

HW11: Expectation Maximization

Hand in at: <http://www.cs.utah.edu/~hal/handin.pl?course=cs726>. Remember that only PDF submissions are accepted. We encourage using L^AT_EX to produce your writeups. See `hw00.tex` for an example of how to do so. You can make a `.pdf` out of the `.tex` by running “`pdflatex hw00.tex`”.

1. Consider the GMM framework, and suppose that each cluster k had its own cluster-specific variance σ_k^2 . What would the updates for these variables look like?
2. I have two coins, A and B. Your job is to figure out π_A and π_B , the probability of heads of each of these coins. However, I'm evil and won't let you flip the coins yourself. What I will do, however, is flip them on my own and tell you the results. In particular I say something like: I picked one of the coins, flipped it 10 times, and it came up heads 7 times and tails 3 times. Then I picked one of the coins (perhaps the same one, perhaps not), flipped it 10 times, and it came up heads 5 times and tails 5 times. I tell you this information N -many times (so you the results of a total of $10N$ coin flips).

Set this up as an EM problem. What is the data, what are the parameters and what are the hidden variables? Write down the *complete data likelihood* for this problem. (You don't need to solve the EM problem: we'll do that in class.)