

Problem set 4

1. If you needed to estimate F0 from a spectrogram, would you be better off using a wide-band or narrow-band spectrogram? Justify your answer.
2. A pair of classifiers have been designed for speaker verification. On 500 trials, one has an error rate of .07 and the other of .05. Plot the errors and the 95% confidence intervals for the two classifiers using Matlab. The function `errorbar` provides a nice way of plotting these. The reduction in error rate from the first classifier to the second is an impressive 25.9%, but can we say with 95% certainty that they are different? Explain your answer.
3. Suppose that two univariate Gaussian classifiers have posterior distributions that are $n(\mu_1, \sigma_1^2)$ and $n(\mu_2, \sigma_2^2)$ respectively. In terms of the means and variances (you may assume that $\mu_1 < \mu_2$), determine at what point the optimal decision boundary will lie.

4. Show that the maximum likelihood estimator for a univariate Gaussian is indeed the sample mean: $\hat{\mu} = \frac{1}{N} \sum_{i=1}^N x_i$. When solving maximum likelihood estimation problems, we assume that observations are independent, hence given data x_1, x_2, \dots, x_N we can represent their joint probability as:

$$f(x_1, x_2, \dots, x_N) = \prod_{i=1}^N f(x_i) \text{ where } f(x_i) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x_i - \mu)^2}{2\sigma^2}}. \text{ Hint: The maximum}$$

of a function and the log of a function are the same. In many cases, it is much easier to maximize the log of the function.

5. Answer the following questions about vector quantizers:
 - a. Suppose we quantize a set of vectors x_1, x_2, \dots, x_N against a vector quantizer codebook. How could we represent the average minimum distortion as a discriminant function?
 - b. What will happen to the average minimum distortion of a vector quantizer on its training set as the number of codewords increase?
 - c. Is there a way to ensure that there is zero distortion on the training set?
 - d. Suppose we have several vector quantizers that are used to classify data as in your last lab. As the distortion on the training set decreases, should we expect the classification performance to increase? Why or why not?