ECE 8110 Homework 1

1. Using a GRV of mean 1 and variance of 1 compared with a GRV of mean -1 and variance of 1, the optimal threshold of a maximum likelihood decoder should be zero (0). During testing with 10,000 data points for each of the two GRVs, the error rate for a 0 threshold was found to be 15.74% compared to an ideal error of 15.87% by using the indexed values from a standard distribution table. The error was found by taking half of the percentage of false low readings from GRV1 and summing them with half the percentage of false high readings from GRV2. In this fashion errors for a threshold of 0.5 [16.02% tested, 18.77% ideal] and -0.5 [15.91% tested, 18.77% ideal] were also calculated. Resulting in errors that fit quite nicely within the ideal values, but not exact given the limited data sets. Increasing the number of points to 100,000 from 10,000 drives the measured error rate to 15.86%, but maintains the 0.5 and -0.5 error rates of 16.07%/16.08%. This suggests I have missed something, but am at a loss for the error.

2. Comparing a GRV of mean -1 with variance of 1 against a GRV of mean 1 with a variance that ranges from 0.1 to 2 produces the following error plots for various threshold levels.

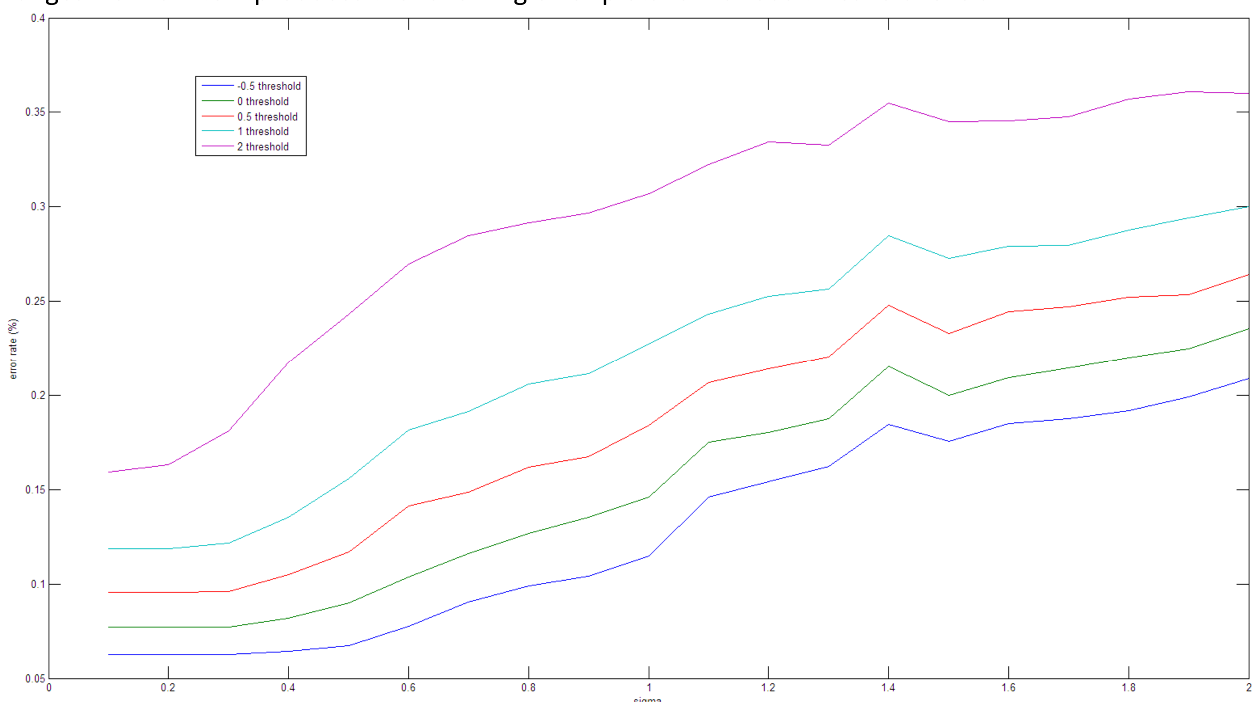

As the variance of the second GRV increases, the error rate begins to grow in a linear fashion. The more interesting plot would be of one that updates the ideal threshold as the second variance increases, but attempting to write something to 'optimize' the threshold giving the changing sigma met with mediocre results as seen below. The hope was to provide at least 1 sigma for each function, until the growing function overtook the control which would result in averaging their Z score.

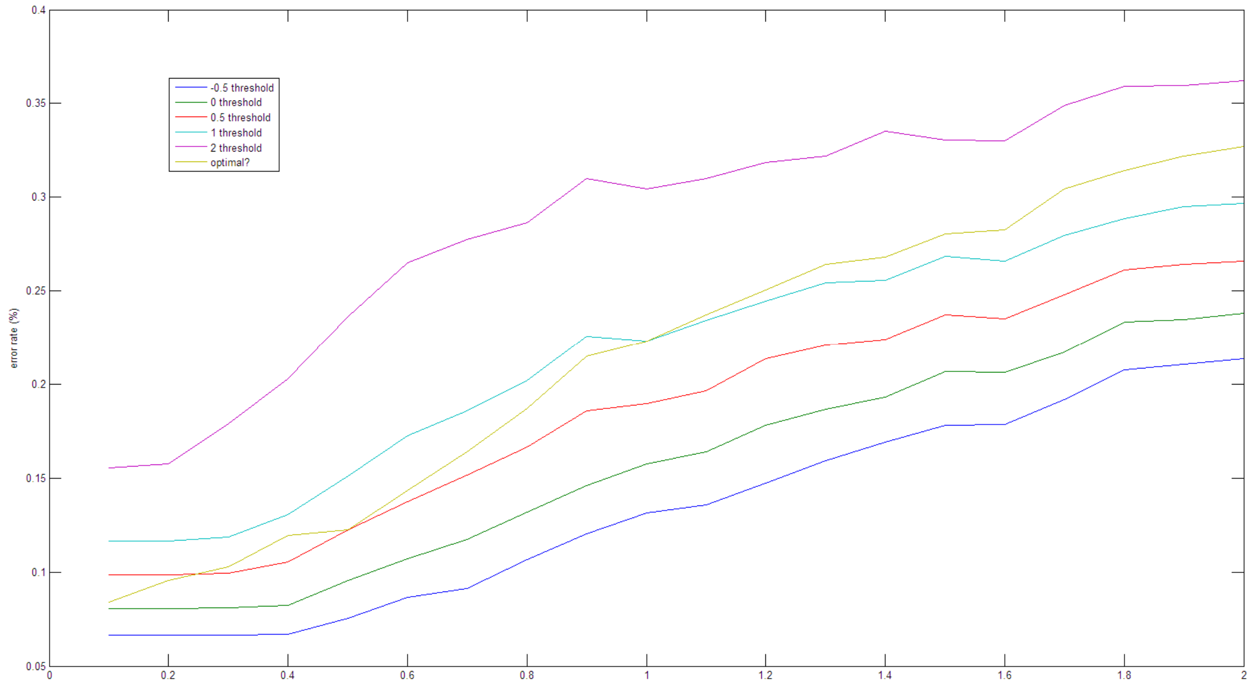

3. Generating two 2D GRV functions, one of mean [-1,1] variance [1 0; 0 1] and the other of mean [1,1] variance [1 0; 0 1], returned a tested error rate of 15.95%. This error is nearly identical to the 1D GRV functions tested in problem 1, which is how it should be as adding the additional dimension does alter the probabilities of the functions. The distance to the threshold must be computed in two dimensions, but the odds remained the same. Ideally the error probably should be 15.87%, but the simulation was run with only 1000 data points for each 2D GRV. Running it again with 100,000 data points for each GRV drives the error rate to 15.87%.

4. To generate plots that resembled the support region Matlab turned somewhat disagreeable with me about how to plot the functions. Using the built in pdf tools I was able to reconstruct various plots from the slides. First the identity covariance matrix, [1 0; 0 1].

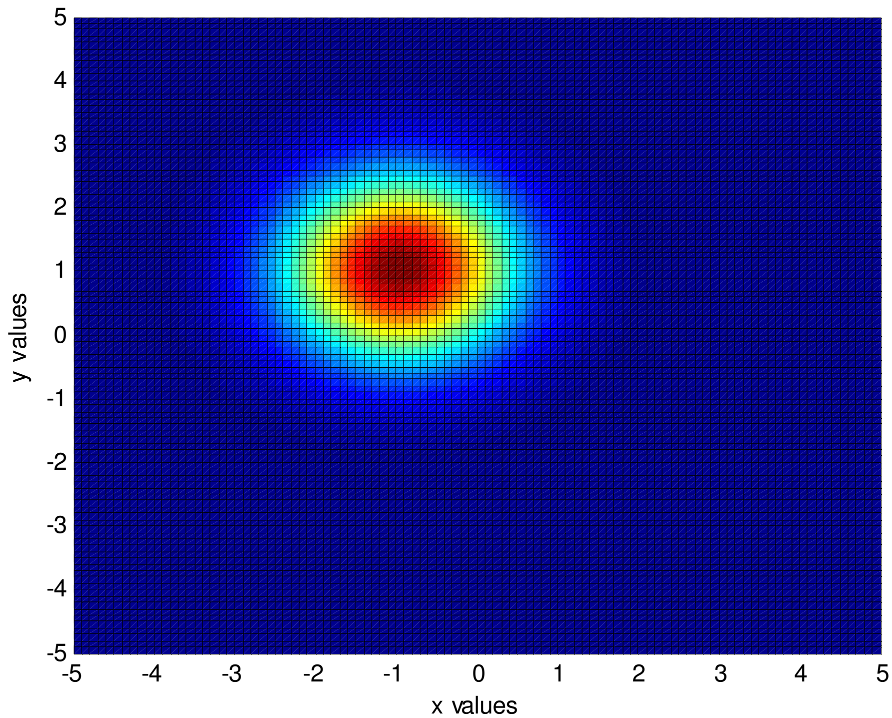

Followed by [5 0; 0 2]:

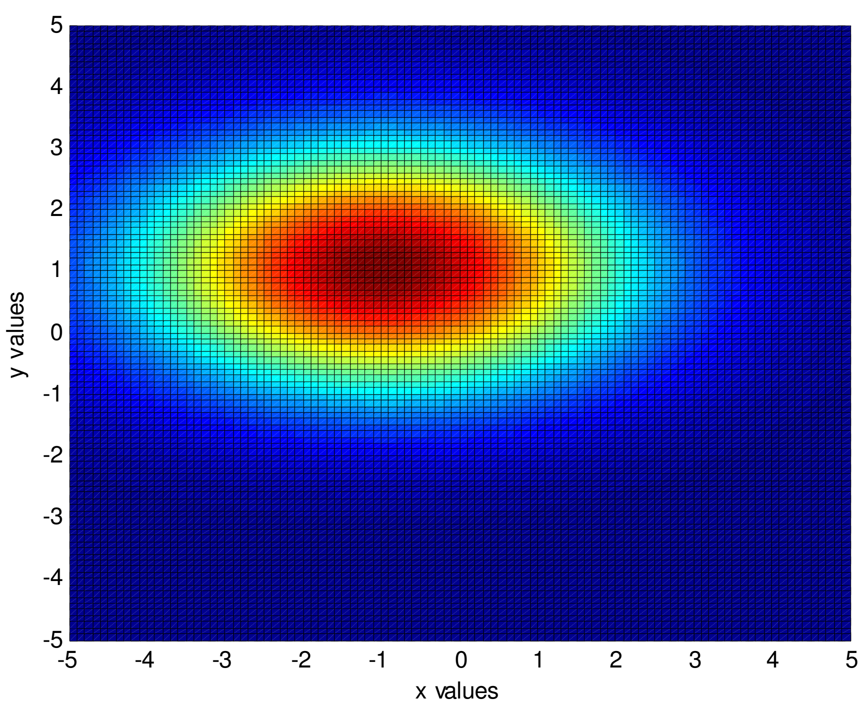

Followed by [1 0.5; 0.5 1]:

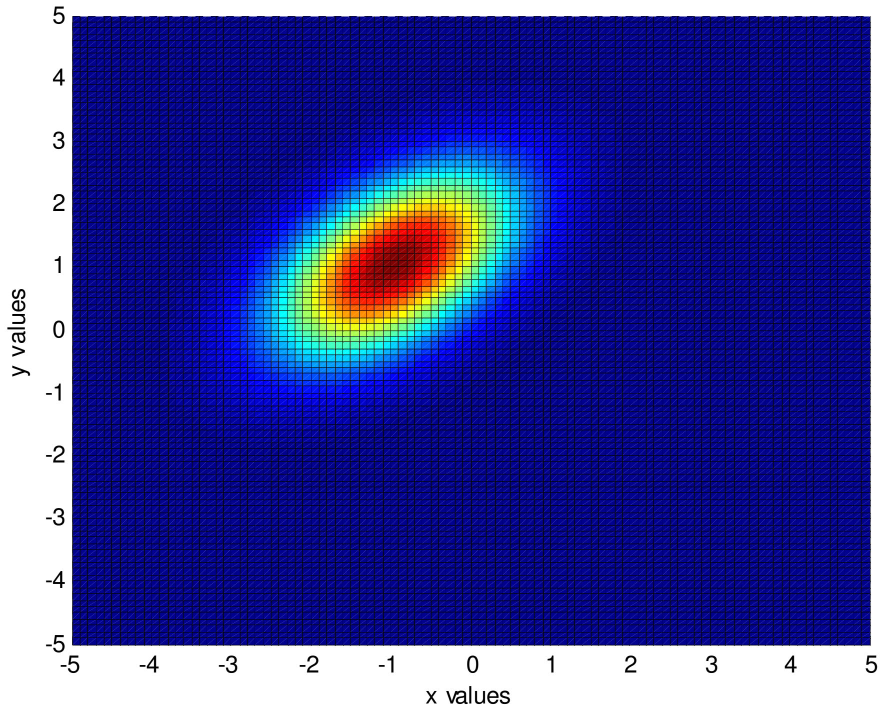

Followed by [5 0.5; 05. 2]: