

Poster Session at Graduate School Information Fair

Affect-guided “sweet spot” exploration using reinforcement learning

Isuru Jayarathne & Michael Cohen
Computer Arts Lab., University of Aizu

Michael Frishkopf & Gregory Mulyk
University of Alberta

Introduction

- Relaxation techniques for reducing effects of stress can improve physical and mental health.
- Among many relaxation techniques, musical relaxation is one of the simplest.
- Nature sounds, instrumental music, voice (chanting), “easy listening” songs, etc. can be played for relaxation.
- Selecting suitable sounds for a particular subject is challenging because of idiosyncratic tastes and circumstances.
- In our approach, computer-guided audition for spatial soundscapes is investigated, automatically exploring a polyphonic area using biosignals as indicators of satisfaction.
- Reinforcement learning (RL) is a technique of programming agents by rewards and penalties to achieve a goal.
- Reward values are calculated using valence and arousal of electroencephalography (EEG) signals fed to a Deep Q-Network (DQN) RL algorithm to adjust the parameters of spatial soundscape.

Methodology

- As shown in Figure 1, a Max/MSP patch was developed as a soundscape interface.
- Headphones present this soundscape featuring six pantophonic sources arranged in a ring around the perimeter of the virtual space, which is parameterized not only by Cartesian position but also by processing attributes such as volume, pitch, reverb, and various parameters (center and cutoff frequencies, and resonance) for low-pass, high-pass, bandpass, and notch filtering.

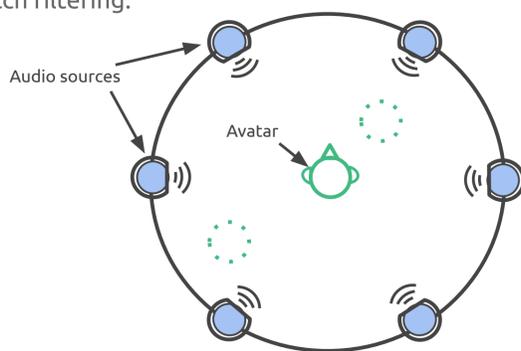


Fig. 1: Horizontal soundscape map showing iconic avatar and distribution of sound sources in Max/MSP interface

- Muse headband with 4 EEG electrodes, was used to collect data.
- At run-time, the subject was instructed to keep eyes closed to minimize distractions, and auditioned the multidimensional soundscape.
- The realtime data stream is broken into windows with 50% overlap and filtered using a Butterworth band-pass filter to extract the relative theta frequency band power.
- Reward was calculated using following equation.

$$\text{Reward} = \frac{\alpha_{AF8}}{\beta_{AF8}} - \frac{\alpha_{AF7}}{\beta_{AF7}}$$

- As shown in Figure 2, the DQN algorithm takes the reward and state parameters from the environment and changes the position of the avatar.
- A simulation (Figure 3) that can feed pre-defined reward and observations was performed to test performance of the DQN model.

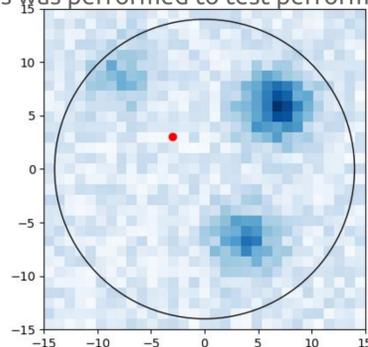


Fig. 3: Visualization of pre-defined rewards. Red dot represents the avatar and each pixel contains reward and observations. Darker pixels indicate higher rewards.

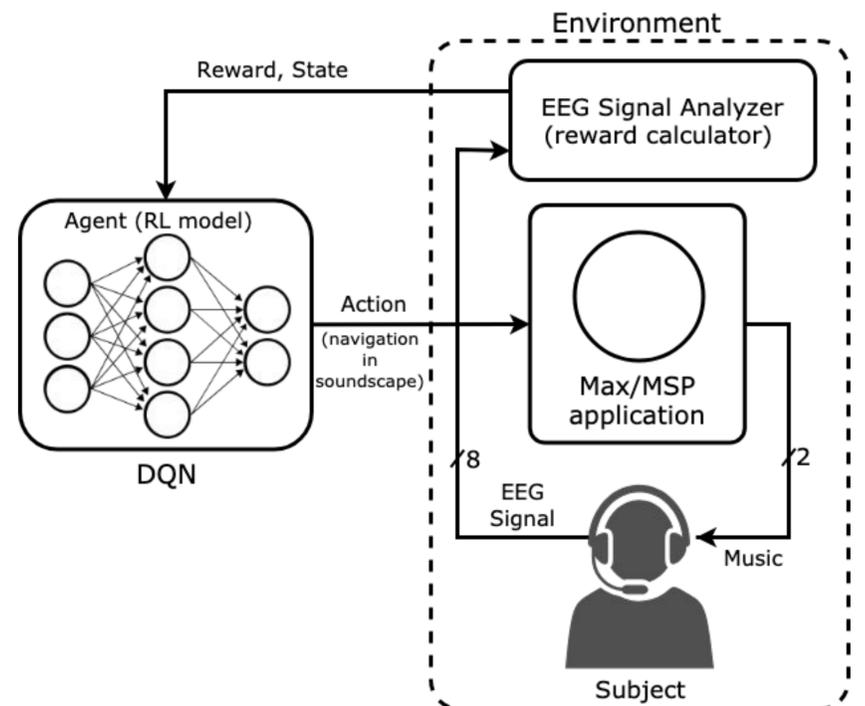


Fig. 2: System architecture

Results

- The simulator was run with several different reward matrices by changing the position of the sweet spot.
- The DQN agent was able to find a sweet spot in around 2000 steps (around 40 – 45 episodes).
- Figure 4a shows the variation of cumulative reward over episodes and Figure 4b shows number of steps for each episode.
- Sometimes the DQN agent was not able to find a global maximum when there is no pixel path to the sweet spot with high reward values.

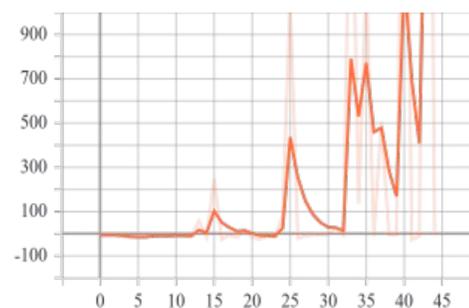


Fig. 4a: Variation of cumulative reward for each episode

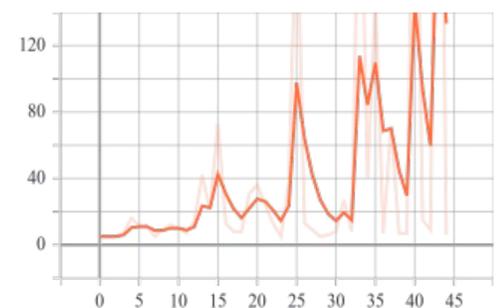


Fig. 4b: Variation of number of steps for each episode

Discussion

- According to the simulation results, DQN agent has a capability find the sweet spot if there is a stable maximum reward.
- Our experiment assumes that relaxation can be described by a stable field across a pantophonic soundscape bounded by the annulus (ring) of a half-dozen musical sources.
- The sweet spot is probably not a point, but an area or collection of locations (perhaps unconnected).
- There are also several reasons the sweet spot probably isn't stable, including memory & habituation, boredom, hysteresis, and fatigue.
- EEG signals might not be the best parameter for measuring relaxation, but we leverage our experience with such biosensing, rather than, for instance, self-reported satisfaction (which could also interfere with such psychological state).

References

- [1] Antoine Hennion. Music lovers: Taste as performance. *Theory, Culture & Society* 18, 5 (2001), 1–22.
- [2] Gregg D. Jacobs and Richard Friedman. EEG spectral analysis of relaxation techniques. *Applied Psychophysiology and Biofeedback* 29, 4 (2004), 245–254.
- [3] Leslie Pack Kaelbling, Michael L Littman, and Andrew W Moore. Reinforcement Learning: A survey. *J. of Artificial Intelligence Research* 4 (1996), 237–285.
- [4] Eun-Jeong Lee, Joydeep Bhattacharya, Christof Sohn, and Rolf Verres. Monochord sounds and progressive muscle relaxation reduce anxiety and improve relaxation during chemotherapy: a pilot EEG study. *Complementary Therapies in Medicine* 20, 6 (2012), 409–416.
- [5] Richard S. Sutton and Andrew G. Barto. 2018. *Reinforcement Learning: An Introduction*. MIT Press. ISBN 0-262-19398-1.
- [6] Hado Van Hasselt, Arthur Guez, and David Silver. 2016. Deep Reinforcement Learning with Double Q-Learning. In *AAAI*, Vol. 2. Phoenix, AZ, page 5.
- [7] Norman M. Weinberger. Music and the brain. *Scientific American* 291, 5 (2004), 88–95. ISBN 978-1-848933-01-8.
- [8] David E. Wolfe, Audree S O'Connell, and Eric G Waldon. Music for relaxation: a comparison of musicians and nonmusicians on ratings of selected musical recordings. *J. of Music Therapy* 39, 1 (2002), 40–55.

