Artificial Intelligence 2 – SS 2020
Assignment 8: Hidden Markov Models
– Given June 5., Due June 14. –

Hint: Exercises need to be handed in via StudOn at 23:59 on the day they are due or earlier. Please use only the exercise group of your tutor to hand in your work.

If any concepts here seem unfamiliar to you or you have no idea how to proceed, consult the lecture materials, ask a fellow student, your tutor, or on the Forum.

If a problem asks for code, comment it or make it otherwise self-explanatory. You do not need to write a lot, but it should be enough to convince your tutor that you understand what the code does. We may deduct up to 30% for uncommented and unclear code, but would prefer not to.

Problems with no points (0pt) will not be graded, but might appear on the exam in a similar form. For these, we will provide a reference solution after the submission deadline. If you find the reference solution unclear, ask about it on the forum or in in a tutorial.

Problem 8.1 (Markov Mood Detection Continued)

In this problem, we again visit your flatmate Moody. It looks like he is now taking a little longer to go between happy and not happy, so you decide to introduce a neutral mood in your model. On any given day d, Moody is in one of three states – either he is happy ($M_d = h$), neutral ($M_d = n$) or he is grumpy ($M_d = g$). He never goes straight from being happy one day to grumpy the next day (or from grumpy to happy). When he's happy, he is likely to stay that way, so $P(M_{d+1} = h|M_d = h) = 0.8$. He is a little less likely to stay grumpy: $P(M_{d+1} = g|M_d = g) = 0.6$. When he's feeling neutral, he's most likely going to be happy the next day, $P(M_{d+1} = h|M_d = n) = 0.4$, and he's as likely to stay neutral as he is to become grumpy. 50pt

Notably, you can hear his music blasting all day which tends to shift depending on his mood. On a good day he often listens to Jazz (i.e. $P(j_d|M_d = h) = 0.6$), on a bad day he usually blasts Death Metal at full volume ($P(\neg j_d|M_d = g) = 0.85$). On a neutral day he is as likely to listen to Jazz as to Death Metal. He has a limited taste in music, so it's always one of the two.

This is a hidden Markov model of the situation, in which the variables M_d describe the state, and the variables J_d describe the evidence.

1. What is the transition matrix of this model? 10 pt
2. You heard your flatmate play Death Metal yesterday and Jazz today. What are the sensor matrices for the first two steps? 10 pt
3. You are not sure what kind of mood your flatmate was in on Wednesday, but it was either good or neutral, with equal probability. Given the music he played today and yesterday, use the matrix version of smoothing to determine the probability distribution of Moody's mood yesterday. 30 pt

   Hints. It can be difficult to get the computation right, so make sure to write out the matrices and vectors you are computing with – they will get you the most points.

Watch out for that transposed transition matrix in the formula! What are the two vectors you need when computing $f_{1:1}$ and $b_{2:2}$? (If you don't know, ask in the tutorial.)