

10-601 Machine Learning, Fall 2009: Homework 4

Due: Wednesday, October 28th, 10:30 am

Instructions There are 2 questions on this assignment worth a total of 100 points. Please hand in a hard copy at the beginning of the class. Please print your code and attach it to the write-up. Refer to the webpage for policies regarding collaboration, due dates, and extensions.

1 Naive Bayes and Logistic Regression [35 points]

1.1 Redundant feature [15 points]

Consider the classification task with Boolean labels, Y , taking values T or F, and three Boolean features, X_1, X_2 and X_3 . The features X_1 and X_2 are conditionally independent given Y and while the feature X_3 is a copy of X_2 (i.e., $X_3=X_2$ always). The conditional probabilities are given by:

$$\begin{aligned}\mathbf{P}(X_1 = T | Y = T) &= p \\ \mathbf{P}(X_1 = T | Y = F) &= 1 - p \\ \mathbf{P}(X_2 = F | Y = T) &= q \\ \mathbf{P}(X_2 = F | Y = F) &= 1 - q \\ \mathbf{P}(Y = T) = \mathbf{P}(Y = F) &= 0.5\end{aligned}$$

You are given a test example with the feature values: $X_1 = T$ and $X_2 = X_3 = F$. We would like to classify this example by predicting the value of Y .

1. Show that if the Naive Bayes assumption is used, that is features are conditionally independent of each other given class labels Y , the decision rule for classifying the test example as $Y = T$ given its feature values is: $p \geq \frac{(1-q)^2}{q^2+(1-q)^2}$

★ **SOLUTION:** To classify the example as $Y = T$, the following should hold:

$$\mathbf{P}(Y = T | X_1 = T, X_2 = F, X_3 = F) \geq 0.5 \tag{1.1}$$

We first calculate this probability:

$$\mathbf{P}(Y = T | X_1 = T, X_2 = F, X_3 = F) = \frac{\mathbf{P}(X_1 = T, X_2 = F, X_3 = F | Y = T) \mathbf{P}(Y = T)}{\mathbf{P}(X_1 = T, X_2 = F, X_3 = F)} \tag{1.2}$$

$$= \frac{\mathbf{P}(X_1 = T | Y = T) \mathbf{P}(X_2 = F | Y = T) \mathbf{P}(X_3 = F | Y = T) \mathbf{P}(Y = T)}{\mathbf{P}(X_1 = T, X_2 = F, X_3 = F)} \tag{1.3}$$

$$= \frac{\mathbf{P}(X_1 = T | Y = T) \mathbf{P}(X_2 = F | Y = T) \mathbf{P}(X_3 = F | Y = T) \mathbf{P}(Y = T)}{\mathbf{P}(X_1 = T, X_2 = F, X_3 = F | Y = T) \mathbf{P}(Y = T) + \mathbf{P}(X_1 = T, X_2 = F, X_3 = F | Y = F) \mathbf{P}(Y = F)} \tag{1.4}$$

$$= \frac{p \times q \times q \times 0.5}{p \times q \times q \times 0.5 + (1 - p) \times (1 - q)^2 \times 0.5}$$

Eq. (1.2) follows from the Bayes' rule. In Eq. (1.3), we applied the Naive Bayes assumption, that is given Y all X_i 's are conditionally independent from each other.

For Eq. (1.1) to hold the following must hold:

$$\begin{aligned}
& \frac{0.5pq^2}{0.5pq^2 + 0.5(1-p)(1-q)^2} \geq 0.5 \\
& \Leftrightarrow \frac{pq^2}{pq^2 + (1-p)(1-q)^2} \geq 0.5 \\
& \Leftrightarrow pq^2 \geq 0.5(pq^2 + (1-p)(1-q)^2) \\
& \Leftrightarrow 0.5pq^2 \geq 0.5(1-p)(1-q)^2 \\
& \Leftrightarrow pq^2 \geq (1-q)^2 - p(1-q)^2 \\
& \Leftrightarrow pq^2 + p(1-q)^2 \geq (1-q)^2 \\
& \Leftrightarrow p(q^2 + (1-q)^2) \geq (1-q)^2 \\
& \Leftrightarrow p \geq \frac{(1-q)^2}{q^2 + (1-p)^2}
\end{aligned}$$

For the p, q values where this condition holds, the test example will be classified as $Y = T$.

2. What would be the optimal decision rule? (The optimal decision rule is the rule when we do not use the Naive Bayes assumption, but use only the given conditional independence assumptions.) Express this decision rule in terms of p and q .

★ **SOLUTION:** Similarly to classify the example as $Y = T$ condition stated in (1.1) should hold. This time in calculating the probability we will use the Naive Bayes assumption:

$$\begin{aligned}
\mathbf{P}(Y = T | X_1 = T, X_2 = F, X_3 = F) &= \frac{\mathbf{P}(X_1 = T, X_2 = F, X_3 = F | Y = T) \mathbf{P}(Y = T)}{\mathbf{P}(X_1 = T, X_2 = F, X_3 = F)} \\
&= \frac{\mathbf{P}(X_1 = T | Y = T) \mathbf{P}(X_2 = F, X_3 = F | Y = T) \mathbf{P}(Y = T)}{\mathbf{P}(X_1 = T, X_2 = F, X_3 = F)} \tag{1.5}
\end{aligned}$$

$$\begin{aligned}
&= \frac{\mathbf{P}(X_1 = T | Y = T) \mathbf{P}(X_2 = F | Y = T) \mathbf{P}(Y = T)}{\mathbf{P}(X_1 = T, X_2 = F)} \tag{1.6} \\
&= \frac{\mathbf{P}(X_1 = T | Y = T) \mathbf{P}(X_2 = F | Y = T) \mathbf{P}(Y = T)}{\mathbf{P}(X_1 = T, X_2 = F | Y = T) \mathbf{P}(Y = T) + \mathbf{P}(X_1 = T, X_2 = F | Y = F) \mathbf{P}(Y = F)} \\
&= \frac{p \times q \times 0.5}{p \times q \times 0.5 + (1-p) \times (1-q) \times 0.5}
\end{aligned}$$

The difference of the above calculation from the previous is in moving from Eq. (1.5) to Eq. (1.6). In arriving Eq. (1.6), this time we used the fact that the joint probabilities of X_2 and X_3 of having some value is basically the same as the probability that of X_2 having that value since X_3 is a copy of X_2 . In order to arrive Eq. (1.6), we used the following: $\mathbf{P}(X_1 = T, X_2 = F, X_3 = F) = \mathbf{P}(X_1 = T, X_2 = F)$ and also $\mathbf{P}(X_2 = F, X_3 = F | Y = T) = \frac{\mathbf{P}(X_2 = F, X_3 = F)}{\mathbf{P}(Y = T)} = \frac{\mathbf{P}(X_2 = F)}{\mathbf{P}(Y = T)} = \mathbf{P}(X_2 = F | Y = T)$.

For a positive classification the following should hold:

$$\begin{aligned}
 & \frac{0.5pq}{0.5pq + 0.5(1-p)(1-q)} \geq 0.5 \\
 \Leftrightarrow & \frac{pq}{pq + (1-p)(1-q)} \geq 0.5 \\
 \Leftrightarrow & pq \geq 0.5pq + 0.5(1-p)(1-q) \\
 \Leftrightarrow & 0.5pq \geq 0.5(1-p)(1-q) \\
 \Leftrightarrow & pq \geq 1-p-q+pq \\
 \Leftrightarrow & pq - pq + p \geq 1-q \\
 \Leftrightarrow & p \geq 1-q
 \end{aligned}$$

3. Plot the two decision rules you derived in 1.1.1 and 1.1.2. The x-axis should be q and y-axis should be p both varying in $[0,1]$. Indicate for which regions of the graph the test example is classified as $Y = T$ and $Y = F$. Show on the graph (shade and label the areas) where the Naive Bayes decision rule makes mistakes relative to the optimal decision rule. See Fig. 1 Naive Bayes decision rule makes mistakes in the shaded areas.

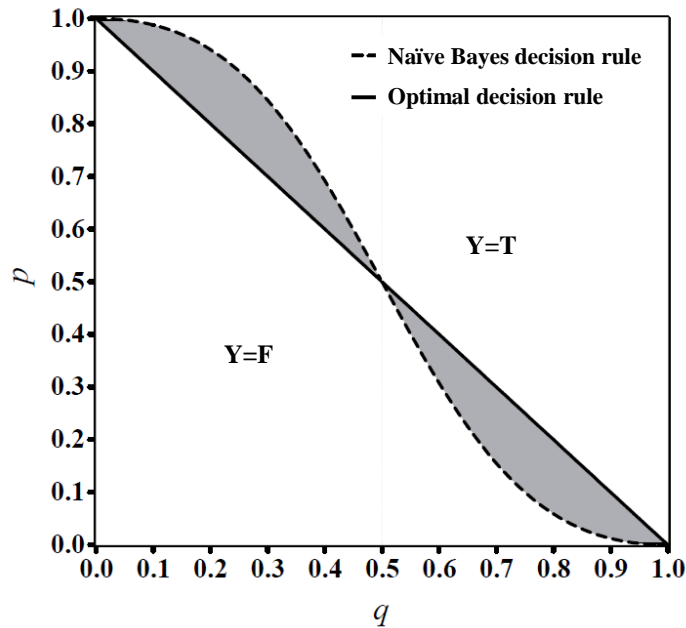


Figure 1: The two decision rules derived in 1.1.1 and 1.1.2.

1.2 Equivalence of NB and LR [10 points]

In section 3 of the Mitchell et al. reading, <http://www.cs.cmu.edu/~7Etom/mlbook/NBayesLogReg.pdf>, it is shown that if Y is Boolean and $X = \langle X_1 \dots X_n \rangle$ is a vector of continuous variables, the form of $\mathbf{P}(Y|X)$

entailed by the assumptions of a Gaussian Naive Bayes (GNB) classifier is precisely the form used by Logistic Regression (LR) with appropriate parameters W . In particular:

$$\mathbf{P}(Y = 1 | X) = \frac{1}{1 + \exp(w_0 + \sum_{i=1}^n w_i X_i)}$$

and

$$\mathbf{P}(Y = 0 | X) = \frac{\exp(w_0 + \sum_{i=1}^n w_i X_i)}{1 + \exp(w_0 + \sum_{i=1}^n w_i X_i)}$$

Now consider the case where Y is Boolean and $X = \langle X_1 \dots X_n \rangle$ is a vector of binary variables instead of continuous variables. Prove for this case also $\mathbf{P}(Y | X)$ follows the same form (and hence that Logistic Regression is also the discriminative counterpart to a Naive Bayes generative classifier over Boolean features).

Hints:

- Using simple notation will help you in the derivation. Since the X_i are Boolean variables, only one parameter is needed to define $\mathbf{P}(X_i | Y = y_k)$. Define $\theta_{i1} \equiv \mathbf{P}(X_i | Y = 1)$, in which case $\mathbf{P}(X_i = 0 | Y = 1) = 1 - \theta_{i1}$. Similarly, use θ_{i0} to denote $\mathbf{P}(X_i = 1 | Y = 0)$.
- Notice with the above notation you can represent $\mathbf{P}(X_i | Y = 1)$ as follows

$$\mathbf{P}(X_i | Y = 1) = (\theta_{i1})^{X_i} (1 - \theta_{i1})^{(1-X_i)}$$

Note when $X_i = 1$ the second term is equal to 1 because its exponent is zero. Similarly, when $X_i = 0$ the first term is equal to 1 because its exponent is zero.

★ **SOLUTION:**

$$\begin{aligned} \mathbf{P}(Y = 1 | X) &= \frac{\mathbf{P}(X | Y = 1) \mathbf{P}(Y = 1)}{\mathbf{P}(X | Y = 1) \mathbf{P}(Y = 1) + \mathbf{P}(X | Y = 0) \mathbf{P}(Y = 0)} \\ &= \frac{1}{1 + \frac{\mathbf{P}(X | Y = 0) \mathbf{P}(Y = 0)}{\mathbf{P}(X | Y = 1) \mathbf{P}(Y = 1)}} \\ &= \frac{1}{1 + \exp \left[\ln \left(\frac{\mathbf{P}(X | Y = 0) \mathbf{P}(Y = 0)}{\mathbf{P}(X | Y = 1) \mathbf{P}(Y = 1)} \right) \right]} \\ &= \frac{1}{1 + \exp \left[\ln \left(\frac{1-\pi}{\pi} \right) + \sum_i \ln \left(\frac{\mathbf{P}(X_i | Y = 1)}{\mathbf{P}(X_i | Y = 0)} \right) \right]} \\ &= \frac{1}{1 + \exp \left[\ln \left(\frac{1-\pi}{\pi} \right) + \sum_i \ln \left(\frac{(\theta_{i0})^{X_i} (1-\theta_{i0})^{(1-X_i)}}{(\theta_{i1})^{X_i} (1-\theta_{i1})^{(1-X_i)}} \right) \right]} \\ &= \frac{1}{1 + \exp \left[\ln \left(\frac{1-\pi}{\pi} \right) + \sum_i \ln((\theta_{i0})^{X_i}) + \ln((1-\theta_{i0})^{(1-X_i)}) - \ln((\theta_{i1})^{X_i}) - \ln((1-\theta_{i1})^{(1-X_i)}) \right]} \\ &= \frac{1}{1 + \exp \left[\ln \left(\frac{1-\pi}{\pi} \right) + \sum_i X_i \ln(\theta_{i0}) + (1-X_i) \ln(1-\theta_{i0}) - X_i \ln(\theta_{i1}) - (1-X_i) \ln(1-\theta_{i1}) \right]} \\ &= \frac{1}{1 + \exp \left[\ln \left(\frac{1-\pi}{\pi} \right) + \sum_i X_i (\ln(\theta_{i0}) - \ln(1-\theta_{i0}) - \ln(\theta_{i1}) + \ln(1-\theta_{i1})) + \ln(1-\theta_{i0}) - \ln(1-\theta_{i1}) \right]} \\ &= \frac{1}{1 + \exp [w_0 + \sum_i w_i X_i]} \end{aligned}$$

where

$$w_0 = \ln\left(\frac{1-\pi}{\pi}\right) + \sum_i \left(\frac{1-\theta_{i0}}{1-\theta_{i1}}\right)$$

$$w_1 = \ln(\theta_{i0}) - \ln(1-\theta_{i0}) - \ln(\theta_{i1}) + \ln(1-\theta_{i1}) = \ln\frac{\theta_{i0}}{\theta_{i1}} + \ln\left(\frac{1-\theta_{i1}}{1-\theta_{i0}}\right)$$

1.3 Relaxing the conditional independence assumption [10 points]

To capture interactions between features, the Logistic Regression model can be supplemented with extra terms. For example, a term can be added to capture a dependency between X_1 and X_2 :

$$\mathbf{P}(Y = 1 | X) = \frac{1}{1 + \exp(w_0 + w_{1,2}X_1X_2 + \sum_{i=1}^n w_iX_i)}$$

Similarly, the conditional independence assumptions made by Naive Bayes can be relaxed so that X_1 and X_2 are not assumed to be conditionally independent. In this case, we can write:

$$\mathbf{P}(Y | X) = \frac{\mathbf{P}(Y) \mathbf{P}(X_1, X_2 | Y) \prod_{i=3}^n \mathbf{P}(X_i | Y)}{\mathbf{P}(X)}$$

Prove that for this case, that $\mathbf{P}(Y | X)$ follows the same form as the logistic regression model supplemented with the extra term that captures the dependency between X_1 and X_2 (and hence that the supplemented Logistic Regression model is the discriminative counterpart to this generative classifier).

Hints:

- Using simple notation will help here as well. You need more parameters than before to define $\mathbf{P}(X_1, X_2 | Y)$. Define $\beta_{ijk} \equiv \mathbf{P}(X_1 = i, X_2 = j | Y = k)$.
- The above notation can be used to represent $\mathbf{P}(X_1, X_2 | Y = k)$ as follows

$$\mathbf{P}(X_1, X_2 | Y = k) = (\beta_{11k})^{X_1X_2} (\beta_{10k})^{X_1(1-X_2)} (\beta_{01k})^{(1-X_1)X_2} (\beta_{00k})^{(1-X_1)(1-X_2)}$$

★ **SOLUTION:**

$$\begin{aligned} \mathbf{P}(Y = 1 | X) &= \frac{\prod_{i=3}^n \mathbf{P}(X_i | Y = 1) \mathbf{P}(X_1, X_2 | Y = 1) \mathbf{P}(Y = 1)}{\prod_{i=3}^n \mathbf{P}(X_i | Y = 1) \mathbf{P}(X_1, X_2 | Y = 1) \mathbf{P}(Y = 1) + \prod_{i=3}^n \mathbf{P}(X_i | Y = 0) \mathbf{P}(X_1, X_2 | Y = 0) \mathbf{P}(Y = 0)} \\ &= \frac{1}{1 + \frac{\prod_{i=3}^n \mathbf{P}(X_i | Y=0) \mathbf{P}(X_1, X_2 | Y=0) \mathbf{P}(Y=0)}{\prod_{i=3}^n \mathbf{P}(X_i | Y=1) \mathbf{P}(X_1, X_2 | Y=1) \mathbf{P}(Y=1)}} \\ &= \frac{1}{1 + \exp\left[\ln\left(\frac{\prod_{i=3}^n \mathbf{P}(X_i | Y=0) \mathbf{P}(X_1, X_2 | Y=0) \mathbf{P}(Y=0)}{\prod_{i=3}^n \mathbf{P}(X_i | Y=1) \mathbf{P}(X_1, X_2 | Y=1) \mathbf{P}(Y=1)}\right)\right]} \\ &= \frac{1}{1 + \exp\left[\ln\left(\frac{1-\pi}{\pi}\right) + \sum_{i=3}^n \left[\ln\left(\frac{\mathbf{P}(X_i | Y=0)}{\mathbf{P}(X_i | Y=1)}\right)\right] + \ln\left(\frac{\mathbf{P}(X_1, X_2 | Y=0)}{\mathbf{P}(X_1, X_2 | Y=1)}\right)\right]} \\ &= \frac{1}{1 + \exp\left[\ln\left(\frac{1-\pi}{\pi}\right) + \sum_{i=3}^n \left[\ln\left(\frac{(\theta_{i0})^{X_i(1-\theta_{i0})^{(1-X_i)}}}{(\theta_{i1})^{X_i(1-\theta_{i1})^{(1-X_i)}}}\right)\right] + \ln\left(\frac{(\beta_{110})^{X_1X_2}(\beta_{100})^{X_1(1-X_2)}(\beta_{010})^{(1-X_1)X_2} + (\beta_{000})^{(1-X_1)(1-X_2)}}{(\beta_{111})^{X_1X_2}(\beta_{101})^{X_1(1-X_2)}(\beta_{011})^{(1-X_1)X_2} + (\beta_{001})^{(1-X_1)(1-X_2)}}\right)\right]} \\ &= \frac{1}{1 + \exp\left[\ln\left(\frac{1-\pi}{\pi}\right) + \sum_{i=3}^n \left[X_i \left(\ln\frac{\theta_{i0}}{\theta_{i1}} + \ln\left(\frac{1-\theta_{i1}}{1-\theta_{i0}}\right)\right)\right] + \ln\left(\frac{1-\theta_{i0}}{1-\theta_{i1}}\right) + \ln\frac{\beta_{000}}{\beta_{001}} + w_{1,2}X_1X_2 + w_1X_1 + w_2X_2\right]} \end{aligned}$$

where

$$\begin{aligned}
w_o &= \ln\left(\frac{1-\pi}{\pi}\right) + \sum_i \ln\left(\frac{1-\theta_{i0}}{1-\theta_{i1}}\right) + \ln\frac{\beta_{000}}{\beta_{001}} \\
w_1 &= \ln\frac{\beta_{100}}{\beta_{101}} + \ln\frac{\beta_{001}}{\beta_{000}} \\
w_2 &= \ln\frac{\beta_{010}}{\beta_{011}} + \ln\frac{\beta_{001}}{\beta_{000}} \\
w_i &= \ln\frac{\theta_{i0}}{\theta_{i1}} + \ln\left(\frac{1-\theta_{i1}}{1-\theta_{i0}}\right) \quad \forall i \in \{3, \dots, n\} \\
w_{1,2} &= \ln\frac{\beta_{110}}{\beta_{111}} + \ln\frac{\beta_{000}}{\beta_{001}} + \ln\frac{\beta_{101}}{\beta_{100}} + \ln\frac{\beta_{011}}{\beta_{010}}
\end{aligned}$$

2 Hot Eigenfaces for Gaussian Naive Bayes [65 Points]

In this problem you will have the opportunity to apply PCA to a moderately sized collection of images scraped from the *Hot-or-Not*¹ website and then use a Naive Bayes classifier to learn to classify facial attractiveness. For this problem we have selected only female faces because they are more consistently labeled and are typically more amenable to basic classifiers. The data for this problem is available at <http://www.cs.cmu.edu/~ggordon/10601/hws/hw4/faces.mat>. The *faces.mat* Matlab file contains the following data:

h: The height of each image.

w: The width of each image.

tr and te: These structs contain the training and test data respectively. The fields in these structs are:

n: The number of images.

hot: A vector with *n* elements. The entry `hot(i)` is 1 if the *i*th image is hot and 0 otherwise.

images: An *n* by *w* * *h* matrix where each row corresponds to a separate image. To view the *j*th image in matlab you would use the following code:

```
load('faces.mat');
figure(1);
j = 3;
% Reshape the row vector into an 86 by 86 pixel image
img = reshape(tr.images(j, :), h, w);
% Show the image
imagesc(img);
% Set the plot to gray-scale
colormap('gray');
```

★ **SOLUTION:** The solution code for this problem can be downloaded from: http://www.cs.cmu.edu/~ggordon/10601/hws/hw4/hw4_code.tar.gz

Use only the data in the `tr` struct for all training and validation. The data in `te` will only be used in the last part to evaluate your classifier. We will explicitly say when to use `te`. Don't use `te` until we tell you.

¹ The images were collected by Ryan White from the Hot-or-Not website <http://www.hotornot.com/>. To learn more about how the data was collected, get the color images, and try your algorithms on the male face data, go to http://www.ryanmwhite.com/research/tr_hot.html.

2.1 Basic Analysis [15 Points]


We will begin with some basic data analysis to become more familiar with the data.

1. What fraction of the faces are hot?

★ SOLUTION: 0.499

2. Using `subplot(1,2,i)` plot both the average hot and not hot faces side by side. To compute the average face simply use the `mean` function on the respective class of face vectors (i.e., `mean(tr.images(tr.hot,:))`).


★ SOLUTION:



3. Based on the average faces, does it seem like their may be sufficient difference to build a classifier for facial attractiveness? Provide a *very* brief explanation.

★ SOLUTION: Yes. There are visually distinguishable differences between the average hot and average not hot face.

4. Compute the covariance between the pixels (i.e., `cov(tr.images)`). Plot the distribution of the values for the off-diagonal terms. Plot the distribution of the diagonal (variance) terms. (That is, make histograms of the given sets of numbers.) Solution



We can see that the off diagonal terms are nonzero and have magnitude similar to the diagonal terms. This indicates strongly correlated features.

5. Based on the above observations, what can you say about the dependence/independence of pixels in these images? Give a *brief* justification why neighboring pixels might be correlated.

★ **SOLUTION:** Pixel values are not independent. Intuitively, neighboring pixels are likely to share similar values since face images are generally smooth. Furthermore, the two bumps in the histogram are likely be due to strongly correlated regions like the hair, cheeks and background.

2.2 Singular Value Decomposition [15 Points]

For this question you will need to use the matlab `svds` function to compute the top 100 eigenvectors. Please attach the code you write for this section to your solution set.


1. Compute the average face `xbar` using the `mean` function. You will need `xbar` throughout the rest of this problem. Center the faces data by subtracting `xbar` from each row of `tr.images` storing the result in `Xcenter`. Run `svds` on `Xcenter` (i.e., `[U, svalues, efaces] = svds(Xcenter, 100);`). The columns of `efaces` are the eigenfaces and the diagonal entries of `svalues` are the singular values. Please include the code you used for the past few steps.

★ **SOLUTION:**

```
function [efaces, xbar, svalues] = eigenfaces(X, dim)
% Compute the number of images (n) and the dimension of the images
% (d)
[n, d] = size(X);
% Compute the average image
xbar = mean(X);
% Compute the svd of the rectentered points
[U, svalues, efaces] = svds(X - repmat(xbar,n,1), dim);
% Transpose for convenience
efaces = efaces';
% Compute the singular values of
svalues = diag(svalues);
end % End of eigenfaces function
```

2. Plot the singular values (i.e., `plot(diag(svalues))`).

★ **SOLUTION:** Notice how the singular values decrease quickly.



3. Do the singular values decrease rapidly?

★ **SOLUTION:** Yes. (See above plot.)

4. What does this suggest about the approximation quality obtained by representing faces as a linear combination of the top few eigenfaces?

★ **SOLUTION:** Because the singular values decrease rapidly the top few eigenfaces will likely provide a reasonable approximation.

5. Plot the top 10 eigenfaces (preferably using `subplot(2,5,i)` to save paper).

★ **SOLUTION:**



6. For two of the eigenfaces provide a *brief* interpretation of what aspect of the images they might encode (i.e., “lighting”, “eyes”, ...).

★ **SOLUTION:** Along the first row from left to right, face 1 likely encodes variations in eye shape and face width. Face 2 encodes variation in face width. Face 3 encodes lighting from the left. Face 4 encodes deviations in lip and eyelashes. Face 5 encodes hair and teeth deviations.

2.3 Dimensionality Reduction [15 Points]

We will now use the eigenfaces computed in the previous part to do dimensionality reduction. We will see how we can transform 7396 dimensional vectors into 100 dimensional vectors while preserving most of the original information. You should not call `svds` again (this will be very slow); instead, use your existing eigenfaces from the previous question.

1. Create a function to project a collection of faces onto a fixed number (`dim`) of eigenfaces. To project the matrix `X` of images (as row vectors) into the low dimensional linear space spanned by the eigenfaces use the following:

```
%> Z = (X - repmat(xbar, n, 1)) * efaces(1:dim,:);
```

where `efaces` are the eigenfaces as row vectors. The `n` by `dim` matrix `Z` contains the low dimensional representation of each image. Provide the code for this function.

★ SOLUTION:

```
function [z, xapprox, mse] = projection(X, efaces, xbar, dim)
[n,d] = size(X);
z = (X - repmat(xbar, n, 1)) * efaces(1:dim, :)';
xapprox = (z * efaces(1:dim, :)) + repmat(xbar, n, 1);
mse = mean((X(:) - xapprox(:)).^2);
end
```

2. Create a function to map points in the low dimensional space back to the original space. You may want to use the following piece of code:

```
%> Xapprox = (Z * efaces(1:dim, :)) + repmat(xbar, n, 1);
```

Provide the code for this function.

★ SOLUTION: My projection function also computes the approximate reconstruction.

```
function [z, xapprox, mse] = projection(X, efaces, xbar, dim)
[n,d] = size(X);
z = (X - repmat(xbar, n, 1)) * efaces(1:dim, :)';
xapprox = (z * efaces(1:dim, :)) + repmat(xbar, n, 1);
mse = mean((X(:) - xapprox(:)).^2);
end
```

3. Project the first 10 faces onto the first 10, 50, and 100 eigenfaces and then compute **Xapprox**. Using a 4 row by 10 column grid (use `\subplot(4,10,i)`), plot:


Row 1: the original images.

Row 2: the **Xapprox** approximation using the first `dim=100` eigenfaces.

Row 3: the **Xapprox** approximation using the first `dim=50` eigenfaces.

Row 4: the **Xapprox** approximation using the first `dim=10` eigenfaces.

★ SOLUTION:



- From the above plot identify characteristics or features that are lost in the low dimensional representations. For example, can you express glasses or mouth gestures in the low dimensional subspaces?

★ **SOLUTION:** Features like glasses are quickly removed in the low dimensional representation. In addition more interesting properties like smile shape (tongue sticking out), jewelery and even facial obstruction by hair are lost in the low dimensional representation.

2.4 Gaussian Naive Bayes [20 Points]

Ultimately, beauty is in the eye of the beholder. In the case of this question the beholder is a simple Gaussian Naive Bayes classifier. Here we will use a Gaussian Naive Bayes classifier with separate variance terms for each feature and class. Thus, the joint probability can be written as:


$$\mathbf{P}(Z_1, \dots, Z_{\text{dim}}, Y | \theta, \mu, \sigma) \propto \theta^Y (1 - \theta)^{(1-Y)} \prod_{i=1}^{\text{dim}} \exp\left(-\frac{(Z_i - \mu_{i,Y})^2}{2\sigma_{i,Y}^2}\right)$$

★ **SOLUTION:** There was a minor bug in the above equation that several people noted. The normalizing constant $1/\sigma_{i,Y}$ was dropped. The correct form of the equation should have been:

$$\mathbf{P}(Z_1, \dots, Z_{\text{dim}}, Y | \theta, \mu, \sigma) \propto \theta^Y (1 - \theta)^{(1-Y)} \prod_{i=1}^{\text{dim}} \frac{1}{\sigma_{i,Y}} \exp\left(-\frac{(Z_i - \mu_{i,Y})^2}{2\sigma_{i,Y}^2}\right)$$

- Project all the training images into the $\text{dim}=100$ linear space using the above function. Compute the covariance matrix for the $\text{dim}=100$ dimensional representation Z and plot the distribution of the diagonal and off diagonal terms.

★ **SOLUTION:**



- If you assume that the rows of Z are Gaussian, what can you say about the independence of the “features” of Z ? How might this be helpful when using a Gaussian Naive Bayes classifier?

★ **SOLUTION:** The definition of SVD ensures that individual pairs of eigenfaces have minimal covariance on the training data. Thus by projecting faces onto the eigenfaces we are minimizing the covariance in the transformed feature space. These transformed features are less “dependent” and therefore closer to the Naive Bayes assumption. Since the original features are unlikely to be Gaussian and we have not addressed higher order dependencies and so it is unreasonable to claim complete independence.

- Write a function that given the data Z and $Y = \text{hot}$ computes the parameters θ and the 2 by dim parameter matrices μ and σ for the Gaussian Naive Bayes classifier. Provide the code in your solution.

★ SOLUTION: See code below:

```
function bn = train_gnb(X, y)
% Get size information
[n, d] = size(X);
% For each class train separate gaussians
bn.mu = zeros(2, d);
bn.sigmaSq = zeros(2, d);
for class = [1,2]
    ind = y == (class - 1);
    bn.mu(class, :) = mean(X(ind, :));
    bn.sigmaSq(class, :) = var(X(ind, :));
end
% Class prior
bn.classP = [mean(y==0), mean(y==1)];
end
```

4. Write a function that given the data Z and the parameters θ , μ , and σ predicts the probability that $Y = 1$ (i.e., the face is hot). Provide the code in your solution.

★ SOLUTION: See code below:

```
function pred = pred_gnb(bn, X)
% Get the size information
[n,d] = size(X);
% Compute the probability for each point in each class
prob = zeros(n,2);
for class = [1,2]
    prob(:,class) = ...
        sum(...
            -(X - repmat(bn.mu(class,:), n, 1)).^2 ./ ...
            repmat(2 * bn.sigmaSq(class, :), n, 1)...
            -log(repmat(bn.sigmaSq(class,:), n, 1)), ...
            2) + ...
        log(bn.classP(class));
end
% Safely normalize (I ensure that at least one entry has value 1)
prob = exp(prob - repmat(min(prob, [], 2),1,2));
prob = prob ./ repmat(sum(prob,2), 1, 2 );
pred = prob(:,2);
end
```

5. We now will train a Gaussian Naive Bayes classifier to predict whether a face is hot using the low dimensional feature space. To determine the best value for `dim` we will use cross validation on `tr`. *Do not* use the `te` data. For each `dim` from 1 to 100 do 100 random splits of `tr` into $\frac{4}{5}$ training and $\frac{1}{5}$ test. For efficiency, you may use the original eigenfaces and `xbar` computed from the earlier sections. You may want to use the following fragment of code:

```
split = round(4/5 * tr.n);
trials = 100;
for p = 1:100
    % Compute the projection
```

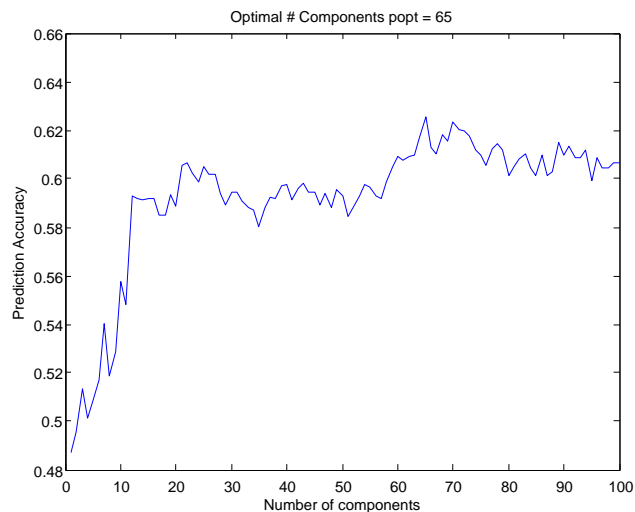
```

X = project(tr.images, efaces, xbar, p);
y = tr.hot;
for t = 1:trials
    % Do Random Split
    ind = randperm(tr.n);
    Xtr = X(ind(1:split), :);      ytr = y(ind(1:split));
    Xval = X(ind((split+1):end),:); yval = y(ind((split+1):end));
    % Train the Gaussian Naive Bayes classifier
    % Predict the y values for the heldout data
    % Compute the classification error
end
end

```

6. Plot the average cross validation prediction accuracy as a function of the number of principal components.

★ **SOLUTION:**



7. What number of components maximizes the average cross validation prediction accuracy?

★ **SOLUTION:** 65

8. You may now use the test data `te`. Using the optimal number of components project the faces from `te` into the low dimensional space. Then using the Gaussian Naive Bayes model trained on all of `tr` compute the prediction accuracy. Compare the prediction accuracy for the Gaussian Naive Bayes classifier to the prediction accuracy for the simple classifier that believes all faces are hot.

★ **SOLUTION:** The test error is 0.65 which is better than the classifier that finds all faces hot which would get an accuracy of 0.50.