# A Study on Contextual Influences on Automatic Playlist Continuation

Anna GATZIOURA [a,1], Miquel SÀNCHEZ-MARRÈ [a] and Alípio Mário JORGE [b,c]

[a] *Knowledge Engineering & Machine Learning Group (KEMLG)*
*Intelligent Data Science and Artificial Intelligence Centre (IDEAI)*
*Universitat Politècnica de Catalunya (UPC), Barcelona, Spain*
[b] *Laboratory of Artificial Intelligence and Decision Support (LIAAD)*
*INESC TEC, Porto, Portugal*
[c] *Faculdade de Ciências da Universidade do Porto (FCUP), Porto, Portugal*

**Abstract.** Recommender systems still mainly base their reasoning on pairwise interactions or information on individual entities, like item attributes or ratings, without properly evaluating the multiple dimensions of the recommendation problem. However, in many cases, like in music, items are rarely consumed in isolation, thus users rather need a set of items, selected to work well together, serving a specific purpose, while having some cognitive properties as a whole, related to their perception of quality and satisfaction, under given circumstances.

In this paper, we introduce the term of *playlist concept* in order to capture the implicit characteristics of joint music item selections, related to their context, scope and general perception by the users. Although playlist consumptions may be associated with contextual attributes, these may be of various types, differently influencing users' preferences, based on their character and emotional state, therefore differently reflected on their final selections. We highlight on the use of this term in HybA, our hybrid recommender system, to identify clusters of similar playlists able to capture inherit characteristics and semantic properties, not explicitly described in them. The experimental results presented, show that this conceptual clustering results in playlist continuations of improved quality, compared to using explicit contextual parameters, or the commonly used collaborative filtering technique.

**Keywords.** hybrid recommender systems, automatic playlist continuation, contextual dimensions, case-based reasoning, latent topic models

## 1. Introduction

*Music items* are rarely consumed in isolation but rather as sequences aiming to create a particular atmosphere [15]. More specific, *playlists* are *sets of music items* designed to be consumed as a sequence, with specific properties as a whole, similar to traditional radio broadcasts [3].

Therefore, in playlist recommendations, and similar domains, more than recommending isolated items, or presenting an ordered list of the most promising alternatives,

[1] Corresponding Author: Anna Gatzioura, Knowledge Engineering & Machine Learning Group (KEMLG), Intelligent Data Science and Artificial Intelligence Centre (IDEAI), Universitat Politècnica de Catalunya (UPC), C/Jordi Girona 1-3, 08034, Barcelona, Spain; E-mail: gatzioura@cs.upc.edu.

like the majority of recommender systems (RSs) do, the underlying structure of *joint item selections* should be evaluated, as item interactions within a set may heavily influence the result [19]. The presence of an item within a concrete concept should be captured, in order to recommend sets of items, addressing quality related attributes, like coherence and diversity, while being relevant to the playlist's purpose and creation moment [18].

Popular recommendation techniques, like collaborative filtering (CF) and content-based (CB), mainly use the user-item matrix to explode past interactions and predict the suitability of an item for a user, and not the suitability of an item for a particular *concept*. More specific, CF techniques are domain independent, focusing only on user ratings on items, while CB approaches, when applied to music base their analysis mainly on sound related attributes [4]. They neither explore the joint item selections nor the circumstances under which those were performed. Nevertheless, users may construct different playlists under different circumstances, not always made of their favourite or the most popular items. Thus, simply predicting if a song or an artist would be liked by a user, without evaluating the whole concept, is not enough and usually results in lower performance, especially in domains where users perform a lot of transactions. Furthermore, the influence even of very similar contextual situations, on users' music perception and preferences, has been found to heavily depend on the user's character and emotional state [12] that need to be captured appropriately.

### *1.1. Motivation*

The motivation of this work arises from the lack of RSs to efficiently recommend sets of items that would fit within a started concept, while addressing additional semantic characteristics, like coherence and diversity, similar to automatic playlist continuation (APC).

More specific, we evaluate the influence of explicit and implicit contextual dimensions on playlist continuation recommendations. *HybA*, a hybrid RS for APC, has been designed with aim to generate recommendations of *"sets" of music items*, not related to a specific user, but to a specific *concept*. This RS uses Case-Based Reasoning (CBR) with entire playlists modelled as cases, to identify their structures, combined with a Latent Dirichlet Allocation (LDA) topic model, to caprure the items' styles appearing within different concepts. In this paper, we highlight on the way that implicit contextual factors are used in this RS to identify additional playlist characteristics. We emphasize mainly on the retrieval step of the CBR cycle and complement our previous work [8]. The basic contributions of this work, can be summarized as:

- The term *playlist concept*, extending the explicit *context*, is introduced.
- We present an approach of capturing *implicit playlist context* through its concept.
- We evaluate the ability of *explicit* and *implicit contextual factors* to identify clusters of similar playlists, and their suitability for the designed hybrid RS.

The rest of the paper is structured as follows: in the next section an overview of the background on music consumption and the related contextual dimensions can be found. Following, HybA is presented, with emphasis on the term of *playlist concept* and its use in the system. Finally, the evaluation results, showing that the proposed conceptual filtering is able to better capture playlist similarities and provide improved recommendations, compared to the use of explicit context and user based CF techniques, can be found.

## 2. Background

### *2.1. Music Consumption*

When referring to *music items* those can be *songs*, *genres*, *artists*, *albums* and *radio stations*. Therefore, music recommendations can be addressed at different levels of abstraction [18]. As music items result from a complicated synthesis process, their analysis in terms of content characteristics, requires deeper domain knowledge. In addition, songs are rarely listened to in isolation [14]. Users rather create playlists/sessions, being sequences of songs, placing more importance on the songs and their relative order [19].

*Playlists* contain the notion of *item sequences and set characteristics*, while being highly affected by the intent and the context in which they were generated and consumed [11]. In addition, music is well known to evoke emotions while at the same time users' music needs are influenced by their actual emotional situation [7]. However, there is still a lack of solid methods combining users' cognitive perception of music with sound characteristics, therefore it becomes even more difficult to capture their perception of a playlist and specify the characteristics that a "good" playlist should have. This notion can be highly subjective, depending on parameters like the user's music knowledge, preferences, personality, emotional state, context and intent [19].

As currently more and more online sites either incorporate some music reproduction into their environment, or focus purely on presenting music sets, *automatic playlist generation (APG)* and recommendation has emerged as among the interesting issues in the music recommendation domain. APG refers to the automatic creation of sequences of music items based on some target characteristics. On the other hand, *automatic playlist continuation (APC)* which is a variation, or sub-case, of APG, consists in adding a set of music items to a playlist in a way that it would match its initial target characteristics. Therefore, APC consists in the selection of the most appropriate music items, and the construction of a sequence of improved quality according to characteristics inferred from a started playlist [3,19].

More precisely, given a started list, the aim is recommending sets of songs able to complete it, while providing a more exciting user experience. These recommendations are generated based on the characteristics of the started list, in terms of music styles and artist variety, independently of the user who made it. Therefore, more than simply predicting whether a music item would be highly rated, the underlying structure of *joint selections* should be evaluated, in order to recommend sets of songs satisfying at the same time relevance, and other *beyond accuracy* dimensions, like coherence and diversity [18].

### *2.2. Contextual Factors*

Incorporating into the recommendation problem contextual information related to the recommendation moment, increases the input data dimensions to three, namely users, items and context. However, when correctly captured, music context, may lead to a significant increase in recommendation accuracy [13].

When referring to music, *context*, initially defined as any *"information describing where you are, whom you are with, and what resources are nearby"* [20], is mapped to the user's situation when consuming the music items, in terms of time, mood, activity and other people's presence. It can be categorized according to several criteria, as:

- Fully observable, partially observable and unobservable
- Primary and secondary
- Environmental and user related

Where primary and environmental dimensions (location, time and weather) can be used to derive the secondary and user related ones (activity, mood, social and cultural) [6,11].

A context-aware recommendation process refers to *the estimation of user contextual preferences, and based on those, on the generation of the most relevant recommendations*. Depending on the part of the reasoning process that the contextual information is taken into account, it can be described as [1]:

- *Contextual pre-filtering*: selection of the relevant data, that will be further used for the recommendations' generation, based on a specific context.
- *Contextual post-filtering*: a recommendation set is generated from the entire data set and is then adjusted based on the contextual information of the active user.
- *Contextual modelling*: contextual parameters are inserted into the recommendation model or are used to transform the items into a different dimension.

Gillhofer and Schedl [10] analyze the relationships between various contextual dimensions to identify whether those permit the accurate prediction of user listening preferences. They evaluate their influence on the different categories of music items, namely songs, genres, artists and mood. Device, task, weather and time have been found as the most important attributes, being able to capture almost the same information as all contextual categories when combined. However, due to the data sparsity in music recommendation problems, still when recommending songs the performance is very low, while when it comes to genre or artist predictions, the additional contextual information indeed improves prediction accuracy. In our previous work [9], we have also found that evaluating the playlist creation time (hour of the day) leads to improved recommendation accuracy for artists and genres.

As explicitly defined context related to playlists may be of different types, or differently reflected on users' selections, Pichl et al. [16] propose the use of playlist names to implicitly extract contextual information. They create contextual clusters based on playlist names, that are then incorporated into the recommendation process. However, the efficiency of this approach heavily depends on the data quality, as highly subjective names like "my favourite", "best", etc., lead to sparse clusters that are not valuable for the recommendation process.

## 3. Recommendation Approach

The designed RS aims to generate recommendations of sets of items able to complete the active user's experience, based on the types of items appearing in similar concepts. HybA aims to address similar semantic concepts, rather than similar users and relies on the basic CBR idea, that *"Similar problems have similar solutions"* [17]. An overview of its reasoning process, that follows the general CBR Cycle, is shown in figure 1.

Given a new playlist, HybA first compares the already selected items with those in the playlists stored in the case base and retrieves the *k* most similar of them. Based on the items found in those, weighted by the similarity degree of the playlist(s) in which they appear with the new one, the continuation is constructed and recommended.

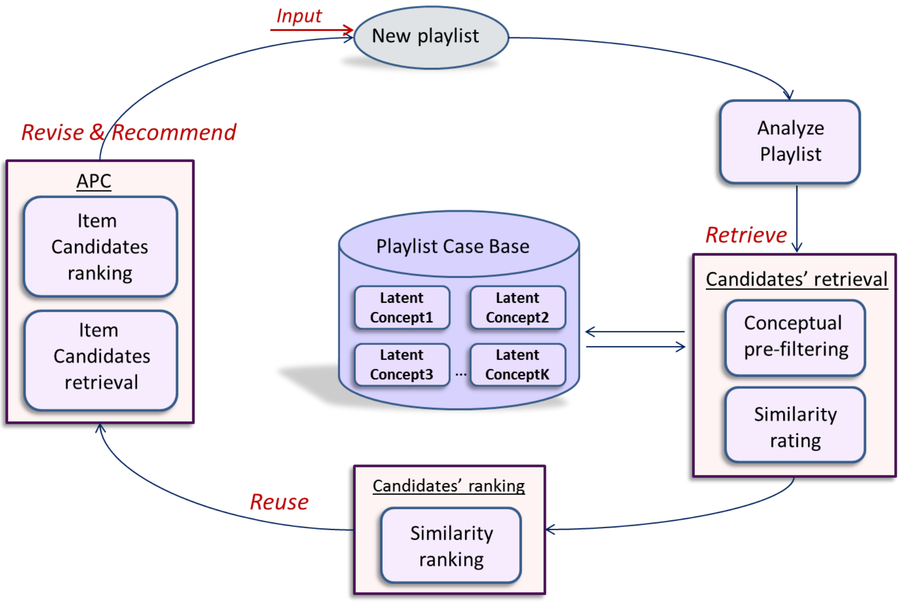


**Figure 1.** HybA Reasoning and Recommendation Model

The problem entities can be described as:

- a set of songs $I = \{i_1, \ldots, i_z\}$
- a set of previously reproduced playlists $L = \{l_1, \ldots, l_k\}$ where each can be written as the set of songs that it consists of, being $l_j = \{i_{j1}, \ldots, i_{jn}\}, i_{jt} \in I, t = 1, \ldots, n$
- a set of users $U = \{u_1, \ldots, u_v\}$ that have formed those playlists

Each of these entities may be associated with additional characteristics, like meta-data or editorial information, temporal or contextual data, preferences and demographic characteristics, respectively, used to compute their similarity degrees, like in [8].

### *3.1. Playlist General Concept*

A *playlist* is a collection of music items, reproduced as a meaningful sequence that should have some special characteristics as a whole [3]. We propose the categorization of playlists' characteristics into:

- *Internal (or content related)*: the exact music items that a playlist consists of. The styles of those items (which may be songs, artists or genres), define to a grand extend the characteristics of a playlist, like its coherence or diversity level.
- *External (or environmental related)*: the parameters related to the playlist but not directly described in it. For instance, the creation purpose, context, etc., and the influence of those factors on the user's emotional state and music perception.

Although not always explicitly reflected in terms of formulation, the context, and other external parameters related to a playlist, have been found to influence its style and the items finally placed in it [10]. However, as music perception is highly subjective, and may be heavily affected by circumstances under which is consumed, the personality and the emotional state of the user, it is hard to establish a direct connection between users and their prefered playlists under a given context. Nevertheless, it has been found [7] that depending on their character, users perform different music selections even under similar circumstances. For instance, more extrovert people tend to rely on “happy” music under sad or stressful situations, in order to “cheer” themselves up, while introverts tend to listen to “sad” or “depressive” music similar to their emotional state.

The term *playlist concept* is introduced as a wider term to describe the general characteristics of a playlist, being the result of the combination of both *internal* and *external parameters*, beyond explicilty defined context. Furthermore, as playlists are formulated to serve a specific semantic concept, they are treated as distributions over music styles. Therefore each playlist can be written as $l_j = \{s_{j1}, \ldots, s_{jn}\}$ where each $s_{jt}$ is the music style of the song $i_{jt}$. In order to capture the general concept of playlists, without focusing on their specific content, or when no additional information, like their names or explicitly defined context, is available, the playlists’ latent topic distribution, based on the *music styles* in them, is proposed. Thus, before analysing the playlist-song distributions, first the playlist-music styles distributions are identified, with aim to capture the tendencies and patterns present in the playlists, rather than finding the exact songs.

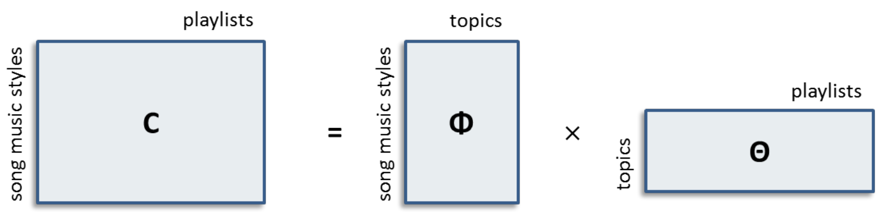


**Figure 2.** Playlists’ latent distributions over music styles

Having playlists described as distributions over music styles, a *Latent Dirichlet Allocation (LDA)* topic model [2] is built to extract the underlying latent topics, as in figure 2. Playlists are then characterized by their *dominant topic(s)* as a representation of their general concept, closely related to their central idea and cognitive perception. We have found this attribute to be capable of capturing the general tendencies in playlists, and their similarities, and to improve recommendation accuracy.

### *3.2. Candidates Retrieval*

An important process of a CBR system that highly affects its performance, is its ability to identify the characteristics of a new case, and retrieve, within an acceptable time, the past cases that seem as most appropriate for its solution. Therefore, in order to identify globally acceptable candidates for a new playlist, while reducing the computational time, first a proper clustering and pre-filtering of the case base is performed.

Especially when applied to large scale datasets the selection of the parameters that mostly characterize the playlists and enable their proper clustering, is of high importance. To this direction, a contextual pre-filtering based on the playlist generation moment (hour of the day) was initially tested, and was found to provide improved results [9]. However, as many times contextual information is not explicitly described, or even if so it may affect differently users’ selections, a more general clustering would be more appropriate.

Therefore, we have extended the contextual to a *conceptual pre-filtering*, evaluating first playlists' conceptual similarities.

More specific, based on the music styles' distribution in a new playlist being $l_N = \{s_{N1}, \ldots, s_{Nn}\}$, its topic distribution, specifying its concept, is identified. Then the set of instances with the same *dominant topic*, $L_C \subseteq L$, that address the same concept, are initially retrieved and then used for further computations, finally leading to the retrieval of the $k$ most similar past playlists.

Given a retrieved playlist $l_R = \{s_{R1}, \ldots, s_{Rm}\}, l_R \in L_C$, its similarity $Sim(l_N, l_R)$ with $l_N$, referred to as *global similarity*, is computed as a function of the *local similarity* levels of the items in them. HybA finally retrieves the $k$ most similar playlist(s), those whose global similarity fulfils equation (1), and constructs the playlist continuation using the songs that initially appeared in them, as in [8].

$$l' \in L_C : \forall l_R \in L_C, l' = argmax\{Sim(l_N, l_R)\} \tag{1}$$

## 4. Evaluation

In this section we present graphically the comparative results of HybA, using the proposed conceptual clustering and pre-filtering, with the use of three different contextual clusterings and a CF approach, on three real music datasets.

More specific, for a new playlist, *cntxT*, uses the cluster $L_H$ of playlists generated in close hours, *cntxM*, the cluster $L_M$ of playlists reproduced in close months and *cntxC* evaluates both the day hour and the month when the playlist was made. Furthermore, a user-based CF[2] using Euclidean similarity and the 10 nearest neighbours for each user has been tested.

The databases used for evaluation come from the Portuguese music social network Palco Principal[3] and contain information on users adding tracks to their listening sessions or their playlists, on a given moment (event). More details can be found in table 1.

| Dataset | Palco Listen1 | Palco Listen2 | Palco Playlists |
|---|---|---|---|
| Events | 1171849 | 295044 | 111942 |
| Songs | 29786 | 22986 | 26117 |
| Users | 21815 | 5543 | 10392 |
| Playlists | 86174 | 22108 | 22132 |

**Table 1.** Palco Principal datasets

To evaluate both recommendation accuracy and quality we have used the recommendations' average precision and coherence. Coherent playlists are generally considered as of "better quality", while the diversity and novelty degree that a user enjoys may be subject to his/her actual music preferences [5,19].

- *Precision* is the ratio of correct recommendations: the missing items that were successfully identified, over the total number of recommendations, calculated as:

$$Precision = \#RelevantRecommendedItems / \#RecommendedItems$$

[2] https://mahout.apache.org/users/recommender/recommender-documentation.html
[3] http://palcoprincipal.com/

- *Coherence* evaluates the homogenity of a playlist $R$, and can be calculated as the average similarity (*sim*) among the pairs of consequent items in the list, being:

$$Coherence = \frac{1}{|R|-1}\sum_{i \in R} sim(i, i+1)$$

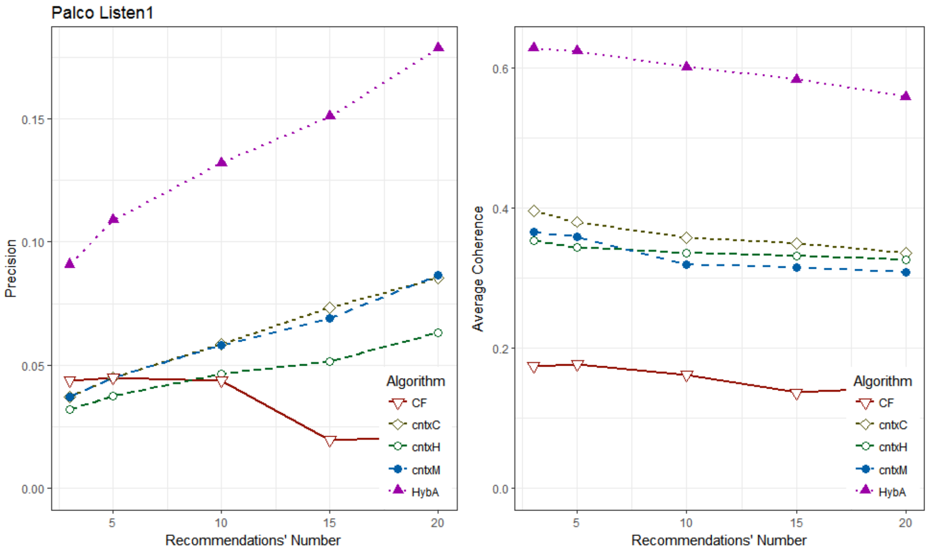


**Figure 3.** Recommendations' precision and coherence on Palco Listen1 dataset

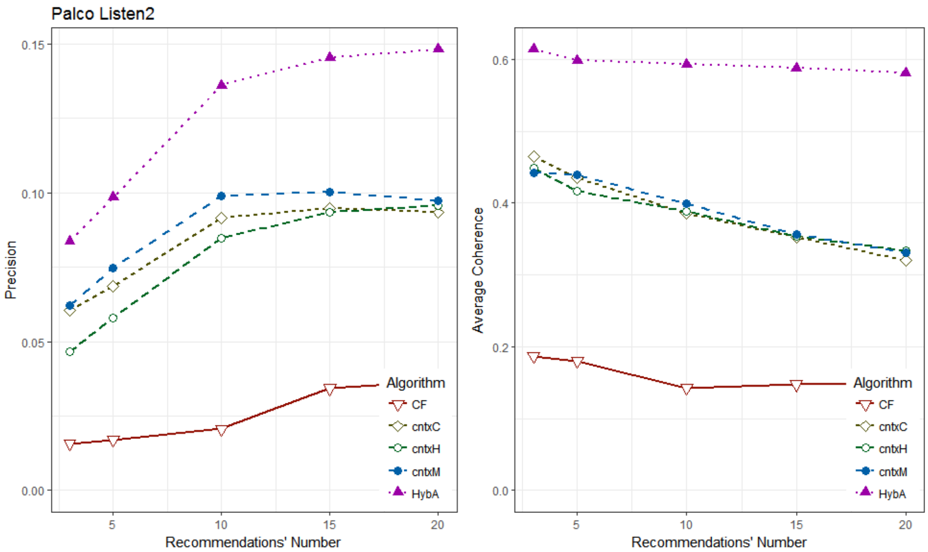


**Figure 4.** Recommendations' precision and coherence on Palco Listen2 dataset

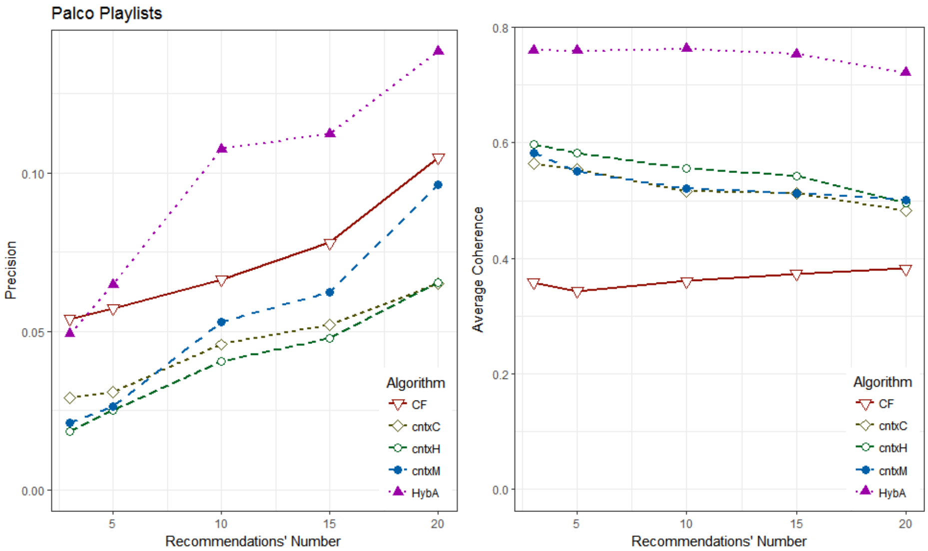


**Figure 5.** Recommendations' precision and coherence on Palco Playlists dataset

As it can be seen from figures 3–5, CF shows a low performance due to its slightly different scope, as it does not take into account any additional information related to joint item selections. In addition, in the majority of cases including information on items' joint selections and basic time context, improves the results over the CF approach, while going beyond explicit context leads to higher improvements. More specific, the proposed conceptual clustering and prefiltering provides an important improvement of both recommendation precision and coherence.

## 5. Conclusions

Although recently context-aware methods have gained space among recommendation techniques, still they focus mainly on *explicitly defined contextual parameters*. However, context may be of different types, differently reflected in users' behaviour, especially in domains related to emotional responses, like in music consumption.

In this paper we have presented the term of *playlist concept*, for *implicitly capturing the context* of playlist generations. This term is used to characterize playlists, constructed under given circumstances, based on the music style combinations in them. Furthermore, it has been found to better identify similar playlists, compared to explicit contextual dimensions, and to provide improved APC recommendations compared with the widely used CF approach. As part of our future work, the exploitation of the latent concepts with aim to perform a proper "translation", enabling the explanation of the resulting playlist clusters, will take place.

## References


[1] Gediminas Adomavicius and Alexander Tuzhilin. Context-aware recommender systems. In *Recommender systems handbook*, pages 217–253. Springer, 2011.

[2] David Blei, Andrew Ng, and Michael Jordan. Latent dirichlet allocation. 3:993–1022, 3 2003.

[3] Geoffray Bonnin and Dietmar Jannach. Automated generation of music playlists: Survey and experiments. *ACM Comput. Surv.*, 47(2):26:1–26:35, 2014.

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

[5] Pablo Castells, Neil J. Hurley, and Saul Vargas. Novelty and diversity in recommender systems. In *Recommender Systems Handbook*, pages 881–918. 2015.

[6] Anind K. Dey, Daniel Salber, Gregory D. Abowd, and Masayasu Futakawa. Providing architectural support for context-aware applications. Technical report, 2000.

[7] Bruce Ferwerda, Marko Tkalcic, and Markus 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, UMAP 2017, Bratislava, Slovakia, July 09 - 12, 2017*, pages 285–288, 2017.

[8] Anna Gatzioura, Alípio Mário Jorge, and Miquel Sànchez-Marrè. A Hybrid Recommendation Algorithm for Improving Automatic Playlist Continuation. *Submitted to: Cognitive Computation, S.I. : Bridging Cognitive Models and Recommender Systems*, 2018.

[9] Anna Gatzioura and Miquel Sànchez-Marrè. Using contextual information in music playlist recommendations. In *Recent Advances in Artificial Intelligence Research and Development - Proceedings of the 20th International Conference of the Catalan Association for Artificial Intelligence, Deltebre, Terres de l'Ebre, Spain, October 25-27, 2017*, pages 239–244, 2017.

[10] Michael Gillhofer and Markus Schedl. Iron maiden while jogging, debussy for dinner? In Xiangjian He, Suhuai Luo, Dacheng Tao, Changsheng Xu, Jie Yang, and Muhammad Abul Hasan, editors, *MultiMedia Modeling*, pages 380–391, Cham, 2015. Springer International Publishing.

[11] Marius Kaminskas and Francesco Ricci. Contextual music information retrieval and recommendation: State of the art and challenges. *Computer Science Review*, 6(2-3):89–119, may 2012.

[12] Youngmoo E. Kim, Erik M. Schmidt, Raymond Migneco, On G. Morton, Patrick Richardson, Jeffrey Scott, Jacquelin A. Speck, and Douglas Turnbull. Emotion recognition: a state of the art review. In *11th International Society for Music Information and Retrieval Conference*, 2010.

[13] Peter Knees and Markus Schedl. A survey of music similarity and recommendation from music context data. *ACM Transactions on Multimedia Computing, Communications, and Applications*, 10(1):1–21, dec 2013.

[14] Brian McFee and Gert R. G. Lanckriet. The natural language of playlists. In *Proceedings of the 12th International Society for Music Information Retrieval Conference, ISMIR 2011, Miami, Florida, USA, October 24-28, 2011*, pages 537–542, 2011.

[15] François Pachet and Pierre Roy. Automatic generation of music programs. In *Principles and Practice of Constraint Programming - CP'99, 5th International Conference, Alexandria, Virginia, USA, October 11-14, 1999, Proceedings*, pages 331–345, 1999.

[16] Martin Pichl, Eva Zangerle, and Günther Specht. Towards a context-aware music recommendation approach: What is hidden in the playlist name? In *IEEE International Conference on Data Mining Workshop, ICDMW 2015, Atlantic City, NJ, USA, November 14-17, 2015*, pages 1360–1365, 2015.

[17] Michael Richter and Rosina Weber. *Case-based reasoning: a textbook*. 09 2013.

[18] Markus Schedl, Peter Knees, Brian McFee, Dmitry Bogdanov, and Marius Kaminskas. Music recommender systems. In *Recommender Systems Handbook*, pages 453–492. 2015.

[19] Markus Schedl, Hamed Zamani, Ching-Wei Chen, Yashar Deldjoo, and Mehdi Elahi. Current challenges and visions in music recommender systems research. *CoRR*, abs/1710.03208, 2017.

[20] B. Schilit, N. Adams, and R. Want. Context-aware computing applications. In *1994 First Workshop on Mobile Computing Systems and Applications*, pages 85–90, Dec 1994.