Brain-Computer Interfaces for Music Recommendation

Kira Radinsky*, **Ashish Kapoor**‡, **Avigad Oron**‡, **Keren Master**‡
‡Microsoft Research and * Technion, Israel
kirar@cs.technion.ac.il, {akapoor, avigado, kerenm}@microsoft.com

Abstract

We explore the opportunity to harness electroencephalograph (EEG) signals for the purpose of music recommendation. The core idea lies on the hypothesis that cortical signals captured by off-the-shelf electrodes carry enough information about mental state of a listener and can be used to build preference models over musical taste for each individual user. We present a reinforcement learning algorithm that aims to build such models over a period of time and then use it effectively to provide recommendations. Our experiments on real users indicate that the recommendation policy learnt via the brain-computer interface provides better recommendations than commercial services such as Pandora.

1 Introduction and Motivation

Recommending music to listeners is a hard problem, that many systems (e.g., Pandora) attempt to solve. However, most of these systems operate on the assumption that the user would first specify the genre or class of the music, and the algorithm would select music corresponding to the specified class (usually by music similarity). For many causal listeners such explicit specification of a genre might not be easy to identify, and they might want to experience multiple genres. Those users might find it easier to specify an emotional/mental state they would like to be in while listening. Given the affective nature of music, it is easy to build upon the assumption that listeners do have a desirable mental state. Consequently, we present a recommendation system that is sensitive to listener's current mental state and makes appropriate recommendations of songs so that the listener can experience the desirable affect.

To solve this problem, we look at the field of brain computer interfaces [2, 3]. The key idea is to measure users' brain signals to provide information about a user's mental state to the system with little conscious effort. This approach is built on the realization that people's subconscious brain activity is indicative of their mental states [1] and can be measured by available brain-sensing methodologies. Note that, the headphone can be easily extended to measure neurophysiological activity and thus, can provide insight into a listener's state. Given the listener's current state and the mental state they desire to be in, we can have an agent select appropriate songs. One of the core contribution of this work is a reinforcement learning methodology to train the agent, which we describe in detail below.

2 Sensing Cortical Activity

The cortical activity is sensed using an Emotiv non-invasive neuroheadset [1] equipped with 14 electrodes. The headset is a low-cost Electroencephalography (EEG) device and provides neurophysiological measurement of brain activity using electrodes placed on the surface of the scalp. The helmet specifically records the activity from the following channels of 32-electrode measurement system [3]: AF3, F7, F3, FC5, T7, P7, O1, O2, P8, T8, FC6, F4, F8 and AF4. The Emotiv set comes with a pre-trained classifier that can be used to classify four emotional states: calm, reflective, meditative and excited.

[1] http://www.emotiv.com

3 Reinforcement Learning of Music Preferences

Given a user's current emotional/mental state, the goal of the algorithm is to perform a song selection that eventually will lead the user to the desired state. We treat this problem as a reinforcement learning problem - consisting of an agent performing actions, and as a consequence of those actions transitioning from state to state, and experiencing rewards. In particular, our state space represents the probability over affective states of: calm, reflective, meditative and excited, formally represented as a vector: $s = (s_1, \ldots, s_4)$, where $s_i \in [0, .., 1]$. Note that a user's emotional states are never directly observed; instead the affect classifier produces probabilities s_i, where $i \in \{1, .., 4\}$, over the four states given the neurophysiological measurement. These probabilities thus constitute our state space. The set of actions consists of possible selection among different genre of songs. At every decision step the algorithm selects a genre, from which it picks a song that has not been played recently. Note that it is fairly important to choose the right action in order to see the desired shift in the affective state of the user. Consequently, the reinforcement learning procedure aims to learn the right mapping of actions given current state s_c and the desired goal state s_g.

The reinforcement learning paradigm uses the notion of reward in order to learn optimal actions for each state. In particular, our proposed algorithm experiences a reward at each step based on the divergence between the reached emotional state and the indicator vector denoting the desired goal:

$$reward(s) = \left(1 + (\sqrt{\sum_{i=1}^{4}(s_i^g - s_i)^2}\right)^{-1}$$

We use a Q-learning algorithm, that calculates the probability of a reward for each state and action $Q(s, a)$, and chooses an action that maximizes this reward. The algorithm operates by a value-operation update:

$$Q(s_t, a_t) = Q(s_t, a_t) + \alpha_t(s_t, a_t)(reward(s_{t+1}) + \gamma max_(a_{t+1})Q(s_{t+1}, a_{t+1}) - Q(s_t, a_t))$$

Here α is the learning rate that influences how much old information on previous states influences the user current song choice, and γ is the discount rate, where higher values encourage the algorithm to take greedy steps. We swiped over these values in the experiments while training on a train subjective not included in the experiments. The values were set to $\alpha = 0.90$ and $\gamma = 0.85$.

4 Results and Conclusions

We collected 35 different songs by first starting with a set of random seed of songs from different genres, and then keeping the record of the suggested songs from Pandora. We had 7 subjects use the system and compare it with Pandora recommendations. In the first stage of the experiment each subject was asked to choose a seed song and listen to the songs recommended by Pandora. In the second stage of the experiment, they were asked to wear the device and describe the emotional state they desire. We asked the user to evaluate both systems after 5 songs had been played. The subjects were asked to rate the quality of recommendations made on a scale of 0-10 (10 is best). We observed that the users preferred the recommendation based by our method (ave. 7.71, std. dev. 0.49) over Pandora (ave. 6.1, std. dev. 0.69). These differences were significant according to the t-test ($p < 0.05$). The results indicate that a personalization based on Neurophysiological measurement is beneficial and can be pragmatically done using off-the shelf devices.

References

[1] EEG-Based Emotion Recognition in Music Listening. Y.P Lin, C.H Wang, T.P Jung et al..IEEE Transactions on Biomedical Engineering 2011.

[2] Complementary Computing for Visual Tasks: Meshing Computer Vision with Human Visual Processing. Ashish Kapoor, Desney Tan, Pradeep Shenoy and Eric Horvitz. IEEE International Conference on Automatic Face and Gesture Recognition 2008.

[3] Special issue on brain-computer interface technology. IEEE Transactions on Neural System Rehabilitation Engineering, 2006.