Support Vector Regression

A) Consider the 1-dimensional SVR problem with two datapoints shown in Figure 1.

1. Estimate the value of b (use equation given in class) and draw the regressive line found by ϵ-SVR, for a kernel width of 0.1 and an $\epsilon = 0.1$.
2. Draw the effect of increasing ϵ to 0.3, 0.5 and 1.
3. Draw the effect of increasing the kernel width.

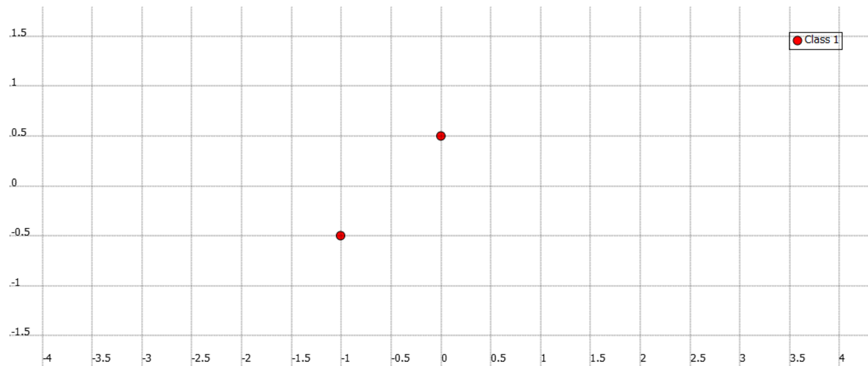

Figure 1: ϵ-SVR, case a

B) Consider the 1-dimensional SVR problem shown in Figure 2, where 2 data points are superimposed on x (i.e. they have the same value for x but a different value for y) .
Discuss the effect of the ϵ parameter on the regression function in this case.

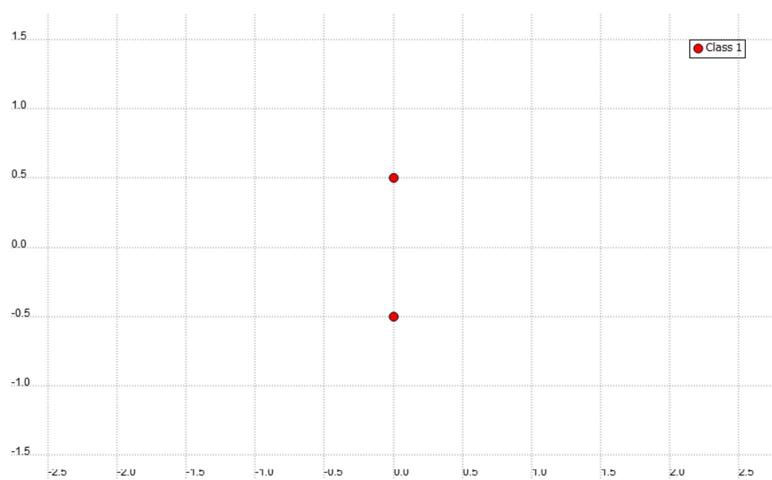

Figure 2: ϵ-SVR, case b

GMR

A) For each of these 2D distributions fitted with Gaussian Mixture Models (GMM), draw approximately the expected result of Gaussian Mixture Regression (i.e. the regressive signal given by GMR) on top of the datapoints (see Figure 3). Consider that, as we did in class, we compute $\mathbb{E}[p(y|x)]$ where x is the horizontal axis and y is the vertical axis. Assume an equal prior for each of the Gaussians:

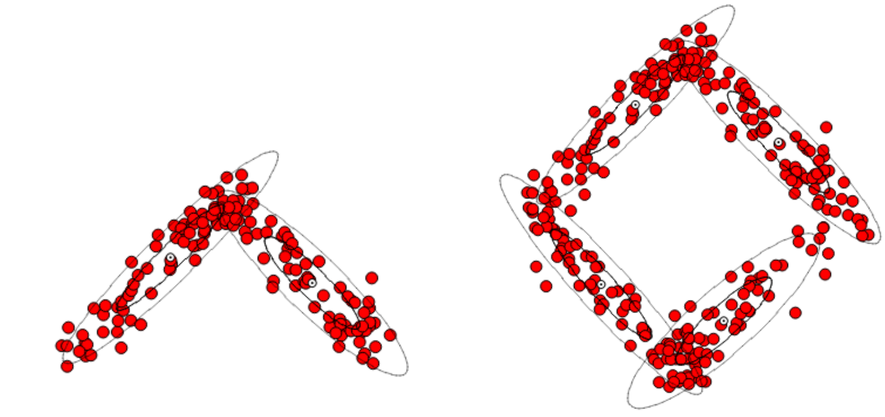

Figure 3: 2D distributions fitted with 2 Gaussians (left) and 4 Gaussians (right). The priors are assumed identical for all Gaussians.

B) As in the previous question, draw approximately the expected result of GMR. However, take this time into account the priors of each Gaussian (see Figure 4). Note that the priors are arbitrary and may not correspond to the real priors found by GMM on this dataset:

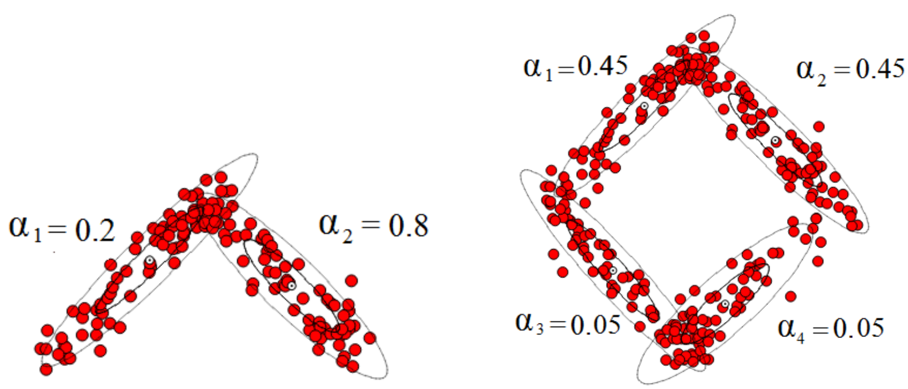

Figure 4: 2D distributions fitted with 2 Gaussians (left) and 4 Gaussians (right). The priors α_i are displayed next to each Gaussian.

C) Observe the original dataset below (Figure 5a) and the solution found by GMR (Figure 5b). How do you explain the poor fit with the two large Gauss functions on both end of the datasets? How could you fix the problem?

D) Consider the dataset shown in Figure 6:

1. Discuss the effect of using one or two components in GMR?

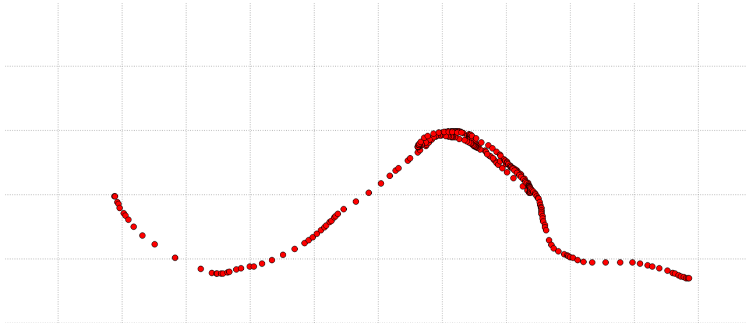

(a)

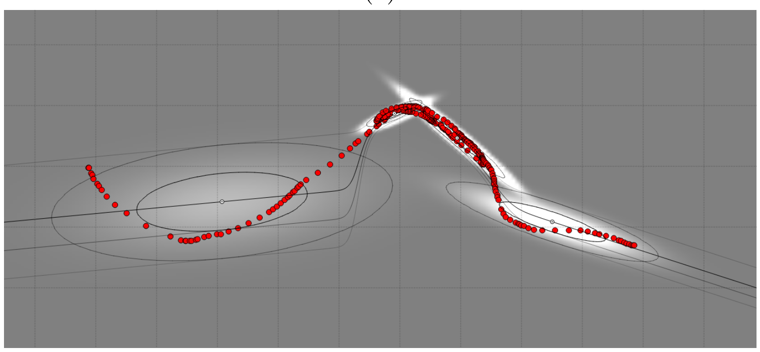

(b)

Figure 5: 2D nonlinear regression problem. (a) Original dataset. (b) Results from the application of GMR. The color shading corresponds to the uncertainty measure (likelihood) of the model.

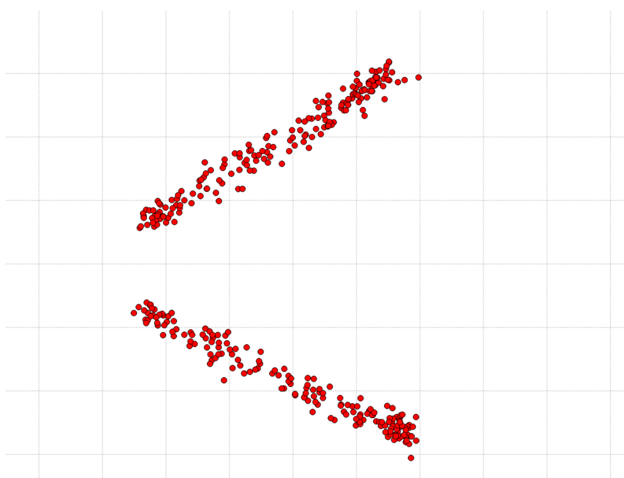

Figure 6: 2D regression problem.

2. GMM models the joint distribution $p(x, y)$ of a dataset. In GMR, one takes the expectation $\hat{y} = \mathbb{E}[p(y|x)]$ of the conditional distribution $p(y|x)$ as the output of the regression. Considering that you use a mixture of two gaussians to model the distribution in Figure 6, do you have suggestions for a different way of obtaining the output from the conditional distribution?

E) Draw a 2D example where two different GMMs give the same GMR output (only draw the Gaussians, you do not have to draw a dataset). Assume equal priors.

F) What shape of regression function can you achieve with only one Gaussian in GMR? Does it depend on the covariance matrix of the Gaussian? (Optional: prove your answer by computing the output of a GMR with a diagonal Gaussian).

GMR & SVR

Consider the distribution of datapoints and assume that you use either:

- SVR with tiny kernel width and very small ϵ
- GMR with a large number of Gauss functions

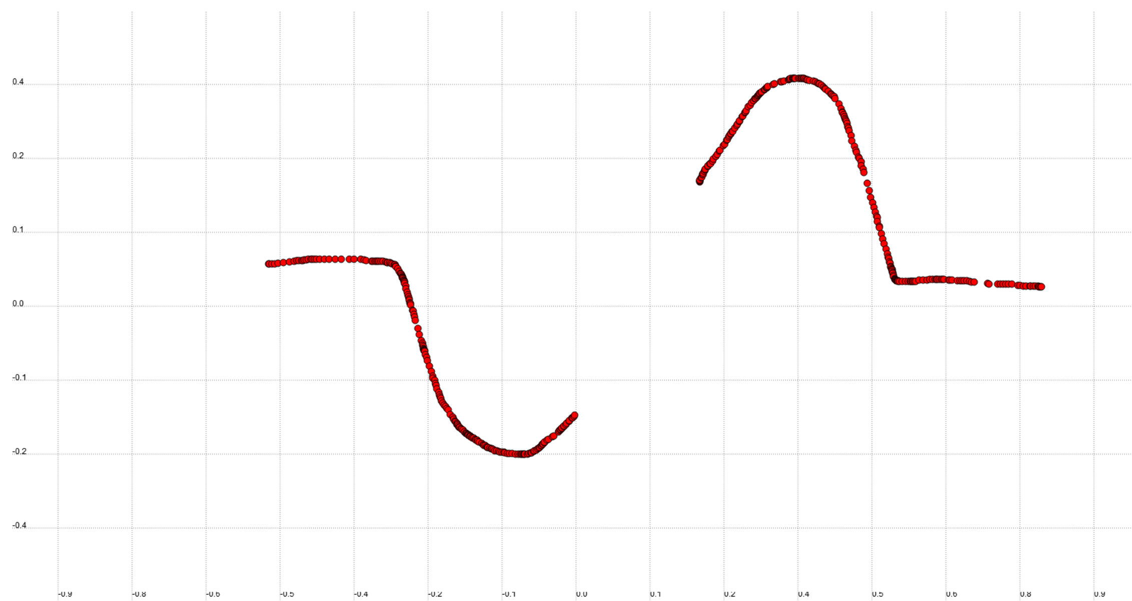

Figure 7: 2D regression problem featuring a "hole" in the dataset