

Introduction to Deep Learning

School of Computing Science

Simon Fraser University

Assignment 3: Conceptual Exercises

Exercise numbers refer to the textbook (Introduction to Deep Learning).

Submit **one** of the questions for grading; you can choose which one. The others are good to solve as preparation for the final exam, but we will grade only one.

1. Solve Exercise 4.1
2. Explain why, if you are using L1 regularization, it is a bad idea to compute the total loss (error), as opposed to computing only the gradient.
3. Solve Exercise 4.4
4. Solve Exercise 4.6

Note: some answers are at the end