

7 Machine Learning and Real-world Data (av308)

You are interested in developing a system to recommend music genres to users according to their moods. As different users have different tastes, each of them provides you with a playlist containing the genre ((C)lassical, (R)ock, (J)azz, (E)lectronica, (H)ip-hop) listened to and their mood (U)pbeat, (D)ownbeat, (N)eutral). Here is an example from a user:

Mood	U	U	N	D	D	N	N	U	U	...
Genre	R	R	J	C	J	J	E	E	H	...

Your manager tells you that you should model this data using a first order hidden Markov model (HMM), with mood as the hidden states and genre as the observed states.

- (a) Define and estimate the parameters of the HMM using the playlist above. Provide equations as needed. Include the start state in your calculations, but not the end state as the example continues in perpetuity. No smoothing should be applied. [4 marks]
- (b) What are the two assumptions made by the HMM? Are they appropriate in this application? [4 marks]
- (c) Using the model you estimated above, predict the most likely music genre to continue this playlist. [4 marks]
- (d) Using the model you estimated above, what is the probability of following Rock immediately with Classical? [2 marks]
- (e) What would happen if you used the HMM you have trained to generate a playlist? Could you generate different playlists? Give reasons for your answer. [2 marks]
- (f) Imagine you now define finer genres, e.g. by splitting rock into classic rock vs heavy metal. Which advantages and disadvantages do you foresee? [2 marks]
- (g) What would happen if you replaced moods with the user's own reactions (e.g. likes, faves, etc.) and learned the parameters of the HMM? Would the modelling assumptions of the HMM still be valid? [2 marks]