

CSC581 Final Spring 2014

due Friday 5/9 @ noon in my office

version 1.0

Name:_____

Short Answer Problems

1. What is meant by a maximum-margin classifier and why are they preferable over perceptrons?

2. What do we mean by the *feasible region* of an optimization problem?

3. What are the KKT conditions and why are they important?

4. How do a Lagrangian and a Lagrangian dual differ?

5. Why do we consider maximum-margin classifiers with large margin less complex than maximum-margin classifiers with a small margin?
6. Explain the basic idea behind the gradient ascent/descent optimization technique.

7. What is *quadratic programming* and how can it be used in the context of SVMs?
8. Given that both SVMs and MLPs can solve nonlinear decision problems, what are some of the advantages of SVMs?

9. What is meant by the term *kernel trick* in the context of SVMs?

10. What is a *support vector*?

Problems

For the final examination you are to build and evaluate support vector machine *regression* models for an appropriate regression data set (same rules for selecting a data set apply as for the midterm).

Part A Perform an exploratory data analysis using summary statistics and histograms. Briefly explain your findings.

Part B Build the best model possible for your data set:

1. Document your grid search/model evaluation process carefully, including the type of kernel you are using, the values of its free parameters, and the value of C .
2. Use the cross-validated mse (or rmse) in order to determine your best model.
3. Select the two best performing models.

Part C Investigate whether the difference in performance of your top two models is statistically significant or not using the bootstrap. You should use 95% confidence intervals for this investigation.

1. What are the 95% error confidence intervals for your two models?
2. Is the performance difference statistically significant? If yes, which model would you pick? if no, which model would you pick and why?

Write a brief report summarizing your findings from Parts A, B, and C.

NOTE: All the work has to be done in R.

NOTE: All work has to be your own, no collaboration is permitted.