

Solution to Homework 1

Question. Why cannot we just use machine learning to estimate the degree of confidence in its predictions – and to estimate how accurate are these predictions?

Answer. When we train a system to predict the result, we can (and do) use thousands of examples. However, when we train it to get the probability of each result, we need about 100 examples of computations to get one probability estimates – so the number of inputs for such training is 100 times smaller, not enough to get good results.