

MISQ Archivist

The Right Music at the Right Time: Adaptive Personalized Playlists Based on Sequence Modeling

Elad Liebman, Maytal Saar-Tsechansky, and Peter Stone

Abstract

Recent years have seen a growing focus on automated personalized services, with music recommendations a particularly prominent domain for such contributions. However, while most prior work on music recommender systems has focused on preferences for songs and artists, a fundamental aspect of human music perception is that music is experienced in *temporal* context and in sequence. Hence, listeners' preferences also may be affected by the sequence in which songs are being played and the corresponding song transitions. Moreover, a listener's sequential preferences may vary across circumstances, such as in response to different emotional or functional needs, so that different song *sequences* may be more satisfying at different times. It is therefore useful to develop methods that can learn and adapt to individuals' sequential preferences in real time, so as to adapt to a listener's contextual preferences during a listening session. Prior work on personalized playlists either considered batch learning from large historical data sets, attempted to learn preferences for songs or artists irrespective of the sequence in which they are played, or assumed that adaptation occurs over extended periods of time. Hence, this prior work did not aim to adapt to a listener's *current* song and sequential preferences in real time, during a listening session. This paper develops and evaluates a novel framework for *online* learning of and *adaptation* to a listener's current song and sequence preferences exclusively by interacting with the listener, during a listening session. We evaluate the framework using both real playlist datasets and an experiment with human listeners. The results establish that the framework effectively learns and adapts to listeners' transition preferences during a listening session, and that it yields a significantly better listener experience. Our research also establishes that future advances of *online* adaptation to listener's temporal preferences is a valuable avenue for research, and suggests that similar benefits may be possible from exploring online learning of temporal preferences for other personalized services.

Keywords: Online preference learning, personalized adaptation, user preferences, music playlist preference, reinforcement learning, data science