Introduction to Machine Learning NPFL054

An exercise on t-test

Assume that the accuracy of your binary classifier (when you measure it on different random test samples) is a normally distributed random variable.

There was a suspicion that the accuracy mean is only 55%.

When you checked the accuracy on 5 random test samples, you obtained the following values: 0.63, 0.66, 0.68, 0.51, and 0.71. -- So you got the average accuracy 63.8%.

However, are you sure about the mean accuracy of your classifier? Is it really better than 55%? --- Compute the statistical test at two different levels:

a) at the confidence level = 95%

b) at the confidence level = 90%

Then make a **brief, clear and exact final statement.**

Explain why you made your decision!

Hints

- use t-test: first compute the t-statistic
- then compare your t-statistic with critical value at the given significance level, and make your decision about the hypothesis
- also, compute the p-value – just to check your decision
- also, compute the confidence intervals at the given levels
- would you reject the hypothesis at the confidence level = 99%?
  – now you should be able to answer immediately!

---

*Critical values of the t-distribution that you *might* use:*

$t_4(0.025) = 3.50$ $t_4(0.05) = 2.78$ $t_4(0.10) = 2.13$ $t_4(0.20) = 1.53$

Additional questions

- How can you compute critical values of the t-distribution in R?
- How would the confidence interval change if you had a sample with different size, having the same sample mean and sample variance?