

ISTA 410/510 Homework V

For contribution to the final grade, due dates, current late policy, and instructions for handing the assignment in, see the assignment web page.

Please create a PDF document with your answers and/or the results of any programs that you write. You should also hand in your programs.

Questions marked by * are required for grad students only. They count as extra problems for undergraduates.

Questions marked by ** are challenge problems. They count as extra problems for both grads and undergraduates.

Any non-challenge problem can be replaced by a challenge problem; please make it clear that this is what you are doing (e.g., for a required problem you could answer “see optional problem #3”). The point here is to enable students to avoid problems that they feel are not instructive.

Extra problems (please indicated in your answer when you are doing an extra problem) are eligible for modest extra credit. The maximum score for an assignment will be capped at 120%. The maximum score for all assignments taken together is capped at 65/60.

For simplicity, problems are generally all worth the same, except ones marked by “+” that are expected to substantively more time consuming, and are worth double (or two no “+” problems). Additional “+”s scale linearly.

1. (+++++) Implement EM for a GMM where the number of clusters is K , the dimensionality of the problem is $D=2$, and the covariance matrices are diagonal. However, the variances for each dimension, for each cluster, should be allowed to vary. You should monitor the log-likelihood while EM iterates. You should also structure your code to hold out every 10th point, and monitor the log-likelihood of the held out data as well. Your program should plot the data it reads in, and once it is done, it should plot the points again, now colored by cluster. For simplicity, your coloring can be based on the maximally likely cluster.

You will probably find it best to generate various kinds of data to test your algorithm. However, some data sets that you should test on are provided on the web side.

Test your code for $D=2$, and $K=2,3,5,7$, and 10, on these data sets using several random initializations.

Report whether or not: a) the log likelihood of the training data always goes up; b) the log likelihood of the held data always goes up; c) the answer is a function of the initialization.

Note that the scale of the two log likelihoods are proportional to the number of points. Hence, the test log likelihood and the held out log likelihood are only comparable if you consider a factor of 10. More generally, we can think in terms of the per-point log-likelihood. In general, report whether the per-point log-likelihood is better for the training data or the held out data. Explain whether this is what you expected or not.

Hand in plots for two examples of clustering that you like, and two examples of clustering that you don't like.

2. (*,++) Modify your program so that the full covariance matrix can be learned. Provide plots for a clustering that you like, and a clustering that you do not like. Is the general behavior of the held out log likelihood consistent with what you found in (1)? Now comment on the absolute values of the fits (likelihoods) on the training and held out data.
3. (**, ++ to ++++++). If you have already implemented EM on toy problems, and are familiar with the vision lab C libraries, you might prefer to use the existing implementations for EM with GMM for (1) and (2) above. Follow this by improving the implementations which are a bit of a mess. In addition to improving modularity, making the feature sets more consistent over the models, rationalizing parameters, improving efficiency, and so on, I suggest making the code multi-threaded, where a sensible number of threads can be set using an option, with the default ideally set based on the platform. (A pair of students made a start on this last year, and if you are interested in continuing this work, I will put you in touch with them). You should do timing tests. Another direction to investigate is the problem of doing EM if the data is too large to fit into memory.

If you chose this assignment, let Kobus know. If more than one person is keen, we will coordinate the effort.

4. (**, ++ to ++++++). If you have already implemented EM, and do not like question (3), propose your own EM related mini-project.