

Instructions

Total points: 50

Submission: Solutions must be typewritten or neatly handwritten and submitted through gradescope. You can submit multiple times, but only the last submission counts. It is your responsibility to make sure that you submit the right things, and we will *not* consider any regrading requests regarding mistakes in making submissions.

Recall that you have a total of three “late days” for the entire semester, and you can use at most one late day for each written assignment.

Notes on notation:

- Unless stated otherwise, scalars are denoted by small letter in normal font, vectors are denoted by small letters in bold font, and matrices are denoted by capital letters in bold font.
- $\|\cdot\|$ means L2-norm unless specified otherwise, i.e., $\|\cdot\| = \|\cdot\|_2$.

Academic integrity: Our goal is to maintain an optimal learning environment. You can discuss the written assignments at a high level with others, but you should not look at any other’s solutions. Trying to find solutions online or from any other sources (including ChatGPT and other similar tools) is prohibited, will result in zero grade and will be reported. To prevent any future plagiarism, uploading any materials from this course to the Internet is also prohibited, and any violations will be reported. Please be considerate and help us help everyone get the best out of this course.

Problem 1 Multiclass Perceptron (17 points)

Recall that a linear model for a multiclass classification problem with C classes is parameterized by C weight vectors $w_1, \dots, w_C \in \mathbb{R}^D$. In the lecture, we derived the multiclass logistic regression by minimizing the multiclass logistic loss. In this problem, you need to derive the multiclass perceptron algorithm in a similar way. Specifically, the multiclass perceptron loss on a training set $(x_1, y_1), \dots, (x_N, y_N) \in \mathbb{R}^D \times [C]$ is defined as

$$F(w_1, \dots, w_C) = \frac{1}{N} \sum_{n=1}^N F_n(w_1, \dots, w_C), \quad \text{where } F_n(w_1, \dots, w_C) = \max \left\{ 0, \max_{y \neq y_n} w_y^\top x_n - w_{y_n}^\top x_n \right\}.$$

1.1 To optimize this loss function, we need to first derive its gradient. Specifically, for each $n \in [N]$ and $c \in [C]$, write down the partial derivative $\frac{\partial F_n}{\partial w_c} \in \mathbb{R}^D$ (and your derivation). For simplicity, you can assume that for any n , $w_1^\top x_n, \dots, w_C^\top x_n$ are always C distinct values (so that there is no tie when taking max over them, and consequently no non-differentiable points needed to be considered). (8 points)

1.2 Similarly to the binary case, multiclass perceptron is simply applying SGD with learning rate 1 to minimize the multiclass perceptron loss. Based on this information, fill in the missing details in the repeat-loop of the algorithm below (your solution cannot contain implicit quantities such as $\nabla F_n(w)$; instead, write down the exact formula based on your solution from the last question). (4 points)

Algorithm 1: Multiclass Perceptron

- 1 **Input:** A training set $(x_1, y_1), \dots, (x_N, y_N)$
- 2 **Initialization:** $w_1 = \dots = w_C = \mathbf{0}$
- 3 **Repeat:**

|

1.3 At this point, you should find that the parameters w_1, \dots, w_C computed by Multiclass Perceptron are always linear combinations of the training points x_1, \dots, x_N , that is, $w_c = \sum_{n=1}^N \alpha_{c,n} x_n$ for some coefficient $\alpha_{c,n}$. Just like kernelized linear regression, this means that one can kernelize multiclass Perceptron as well for any given kernel function $k(\cdot, \cdot)$ (with a corresponding feature mapping ϕ).

Specifically, fill in the missing details in the repeat-loop of the algorithm below which maintains and updates weights $\alpha_{c,n}$, $\forall c \in [C], n \in [N]$ such that $w_c = \sum_{n=1}^N \alpha_{c,n} \phi(x_n)$ is always the same as what one would get by running Algorithm 1 with x_n replaced by $\phi(x_n)$ for all n . No reasoning is required. Keep in mind that in your solution $\phi(x_n)$ should never appear. (5 points)


Algorithm 2: Multiclass Perceptron with kernel function $k(\cdot, \cdot)$

- 1 **Input:** A training set $(x_1, y_1), \dots, (x_N, y_N)$
- 2 **Initialize:** $\alpha_{c,n} = 0$ for all $c \in [C]$ and $n \in [N]$
- 3 **Repeat:**

└

Problem 2 Backpropagation for CNN (18 points)

Consider the following mini convolutional neural net, where (x_1, x_2, x_3) is the input, followed by a convolution layer with a filter (w_1, w_2) , a ReLU layer, and a fully connected layer with weight (v_1, v_2) .



More concretely, the computation is specified by

$$a_1 = x_1 w_1 + x_2 w_2$$

$$a_2 = x_2 w_1 + x_3 w_2$$

$$o_1 = \max\{0, a_1\}$$

$$o_2 = \max\{0, a_2\}$$

$$\hat{y} = o_1 v_1 + o_2 v_2$$

For an example $(x, y) \in \mathbb{R}^3 \times \{-1, +1\}$, the logistic loss of the CNN is

$$\ell = \ln(1 + \exp(-y\hat{y})),$$

which is a function of the parameters of the network: w_1, w_2, v_1, v_2 .

2.1 Write down $\frac{\partial \ell}{\partial v_1}$ and $\frac{\partial \ell}{\partial v_2}$ (show the intermediate steps that use chain rule). You can use the sigmoid function $\sigma(z) = \frac{1}{1+e^{-z}}$ to simplify your notation. (4 points)

2.2 Write down $\frac{\partial \ell}{\partial w_1}$ and $\frac{\partial \ell}{\partial w_2}$ (show the intermediate steps that use chain rule). The derivative of the ReLU function is $H(a) = \mathbb{I}[a > 0]$, which you can use directly in your answer. (6 points)

2.3 Using the derivations above, fill in the missing details of the repeat-loop of the Backpropagation algorithm below that is used to train this mini CNN. (8 points)

Algorithm 3: Backpropagation for the above mini CNN

- 1 **Input:** A training set $(x_1, y_1), \dots, (x_N, y_N)$, learning rate η
 - 2 **Initialize:** set w_1, w_2, v_1, v_2 randomly
 - 3 **Repeat:**
 - 4 randomly pick an example (x_n, y_n) (the three features of x_n are denoted by x_{n1}, x_{n2} , and x_{n3} .)
 - 5 Forward propagation:
 - 6 Backward propagation:
-

Problem 3 Support Vector Machines (15 points)

Consider a dataset consisting of points in the form of (x, y) , where x is a real value, and $y \in \{-1, 1\}$ is the class label. There are only three points $(x_1, y_1) = (-1, -1)$, $(x_2, y_2) = (1, -1)$, and $(x_3, y_3) = (0, 1)$, shown in Figure 1.




Figure 1: Three data points considered in Problem 1

3.1 Can these three points in their current one-dimensional feature space be perfectly separated with a linear classifier of the form $\text{SGN}(wx + b)$ for some $w, b \in \mathbb{R}$? Why or why not? (2 points)

3.2 Now consider the feature mapping $\phi(x) = [x, x^2]^T$ (convince yourself that the data are now linearly separable in this new feature space). Write down the 3×3 kernel/Gram matrix \mathbf{K} for this dataset. (2 points)

3.3 Next, write down the dual formulations of SVM for this dataset in the two-dimensional feature space. Note that when the data is separable, we set the hyperparameter C to be $+\infty$. (You have to plug in the actual data instead of just showing the generic dual formulation.) (2 points)

3.4 Next, solve the dual formulation exactly (note: while this is not generally feasible as discussed in the lecture, the simple form of this dataset makes it possible). (3 points)

3.5 Given a test point $x = 2$, what is the prediction of this SVM? Follow the steps below to derive your answer: first, compute the primal solution b^* ; then, compute $k(x_n, x)$ for all $n = 1, 2, 3$ and the test point $x = 2$; finally, calculate the prediction. Note that to get full credits, you have to avoid directly operating in the feature space that corresponds to the kernel function. (6 points)