

6.867 Machine Learning

Problem set 6

Due Wednesday, November 27, in tutorial

What and how to turn in?

Turn in short written answers to the questions explicitly stated, and when requested to explain or prove. Do **not** turn in answers when requested to “think”, “consider”, “try” or “experiment” (except when specifically instructed). You may turn in answers to questions marked “optional”— they will be read and corrected, but a grade will not be recorded for them.

Turn in all MATLAB code explicitly requested, or that you used to calculate requested values. It should be clear exactly what command was used to get the answer to each question.

To help the graders (including yourself...), please be neat, answer the questions briefly, and in the order they are stated. Staple each “Problem” separately, and be sure to write your name on the top of every page.

Problem 1: Clustering

In this problem we will experiment with the spectral clustering algorithm, and explore its properties and the influence of its parameters. A Matlab implementation of spectral clustering is provided in `spectral.m`. You can call it as

```
labels = spectral(X,k,beta);
```

where `X` is the 2D data, `k` is the number of neighbors to use in the neighborhood graph, and `beta` is the weight falloff parameter. The default values for `k` is 3, for `beta` 1. Use `spectral` to experiment with the algorithm, and to test your hypotheses regarding the questions.

You may also find use for the little function

```
X = mkdata(m,n);
```

that generates random data points, from m Gaussian distributions, with n points from each. Such data typically has m sets of points that form natural clusters.

The code is written to perform binary clustering; let us first see how the algorithm deals with the data where there are indeed two clusters.

1. [5pt] What will typically happen if only a single nearest neighbor is used to create the graph (i.e. $k=1$)? Explain why that happens.
2. [5pt] What happens if `beta` is set to a very large value (say, 100), when k is reasonable - say, 3? Explain.

Now assume that there are in fact three clusters in the data. In the form given to you, the clustering algorithm of course can not discover the three clusters (since it only looks for two).

3. [5pt] How will the algorithm, as given to you, behave on the data with three clusters? Explain.
4. [10pt] Modify the algorithm in order to find the three clusters. Note that deciding how many clusters are in fact present in the data is an important problem in clustering; you do not have to solve it, however - the modified algorithm should simply assume there are three clusters. Explain your modifications, and turn in your code.

Problem 2: HMM

One of the most notable contributions of HMM has been is speech recognition. In this problem we will take a look at an application of HMM to a relatively simple problem in speech – recognition of isolated words. We will focus on a “toy” version of this problem, namely, we are only interested in understanding whether the input utterance (acoustic speech signal) represents the word “four” or “five”.

We will be working with the Mel-frequency cepstral coefficients (MFCC) of the signal, which are essentially the Fourier coefficients transformed in a way consistent with how the biological systems process acoustic information. You can find more about this representation in the literature devoted to speech recognition; one reference is

L. Rabiner and B.-H. Juang, *Fundamentals of Speech Recognition*. Englewood Cliffs, NJ: Prentice Hall, 1993.

The 13-element MFCC vectors are the observations \mathbf{x}_t in our model. The MFCC are computed within sequential time windows of fixed size, and since the length of the utterance varies from sequence to sequence, so does the number of observations. All the data used in this problem is found in the file `data45.mat`. The data was collected from 8 male speakers, each uttering the two words a number of times. The cell arrays `train4`, `train5` contain in each element (e.g. `train4{j}`) a $13 \times T$ matrix of normalized MFCC, where T is the length of the corresponding observation sequence.

Our HMM for each spoken digit will contain 5 hidden states. The states roughly correspond to the phonemes (sounds) in the word; in our case, both “four” and “five” have three phonemes, and we expect the states to model these phonemes, plus the “silent” segments before and after the word. The transitions between states correspond to a sequential process, and are constrained in the following way: from a state s , the model can only transition to the state $s + 1$, no “skipping states” or going “backwards” is allowed. Figure 1(b) schematically shows the state transition diagram for this model.

Finally, we will model the probability density of observation (MFCC vectors) for each state by a single multivariate (13-dimensional) Gaussian with full covariance. The model is summarized in Figure 1. Before we proceed, it is important to note that this is an over-simplified model, and the HMM typically used for similar task (isolated word recognition) in speech would usually be more complex (larger number of states, observations modeled by a mixture of Gaussians, etc.)

We will use a (slightly modified) subset of the Matlab HMM toolkit written by Dr. Kevin Murphy; the original toolkit and many other useful software tools can be found at <http://www.ai.mit.edu/~murphyk/Software>.

Extracting the archive `HMM.zip` or `HMM.tar.gz` will create a directory `HMM`; make sure you include it in Matlab working path (use the `addpath` command).

1. [2pt] What aspect of the HMM captures the estimated durations of the word parts (i.e., if one part usually takes longer time than another, how will this affect the HMM

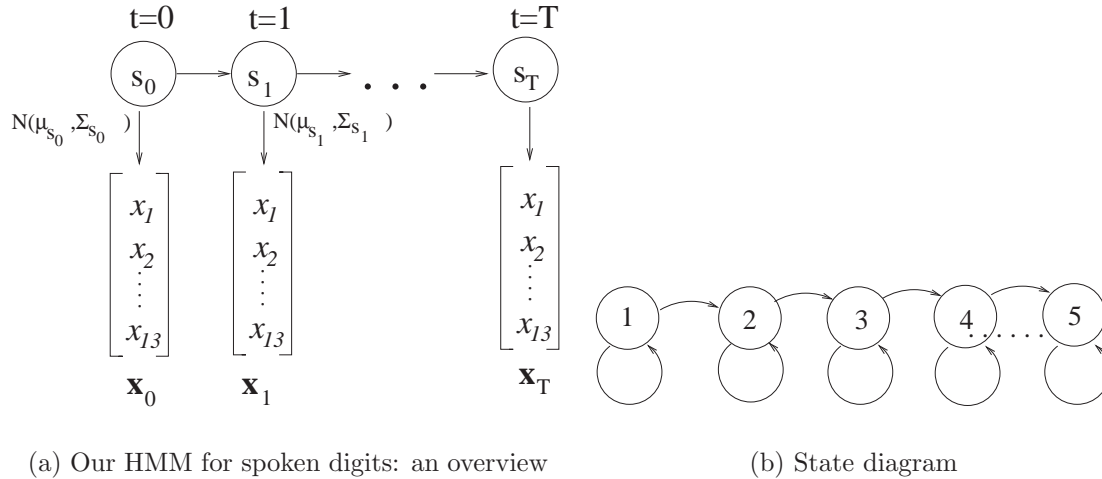


Figure 1: The HMM used in Problem 2

parameters)?

2. [3pt] Prove that if an entry i, j in the initial “guess” of the transition probability matrix provided to the EM is set to zero, it will remain zero in the eventual result obtained with the EM.
3. [2pt] Precisely what constraints do the 5×5 transition probabilities matrix $P(s_{t+1}|s_t)$ and the 5-element initial state probabilities vector P_0 have to satisfy in our model?
4. [3pt] Using the implementation of the EM algorithm provided to you, estimate the parameters of the two HMMs, based on the data in `train4` and `train5` respectively. The function you should use is

```
[p0, pt, mu, sigma] = learn_ghmm(X, initP0, initPt);
```

The inputs to `learn_ghmm` are the cell array of observation sequences \mathbf{X} and the values, respectively, of $P_0(i)$ and $P(s_{t+1}|s_t)$ that the EM should use as the initial guess. In the output, `p0` is the estimated P_0 , `pt` the estimated transition probabilities matrix $P(s_{t+1}|s_t)$, and `mu` and `sigma` are the estimated parameters of the Gaussian distributions of observations for each state (the Gaussian for state k is determined by `mu(:,k)` and `sigma(:, :, k)`). Make sure you appropriately constrain the transition matrix and the starting state probabilities, according to your answers to the previous two questions.

Turn in the resulting transition probability matrix for digit “four”.

Having trained an HMM (estimated its parameters) we can apply it to various tasks related to the spoken digit recognition. One of the main goals is, of course, classification - given an unlabeled utterance, we would like to tell whether it is a “four” or a “five”. The cell array `test45` contains 252 MFCC representations of spoken digits; the true labels are given in the array `labels`. We will now test how well our HMMs can classify these utterances. To do so, we will need to use the function

```
loglikes = log_lik_ghmm(data, prior, transmat, mu, Sigma);
```

which takes the cell array `data`, the probabilities of the initial state `prior`, the transition probability matrix `transmat`, and the parameters of observation densities `mu`, `sigma`, and returns in `loglikes(k)` the log-likelihood of the observation sequence `data{k}` given the HMM.

5. [5pt] Classify the utterances in `test45` based on their log-likelihoods under the two models. Explain in one or two sentences how you perform the classification and why it makes sense (no need to turn in code). What is the misclassification rate on `test45`?

Another possible use for HMMs can be to align two observation sequences. In the case of speech analysis, this could mean marking the segments of the signal corresponding to the same hidden state. For this part, you will need to use the function

```
path= viterbi_path(data, prior, transmat, mu, sigma);
```

which for the observation sequence `data` returns the most probable sequence of states, found by Viterbi algorithm with the HMM described by `prior`, `transmat`, `mu` and `sigma`. We will try this on just one example - a pair of utterances of the word “four”. The original sampled audio signal is given in the vectors `signal1`, `signal2`, and the MFCC representation in `mfcc1`, `mfcc2`. Note that these utterances (by different speakers) have quite different lengths.

6. [3pt] Explain how the most probable sequence of states helps us to align the sequences.
7. [7pt] Using Matlab function `specgram`, plot the spectrogram of each signal, and under it plot the graph showing the states found by Viterbi algorithm (use the `subplot` command). Assuming that the states roughly correspond to sounds, mark the areas on the two spectrograms that probably correspond to the sound “f”. (Recall that the observed sequences typically have a “silence” period in the beginning, which is modeled by a separate hidden state).

We now turn to a more interesting task: separating words in a sequence. We will do it in a simplified setting, for a sequence of words known to be either “four” or “five”. In order

to separate the words, we must model not only the separate words, but also the transition between the words. Normally, these would be estimated based on a training set of word sequences; here, we shall simply assume that the probability of leaving the final state of the word and going to the initial state of the next word is 0.2, and that the probabilities of “four” and “five” are equal.

8. [5pt] Explain how a single HMM would combine the two individual HMMs; turn in the initial state probabilities and the transition matrix of this combined HMM.
9. [5pt] We will work with 4 sequences of two spoken digits `seq1, . . . , seq4`. Using the combined HMM, and the function `viterbi_path`, find for each sequence what are the two digits, and where is the boundary, according to the HMM estimate. Along with your answers turn in the plots, or numerical results, on which you base your estimates.

C	$P(C)$
g	.3
a	.5
b	.2

(a)

C	$P(T_1 C)$		$P(T_2 C)$	
	1	0	1	0
g	.7	.3	.8	.2
a	.5	.5	.4	.6
b	.2	.8	.1	.9

(b)

T_1	T_2	$P(D T_1, T_2)$	
		0	1
0	0	1	0
0	1	1	0
1	0	0	1
1	1	0	1

(c)

Table 1: Settings for Problem 3. (a) $P(C)$. (b) $P(T_1|C)$ and $P(T_2|C)$. (c) $P(D|T_1, T_2)$.

Problem 3: Graphical models

Let us start by warming up with a simple belief network. Consider the following scenario. A person (hereafter called the “buyer”) wants to buy a used car; the condition of the car is a random variable C , with the value in $\{g, a, b\}$ corresponding, respectively, to good, average and bad quality. Our belief about the frequency of each kind in the market is expressed by the distribution of C , shown in Table 1(a). The buyer bases his decision whether to buy a car on the outcome of two tests which he runs on the car, T_1 and T_2 . The outcomes, t_1 and t_2 respectively, are binary - i.e. the car either passes a test (1) or fails (0). Unfortunately, the results of the tests are non-deterministic, even given the condition of the car. We do know the distribution of this results, conditioned on the car quality - see Table 1(b). Based on the test outcomes, the buyer decides whether to buy a car; for the sake of convenience, let us for now treat this decision as a random variable, even though its distribution (conditioned on the test outcomes) shown in Table 1(c) shows that it is in fact a deterministic decision. We encode the decision to buy as 1 and not to buy as 0.

1. [2pt] Draw the graphical model corresponding to the described settings. Your model should contain 4 nodes (C, T_1, T_2, D) – mark them clearly.
2. [2pt] Write down the joint probability of observing the quadruple ($C = c, T_1 = t_1, T_2 = t_2, D = d$) as a product of known distributions.
3. [2pt] What are the marginal distributions of test outcomes, $P(T_1)$ and $P(T_2)$? What is their joint distribution? Are these two random variables marginally independent?
4. [2pt] You are told that the car in question is actually in average condition. What is the joint distribution of T_1, T_2 given that knowledge? Are T_1 and T_2 independent given C ?




Figure 2: Simple influence diagram. Semantics of the nodes: C – condition of the car, T – the outcome of the test, D - the decision whether to buy or not, U - the utility to be maximized by the decision.

5. [2pt] Suppose that in addition to knowing that the car is in average condition, you also know that the buyer purchased the car. What is the joint distribution of T_1, T_2 given that knowledge? Are T_1 and T_2 independent given C and D ?

Treating D as a random variable is somewhat of an abuse - it is, in fact, a special entity – a *decision node*. Such nodes are usually shown in a graphical model as squares, to distinguish them from the random variable nodes (circles). A decision node represent choices available to a decision-maker. The arcs entering a decision node have different meaning from the arcs entering a random variable node; rather than denoting probabilistic dependence, they show what informational input is available to the decision maker at that node.

In addition, we can define yet another special kind of nodes in a graphical model – *value nodes*. A value node, which is shown as a diamond, represents a utility (function) to be maximized by the decision maker. Arcs entering a value node show functional dependence (as in Bayesian network, the exact form of the dependence must be stored in the node). The belief network, augmented by decision and value nodes, is called an *influence diagram*.

Before we proceed, let us look at an example. Suppose we have a simplified version of our car purchase scenario, in which there is a single test, T , which the car passes with probability 0.8, 0.4 or 0.1 if its condition is respectively g , a or b . The distribution of C is the same as in Table 1. Based on the test outcome, the buyer decides whether to buy the car or not. The influence diagram for this scenario is shown in Figure 2.

In order to determine the utility function $U(c, d)$, we will also need the information on losses and gains associated with each decision. Let the price of the car be \$1000, and assume that the market value for a good car of this model is $V(g)=\$1,500$, for a bad car it is $V(b)=\$300$, and for an average car $V(a)=\$1100$. The buyer must also pay the fee for

the test - say, \$50. The utility function then is

$$U(c, d) = \begin{cases} u^0 = -\$50, & \text{if } d = 0 \\ u^1(c) = V(c) - \$1050, & \text{if } d = 1. \end{cases} \quad (1)$$

The task typically associated with an influence diagram is to select the strategy at the decision nodes that would maximize the value of the utility. A strategy $S(X)$ is a function associated with a decision node, which given the input X to that node, computes the decision value. There are four possible strategies in our case:

$$S_1(t) : \quad t = 0 \rightarrow d = 0, \quad t = 1 \rightarrow d = 0, \quad (2)$$

$$S_2(t) : \quad t = 0 \rightarrow d = 0, \quad t = 1 \rightarrow d = 1, \quad (3)$$

$$S_3(t) : \quad t = 0 \rightarrow d = 1, \quad t = 1 \rightarrow d = 0, \quad (4)$$

$$S_4(t) : \quad t = 0 \rightarrow d = 1, \quad t = 1 \rightarrow d = 1. \quad (5)$$

We can compute the expected utility for a strategy S , given the distribution of the random variables in the diagram:

$$E_S = \mathbf{E}_{C,T} [U(c, S(t))] = \mathbf{E}_C [\mathbf{E}_{T|C} [u(c, S(t))]] = \sum_{c=g,a,b} P(c) \sum_{t=0,1} P(t|c) u(c, S(t)), \quad (6)$$

which becomes, for our four strategies,

$$E_{S_1} = -50,$$

$$E_{S_2} = .3(.8 \times 450 - .2 \times 50) + .5(.4 \times 50 - .6 \times 50) + .2(-.1 \times 750 - .9 \times 50) = 76,$$

$$E_{S_3} = .3(.2 \times 450 - .8 \times 50) + .5(.6 \times 50 - .4 \times 50) + .2(-.9 \times 750 - .1 \times 50) = -116,$$

$$E_{S_4} = .3 \times 450 + .5 \times 50 - .2 \times 750 = 10.$$

Thus, perhaps not surprisingly, the strategy that maximizes the expected utility value is S_2 - buy the car only if it passes the test. Of course, a different cost, or a different probability distribution of a random variable could change that.

Let us now proceed and develop a somewhat more interesting framework for the car buying problem. The meaning of C and its distribution are same as before. Test 1, with probability distribution of the outcome T_1 as described in Table 1, consists of simply looking at the car, and costs nothing. Its outcome can be interpreted as whether the buyer liked the car (1) or did not (0). We can make a decision whether to buy the car (perhaps based on the outcome of T_1), or ask a mechanic to check the car first - test 2, with outcome T_2 , for which there is a fee of \$50 - and then decide. The car price and the market value are the same as described above. Note that different decision strategies involve decisions based on both tests, one of the tests, or no tests at all. The utility of the decision is again measured by the financial gain of the buyer.

6. **[10pt]** Draw an influence diagram corresponding to this setting. Include the value node U . For each node clearly show its kind (by drawing it as a circle/square/diamond), name it, and describe its semantics (meaning) in no more than one sentence.
7. **[3pt]** How many possible strategies are implementable in this influence diagram? Explain what they are (do not enumerate all the strategies!). Note that the strategy is described by the functions associated with *all* the decision nodes in the diagram.
8. **[2pt]** Write down the expected value of the utility as a function of the strategy.
9. **[5pt]** Find the strategy with the maximal expected utility, and the maximum value attained by that strategy. *Advice: You may find it appropriate to run a short script in Matlab or another programming language that will compute the expected utility for different strategies; no need to turn it in.*