

**Homework # 3 - Due Friday 30th of November**

NAME: \_\_\_\_\_

Signature: \_\_\_\_\_

STD. NUM: \_\_\_\_\_

1. Implement a 1-hidden-layer neural network for binary classification. You have to apply this neural net to the twitter dataset. Submit the classifier following the same procedure as in for the naive Bayes classifier of homework 5, with the exception that you need to set the subject of the email to CPSC340\_HW6 instead of CPSC340\_HW5. . Hand in all your code. You will have to decide which form of optimizer (Newton's, on-line gradient descent, mini-batches, etc.) to use. You will also have to decide how many neurons in the hidden layer (try cross-validation for this). Finally, you may want to try a weight decay (ridge,  $L_2$ ) regularizer. The classifier must do better than random of course.
2. Derive the expressions for the gradient vector and Hessian matrix of logistic regression, which appear on the slides.