

CSCI-UA 9473. Machine Learning

Material for the Midterm, Fall 2021

Augustin Cosse

October 16, 2021

1 Material covered

1. You must know and understand the [residual sum of squares criterion](#) and derive the [gradient iterations](#) on that function as well as the [closed form solution](#) for the vector of weights (obtained by setting the derivatives to zero).
2. You must be able to list and explain the three main [regularization approaches](#) (Ridge, Lasso, Best Subset Selection)
3. You must be able to explain the notion of [cross validation](#) and its use within best subset selection
4. You must be able to explain the [statistical assumptions](#) leading to the residual sum of squares model, the Ridge and the Lasso. You must be able to describe the distributions involved in each regularization approach (Gaussian and Laplace).
5. You must be able to [compare the regularization approaches](#) in terms of their respective efficiency and complexity
6. You must be able to explain the [bias variance](#) decomposition of the Mean Squared Error (MSE)
7. You must be able to explain how to use the Residual Sum of Squares criterion to [learn a binary classifier](#)
8. You must be able to explain how the [binary classifier can be extended into a multiclass classifier](#) (i.e. one vs rest, one vs one, and one hot encoding)
9. You must be able to explain the [distinction between generative and discriminative classifiers](#) and give an example from each family.
10. You must be able to explain and derive the expression for the [logistic regression classifier](#)
11. You must be able to explain the [perceptron model](#) as well as the [perceptron learning rule](#) and the associated [convergence theorem](#)