PAPA: Physiology and Purpose-Aware Automatic Playlist Generation

Nuria Oliver
Microsoft Research
One Microsoft Way
Redmond, WA 98052, USA
nuria@microsoft.com

Lucas Kreger-Stickles
Microsoft Research
One Microsoft Way
Redmond, WA 98052
lucasks@cs.washington.edu

Abstract

In this paper we present PAPA, a novel approach for automatically generating playlists. The proposed framework utilizes the user's physiological response to music, together with traditional song meta-data to generate a playlist the user will not only enjoy, but which will assist him or her in achieving various user-defined goals ("purpose"). In addition to outlining the generic framework, we present an exemplary application named MPTrain that (1) creates a playlist in real-time to assist users in achieving specific exercise goals; and (2) incorporates the user's physiological response to the music to determine the next song to play.

Keywords: Automatic Playlist Generation, Physiological Monitoring, User Modeling

1. Introduction and Related Work

In recent years, there has been increasing interest in the automatic generation of playlists, partly due to the broad adoption of digital music and personal digital music players. Common approaches to creating playlists range from randomly shuffling a collection (*e.g.* iPod shuffle) to manually creating the playlist by selecting the order of the songs. The first approach does not give any control to the user in terms of which songs will be selected. The second approach, on the other hand, gives the user complete control over the selection, but is typically time consuming, tedious, static and potentially hard to implement when the music collection is large (*e.g.* tens of thousands of songs). Moreover, as it entirely relies on the user's knowledge of his or her collection, it does not allow for the discovery of "forgotten" songs in the collection.

Consequently, there have been multiple efforts in the research community to assist users creating playlists. Most of the previous work to date has focused on efficient algorithms to automatically create a playlist that satisfies some given constraints [2, 3, 1], and on collaborative filtering [5, 6] algorithms. Typically, the user provides the system with a *seed* song for the playlist to be generated and the algorithms find similar sounding songs.

There are two projects that are particularly relevant to the approach described in this paper: (1) The PATS system [8], where the authors introduce the concept of *context-of-use* as an important input variable when automatically creating a playlist. The *context-of-use* is defined in [8] as *the real-world environment in which the music is heard, being it a party, a romantic evening or traveling in the car or train*; and (2) the HPDJ system [4], which uses sensors to determine physical and physiological responses of a crowd to the music and uses this feedback to automatically sequence and mix the music in nightclubs. The PATS system is similar to the approach proposed in this paper in that it takes into account the context-of-use as an input variable when creating a playlist. The HPDJ system shares with our approach the addition of real-time biofeedback to inform the song selection algorithms.

In this paper we combine information about the user's physiological response and purpose in order to automatically generate playlists. We refer to this framework as PAPA, because our approach takes into account both the purpose and the physiological response of the user to the music. Whereas other systems employ context in song selection so as to better find an asthetically appropriate song, we believe we are the first to propose utilizing the music's potential impact not only on the user's enjoyment, but also on helping him/her achieve a particular goal.

We illustrate our approach by means of an application named MPTrain that (1) selects the music to assist users in achieving specific exercising goals; (2) incorporates the user's physiological response to the music to determine the next song to play.

The paper is organized as follows: In Section 2 we present PAPA, the framework for automatically generating playlists. Next, Section 3 is devoted to describing the MPTrain system, an exemplary application of the proposed framework. Finally, Section 4 contains some conclusions and future lines of work.

2. Automatic Generation of Playlists from Purpose and Physiological Response

Figure 1 illustrates the basic components of PAPA.

Typically, the user (1) is listening to music from his personal digital music library (DML) (2) by means of a portable digital music player. Let the DML consist of a collection of

N songs, $\{S_1, S_2, ..., S_N\}$. In the Figure, the portable device is a mobile phone with MP3 playback capabilities (3). The system has access to: (a) the user's profile (10) containing personal information, P^u, such as the user's age, weight, gender, and musical preferences, together with (b) historic data in the form of logs of previous interactions with the system, H^u.

The user also wears a set of K physiological and environmental sensors (2), such as sensors for determining the user's heart-rate, galvanic skin response, respiration rate, movement, position, etc. Let $F^u = \{f_1^u, f_2^u, ..., f_K^u\}$ denote the user's physiology and environmental data, where f_i^u is the feature vector for sensor i, with $i = 1, ..., K$. These sensors continuously send data wirelessly to the mobile phone.

The system takes into account the context in which the user is listening to the music (*e.g* at work, in the car, at home, with friends, etc) together with the purpose of listening to the music (*e.g.* relaxing, concentrating, exercising, driving, etc.) (4). The context and purpose can be expressed as a function G of the user's: (a) real-time physiological and contextual data, F^u; (b) profile, P^u and (c) historic data, H^u. The user's DML is augmented with relevant metadata M^s (5), such as the songs' tempo, average energy, duration, genre, etc. Let $M^{si} = \{m_1^{si}, m_2^{si}, ..., m_L^{si}\}$ denote the set of L features extracted from song i, with $i = 1...N$.

The system utilizes the user's biofeedback and explicit feedback (7) to learn a model of the set of features in the music, M^s – *e.g* tempo, average energy, etc – and the user's response to it, F^u – *e.g.* increased heart-rate, decreased respiration rate, etc. The system will select the next song to play based on the deviation between the desired goal G and the user's current state $\{F^u, P^u, H^u\}$.

For example, let us assume that the user's goal is to relax. This goal corresponds to certain values in the user's physiology, *e.g.* low heart-rate, galvanic skin response, respiration rate and movement. The user starts listening to the first musical piece as selected by the system. As time progresses, the system constantly monitors the user's response to the music and uses this information to select the next "optimal" song to play. If, for example, the user's current physiology does not correspond to the system's model of a "relaxed" state, the system will select a more "relaxing" song to play next. The model of the user's state is learned from data.

Note that there are two different mappings that need to be estimated from data: (1) The mapping Map_{MU} between musical features M^s and the user's state (F^u, P^u, H^u); and (2) the mapping Map_{GU} between the goal G and the user's state (F^u, P^u, H^u). Those two mappings, together with the user's profile (P^u) and historic data (H^u) are part of the user model (6).

Finally, there are two important aspects in evaluating this paradigm: (1) How well the music helps users in achieving their desired goal, and (2) how enjoyable the music selection is to the user, given the context.

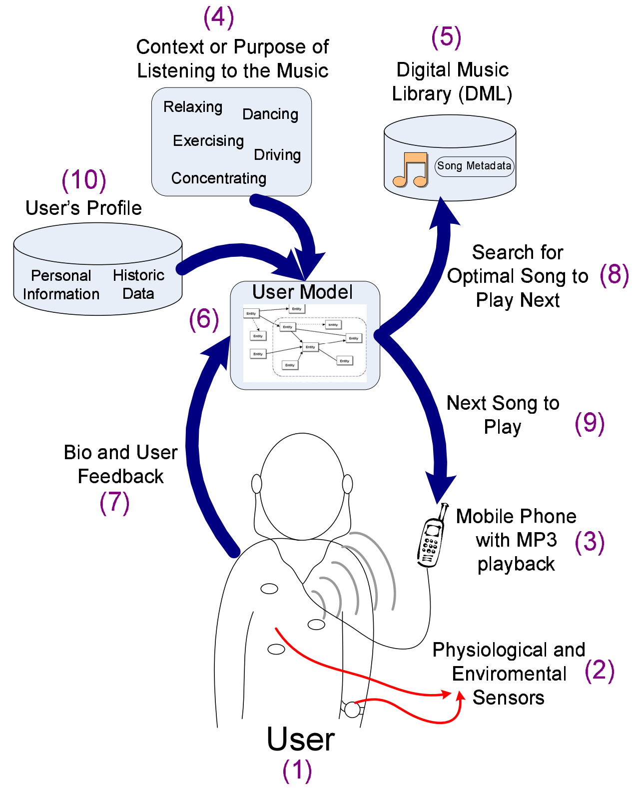

Figure 1. PAPA: Physiology and Purpose-Aware Playlist Generation.

We shall illustrate the proposed approach with a prototype named MPTrain that creates a playlist to assist users in achieving specific exercise goals.

3. MPTrain: Automatic Playlist Generation for Optimal Exercise Performance

MPTrain is a mobile phone based system that takes advantage of the influence of music in exercise performance enabling users to more easily achieve their exercise goals.

MPTrain is designed as a mobile and personal system (hardware and software) that users wear while exercising (walking, jogging or running). MPTrain's hardware includes a continuous heart-rate and acceleration monitor wirelessly connected to a mobile phone carried by the user. MPTrain's software allows the user to enter a desired workout in terms of desired heart-rate stress over time. It then assists the user in achieving the desired exercising goals by: (1) constantly monitoring his/her physiology (heart-rate in number of beats per minute) and movement (speed in number of steps per minute); and (2) selecting and playing music (MP3s) with specific features that will guide him/her towards achieving the desired exercising goals.

MPTrain uses algorithms that learn the mapping between musical features (*e.g.* beat), the user's current exercise level (*e.g.* running speed or gait) and the user's current physiological response (*e.g.* heart-rate). The goal is to automatically choose and play the "right" music to encourage the user to speed up, slow down or maintain their pace while keeping him/her on track with the desired workout.

Figure 2 illustrates MPTrain's data flow. The user is listening to digital music on his/her mobile phone while jogging. At the same time, the user's hear-rate and speed are monitored and stored on the mobile phone. A few seconds (typically 10 s) before the conclusion of the song currently being played, MPTrain compares the user's current heart-rate with the desired one according to the current pre-selected workout. The user's model is composed of two elements: (1) The *next action module*, which determines if the user needs to speed up, slow down or keep their pace of jogging, based on whether his/her heart-rate needs to increase, decrease or stay the same. With this information, (2) the *music finding module* identifies the next song to be played from the music database. The current implementation of the music finding algorithm determines the next song by identifying a song (1) that hasn't been played yet; and (2) has a tempo (in beats per minute) that is similar to the user's current gait, but increasing or decreasing an amount inversely related to the deviation between the actual heart-rate and the desired heart-rate from the preset workout.

At any instant of time, the user can check how well (s)he is doing with respect to the desired exercise level, modify the exercising goals or change the music track from the one automatically selected by MPTrain. We are currently working on including additional user adaptation to MPTrain by incorporating information about that user's past sessions and by constantly monitoring the user's actions and learning from them. As the user interacts with the MPTrain system, we would like for its music selection algorithm to become progressively better suited for that particular user.

We refer the reader to [7] for a detailed description of the MPTrain system.

3.1. Automatic Playlist Generation

MPTrain acts as a personal trainer that uses music to encourage the user to accelerate, decelerate or keep the running speed. The key element is that music improves gait regularity due to the use of the beat, which helps individuals to anticipate the desired rate of movement [9]. The rhythmic structure of the music and the rhythmic actions performed by the body are believed to combine and sychronize. We shall describe next MPTrain's algorithm for selecting the music to play.

In its current implementation, MPTrain does not take any action until one of these three conditions are true:

1. There are a few seconds left (*e.g.* 10) before the end of the current song: In this case, MPTrain determines whether the user needs to increase, decrease or keep their running pace, by comparing the user's current heart-rate with the desired heart-rate from the *desired workout* for that day. Once it has determined the action to take, it searches the user's

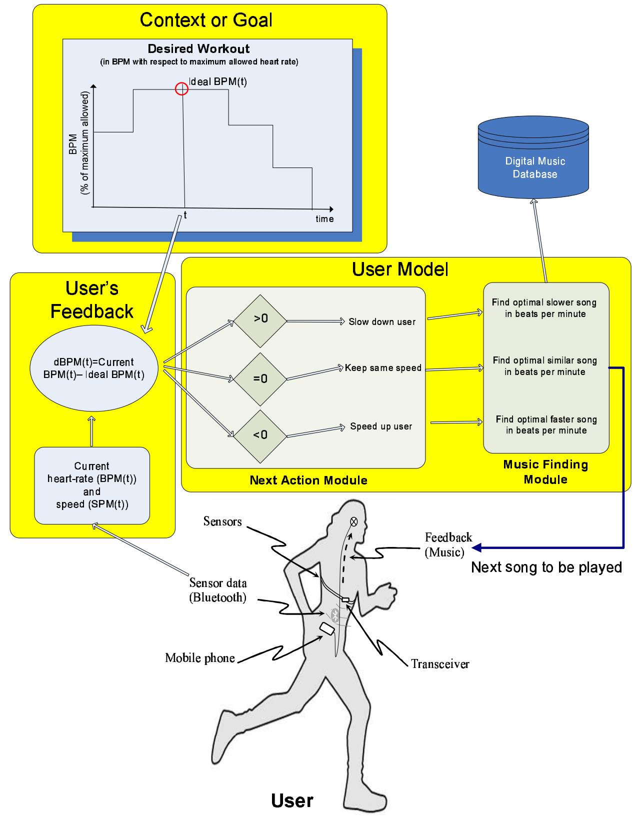

Figure 2. MPTrain's dataflow.

digital music library (DML) for the optimal song to play. Note that the DML contains not only the user's personal collection of MP3 songs, but also additional information about each song, such as its beat [1] and average energy. Depending on the situation, MPTrain will look for a song whose beat is similar, higher or lower than that of the song currently being played, according to the difference between the actual and desired heart-rates. In the event that the workout target is about to change, MPTrain selects a song appropriate to the *next* workout target.

For example, if the user's current heart-rate is at 55% of the maximum heart-rate reserve [2], but the desired heart-rate at that point is at 65%, MPTrain will find a song that has faster beat than the one currently being played. The increase in beat in the song is proportional to the percentage of error between the user's actual and desired heart-rate.

2. There is a discontinuity in the desired workout pattern, such as moving from a warming-up (about 60% of maximum heart-rate reserve) to a weight management section (about 70% of maximum heart-rate reserve) in the desired workout. In this case, MPTrain interrupts the song that is currently playing, unless the song has been playing for a

[1] Determined automatically or manually.

[2] Heart-rate reserve is defined as the user's maximum heart-rate minus the user's resting heart-rate.

very short time (*e.g.* less than 20s) [3]. MPTrain then selects the next song to play as described in case (1) above.

3. The user explicitely requests a change of song. In this case, MPTrain selects a different song from the DML whose features still satisfy the constraints given the situation.

MPTrain's current implementation uses an empirically learned function to map the physiological response to the music's beat. This model is used to make a statistically accurate track selection. Further versions will also incorporate information about the user's past performance and specific response to each song.

3.2. Evaluation

We are currently carrying out a comprehensive user study with 20 runners. The study consists of three to four running sessions. In each of the sessions the runners know what is the desired exercise pattern for the session and their heart-rate and speed are monitored. The first running session is without music (*mute* condition); the second one with random music from the digital music library that is stored in the phone (*random* condition); and the third one with the music as it is selected by the MPTrain system (*MPTrain* condition) and the fourth and optional final session is with a *Metronome* as selected by the system.

All runners fill out a pre-run questionnaire and a post-run questionnaire after each of the sessions. With the study we plan to evaluate, for each of the sessions: (1) how well each runner achieves the pre-defined exercise goal; (2) their perception of the workout; (3) their level of enjoyment of the music selected by MPTrain, especially when compared to the rest of the conditions.

Before deploying the comprehensive user study, we have extensively tested MPTrain with two runners. The results have been very positive and encouraging. The focus of these suite of tests has been to evaluate and refine MPTrain's music selection algorithms. In our tests, the Digital Music Library contained up to 57 songs with durations and tempos ranging from 2:02 to 5:55 minutes and 65 to 185 beats per minute, respectively. The songs belonged to a variety of music genres (*e.g.* pop, rock, soul, hip-hop, etc), both instrumental and vocal. Note that MPTrain's metadata about each song includes the tempo and energy of the song in 20 s window intervals and the average tempo and energy. The music selection algorithm described in MPTrain uses the song's tempo.

In the outdoor running experiments that we have carried out so far, the runners were able to achieve a workout significantly more similar to the desired workout when using MPTrain than in the *mute* and *random* conditions. We have also observed a *learning curve* phenomenon. It typically took the runners a few sessions to get used to both the music in the DML and to MPTrain's style of "coaching". In all cases, the users became better at achieving the desired workout as their experience with the system was increased. Finally, the runners reported that (1) running with music made their workout more enjoyable than running without it; (2) running with MPTrain's music selection was not only fun, but also an efficient way to achieve their workout goal.

[3] In this case a song has already been selected based on the new target.

4. Conclusions and Future Work

We have presented a novel approach to the automatic generation of playlists that incorporates the user's physiological response and purpose in the music selection algorithms. We have illustrated our approach with MPTrain, a mobile phoned based system that utilizes the user's biofeedback to select the "right" music to play to assist the user in achieving a specific exercise goal.

In addition to finishing the user study and analyzing the data from the study, there are several lines of future research that we would like to pursue with the proposed approach, including (1) exploring other domains where our approach could be applied; (2) enriching the user's model by adding information about the user's past interactions with the system to enhance the automatic playlist generation algorithm; (3) incorporating new musical features to the music selection algorithm, such as the song's average energy and the volume at which is being played; (4) altering the speed and volume with which each song is being played, to better direct users towards achieving their goal; (5) incorporating additional contextual information, such as GPS data, body and external temperature, barometric pressure to determine incline, etc.; and (6) working on different user interfaces to allow users to rate the playlist, review their past interactions with the system, and identify trends and deviations from those trends. We would also like to include lifestyle variables, such as diet, overall mood, stress levels, etc, and find correlations between them and the efficacy of the system, both in terms of how well it helps users achieve their goal and how enjoyable the music selection is.

References

[1] http://www.gracenote.com/gn_products/playlists.html.

[2] M. Alghoniemy and A. Tewfik. A network flow model for playlist generation. In *Proc. ICME'01*, 2001.

[3] M. Aucouturier and F. Pachet. Scaling up music playlist generation. In *Proc. ICME'02*, 2002.

[4] D. Grahan-Rowe. Computer dj uses biofeedback to pick tracks. http://www.newscientist.com/article.ns?id=dn1563, 2001.

[5] D. Hauver and J. French. Flycasting: using collaborative filtering to generate a playlist for online radio. In *Proc. Int. Conf. on Web Delivering of Music*, pages pp. 123–130, 2001.

[6] C. Hayes and P. Cunningham. Smart radio: Building music radio on the fly. In *Proc. SGES Conf. on Knowledge based Systems and Applied Artificial Intelligence*, 2000.

[7] N. Oliver and F. Flores-Mangas. Mptrain: A mobile, music and physiology-based personal trainer. In *Proc. Mobile-HCI'06 (To Appear)*, 2006.

[8] S. Pauws and B. Eggen. Pats: Realization and user evaluation of an automatic playlist generator. In *Proc. ISMIR'02*, 2002.

[9] M. Staum. Music and rhythmic stimuli in the rehabilitation of gait disorders. *Journal of Music Therapy*, 20:69–87, 1983.