

Supervised learning – Part 2

1. Model selection and training [Old exam question]


Consider a classifier trained till convergence on some training data \mathcal{D}_{train} , and tested on a separate test set \mathcal{D}_{test} . You look at the test error, and find that it is very high. But when you computed the training error it was close to 0.

(a) Which of the following is expected to help? Select all that apply.


- ☐ Increase the training data size.
- ☐ Decrease the training data size.
- ☐ Increase model complexity (For example, if your classifier is an SVM, use a more complex kernel. Or if it is a decision tree, increase the depth).
- ☐ Decrease model complexity.
- ☐ Train on a combination of \mathcal{D}_{train} and \mathcal{D}_{test} and test on \mathcal{D}_{test} .

(b) Explain your choices.

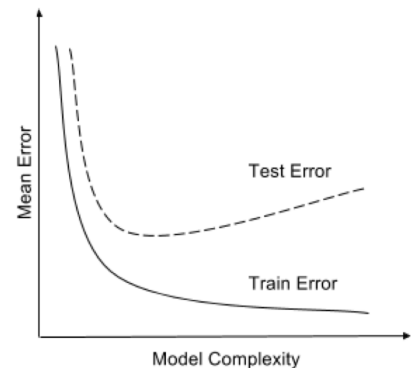
(c) Say you plot the train and test errors as a function of the model complexity. Which of the following three plots is your plot expected to look like? Explain your choice briefly.



(i) ☐



(ii) ☐



(iii) ☐

(d) In machine learning, we typically need a separate validation set next to the training and test sets. What is its role?

2. Consider a binary classification problem whose features are in \mathbb{R}^2 . Suppose the predictor learned by logistic regression is $\sigma(w_0 + w_1x_1 + w_2x_2)$, where $w_0 = 4$, $w_1 = -1$ and $w_2 = 0$. Find and plot the curve along which $P(\text{class} = 1) = 0.5$ and the curve along which $P(\text{class} = 1) = 0.95$.
3. Consider a 3-class classification problem. You have trained a predictor whose input is $\mathbf{x} \in \mathbb{R}^2$ and whose output is $\text{softmax}(x_1 + x_2 - 1, 2x_1 + 3, x_2)$. Find and sketch the three regions in \mathbb{R}^2 that get classified as class 1, 2 and 3.