

# 6.867 Machine learning and neural networks

## Problem set 4

Due: November 28th, in recitation

### 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).

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.

## Density Estimation using Mixtures of Gaussians

In this problem, we will fit mixture of Gaussians models to traced handwritten digits using the EM algorithm. Most of the code and utilities that you will need will be provided.

`EM.m` fits a mixture of  $k$  Gaussian distributions to the training data. The simplest way to call this function is `[params,loglikes] = EM(X,k)`, where `params` contains the estimated parameters (mixing proportions and the means and covariances of each of the component Gaussian); `loglikes` stores the log-data likelihoods of the mixture models we produce in the course of the EM iterations. The following functions are called inside `EM.m`:

**Estep** Performs the *expectation* step: for each data point, we calculate the posterior probability that it was drawn from each of the component Gaussians.

**Mstep** Performs the *maximization* step: finds the maximum likelihood settings of the parameters assuming we already have the posterior probabilities computed in the E-step.

**mixtureLL** Calculates the log-likelihood of the data given the current setting of the mixture model parameters (component Gaussians as well as the mixing proportions).

1. Complete the function `Mstep.m` in terms of updating the means and the covariances of the Gaussian components (use the skeleton we have provided).

**Answer:**

```
function compParams = Mstep(X, p)
    [n,d] = size(X) ;
    k = size(p,2) ;
    compParams = [] ;
    for i=1:k
        reped_p = repmat(p(:,i),1,d) ; % This can come in usefull
        ni = sum(p(:,i)); % Total 'weight' of ith component
        prior = ni/n;
        mu = sum( reped_p.*X )/ni; % weighted mean
        XminusMu = op(X,'-',mu);
        Sigma = (reped_p.*XminusMu)'.*XminusMu/ni; % weighted covariance
        compParams = [compParams, struct('prior',prior,'mean',mu,'cov',Sigma)];
    end;
```

**Scoring:** 7 points

Our implementation of the E-step (`EstepX.m`) deviates slightly from a “straight-forward” implementation. Modifications of this type are necessary in practice to avoid numerical problems.

2. By looking through `EstepX.m`, `mixtureLLX.m` and `evalmixtureX.m` can you suggest what the problem is that these routines are trying to avoid?

**Answer:** The numeric values of the densities are extremely large. Multiplying many of these probabilities quickly leads to numeric overflow.

**Scoring:** 1 point

In the following questions, we study a dataset of handwritten digits given in the file `pendigits.mat` (type `load pendigits` to load it into `MATLAB`). The digits were drawn with a stylus input device. The raw trajectories of the digits can be found in the files `pendigits-orig.(tra|tes|names)`. The trajectories were then scaled to normalize them, and eight equi-distant points along the trajectory were extracted from each digit sample. We will use versions of the resulting 16-dimensional dataset (two dimensions for each of the eight points). You can find more information in the file `pendigits.names`.

The provided routine `showdigit.m` can be used to plot the digits:

```
>> showdigit(pentrain8(42,:))
>> showdigit(pentrain4(42,:))
```

The green circle is the beginning of the trajectory. Each row of the data matrix `pentrain4` contains only four points, the means of every two consecutive points.

In order to visualize our estimation results, we will simplify the problem a bit by considering only the second point in `pentrain4`, given by `pentrain4(:,3:4)`. Moreover, we will restrict ourselves to just digits three, four, and five. The resulting  $2219 \times 2$  matrix `second345` contains the relevant points.

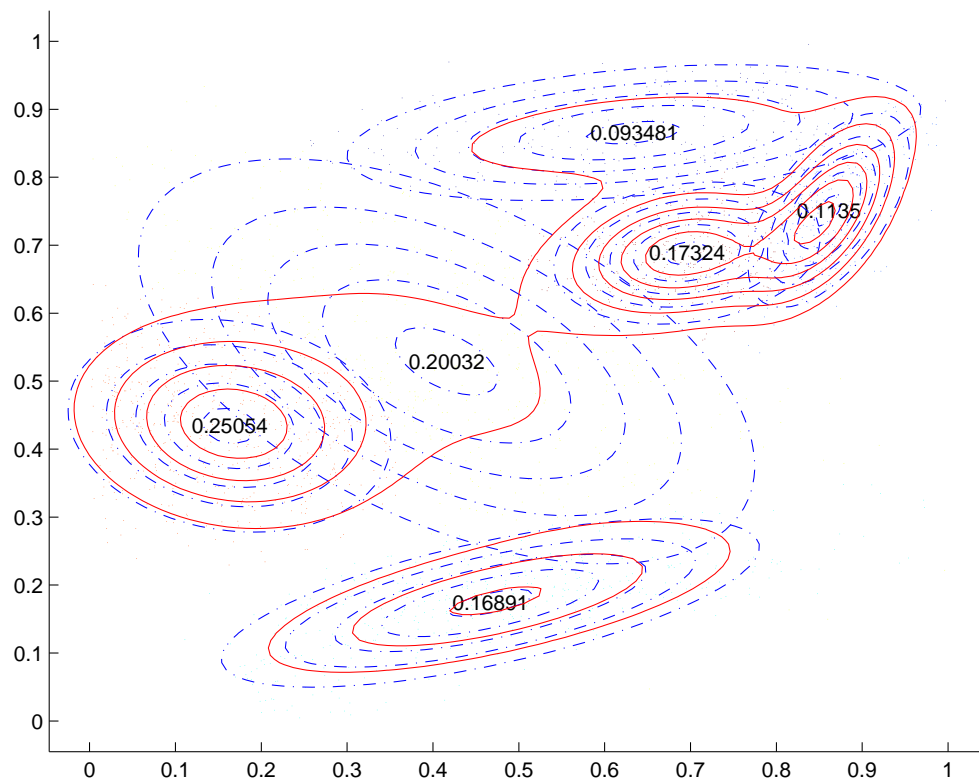
3. Use `EM.m` to fit the distribution of points in `second345` as a mixture of six Gaussian distributions. Plot the resulting distribution, and the resulting components, using the routine `plotmixture2D.m` that we have provided. You can call `plotmixture2D.m`, e.g., as follows:

```
>> plotmixture2D(second345,'evalmixtureX',params)
```

where `params` specify the 6-component mixture model. The blue lines in the contour plot correspond to each of the component Gaussians in the mixture and the red lines provide the contours of the 6-component mixture model. At the center of each Gaussian component, you also see the numerical value of the mixing proportion allocated to that component.


Turn in the contour plot (no need to turn in the 3D plot).

**Answer:** The resulting mixture might be (you might converge to a different mixture):



**Scoring:** 1 point

It is also interesting to note the improvement in log likelihood over the progress of EM. Note that we actually got extremely close to the solution after about 150 iterations:



The resulting log-likelihood should be a bit over 1000, and EM should converge after a few hundred iterations.

4. Rerun EM on the same data-set a few times, looking at the resulting plots and the log-likelihoods. Explain why you do not always get close to the same distribution, and why you sometimes do get the exact same distribution (and not just a similar

one).

**Answer:** The EM algorithm converges to a local optimum. Initializing the mixture differently might lead to convergence to a very different local optimum. However, once we are close enough to a local optimum, we will always end up in it, and so we get the same mixture with different starting points.

In some way, the space of mixtures is partitioned to parts of the space that all lead to the same local optimum. This is similar to slopes of a mountain— think of partitioning a mountain range according to which peak you would get to if you just went up. If you started in different places, you might end up on different peaks. But for each peak, there is a large area which leads to it.

**Scoring:** 3 points

Plot the log-likelihood at each EM iteration (this is returned as the second output value of `EM.m`). Notice (sometimes) the step-like behavior of the log-likelihood: the log-likelihood is almost unchanged throughout many iterations, and then increases dramatically after a few additional iterations. You can use the history of mixture parameters (returned as the third output value of `EM.m`) to see how the mixture distribution itself changes in relation to the log-likelihood. It takes time to resolve how to best allocate the mixture components, especially if (some of) the Gaussian components are only slightly different from each other in the current model.

## Initialization

The EM algorithm describes how to *update* the mixture parameters in order to improve the log-likelihood of the data. This search process must start somewhere. An optional parameter to `EM.m` can be used to control the way the mixture parameters are initialized.

5. The provided routine `initSame.m` sets all mixture components to have zero mean, unit (spherical) covariances, and equal mixing proportions. Try using it to fit `second345`. How many iterations does it take for EM to converge? Why? Describe what happens at each iteration (you might want to set the `histfreq` parameter to one, and plot the mixture at each iteration).

**Answer:** It takes only two iterations. In the E step of the first iteration, every point is soft-assigned with equal weights to all the components (since they are all identical). In the M step, all components will be updated in the same way, since they all have the same weighted set of points associated with them. In the second iteration, the E step will not change the assignments, and so the M step will not change the parameters, and so we have converged to a stationary point.

Note that this stationary point is very different from the local optima we usually converge to. This is an *unstable* stationary point— any slight perturbation of the parameters will cause us to start getting farther away from it (if we run EM).

**Scoring:** 2 points

6. What would happen if the means and covariances were all equal, but the initial mixing proportions were chosen randomly? If the means and mixing proportions were equal but the initial covariance matrices were chosen randomly?

**Answer:** If only the mixing proportions were different, we would still get the same behavior. All the points will have weights equal to the mixing proportions, and all the components will remain identical.

If the covariances were different, farther away points will get assigned more heavily to wider components. The components would then start changing to reflect that, and EM would proceed normally.

**Scoring:** 3 points

The default initialization routine used by `EM.m` is `initRandPoints.m`, which initializes the means to random data points, and sets the covariance matrices to diagonal matrices with the same overall variance parameter computed from the overall covariance of the data.

### Data anomalies and regularization

The actual data was given as integers in the range 0..100 in each coordinate. The matrix `pentrain8`, and its derivatives `pentrain4` and `second345` represent the original data scaled to [0,1] with uniform noise added.

The matrix `cleanpentrain8` contains the 0..100 integer valued data.

7. Try estimating a 5-component mixture model to the first point `cleanpentrain8(:,1:2)`. You can also try the rest of the time points `cleanpentrain8(:,2t-1:2t)`. Why does EM fail? What happens? Hint: plot the distribution using

```
>> plot(cleanpentrain8(:,1),cleanpentrain8(:,2),'.')
```

**Answer:** EM fails since one of the components quickly converges to a degenerate Gaussian. Since about a fifth of the points have the first coordinate equal to exactly zero, we try to fit these points with a Gaussian with zero variance in the first coordinate. But the 2-dimensional density of such Gaussian is infinite (its covariance matrix is not invertible).

**Scoring:** 3 points

8. Would we be able to resolve the above problem by introducing regularization to the M-step? Regularization would modify how we update the covariance matrix of each

of the component Gaussians in the M-step. Let  $\hat{n}_i$  is the number of “soft” counts assigned to component  $i$  in the E-step and  $\hat{\Sigma}_i$  the corresponding (unregularized) estimate of the covariance matrix. We can now imagine having an additional set of  $n'$  points with covariance  $S$  (the prior covariance matrix) also assigned to component  $i$ . These additional points would change the covariance update according to

$$\hat{\Sigma}'_i = \left(\frac{\hat{n}_i}{\hat{n}_i + n'}\right) \hat{\Sigma}_i + \left(\frac{n'}{\hat{n}_i + n'}\right) S \quad (1)$$

where  $\hat{\Sigma}'_i$  is the new regularized estimate of the covariance matrix for component  $i$ . (We do not expect you to try this out but answer the question based on what you think would happen).

**Answer:** This would help, since even if  $\hat{\Sigma}_i$  is degenerate, since  $S$  is not degenerate, the resulting  $\hat{\Sigma}'_i$  will not be degenerate. By using this regularization we are avoiding getting close to problematic areas of the space of mixture parameters.

**Scoring:** 1 point

## Unsupervised Learning Using Mixture Models


We continue to use the trajectories in an unsupervised manner, i.e., not paying attention to the accompanying labels. We can use the mixture models to try to partition the set of input samples (trajectories) into *clusters*. We will also assess how the emerging cluster identities will correlate with the actual labels.

The first question we have to answer is how many Gaussian components we should use in the mixture model. How could we decide the issue? We could proceed similarly to the case of setting the kernel width parameter in non-parametric density estimation, i.e., via cross-validation. The problem here, however, is that estimating any single mixture model with say 5 components already takes a fair amount of computing time and repeating this process  $n$  times ( $n$  is the number of training examples) does not appear very attractive. Instead, we will use a simple approximate criterion or *Bayesian information criterion* to find the appropriate number of clusters. This involves adding a complexity penalty to the log-data likelihoods. In other words, if  $LL(k)$  is the log-likelihood that our  $k$ -component mixture model assigns to the training data after training, we find  $k$  that maximizes

$$score(k) = LL(k) - \frac{1}{2} D_k \log(n) \quad (2)$$

where  $D_k$  denotes the number of parameters in the  $k$ -component mixture model. For a Gaussian mixture mode,  $D_k = k * (d + d * (d + 1)/2) + k - 1$ , where  $d$  is the dimensionality of the examples (e.g., 2). Note that the penalty is a function of the number of parameters and the number of training examples  $n$ .

9. Use different numbers of mixture components with `second345` and select the “optimum” number of components. Use `mixtureModSel.m` to plot the scores for  $k = 1, \dots, 8$ . Return the plot. Your choice of  $k$  may change if you re-estimate the mixture models (make sure you understand why).



**Scoring:** 1 point

Note that there may be more clusters than there are different types of digits. Some of the digits may have sub-types or our assumption that the clusters look like Gaussians may be invalid and we need more than one Gaussian to account for any single real cluster in the data.

Now, given the “best” mixture model, we would like to correlate the cluster identities with the available labels. We can do this to find out how well we could capture relevant structure in the data without actually knowing anything about the labels.

10. The labels are provided in `labels345`. Write a simple MATLAB routine to get hard assignments of training examples to clusters. You’ll probably want to use `mixtureScaledPiX` for this purpose. Use `correlate.m` to correlate the cluster identities with the available labels `labels345`. Return your MATLAB code and the correlation matrix. Briefly discuss whether the approach was successful, i.e., whether we could associate the clusters with specific labels.

**Answer:** Using the functions:

```
function c = hard(X, compParams)
```



```
scaledpix = mixtureScaledPiX(X,compParams);  
[dummy,c] = max(scaledpix,[],2);
```

We get:

```
>> c = hard(second345,mm);  
>> correlate(c,labels345)  
ans =  
    0    15   389  
   26   140   184  
    0   378    2  
    0   228    8  
  284    3    2  
  352   16    2  
   57    0   133
```

From the above correlation matrix, we observe that most components do correspond to separate digits. Up to a fairly low number of “errors”, the digit “3” was almost completely separated, while the digits “4” and “5” have one component in common.

The main problem is that by just looking at the seven components, we could not know that the data actually consists of a mixture of three digits. We would have suspected there are seven distinct classes in the data.

**Scoring:** *2 points*

## Using Mixture Models for Classification

We will now try to use the Gaussian mixture models in a classification setting. We can estimate class-conditional density models (mixture models) separately for each class and classify any new test example on the basis of which class-conditional model assigns the highest likelihood to the example.

11. Using the provided `fitAllDigits.m`, fit a 3-component Gaussian mixture model to each of the digits in `pentrain8,trainlabels`.

**Answer:**

```
>> dcp = fitAllDigits(pentrain8, trainlabels,3);
```

**Scoring:** *No score*

12. Complete the routine `classify.m` that takes as input the examples to be classified and a cell-array of mixture parameters (as returned by `fitAllDigits.m`) to return the maximum likelihood assignment of labels to the examples. Provide the code for `classify.m`.

**Answer:**

```
function c = classify(X, classes_compParams)
    m = length(classes_compParams);
    n = size(X,1);
    p = zeros(n,m);
    for i=1:m
        p(:,i) = evalmixtureX(X, classes_compParams{i});
    end
    [dummy, c] = max(p, [], 2);
```

**Scoring:** 5 points

13. Classify the `pentrain8` and `pentest8`, and use `correlate.m` to compare the results to the true labeling in `trainlabels` and `testlabels`. Turn in the two correlation matrices.

**Answer:**

```
>> trainc = classify(pentrain8, dcp);
>> testc = classify(pentest8, dcp);
>> correlate(trainc, trainlabels)
ans =
    0    775     0     1     0     0     0     5     0     0
    0     2    779     2     0     0     0     0     0     0
    0     0     1    714     0     0     0     0     0     0
    0     0     0     0    780     0     3     0     0     0
    0     1     0     0     0    720     1     0     0     0
    0     0     0     0     0     0    715     0     0     0
    0     0     0     0     0     0     0    772     0     0
    2     0     0     1     0     0     1     1    719     0
    1     1     0     1     0     0     0     0     0    719
   777     0     0     0     0     0     0     0     0     0

>> correlate(testc, testlabels)
ans =
    0    357     4     3     0     6     3    11     0     2
    0     5    359     0     0     0     0     0     0     0
```

0	0	0	330	0	2	0	0	0	0
0	1	0	0	362	1	0	0	0	0
0	0	0	0	0	289	0	0	0	0
1	0	0	0	0	0	325	0	0	0
0	0	1	0	0	0	0	330	0	2
16	0	0	2	0	5	8	0	336	1
0	1	0	1	2	32	0	23	0	331
346	0	0	0	0	0	0	0	0	0

**Scoring:** 1 point

For comparison, you probably want to repeat the above for 1-component Gaussian mixtures. You could also use the model selection routine we have provided to find a different number of mixture components for each class-conditional density. Increasing the number of components per class would eventually force you to use regularized estimates (see above).

## Markov and Hidden Markov Models

In this problem, we will use Markov and Hidden Markov models to identify the language of written sentences. For simplicity our representation of text will include only 27 symbols—the 26 letters of the Latin alphabet, and the space symbol. Any accented letter is represented as a non-accented letter, none-Latin letters are converted to their closest Latin letters, and punctuation is removed. This representation naturally loses quite a bit of information compared to the original ASCII text. This 'handicap' is in part intentional so that the classification task would be a bit more challenging.

Most of the MATLAB code you will need here will be given. You will find the following routines useful (here and perhaps in some of your projects as well):

`readlines.m` Reads a named text file, returning a cell array of the lines in the file. To get line `i` of cell-array `lines` returned from, e.g., `lines = readlines('cnn.eng')`, use `lines{i}`.

`text2stream.m` Converts a string (a line of text) into a row vector of numbers in the range  $\{1, \dots, 27\}$ , according to the representation discussed above. So, for example, `numberline = text2stream(lines1)` would convert the first line of text from `lines` into a row vector of numbers. The conversion of the full output of `readlines` would have to be done line by line.

`count.m` Given text in a row vector representation and a width  $k$ , the function computes the count of all  $k$ -grams in the array. In other words, the function returns a  $k$ -dimensional array representing the number of times each configuration of  $k$  successive letters occurs in the text.

`totalcount.m` This function allows you to compute the accumulated counts from each of the lines of text returned by `readlines`. Use this function to find the training counts for the different languages.

## Language identification using Markov models

Here we will construct a language classifier by using Markov models as class-conditional distributions. In other words, we will separately train a Markov model to represent each of the chosen languages: English, Spanish, Italian and German. The training data is given in the files `cnn.eng`, `cnn.spa`, `cnn.ita`, `cnn.ger`, which contain several news articles (same articles in different languages), one article per line.

We will first try a simple independent (zeroth-order Markov) model. Under this model, each successive symbol in text is chosen independently of other symbols. The language is in this case identified based only on its letter frequencies.

1. Write a function `naiveLL(stream,count1)` which takes a 1-count (frequency of letters returned by `count.m`) and evaluates the log-likelihood of the text stream (row vector of numbers) under the independent (zeroth-order Markov) model.

**Answer:**

```
function ll = naiveLL(stream,c1)
    lp = log(c1/sum(c1));
    ll = sum(lp(stream));
```

**Scoring:** 5 points

Extract the total 1-counts from the language training sets described above. Before proceeding, let's quickly check your function `naiveLL`. If you evaluate the log-likelihood of 'This is an example sentence' using the English 1-counts from `cnn.eng`, you'll get -76.5690, while the Spanish log-likelihood of the same sentence is -77.2706.

2. Write a short function `naiveC` which takes a stream, and several 1-counts corresponding to different languages, and finds the maximum-likelihood language for the stream. You could assume, e.g., that the 1-counts are stored in an array, where each column corresponds to a specific language. The format of the labels should be in correspondence with the `test_labels` described below.

**Answer:**

```
function [c,ll] = naiveC(stream,counts1)
    for i=1:size(counts1,2)
        ll(i) = naiveLL(stream,counts1(:,i));
    end
    [dummy,c]=max(ll);
```

**Scoring:** *3 points*

The files `song.eng`, `song.spa`, `song.ita`, `song.ger` contain additional text in the four languages. We will use these as the test set. For your convenience, we have provided you with script `generate_test.m` :

```
test_sentences = [ readlines('song.eng') ; ...
    readlines('song.ger') ; ...
    readlines('song.spa') ; ...
    readlines('song.ita') ] ;
test_labels = [ ones(17,1) ; ones(17,1)*2 ; ones(17,1)*3 ; ones(17,1)*4 ]
```

In order to study the performance of the classifier as a function of the length of test strings, we will classify all prefixes of the lines in the test files. The provided routine `testC.m` calculates the success probability of the classification, for each prefix length, over all the streams or strings in a given cell-array. You can call this function, e.g., as follows

```
successprobs = testC(test_sentences,test_labels,'naiveC',count1s)}
```


where `count1s` provides the array of training counts that your function `naiveC` should accept as an input.

3. Plot the success probability as a function of the length of the string. What is the approximate number of symbols that we need to correctly assign new piece of text to one of the four languages?

**Answer:** We carry out the following operations:

```
>> eng = readlines('languages/cnn.eng');
>> ita = readlines('languages/cnn.ita');
>> ger = readlines('languages/cnn.ger');
>> spa = readlines('languages/cnn.spa');
>> counts1 = [totalcount(eng,1) , totalcount(ger,1) , totalcount(spa,1) , totalcount(ita,1)];
>> p = testC(test_sentences,test_labels,'naiveC',counts1);
>> plot(p)
```

Resulting in the plot:



From the plot we see that approximately 100 characters are needed for reliable language classification.

**Scoring:** *1 point*

In order to incorporate second order statistics, we will now move on to modeling the languages with first-order Markov models.

4. Write a function `markovLL(stream,count2)` which returns the log-likelihood of a stream under a first-order Markov model of the language with the specified 2-count. For the initial probabilities, you can use 1-counts calculated from the 2-counts.

**Answer:**

```
function ll = markovLL(stream,c2)
    c1 = sum(c2,2); % 1-count calculated from 2-count
    p0 = c1/sum(c1); % Initial probs
    p1 = op(c2, './',c1); % Transition probs
    lp0 = log(p0);
    lp1 = log(p1);
    ll = lp0(stream(1));
    for i=2:length(stream)
        ll = ll+lp1(stream(i-1),stream(i));
    end
```

**Scoring:** *5 points*

Quick check: The English log-likelihood of 'This is an example sentence' is -63.0643, while its Spanish log-likelihood is -65.4878. We are again assuming that you are using the training sets described above to extract the 2-counts for the different languages.

5. Write a corresponding function `markovC.m` that classifies a stream based on Markov models for various languages, specified by their 2-counts.

**Answer:**

```
function [c,ll] = markovC(stream,counts2,zamir)
    if nargin<3, zamir=0; end;
    if ischar(stream), stream=preproc(stream); end;
    for i=1:size(counts2,3)
        ll(i) = markovLL(stream,counts2(:,:,i)+zamir);
    end
    [dummy,c]=max(ll);
```

The above function can also incorporate a pseudo-count (as required later on) and can deal with either streams or strings as inputs. These features were not required for a solution to be complete.

**Scoring:** *1 point*

6. Try to classify the sentence 'Why is this an abnormal English sentence'. What is its likelihood under a Markov model for each of the languages ? Which language does it get classified as ? Why does it not get classified as English ?

**Answer:** After calculating the two-counts:

```
>> counts2 = cat(3,totalcount(eng,2) , totalcount(ger,2) , totalcount(spa,2) , tota
```

We can calculate:

```
>> [c,ll] = markovC(text2stream('Why is his is an abnormal English sentence'),count
c =      3
ll =      -Inf      -Inf -115.9894      -Inf
```

(We also get a few warnings).

The sentence gets classified as Spanish with a log-likelihood of -116. In all other languages, including English, it has infinitely negative log-likelihood. This is because the letter pair "bn" does not appear in the training text, and so no sentence with this letter combination will be generated by the learned model. The probability of generating the above sentence under our English model is thus zero, and its log-likelihood is negative infinity.

**Scoring:** *3 points*

When estimating discrete probability models, we often need to regularized the parameter estimates to avoid concluding that any configuration (e.g., two successive characters) have

probability zero unless such the configuration appeared in the limited training set. Let  $\{\theta_i\}$ , where  $\sum_{i=1}^m \theta_i = 1$ , define a discrete probability distribution over  $m$  elements  $i \in \{1, \dots, m\}$ . For the purposes of estimation, we treat  $\theta_i$  as parameters. Given now a training set summarized in terms of the number of occurrences of each element  $i$ , i.e.,  $\hat{n}_i$ , the maximum likelihood estimate of  $\{\theta_i\}$  would be

$$\hat{\theta}_i = \frac{\hat{n}_i}{\sum_{j=1}^m \hat{n}_j} \quad (3)$$

This is zero for all elements  $i$  that did not occur in the training set, i.e., when  $\hat{n}_i = 0$ . To avoid this problem, we introduce a prior distribution over the parameters  $\{\theta_i\}$ :

$$P(\theta) = \frac{1}{Z} \prod_{i=1}^m \theta_i^{\alpha_i} \quad (4)$$

where  $Z$  is a normalization constant and  $\alpha_i \geq 0$ 's are known as *pseudo-counts*. This prior, known as a *Dirichlet* distribution, assigns the highest probability to

$$\tilde{\theta}_i = \frac{\alpha_i}{\sum_{j=1}^m \alpha_j} \quad (5)$$

Thus  $\alpha_i$ 's behave as if they were additional counts from some prior (imaginary) training set.

We can combine the prior with the maximum likelihood criterion by maximizing instead the penalized log-likelihood of the data (expressed here in terms of the training set counts  $\hat{n}_i$ ):

$$J(\theta) = \underbrace{\sum_{i=1}^m \hat{n}_i \log \theta_i}_{\text{log-probability of data}} + \underbrace{\log P(\theta)}_{\text{log-prior}} \quad (6)$$

$$= \sum_{i=1}^m \hat{n}_i \log \theta_i + \sum_{i=1}^m \alpha_i \log \theta_i + \text{constant} \quad (7)$$

$$= \sum_{i=1}^m (\hat{n}_i + \alpha_i) \log \theta_i + \text{constant} \quad (8)$$

The maximizing  $\{\theta_i\}$  is now

$$\hat{\theta}_i = \frac{\hat{n}_i + \alpha_i}{\sum_{j=1}^m (\hat{n}_j + \alpha_j)} \quad (9)$$

which will be non-zero whenever  $\alpha_i > 0$  for all  $i = 1, \dots, m$ . Setting  $\alpha_i = 1/m$  would correspond to having a single prior observation distributed uniformly among the possible elements  $i \in \{1, \dots, m\}$ . Setting  $\alpha_i = 1$ , on the other hand, would mean that we had  $m$  prior observations, observing each element  $i$  exactly once.



7. Add pseudocounts (one for each configuration) and reclassify the test sentence. What are the likelihoods now. Which language does the sentence get classified as ?

**Answer:** We now have:

```
>> [c,ll] = markovC(text2stream('Why is this is an abnormal English sentence'),counts2,1);
c =      1
ll = -106.1994 -123.9184 -114.7622 -128.4992
```


And the sentences gets correctly classified as English.

**Scoring:** 2 points

8. Use `testC.m` to test the performance of Markov-based classification (with the corrected counts) on the test set. Plot the correct classification probability as a function of the text length. Compare the classification performance to that of `naiveC.m`. (Turn in both plots).

**Answer:**

```
>> p = testC(test_sentences,test_labels,'markovC',counts2,1);
>> plot(p)
```



The performance using Markov models is much better— accurate performance is achieved with less than 40 characters.

**Scoring:** 1 point

## Hidden Markov Models

We will now turn to a slightly more interesting problem of language segmentation: given a mixed-language text, we would like to identify the segments written in different languages. Consider the following simple approach: we will first find the character statistics for each language by computing the 1-counts from the available training sets as before. We will then go through the new text character by character, classifying each character to the most likely language (language whose independent model assigns the highest likelihood to the character). In other words, we would use `naiveC` to classify each character in the new text. To incorporate higher-order statistics, we could train a Markov model for each language (as before), and again assign the characters in the text to the most likely language.

9. Why would we expect the resulting segmentation not to agree with the true segmentation? What would the resulting segmentation look like? What is the critical piece of information we are not using in this approach ?

**Answer:** We would classify each letter according to the language in which it is most frequent. And so, all 't's will get classified as English, while all 'a' will get classified as Spanish. Letters in the same word would get classified as different languages. This is of course unreasonable.

The crucial information we are not using is that adjacent characters tend to come from the same language.

**Scoring:** *4 points*

We would rather model the multi-lingual text with a hidden Markov model. For simplicity, we will focus on segmenting text written in only two languages.

10. Suggest how a hidden Markov model can solve the problem you identified above. Provide an annotated transition diagram (heavy lines for larger transition probabilities) and describe what the output probabilities should be capturing.

**Answer:** We can use the hidden variable to designate which language the characters come from. The transition probabilities for this hidden variable will encode the information that transitions between languages are much less frequent than remaining in the same language. The output probabilities given that we are in a specific language would correspond to the single-character statistics of that language.

**Scoring:** *4 points*

For some setting of the parameters, this HMM probably degenerates to the independent model. Make sure you understand when this happens.

Now, load the example text from the file `segment.dat`. The text, mixed German and Spanish, is given in the variable `gerspa` as numbers in the range 1..27 (as before). You can use the provided function `stream2text.m` to view the numbers as letters.

The routine `[hmm, ll] = esthmm(data, hmm)` uses EM to find the maximum likelihood parameters of a hidden Markov model, for a given output sequence(s). As discussed earlier, the EM algorithm has to start its search with some parameter setting. The `hmm` input argument to `esthmm` provides this initial parameter setting.

The routine `hmm = newhmm(numstates, numout)` creates random HMM parameters for an HMM with `numstates` hidden states, and `numout` output symbols. An optional third argument controls the “non-uniformity” of the parameters.

**Note:** The orientation of the matrices `Po` and `Pt` is different from those in recitation:  $Pt(i, j) = P(state_t = j | state_{t-1} = i)$  and  $Po(i, a) = P(O_t = a | state_t = i)$ .

With two different initializations, you will probably find two different HMMs. It is a good idea to run the EM algorithm multiple times, starting from slightly different initial settings. Every such run can potentially lead to a different result. Make sure you understand how you would choose among the alternative solutions.

11. Estimate a two-state hidden Markov model with two different types of initial settings of the state transition probabilities. First, set the transition probabilities close to uniform values. In the second approach, make the “self-transitions” quite a bit more likely than the transitions to different states. Which initialization leads to a better model?

**Answer:** We can compare the two estimated models by comparing their likelihood. We are seeking a maximum likelihood model, and so we will select the one with a higher likelihood.

```
>> randhmm = newhmm(2,27);
>> [hmm1,ll1] = esthmm(gerspa,randhmm);
>> ll1(end)
ans = -7.8976e+03
>> biasedhmm = newhmm(2,27);
>> biasedhmm.Pt = [0.9,0.1;0.1,0.9];
>> [hmm2,ll2] = esthmm(gerspa,biasedhmm);
>> ll2(end)
ans = -7.8115e+03
```

The second HMM has a higher log-likelihood, and thus will be selected.

We can also look at the resulting HMM parameters and notice that the second model better fits what we were looking for: the transition probabilities are very biased, the initial probabilities are almost deterministic, and the output probabilities resemble the language statistics we are already familiar with from the first part of the problem.

**Scoring:** 2 points


Now try to segment `gerspa` into German and Spanish, using a two-state hidden Markov model. You can then use the routine `viterbi(sequence,hmm)` to examine the most likely (maximum a posteriori probability or MAP) state sequence. The correct segmentation of `gerspa` is given in `gerspa_lang`.

12. Examine the difference between your segmentation and the correct one. Do the errors make sense?

**Answer:** There are sixty-four miss-labeled characters. Note that all of them occur on the boundary between language transitions. All the transitions were identified, but the exact location was sometimes missed. This is very reasonable.

**Scoring:** 1 point

13. Use the provided routine `hmmposteriors(sequence,hmm)` to calculate the per character posterior probabilities over the states. These are the  $\gamma_t(i)$  probabilities described in lectures ( $t$  gives the character position in text and  $i$  specifies the state). Plot these probabilities as a function of the character position in the text sequence and turn in the plot. You might want to re-scale the axis using `axis[0 3000 0 1.1]`.



**Scoring:** 1 point

14. Find the sequence of maximum a posteriori probability states. That is, for each time point  $t$ , find the state  $i$  with the maximum a posteriori probability  $P(state_t = i | observedseq)$ . Compare this state sequence with the maximum a posteriori state sequence. Are they different?

**Answer:** They differ in one position (position 631).

**Scoring:** 2 point

15. Give an example of a hidden Markov model, with two hidden states and two output symbols, and a length two output sequence, such that the MAP state sequence differs from the sequence of MAP states. Be sure to specify all the parameters of the HMM (these need not be the maximum likelihood parameters for this specific length two output sequence).

**Answer:**

$$\begin{aligned}P(X_0 = a) &= 0.6, P(X_0 = b) = 0.4 \\P(X_1 = a|X_0 = a) &= 0.5, P(X_1 = b|X_0 = a) = 0.5 \\P(X_1 = a|X_0 = b) &= 0, P(X_1 = b|X_0 = b) = 1\end{aligned}$$

. All output probabilities are uniform. For any output sequence, the MAP state sequence is  $b, b$  with posterior probability 0.4. But the MAP state of  $X_0$  is  $a$ , with posterior probability 0.6 (each of the state sequences  $a, a$  and  $a, b$  have posterior probability 0.3).

**Scoring:** 3 points

In the above example, we assumed no knowledge about the languages. We would like to improve the segmentation by incorporating some additional information about the languages such as single-character statistics.

16. What parameters of the HMM still need to be estimated ? Modify the routine `esthmm.m` accordingly. (Turn in the modifications)

**Answer:** The output probabilities are now known and fixed. We only search for the maximum likelihood initial and transition probabilities. We make one change to `esthmm.m`: delete the following line:

```
hmm.Po = op(Nsx, './', sum(Nsx,2));
```

It would also be appropriate, though not necessary, to delete the calculation of `Nsx`, since it is now unused.

**Scoring:** 3 points

17. (*Optional*) Use the single-character statistics you calculated from `cnn.spa` and `cnn.ger` to segment the text. What is the maximum likelihood estimate of the remaining parameters ? Compare the resulting most likely state sequence to the one you found without this prior information. Are they different ?

18. Hidden Markov models can also be useful for classification. Suggest how you would use HMMs in place of the Markov models described above to identify the language of any specific (single-language) document. Specifically, list the routines we have provided that you could make use of and describe which additional routines you would need.

**Answer:** For each language, we would train a HMM (using `esthmm.m`) on the corpus of texts for that languages. We would have to choose the number of hidden states in each of those HMMs. We can set a fixed number of hidden states (e.g. any number of hidden states above 27 would give us a model with a training accuracy at least as good as a first order Markov model, but good training accuracy might not lead to good generalization). We can also use cross-validation to choose the number of states for each language model.

After training the models, in order to classify a sentence, we would compare the likelihood of the query sentence under each of the models. This is given as the third output parameter of `hmmPosteriors.m`. However, all the other calculations in `hmmPosteriors.m` are unnecessary (and take significant time)– it is enough to run only the forward calculation as follows:

```
function logPx = hmmloglikelihood(x,hmm)
    loga = logalpha(x,hmm);
    logPx = logsumexp(loga(:,n));
```

**Scoring:** 5 points

19. (*Optional*) Use hidden Markov models, with a varying number of states, for the language classification task investigated in the first part of this problem. Using the same training set and test set as before, create a graph of the probability of correct classification as a function of the length of the text. Discuss how the results compare to using naive classification and first-order Markov models, and how changing the number of hidden states affects the results.