):1236, 2003.

[111] J. A. Russell. A Circumplex Model of Affect. *Journal of Personality and Social Psychology*, 39(6):1161–1178, 1980.

[112] S. Sasaki, T. Hirai, H. Ohya, and S. Morishima. Affective music recommendation system based on the mood of input video. In X. He, S. Luo, D. Tao, C. Xu, J. Yang, and M. A. Hasan, editors, *MultiMedia Modeling*, pages 299–302, Cham, 2015. Springer International Publishing.

[113] T. Schäfer, P. Sedlmeier, C. Städtler, and D. Huron. The psychological functions of music listening. *Frontiers in Psychology*, 4(511):1–34, 2013.

[114] M. Schedl, G. Breitschopf, and B. Ionescu. 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, 2014.

[115] M. Schedl and B. Ferwerda. Large-scale analysis of group-specific music genre taste from collaborative tags. In *The 19th IEEE International Symposium on Multimedia (ISM2017), Taichung, December 11-13, 2017.*, 2017.

[116] M. Schedl, A. Flexer, and J. Urbano. The neglected user in music information retrieval research. *Journal of Intelligent Information Systems*, July 2013.

[117] M. Schedl, E. Gómez, E. Trent, M. Tkalčič, H. Eghbal-Zadeh, and A. Martorell. On the Interrelation between Listener Characteristics and the Perception of Emotions in Classical Orchestra Music. *IEEE Transactions on Affective Computing*, PP(99), 2017.

[118] M. Schedl, P. Knees, B. McFee, D. Bogdanov, and M. Kaminskas. Music Recommender Systems. In F. Ricci, L. Rokach, B. Shapira, and P. B. Kantor, editors, *Recommender Systems Handbook*, chapter 13, pages 453–492. Springer, 2nd edition, 2015.

[119] M. Schedl, F. Lemmerich, B. Ferwerda, M. Skowron, and P. Knees. Indicators of country similarity in terms of music taste, cultural, and socio-economic factors. In *The 19th IEEE International Symposium on Multimedia (ISM2017), Taichung, December 11-13, 2017.*, 2017.

[120] M. Schedl, H. Zamani, C.-W. Chen, Y. Deldjoo, and M. Elahi. Current challenges and visions in music recommender systems research. *International Journal of Multimedia Information Retrieval*, April 2018.

[121] K. Scherer. What are emotions? And how can they be measured? *Social Science Information*, 44(4):693–727, 2005.

[122] M. Skowron, F. Lemmerich, B. Ferwerda, and M. Schedl. Predicting genre preferences from cultural and socio-economic factors for music retrieval. In *Proceedings of the 39th European Conference on Information Retrieval (ECIR)*, 2017.

[123] M. Skowron, M. Tkalčič, B. Ferwerda, and M. Schedl. Fusing social media cues: personality prediction from twitter and instagram. In *Proceedings of the 25th international conference companion on world wide web*, pages 107–108. International World Wide Web Conferences Steering Committee, 2016.

[124] M. Slaney and W. White. Similarity Based on Rating Data. In *Proceedings of the 8th International Conference on Music Information Retrieval (ISMIR)*, Vienna, Austria, 2007.

[125] L. Spinelli, J. Lau, L. Pritchard, and J. H. Lee. Influences on the social practices surrounding commercial music services: A model for rich interactions. In *Proceedings of the 19th International Society for Music Information Retrieval Conference (ISMIR)*, Paris, France, 2018.

[126] H. Steck, R. van Zwol, and C. Johnson. Interactive recommender systems: Tutorial. In *Proceedings of the 9th ACM Conference on Recommender Systems*, RecSys '15, pages 359–360, New York, NY, USA, 2015. ACM.

[127] S. J. Tepper and E. Hargittai. Pathways to music exploration in a digital age. *Poetics*, 37(3):227–249, 2009.

[128] M. Tkalčič, B. Ferwerda, D. Hauger, and M. Schedl. Personality correlates for digital concert program notes. In *International Conference on User Modeling, Adaptation, and Personalization*, pages 364–369. Springer, 2015.

[129] A. van den Oord, S. Dieleman, and B. Schrauwen. Deep Content-based Music Recommendation. In C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Weinberger, editors, *Advances in Neural Information Processing Systems 26 (NIPS)*. Curran Associates, Inc., 2013.

[130] S. Vargas, L. Baltrunas, A. Karatzoglou, and P. Castells. Coverage, redundancy and size-awareness in genre diversity for recommender systems. In *Proceedings of the 8th ACM Conference on Recommender Systems*, RecSys '14, pages 209–216, New York, NY, USA, 2014. ACM.

[131] S. Volokhin and E. Agichtein. Understanding music listening intents during daily activities with implications for contextual music recommendation. In *Proceedings of the 2018 Conference on Human Information Interaction & Retrieval*, CHIIR '18, pages 313–316, New York, NY, USA, 2018. ACM.

[132] W. Wang, A. X. Liu, and M. Shahzad. Gait recognition using wifi signals. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp '16, pages 363–373, New York, NY, USA, 2016. ACM.

[133] X. Wang, D. Rosenblum, and Y. Wang. Context-aware Mobile Music Recommendation for Daily Activities. In *Proceedings of the 20th ACM International Conference on Multimedia*, pages 99–108, Nara, Japan, 2012. ACM.

[134] B. Whitman and S. Lawrence. Inferring Descriptions and Similarity for Music from Community Metadata. In *Proceedings of the 2002 International Computer Music Conference (ICMC)*, Göteborg, Sweden, 2002.

[135] K. Yadati, C. C. Liem, M. Larson, and A. Hanjalic. On the automatic identification of music for common activities. In *Proceedings of the 2017*

ACM on International Conference on Multimedia Retrieval, ICMR '17, pages 192–200, New York, NY, USA, 2017. ACM.

[136] H. Yakura, T. Nakano, and M. Goto. Focusmusicrecommender: A system for recommending music to listen to while working. In *23rd International Conference on Intelligent User Interfaces*, IUI '18, pages 7–17, New York, NY, USA, 2018. ACM.

[137] Y.-H. Yang and H. H. Chen. *Music Emotion Recognition*. CRC Press, 2011.

[138] Y.-H. Yang and H. H. Chen. Machine recognition of music emotion: A review. *ACM Transactions on Intelligent Systems and Technology*, 3(3):40:1–40:30, May 2012.

[139] M. Zentner, D. Grandjean, and K. R. Scherer. Emotions evoked by the sound of music: characterization, classification, and measurement. *Emotion*, 8(4):494, 2008.