

CS181 Independent Assignment 2020

Dates: This assignment is open from April 29 through May 4 (anywhere on earth). We expect that students should be able to achieve a SAT-level of performance in 3-5 hours of a single sitting, but are providing plenty of time to avoid time pressure and allow you to reflect on the questions if you wish.

Collaboration Policy: You may not collaborate with any people for this assignment (students, staff, or anyone else); doing so is an Honor Code violation. You may use notes, books, and internet resources. If you have clarification questions, then use *private* posts on Piazza. Do *not* post public posts; the staff will decide whether to share any response broadly.

Grading: To get a SAT, you must make a meaningful effort on every part of every question. Do not leave questions blank!

Length of Responses: We expect a typical response to be 2-3 paragraphs, potentially with accompanying figures or graphs. Please keep your responses to less than a page, maximum! Note: Some problems may take longer than others.

Submission Instructions: You will submit via Gradescope. Please use this latex template *and* annotate where each question starts in Gradescope when you submit to make it easier for grading.

Question 1: Generative vs. Discriminative Models. Explain the difference between generative and discriminative models: What distributions do they each model? What is the difference between them conceptually/in words?

Next, give two examples of real-world settings where having a generative model would be more useful than having a discriminative model. Be sure to describe how each model would be set up and justify your answer.

Solution

Question 2: Bias and Variance. In this problem, you will demonstrate situations in which bias is the dominating factor for generalization error and in which variance dominates. To do so, first define a data generating process (it can be very simple/one-dimensional). Next, choose two model classes, one which will have high bias but small variance on these data and one which will have low bias but high variance.

For your answer: On the first page, describe the data generating process, the models, and why you chose those. Include plots and/or numbers to demonstrate that bias or variance are the cause of the errors. On later pages, append your code. We're not expecting this problem will take more than a page or so of code.

Solution

Code for Solution

Question 3: Loss functions. You're a CS181 TF next year tasked with coming up with a concept check or section problem to help them understand the many different options for loss functions for classification (L^2 loss, 0-1 loss, hinge loss, SVM loss, log loss...). Write your question as well as a full solution. The question need not consider all the loss functions we covered, but should point to some insights on the different properties of losses and the impact of those properties on the boundaries learned and performance.

Solution

Question 4: Clustering vs. Embeddings. Provide a pair of examples when it would be sensible to use clustering methods (e.g., k-means or Gaussian mixture models) vs. an embedding method such as principal component analysis (PCA). Be sure to clearly explain what criteria you are using to determine “sensible.”

For the same data set, is it possible for clustering or embedding methods to be preferable depending on the criteria? If so, provide an example. If not, explain why.

Solution

Question 5: Inference and Estimation with Graphical Models. Describe the similarities between the graphical models for topic models (e.g. latent Dirichlet allocation, LDA), Gaussian Mixture Models (GMMs), and a Naive Bayes classifier. In your description, include correspondences between input and output variables, local and global variables, observed and hidden variables.

For the LDA and GMMs, what makes inference harder than Naive Bayes? How does EM try to make that simpler? Suppose instead of the EM strategy we discussed in class, we randomly assigned each hidden variable (global and local) into one of two groups, and then performed EM with respect to those two (randomly-created) groups. Would that have the same good properties? Why or why not?

Solution

Question 6: Reinforcement Learning. Name an on-policy and an off-policy RL algorithm. What makes them different, and how does this relate to being on or off-policy?

Now, consider the challenge of learning to control an automated vehicle. The vehicle controller starts off with no knowledge about how to drive. Let us consider three different scenarios:

- You can only learn (that is, train the controller) by driving around the city.
- You can train the controller by both driving the vehicle around the city and (simulated) driving the controller through a very accurate simulator of the city.
- You can train the controller by both driving the vehicle around the city and (simulated) driving the controller through a low-quality simulator of the city.

How might these different scenarios impact how you explore and how you exploit? Your choice of on or off-policy RL methods?

Solution