

10-701 Introduction to Machine Learning

Homework 3, **version 1.4**

Due Oct 30, 11:59 am

Rules:

1. Homework submission is done via CMU Autolab system. Please package your writeup and code into a zip or tar file, *e.g.*, let `submit.zip` contain `writeup.pdf` and your code. Submit the package to <https://autolab.cs.cmu.edu/courses/10701-f15>.
 2. Like conference websites, repeated submission is allowed. Please feel free to refine your answers since we will only grade the latest version. Submitting incomplete solutions early will be helpful in preventing last minute panic as well.
 3. Autolab may allow submission after the deadline, note however it is because of the late day policy. Please see course website for policy on late submission.
 4. We recommend that you typeset your homework using appropriate software such as L^AT_EX. If you are writing please make sure your homework is cleanly written up and legible. The TAs will not invest undue effort to decrypt bad handwriting.
 5. You are allowed to collaborate on the homework, but you should write up your own solution and code. Please indicate your collaborators in your submission.
-

1 Neural Networks (50 Points) (Zhiting)

1.1 Neural network for regression

Figure.1 shows a two-layer neural network which learns a function $f : X \rightarrow Y$ where $X = (X_1, X_2) \in \mathbb{R}^2$. The weights $\mathbf{w} = \{w_1, \dots, w_6\}$ can be arbitrary. There are two possible choices for the function implemented by each unit in this network:

- S: signed sigmoid function $S(a) = \text{sign}[\sigma(a) - 0.5] = \text{sign}[\frac{1}{1+\exp\{-a\}} - 0.5]$
- L: linear function $L(a) = ca$

where in both cases $a = \sum_i w_i X_i$.

1. Assign proper activation functions (S or L) to each unit in Figure.1 so this neural network simulates a linear regression: $Y = \beta_1 X_1 + \beta_2 X_2$.

Answer:

L, L, L

2. Assign proper activation functions (S or L) for each unit in Figure.1 so this neural network simulates a binary logistic regression classifier: $Y = \arg \max_y P(Y = y|X)$, where $P(Y = 1|X) = \frac{\exp(\beta_1 X_1 + \beta_2 X_2)}{1 + \exp(\beta_1 X_1 + \beta_2 X_2)}$, and $P(Y = -1|X) = \frac{1}{1 + \exp(\beta_1 X_1 + \beta_2 X_2)}$. Derive β_1 and β_2 in terms of w_1, \dots, w_6 .

Answer:

L, L, S

$$\beta_1 = c(w_1 w_5 + w_2 w_6)$$

$$\beta_2 = c(w_3 w_5 + w_4 w_6)$$

3. Assign proper activation functions (S or L) to each unit in Figure.1 so this neural network simulates a boosting classifier which combines two logistic regression classifiers, $f_1 : X \rightarrow Y_1$ and $f_2 : X \rightarrow Y_2$, to produce its final prediction: $Y = \text{sign}[\alpha_1 Y_1 + \alpha_2 Y_2]$. Use the same distribution in problem 1.1.2 for f_1 and f_2 . Derive α_1 and α_2 in terms of w_1, \dots, w_6 .

Answer:

S, S, S

$$\alpha_1 = w_5$$

$$\alpha_2 = w_6$$

1.2 Convolutional neural networks

1. Count the total number of parameters in LeNet (pp.46, slides of Lecture.8). How many parameters in all of the convolutional layers? How many parameters in all of the fully-connected layers?

Note:

- (a) The filter size of each convolutional and pooling(subsampling) layer:

C1: 5×5 (i.e., each unit of C1 has a 5×5 receptive field in its preceding layer);

S2: 2×2 ;

C3: 5×5 ;

S4: 2×2 ;




Figure 1: A two-layer neural network.

(b) Fully-connected layers in LeNet include C5, F6, and OUTPUT

Answer:

ConvLayers

C1: $(5 * 5 + 1) * 6 = 156$

S2: $0 \text{ or } 2 * 6 = 12$

C3: $(5 * 5 * 3 + 1) * 6 + (5 * 5 * 4 + 1) * 9 + (5 * 5 * 6 + 1) = 1516$
(Or: $(16 * (6 * 5 * 5 + 1)) = 2416$)

S2: $0 \text{ or } 2 * 16 = 32$

FCLayers

C5: $(5 * 5 * 16 + 1) * 120 = 48120$

F6: $(120 + 1) * 84 = 10164$

OUTPUT: $(84 + 1) * 10 = 850$

2. In a convolutional layer the units are organized into planes, each of which is called a feature map. The units within a feature map (indexed q) have different inputs, but all share a common weight vector, $\mathbf{w}^{(q)}$. A convolutional network is usually trained through backpropagation. Let $J^{(q)}$ be the number of units in the q th feature map, $z_j^{(q)}$ the activation of the j th unit, $x_{ji}^{(q)}$ the i th input for the j th unit, $w_i^{(q)}$ the i th element of $\mathbf{w}^{(q)}$, L the training loss. Derive the gradient of $w_i^{(q)}$.

Answer:

$$\frac{\partial L}{\partial w_i^{(m)}} = \sum_j \frac{\partial L}{\partial z_j^{(m)}} \frac{\partial z_j^{(m)}}{\partial w_i^{(m)}} = \sum_j \delta_j^{(m)} x_{ji}^{(m)}$$

1.3 Gradient vanishing/explosion

In this problem we will study the difficulty of back-propagation in training deep neural networks. For simplicity, we consider the simplest deep neural network: one with just a single neuron in each layer, where the output of the neuron in the j th layer is $z_j = \sigma(a_j) = \sigma(w_j z_{j-1} + b_j)$. Here σ is some activation function whose derivative on x is $\sigma'(x)$. Let m be the number of layers in the neural network, L the training loss.

1. Derive the derivative of L w.r.t. b_1 (the bias of the neuron in the first layer).

Answer:

$$\frac{\partial L}{\partial b_1} = \frac{\partial L}{\partial z_m} \sigma'(a_1) \prod_{k=2}^m \sigma'(a_k) w_k$$

2. Assume the activation function is the usual sigmoid function $\sigma(x) = 1/(1 + \exp\{-x\})$. The weights \mathbf{w} are initialized to be $|w_j| < 1$ ($j = 1, \dots, m$).

- (a) Explain why the above gradient $(\partial L / \partial b_1)$ tends to vanish ($\rightarrow 0$) when m is large.

Answer:

The derivative of sigmoid function reaches a maximum at $\sigma'(0) = \frac{1}{4}$. Since $|w_j| < 1$, we have $|w_j \sigma'(a_j)| < \frac{1}{4} < 1$

- (b) Even if $|w|$ is large, the above gradient would also tend to vanish, rather than explode ($\rightarrow \infty$). Explain why. (A rigorous proof is not required.)

Answer:

To avoid the vanishing gradient problem we need $|w \sigma'(a)| \geq 1$. But, the $\sigma'(a)$ term also depends on w : $\sigma'(a) = \sigma(wz + b)$. If we make w large we tend to make $wz + b$ very large, and $\sigma'(a)$ very small.

3. One of the approaches to (partially) address the gradient vanishing/explosion problem is to use the rectified linear (ReLU) activation function instead of the sigmoid. The ReLU activation function is $\sigma(x) = \max\{0, x\}$. Explain why ReLU can alleviate the gradient vanishing problem as faced by sigmoid.

Answer:

As long as $a > 0$, $\sigma'(a) = 1$. So we don't have the issue as in 2(b).

4. A second approach to (partially) address the gradient vanishing/explosion problem is layer-wise pre-training. Restricted Boltzmann machine (RBM) is one of the widely-used models for layer-wise pre-training. Figure 2 shows an example of RBM which includes K hidden units \mathbf{h} , and J input units \mathbf{v} . Let us define the joint distribution as the following general form:

$$P(\mathbf{v}, \mathbf{h}) = \frac{1}{Z} \exp \left(\sum_i \theta_i \phi_i(\mathbf{v}, \mathbf{h}) \right), \quad (1)$$

where $Z = \sum_{\mathbf{v}, \mathbf{h}} \exp(\sum_i \theta_i \phi_i(\mathbf{v}, \mathbf{h}))$ is the normalization term; $\phi_i(\mathbf{v}, \mathbf{h})$ are some features; θ_i are the parameters corresponding to the weights in the RBM. Consider the simplest learning algorithm, gradient descent. Show that

$$\frac{\partial \log P(\mathbf{v})}{\partial \theta_i} = \sum_{\mathbf{h}} \phi_i(\mathbf{v}, \mathbf{h}) P(\mathbf{h}|\mathbf{v}) - \sum_{\mathbf{v}, \mathbf{h}} \phi_i(\mathbf{v}, \mathbf{h}) P(\mathbf{v}, \mathbf{h}). \quad (2)$$

Answer:

We have that:

$$\begin{aligned}
P(\mathbf{x}, \mathbf{y}) &= \frac{1}{\sum_{\mathbf{x}, \mathbf{y}} \exp(\sum_k \theta_k \phi_k(\mathbf{x}, \mathbf{y}))} \exp\left(\sum_k \theta_k \phi_k(\mathbf{x}, \mathbf{y})\right) \Rightarrow \\
P(\mathbf{x}) &= \frac{1}{\sum_{\mathbf{x}, \mathbf{y}} \exp(\sum_k \theta_k \phi_k(\mathbf{x}, \mathbf{y}))} \sum_{\mathbf{y}} \exp\left(\sum_k \theta_k \phi_k(\mathbf{x}, \mathbf{y})\right) \Rightarrow \\
\log P(\mathbf{x}) &= \log \sum_{\mathbf{y}} \exp\left(\sum_k \theta_k \phi_k(\mathbf{x}, \mathbf{y})\right) - \log \sum_{\mathbf{x}, \mathbf{y}} \exp\left(\sum_k \theta_k \phi_k(\mathbf{x}, \mathbf{y})\right) \Rightarrow \\
\frac{\partial \log P(\mathbf{x})}{\partial \theta_l} &= \sum_{\mathbf{y}} \frac{\exp(\sum_k \theta_k \phi_k(\mathbf{x}, \mathbf{y}))}{\sum_{\mathbf{y}} \exp(\sum_k \theta_k \phi_k(\mathbf{x}, \mathbf{y}))} \phi_l(\mathbf{x}, \mathbf{y}) - \sum_{\mathbf{x}, \mathbf{y}} \frac{\exp(\sum_k \theta_k \phi_k(\mathbf{x}, \mathbf{y}))}{\sum_{\mathbf{x}, \mathbf{y}} \exp(\sum_k \theta_k \phi_k(\mathbf{x}, \mathbf{y}))} \phi_l(\mathbf{x}, \mathbf{y}), \\
&= \sum_{\mathbf{y}} \frac{P(\mathbf{x}, \mathbf{y})}{P(\mathbf{x})} \phi_l(\mathbf{x}, \mathbf{y}) - \sum_{\mathbf{x}, \mathbf{y}} P(\mathbf{x}, \mathbf{y}) \phi_l(\mathbf{x}, \mathbf{y}), \\
&= \sum_{\mathbf{y}} P(\mathbf{y} | \mathbf{x}) \phi_l(\mathbf{x}, \mathbf{y}) - \sum_{\mathbf{x}, \mathbf{y}} P(\mathbf{x}, \mathbf{y}) \phi_l(\mathbf{x}, \mathbf{y}).
\end{aligned}$$




Figure 2: A restricted Boltzmann machine.

2 Support Vector Machines (50 Points) (Yuntian)

2.1 Support Vector Regression (25 Points)

We now extend support vector machines (SVM) to regression problems. Recall that in regression problems, we have n data points $(x_i, y_i)_{i=1}^n$ where $x_i \in \mathbb{R}^m$ and $y_i \in \mathbb{R}$. Given a function class \mathcal{F} (e.g. linear or quadratic functions), we want to fit a function $f \in \mathcal{F}$ on the training set:

$$f^* = \operatorname{argmin}_{f \in \mathcal{F}} C \sum_{i=1}^n l(f(x_i), y_i) + R(f) \quad (3)$$

where $l(\cdot, \cdot)$ is the loss function, $R(f)$ is the regularization term, C controls the regularization strength. The first part tries to fit data, and the second part penalizes complex f to avoid over-fitting.

In the support vector regression (SVR) framework, we consider linear function class $\mathcal{F} = \{x \rightarrow w^T x\}$ (we do not consider interception term for simplicity). We use ℓ_2 -regularizer $R(f) = \frac{1}{2} \|w\|_2^2$ for $f(x) = w^T x$. For the loss function l , similar to the hinge-loss function in SVM classification, we employ an ϵ -insensitive error function

$$l_\epsilon(f(x), y) = \begin{cases} 0 & \text{if } |f(x) - y| < \epsilon \\ |f(x) - y| - \epsilon & \text{otherwise.} \end{cases} \quad (4)$$

Then we get the following optimization problem:

$$w^* = \operatorname{argmin}_w C \sum_{i=1}^n l_\epsilon(w^T x_i, y_i) + \frac{1}{2} \|w\|_2^2. \quad (5)$$

1. Write down the dual problem of SVR. (Hint: follow the derivations for SVM)

Answer:

We introduce slack variables ξ_i, ξ_i^* and write the original problem as

$$\begin{aligned} \min \quad & \frac{1}{2} \|w\|_2^2 + C \sum_{i=1}^n (\xi_i + \xi_i^*) \\ \text{s.t.} \quad & \xi_i \geq y_i - w^T x_i - \epsilon \\ & \xi_i^* \geq w^T x_i - y_i - \epsilon \\ & \xi_i, \xi_i^* \geq 0 \end{aligned} \quad (6)$$

By introducing slack variables $\eta_i, \eta_i^*, \alpha_i, \alpha_i^*$, the Lagrangian is

$$\begin{aligned} L = \frac{1}{2} \|w\|_2^2 + C \sum_{i=1}^n (\xi_i + \xi_i^*) - \sum_{i=1}^n \eta_i \xi_i - \sum_{i=1}^n \eta_i^* \xi_i^* \\ + \sum_{i=1}^n \alpha_i (y_i - w^T x_i - \epsilon - \xi_i) + \sum_{i=1}^n \alpha_i^* (w^T x_i - y_i - \epsilon - \xi_i^*) \end{aligned} \quad (7)$$

The primal problem is

$$\begin{aligned} \min_{w, \xi_i, \xi_i^*} \max_{\alpha_i, \alpha_i^*, \eta_i, \eta_i^*} L \\ \text{s.t.} \quad \alpha_i, \alpha_i^*, \eta_i, \eta_i^* \geq 0 \end{aligned} \quad (8)$$

So the dual problem is

$$\begin{aligned} \max_{\alpha_i, \alpha_i^*, \eta_i, \eta_i^*} \min_{w, \xi_i, \xi_i^*} L \\ \text{s.t.} \quad \alpha_i, \alpha_i^*, \eta_i, \eta_i^* \geq 0 \end{aligned} \quad (9)$$

We can get $\min_{w, \xi_i, \xi_i^*} L$ by taking the derivatives and solve for w, ξ_i, ξ_i^* , so the dual problem is

$$\begin{aligned} \max_{\alpha_i, \alpha_i^*, \eta_i, \eta_i^*} \quad & -\frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n (\alpha_i - \alpha_i^*)(\alpha_j - \alpha_j^*) x_i^T x_j + \sum_{i=1}^n (\alpha_i - \alpha_i^*) y_i - \sum_{i=1}^n (\alpha_i + \alpha_i^*) \epsilon \\ \text{s.t.} \quad & 0 \leq \alpha_i \leq C \\ & 0 \leq \alpha_i^* \leq C \end{aligned} \quad (10)$$

2. Write down the KKT conditions, and explain what are the “support vectors”.

Answer: The KKT conditions are:

$$\frac{\partial L}{\partial w} = w - \sum_{i=1}^n (\alpha_i - \alpha_i^*) x_i = 0 \quad (11)$$

$$\frac{\partial L}{\partial \xi_i} = C - \alpha_i - \eta_i = 0 \quad (12)$$

$$\frac{\partial L}{\partial \xi_i^*} = C - \alpha_i^* - \eta_i^* = 0 \quad (13)$$

$$\alpha_i, \alpha_i^*, \eta_i, \eta_i^* \geq 0 \quad (14)$$

$$\eta_i \xi_i = 0, \eta_i^* \xi_i^* = 0 \quad (15)$$

$$\alpha_i (y_i - w^T x_i - \epsilon - \xi_i) = 0 \quad (16)$$

$$\alpha_i^* (w^T x_i - y_i - \epsilon - \xi_i^*) = 0 \quad (17)$$

The support vectors are those with $\alpha_i - \alpha_i^* \neq 0$.

3. Derive a kernelized version of SVR. For a test point x , write down the prediction rule.

Answer: In a kernelized version, for any kernel function $K(\cdot, \cdot)$, just replace the $x_i^T x_j$ with $K(x_i, x_j)$. For a test point x , the prediction rule is

$$f(x) = \sum_{i=1}^n (\alpha_i - \alpha_i^*) K(x_i, x) \quad (18)$$

4. Give one reason why do we usually solve the dual problem of SVR and SVM instead of the primal.

Answer: We can introduce kernel functions here.

5. Implement SVR on a **1-D toy dataset**. Each line of the dataset contains a training instance (x_i, y_i) (separated by a tab). For this problem, you need to

- Use RBF kernel $k(x_i, x_j) = \exp(-\frac{\|x_i - x_j\|_2^2}{2h^2})$, and take $h = 0.5$, $C = 4$, $\epsilon = 0.1$.
- Plot the prediction curve for $x \in [0, 1]$ and show the support vectors versus other training points in the training dataset.

(Hint: You are allowed to use optimization toolkits such as CVX or Matlab’s inbuilt function quadprog to solve the dual problem.)

Answer:

```
clc;
clear;

data = dlmread('SVR_dataset.txt');
```




Figure 3: SVR results.

```

x = data(:, 1);
y = data(:, 2);

h = 0.5;
C = 4;
eps = 0.1;

n = size(x, 1);

% Parameters are 2n by 1 vector, whose first n elements are
% p = alpha - alpha^star and next n elements are
% q = alpha + alpha^star

H = zeros(2 * n, 2 * n);
for i = 1:n
    for j = 1:n
        temp = norm(x(i) - x(j), 2);
        H(i, j) = exp(- temp ^ 2 / 2 / h ^ 2);
    end
end

f = - eps * ones(2 * n, 1);
f(1:50, 1) = y;

A = zeros(4 * n, 2 * n);
for i = 1:n
    A(i, i) = 0.5;
    A(i, i + n) = 0.5;

    A(i + n, i) = -0.5;
    A(i + n, i + n) = 0.5;

    A(i + 2 * n, i) = -0.5;

```



```

        A(i + 2 * n, i + n) = -0.5;

        A(i + 3 * n, i) = 0.5;
        A(i + 3 * n, i + n) = -0.5;
    end

    b = zeros(4 * n, 1);
    b(1:2 * n, 1) = C;

    params = quadprog(H, -f, A, b);

    %% Plot

    p = params(1:n, 1);
    q = params(n + 1:100, 1);

    limit = 0.00001;
    sv_idx = find(abs(p) > limit);
    nsv_idx = find(abs(p) <= limit);

    x_sv = x(sv_idx, 1);
    x_nsv = x(nsv_idx, 1);
    y_sv = y(sv_idx, 1);
    y_nsv = y(nsv_idx, 1);

    cut = 99;
    x_axis = 0:1/cut:1;
    y_axis = zeros(size(x_axis));

    for i = 1:cut + 1
        K = zeros(n, 1);
        for j = 1:n
            K(j, 1) = exp( - norm(x_axis(i) - x(j), 2) ^ 2 / 2 / h ^ 2);
        end
        y_axis(i) = dot(K, p);
    end

    figure % opens new figure window
    plot(x_axis, y_axis);
    hold on;
    scatter(x_sv, y_sv, '*');
    hold on;
    scatter(x_nsv, y_nsv);
    legend('prediction_curve', 'support_vectors');

```

2.2 Support Kernel Machines (20 Points)

In SVM, the kernel function can be viewed as a similarity measure between data points. In some classification scenarios, features may come from different sources or modalities, e.g. in some tasks the data may contain both image features and text features. In that case, since these are different representations, they have different measures of similarity corresponding to different kernels. In such a case, we want to learn a

combination of kernels instead of using a single kernel. There is significant amount of work in combining kernels, here we adapt the notations in [1].

We begin by considering a linear case of Support Kernel Machine (SKM). Suppose the data points $x_i \in \mathcal{X} = \mathbb{R}^k$. We also assume we are given a decomposition of $\mathbb{R}^k = \mathbb{R}^{k_1} \times \dots \times \mathbb{R}^{k_m}$, so that each data point x_i can be decomposed into m block components, i.e. $x_i = (x_{1i}, \dots, x_{mi})$ where each x_{ji} is in general a vector. In real tasks, each block may correspond to a certain kind of representation, e.g. x_{1i} may correspond to image features and x_{2i} may be text features.

Our goal is to find a linear classifier of the form $y = \text{sign}(w^T x + b)$ where w has the same block decomposition $w = (w_1, \dots, w_m) \in \mathbb{R}^{k_1 + \dots + k_m}$. Recall that in linear SVM the objective is:

$$\underset{\substack{w \in \mathbb{R}^{k_1} \times \dots \times \mathbb{R}^{k_m} \\ \xi_i \geq 0, b \in \mathbb{R}}}{\text{minimize}} \quad \frac{1}{2} \|w\|_2^2 + C \sum_{i=1}^n \xi_i \quad (19)$$

$$\text{subject to} \quad y_i \left(\sum_j w_j^T x_{ji} + b \right) \geq 1 - \xi_i, \quad \forall i \in \{1, \dots, n\} \quad (20)$$

In SKM, we encourage the sparsity of the vector w at the level of blocks. The primal problem for the SKM is defined as:

$$\underset{\substack{w \in \mathbb{R}^{k_1} \times \dots \times \mathbb{R}^{k_m} \\ \xi_i \geq 0, b \in \mathbb{R}}}{\text{minimize}} \quad \frac{1}{2} \left(\sum_{j=1}^m d_j \|w_j\|_2 \right)^2 + C \sum_{i=1}^n \xi_i \quad (21)$$

$$\text{subject to} \quad y_i \left(\sum_j w_j^T x_{ji} + b \right) \geq 1 - \xi_i, \quad \forall i \in \{1, \dots, n\} \quad (22)$$

where $d_j > 0$ can be seen as constant.

1. By introducing dual variables $\alpha_i \geq 0$ and $\beta_i \geq 0$, we get the Lagrangian function

$$\mathcal{L} = \frac{1}{2} \left(\sum_{j=1}^m d_j \|w_j\|_2 \right)^2 + C \sum_{i=1}^n \xi_i - \sum_{i=1}^n \alpha_i \left(y_i \left(\sum_{j=1}^m w_j^T x_{ji} + b \right) - 1 + \xi_i \right) - \sum_{i=1}^n \beta_i \xi_i \quad (23)$$

Denote $\gamma = \sum_{j=1}^m d_j \|w_j\|_2$.

- (a) Show that at the minimum of the Lagrangian function, i.e. $w = \text{argmin}_w \mathcal{L}$ for this and the following questions,

$$\|w_j\|_2 d_j \gamma = w_j^T \sum_{i=1}^n \alpha_i y_i x_{ji}, \quad \forall j \in \{1, \dots, m\} \quad (24)$$

Answer: It is apparent that the conclusion holds when $w_j = 0$. When $w_j \neq 0$, we have

$$\frac{\partial \mathcal{L}}{\partial w_j} = \gamma d_j \frac{w_j}{\|w_j\|_2} - \sum_{i=1}^n \alpha_i y_i x_{ji} \quad (25)$$

At the minimum, $\frac{\partial \mathcal{L}}{\partial w_j} = 0$, we have

$$\gamma d_j \frac{w_j}{\|w_j\|_2} = \sum_{i=1}^n \alpha_i y_i x_{ji} \quad (26)$$

Multiply both sides by w_j^T , we have

$$\|w_j\|_2 \gamma d_j = w_j^T \sum_{i=1}^n \alpha_i y_i x_{ji} \quad (27)$$

- (b) Show that $\|\sum_{i=1}^n \alpha_i y_i x_{ji}\|_2 \leq d_j \gamma$, $\forall j \in \{1, \dots, m\}$.

Note: for (b) you can get full credit if you only consider $w_j \neq 0$, but you can get 5 extra points if you include $w_j = 0$ case in your proof. A hint is that \mathcal{L} is not differentiable w.r.t. w_j if $w_j = 0$, and you may refer to [this link](#) for how to deal with that case by using $\partial\|x\|_2 = \{g : \|g\|_2 \leq 1\}$ if $x = 0$.

Answer: If $w_j \neq 0$. Use the intermediate results of (a), we have

$$\gamma d_j \frac{w_j}{\|w_j\|_2} = \sum_{i=1}^n \alpha_i y_i x_{ji} \quad (28)$$

Take the l_2 norm of both sides, we have

$$\gamma d_j = \left\| \sum_{i=1}^n \alpha_i y_i x_{ji} \right\|_2 \quad (29)$$

If $w_j = 0$, $0 \in \partial L$, we have that for some g with $\|g\|_2 \leq 1$,

$$\gamma d_j g = \sum_{i=1}^n \alpha_i y_i x_{ji} \quad (30)$$

Take the l_2 norm of both sides, we have

$$\gamma d_j \|g\|_2 = \left\| \sum_{i=1}^n \alpha_i y_i x_{ji} \right\|_2 \quad (31)$$

As $\|g\|_2 \leq 1$, we have $d_j \gamma \geq \left\| \sum_{i=1}^n \alpha_i y_i x_{ji} \right\|_2$.

- (c) Show that

- if $\left\| \sum_i \alpha_i y_i x_{ji} \right\|_2 < d_j \gamma$, then $w_j = 0$,
- if $\left\| \sum_i \alpha_i y_i x_{ji} \right\|_2 = d_j \gamma$, then $\exists \eta_j \geq 0$, such that $w_j = \eta_j \sum_i \alpha_i y_i x_{ji}$.

Answer: If $w_j \neq 0$, using the intermediate results of (b) we have

$$\gamma d_j = \left\| \sum_{i=1}^n \alpha_i y_i x_{ji} \right\|_2 \quad (32)$$

So if $\left\| \sum_{i=1}^n \alpha_i y_i x_{ji} \right\|_2 < d_j \gamma$, $w_j = 0$.

On the other hand, we have proved in (a) that if $w_j \neq 0$, $\gamma d_j \frac{w_j}{\|w_j\|_2} = \sum_{i=1}^n \alpha_i y_i x_{ji}$. Let η_j be $\frac{\|w_j\|_2}{\gamma d_j}$, then $w_j = \eta_j \sum_{i=1}^n \alpha_i y_i x_{ji}$.

2. Recall from homework 2 that ℓ_1 norm can encourage sparsity. Explain the effect of the regularization term $\frac{1}{2}(\sum_{j=1}^m d_j \|w_j\|_2)^2$.

Answer: Recall that l_1 regularizer has the effect of putting some coefficients to zero, and if $d_j \|w_j\|_2 = 0$, then $w_j = 0$. Therefore some blocks w_j tend to zero due to this regularizer (group lasso).

3. Now we extend the above analysis to a kernelized version. Assume that we have a mapping $\phi : \mathcal{X} \rightarrow \mathbb{R}^f$ which is generally a non-linear function. We assume that $\phi(x)$ has m block components $\phi(x) = (\phi_1(x), \dots, \phi_m(x))$, and we also assume w has the same decomposition $w = (w_1, \dots, w_m)$. Show that at the minimum of the Lagrangian function, $\exists \eta_j \geq 0$ such that $w_j = \eta_j \sum_{i=1}^n \alpha_i y_i \phi_j(x_i)$.

Answer: Since it is a mapping for x , we can simply substitute x_{ji} with $\phi_j(x_i)$ and every conclusions above still hold. The proof completes by using the conclusion in 1(c) and replace x_{ji} with $\phi_j(x_i)$.

2.3 SVM Error Analysis (5 Points)

In this problem, we want to analyze the error of SVM classification. Assume that we have n data points $(x_i, y_i)_{i=1}^n$ where $x_i \in \mathbb{R}^m$ and $y_i = 1, \dots, K$. Assume that we train an SVM classifier $f_{(x_1, y_1), \dots, (x_n, y_n)}$ on these n data points.

For a randomly drawn test data point (x_{n+1}, y_{n+1}) , the prediction is $y_{n+1}^{\text{pred}} = f_{(x_1, y_1), \dots, (x_n, y_n)}(x_{n+1})$. We assume that the n training data points and the test data point (x_{n+1}, y_{n+1}) are drawn i.i.d from some unknown underlying distribution. The expected error rate is defined as:

$$\text{err} = \mathbb{E}_{(x_1, y_1), \dots, (x_n, y_n)} \mathbb{E}_{(x_{n+1}, y_{n+1})} [\mathbf{1} \{f_{(x_1, y_1), \dots, (x_n, y_n)}(x_{n+1}) \neq y_{n+1}\}] \quad (33)$$

where the indicator function $\mathbf{1} \{A\} = 1$ if A is true, otherwise 0.

1. Show the the expected error rate is equal to the expectation of leave-one-out cross validation error for $n + 1$ data points.

Answer: As the points are i.i.d, we have for any i ,

$$\begin{aligned} & \mathbb{E}_{(x_1, y_1), \dots, (x_n, y_n)} \mathbb{E}_{(x_{n+1}, y_{n+1})} [\mathbf{1} \{f_{(x_1, y_1), \dots, (x_n, y_n)}(x_{n+1}) \neq y_{n+1}\}] \\ &= \mathbb{E}_{(x_1, y_1), \dots, (x_{n+1}, y_{n+1})} [\mathbf{1} \{f_{(x_1, y_1), \dots, (x_{i-1}, y_{i-1}), (x_{i+1}, y_{i+1}), \dots, (x_{n+1}, y_{n+1})}(x_i) \neq y_i\}] \end{aligned} \quad (34)$$

By looping through $i = 1$ to $i = n + 1$ and averaging, we have

$$\begin{aligned} & \mathbb{E}_{(x_1, y_1), \dots, (x_n, y_n)} \mathbb{E}_{(x_{n+1}, y_{n+1})} [\mathbf{1} \{f_{(x_1, y_1), \dots, (x_n, y_n)}(x_{n+1}) \neq y_{n+1}\}] \\ &= \frac{1}{n+1} \sum_{i=1}^{n+1} \mathbb{E}_{(x_1, y_1), \dots, (x_{n+1}, y_{n+1})} [\mathbf{1} \{f_{(x_1, y_1), \dots, (x_{i-1}, y_{i-1}), (x_{i+1}, y_{i+1}), \dots, (x_{n+1}, y_{n+1})}(x_i) \neq y_i\}] \\ &= \mathbb{E}_{(x_1, y_1), \dots, (x_{n+1}, y_{n+1})} \left[\frac{1}{n+1} \sum_{i=1}^{n+1} \mathbf{1} \{f_{(x_1, y_1), \dots, (x_{i-1}, y_{i-1}), (x_{i+1}, y_{i+1}), \dots, (x_{n+1}, y_{n+1})}(x_i) \neq y_i\} \right] \end{aligned} \quad (35)$$

The right hand side is the expectation of the leave on out error rate.

2. In the lecture, we have the statement that “the leave-one-out cross-validation error does not depend on the dimensionality of the feature space but only on the number of support vectors”. Show that this statement is true by explaining why

$$\text{err}_{\text{loocv}} \leq \frac{n_s}{n+1} \quad (36)$$

where $\text{err}_{\text{loocv}}$ is the leave-one-out cross-validation error for training set $(x_i, y_i)_{i=1}^{n+1}$, n_s is the number of support vectors.

Answer: If we remove a non-support vector and retrain the SVM, the decision boundary before removing it is still the optimum, and if we use that decision boundary, the previous non-support vector will not be misclassified. So the leave one out error possibly occurs when we remove the support vectors, hence

$$\text{err}_{\text{loocv}} \leq \frac{n_s}{n+1} \quad (37)$$

Acknowledgements

Thanks for Desai's homework for providing this solution.

References

- [1] Francis R Bach, Gert R.G. Lanckriet, and Michael I Jordan. Multiple kernel learning, conic duality, and the smo algorithm. In *ICML*, 2004. 10