

This homework is due **Monday, November 13 at 10pm.**

## 1 Getting Started

You may typeset your homework in latex or submit neatly handwritten and scanned solutions. Please make sure to start each question on a new page, as grading (with Gradescope) is much easier that way! Deliverables:

1. Submit a PDF of your writeup to assignment on Gradescope, “HW[n] Write-Up”
2. Submit all code needed to reproduce your results, “HW[n] Code”.
3. Submit your test set evaluation results, “HW[n] Test Set”.

After you’ve submitted your homework, be sure to watch out for the self-grade form.

- (a) Before you start your homework, write down your team. Who else did you work with on this homework? List names and email addresses. In case of course events, just describe the group. How did you work on this homework? Any comments about the homework?

- (b) Please copy the following statement and sign next to it:

*I certify that all solutions are entirely in my words and that I have not looked at another student’s solutions. I have credited all external sources in this write up.*

## 2 Kernel SVM and Kernel Ridge Regression

In lecture, kernels were discussed in a way that avoided formality.

In this problem, we will give a derivation of kernel SVM and kernel ridge regression from the view of function approximation with penalization. The objective function of a linear SVM can be interpreted as Hinge Loss combined with  $L_2$  regularization over the space of linear functions. The objective function of a kernel SVM can be interpreted in the same way: Hinge Loss plus  $L_2$  regularization over the Hilbert space of functions defined by a kernel function.

Assume we are doing classification or regression over  $\mathbb{R}^d$ . We first introduce the following abstract vector space:

$$H = \{f : \mathbb{R}^d \rightarrow \mathbb{R} : f(x) = \sum_{m=1}^M \alpha_m k(x, y_m) : \alpha_i \in \mathbb{R}, M \in \mathbb{N}, y_m \in \mathbb{R}^d\}, \quad (1)$$

where  $k(x, y) : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$  is a kernel function that satisfies the property  $k(x, y) = k(y, x)$  for any  $x, y \in \mathbb{R}^d$  and for any distinct  $y_1, y_2, \dots, y_M \in \mathbb{R}^d$ , the matrix  $K \in \mathbb{R}^{M \times M}$  defined by  $K_{ij} = k(y_i, y_j)$  is positive definite.

This is a vector space since it has a zero and linear combinations make sense. It can be infinite-dimensional however since there are lots of possible  $y_m$ .

- (a) Now we introduce an inner product on the above vector space  $H$ . We define the inner product between any two functions

$$f(x) = \sum_{m=1}^M \alpha_m k(x, y_m), g(x) = \sum_{s=1}^S \beta_s k(x, x_s) \in H,$$

in  $H$  as

$$\langle f, g \rangle_H = \sum_{m=1}^M \sum_{s=1}^S \alpha_m \beta_s k(y_m, x_s). \quad (2)$$

**Show that the defined inner product is valid.** That is, it satisfies the symmetry, linearity and positive-definiteness properties stated below: For any for any  $f, g, h \in H$  and any  $a \in \mathbb{R}$ , we have

- $\langle f, g \rangle_H = \langle g, f \rangle_H$ .
- $\langle af, g \rangle_H = a \langle f, g \rangle_H$  and  $\langle f + h, g \rangle_H = \langle f, g \rangle_H + \langle h, g \rangle_H$ .
- $\langle f, f \rangle_H \geq 0$ ;  $\langle f, f \rangle_H = 0$  if and only if  $f$  is a constant zero function.

**What is the norm of the function  $f$ ?** (The natural norm in an inner product space is defined as  $\|f\|_H = \sqrt{\langle f, f \rangle_H}$ .)

- (b) **Show that the defined inner product has the reproducing property**  $\langle k(x, \cdot), k(y, \cdot) \rangle_H = k(x, y)$ , where we take  $\cdot \rightarrow k(x, \cdot)$  and  $\cdot \rightarrow k(y, \cdot)$  as two functions in  $H$ . In other words, the inner-product is natural for the vector space as defined. **Conclude that**

$$\langle k(\cdot, x_i), f \rangle_H = f(x_i). \quad (3)$$

(For those who have taken signal processing courses, you can see here that what this family of definitions is trying to do is to parallel the example of the sifting property that we know from signal processing.)

- (c) The completion of this inner product space is a Hilbert space called a Reproducing Kernel Hilbert Space (RKHS), but in this problem, this knowledge is not required. From now on, we will work in the completion of  $H$ . We will also call it  $H$  for notational simplicity. We assume we have the following knowledge of a Hilbert space: Suppose  $M$  is a finite-dimensional subspace of a Hilbert space  $H$ . Then any element  $f$  in the full space  $H$  has a unique representation as the sum of an element of  $M$  and an element that is orthogonal to any element in  $M$ . That is,

$$f = m + g, \quad (4)$$

for some  $m \in M$  and some  $g$  such that  $\langle m', g \rangle_H = 0$  for all  $m' \in M$ . (In other words, it behaves exactly like the vector spaces that you are used to.)

Below we introduce a general optimization problem over the Hilbert space  $H$ . We will see many kernelized machine learning algorithms, including kernel SVM, can be written in the following form: Given a data set with  $N$  points  $x_i, y_i, i = 1, \dots, N$ , the optimization problem is

$$\min_{f \in H} \frac{1}{N} \sum_{i=1}^N L(y_i, f(x_i)) + \lambda \|f\|_H^2, \quad (5)$$

where  $L$  is any loss function on pairs of real numbers. (Remember, the  $y_i$  here in this part are real numbers. They are not training points in  $d$ -dimensional space. Those are the  $x_i$ .)

**Show that the minimizing solution to the problem has the form**

$$f(x) = \sum_{i=1}^N \alpha_i k(x, x_i).$$

That is, the solution can be expressed as a weighted sum of kernel functions based on the training points. This phenomenon that reduces an infinite-dimensional optimization problem to be finite-dimensional is called the *kernel property*. Hint: Define  $M = \{\sum_{n=1}^N \alpha_n k(x, x_n) : \alpha_i \in \mathbb{R}\}$  to be the subspace of interest.

- (d) (Kernel SVM) The kernel SVM, is nothing but defining the loss function  $L$  concretely as a Hinge loss:

$$L(y, f(x)) = \max(0, 1 - yf(x)). \quad (6)$$

In other words, kernel SVM is

$$\min_{f \in H} \frac{1}{N} \sum_{i=1}^N \max(0, 1 - y_i f(x_i)) + \lambda \|f\|_H^2. \quad (7)$$

**Show kernel SVM is of the form**

$$\min_{\alpha \in \mathbb{R}^d} \frac{1}{N} \sum_{i=1}^N \max(0, 1 - y_i \sum_{j=1}^N \alpha_j k(x_i, x_j)) + \lambda \alpha^T K \alpha. \quad (8)$$

- (e) (Kernel ridge regression) Take  $L$  as  $l_2$  loss, that is,  $L(a, b) := \|a - b\|_2^2$ . **Show that optimization problem of kernel ridge regression has the following form**

$$\min_{\alpha \in \mathbb{R}^d} \frac{1}{N} \|Y - K\alpha\|_2^2 + \lambda \alpha^T K \alpha, \quad (9)$$

where  $Y \in \mathbb{R}^n$  with the  $i$ th element of  $Y$  being  $y_i$ . **Derive a closed form solution for the kernel ridge regression:**

$$\alpha = (K + \lambda N I_N)^{-1} Y, \quad (10)$$

where  $I_N$  is an  $n$ -dimensional identity matrix.

- (f) (Polynomial Regression from a kernelized view) In this part, we will show that polynomial regression with a particular Tikhonov regularization is the same as kernel ridge regression with a polynomial kernel for second-order polynomials. Recall that a polynomial kernel function on  $\mathbb{R}^d$  is defined as

$$k(x, y) = (1 + x^T y)^2, \quad (11)$$

for any  $x, y \in \mathbb{R}^d$ . Given a dataset  $(x_i, y_i)$  for  $i = 1, 2, \dots, n$ . **Show the solution to kernel ridge regression is the same as the least square solution to polynomial regression for  $d = 2$  given the right choice of Tikhonov regularization for the polynomial regression.** That is, show for any new point  $x$  given in the prediction stage, both methods give the same  $y$ . What is the Tikhonov regularization matrix here?

Hint: You may or may not use the following matrix identity:

$$A(aI_d + A^T A)^{-1} = (aI_N + AA^T)^{-1} A, \quad (12)$$

for any matrix  $A \in \mathbb{R}^{n \times d}$  and any positive real number  $a$ .

- (g) (Bonus) In general, for any polynomial regression with  $p$ th order polynomial on  $\mathbb{R}^d$ , with a selected Tikhonov regression, we can show the equivalence between it and kernel ridge regression with a polynomial kernel of order  $p$ . (You are not required to show this.) **Comment on the computational complexity of doing least squares for polynomial regression with a Tikhonov regression directly and that of doing kernel ridge regression in the training stage.** (That is, the complexity of finding  $\alpha$  and finding  $w$ .) **Compare with the computational complexity of actually doing prediction as well.**

### 3 Nearest Neighbors, from A to Z

For this problem, we will use data from the UN to have some fun with the nearest neighbors algorithm. A lot of the code you will need has been provided for you.

The data we are using is called the “World Values Survey.” It consists of survey data collection over several years from almost all countries. The survey asked “Which of these are most important for you and your family?” There were 16 possible responses, including needs like “Freedom from Discrimination and Persecution” and “Better Transport and Roads.” The data reported is the fraction of responses in each country that chose each option.

We would like to use these 16 features of each country (the citizen's responses to the survey) to predict that country's HDI (Human Development Index). In reality, the HDI is a complex measure which takes into account lots of data about a country, including factors like life expectancy, education, per capita income, etc. Intuitively, though, you would expect citizens of countries with different HDI to have different priorities. For that reason, predicting the HDI from survey data might be a reasonable endeavor.

Note that throughout the problem we will be using RMSE, which stands for Root Mean Squared Error.

- (a) (Bonus): **Fill out the “Berkeley F2017 Values Survey.”** The purpose of this is so that you have a sense of how the data was generated, a useful first step in any ML problem. Just for fun, at the end of this problem we will attempt to predict what the HDI of Berkeley would be if it were its own country.
- (b) First, we should do some basic data exploration. **Compute the correlation of each feature with HDI. Which feature is the most positively correlated with HDI? Which feature is the most negatively correlated with HDI? Which feature is the least correlated with HDI (closest to 0)?**
- (c) **For each of these three features identified in (b) (most positively correlated, most negatively correlated, least correlated), plot “HDI versus [Feature].”** You will create three plots in total. **What do you observe?**
- (d) Let's visualize the data a bit more. **Plot the data in its first two PCA dimensions, colored by HDI.** The code to do this has been provided for you.
- (e) Now, let's use our first ML technique. **Use the code provided to train and cross-validate ridge regression to predict a country's HDI from its citizens' world values survey responses.** You may need to modify the hyper-parameters. **What is the best RMSE?**
- (f) Let's try another ML technique. **Use the code provided to train and cross-validate lasso regression to predict a country's HDI from its citizens' world values survey responses.** You may need to modify the hyper-parameters. **What is the best RMSE?**
- (g) **Examine the model returned by lasso regression (that is, the 16 feature weights). Does lasso regression indeed give more 0 weights?**
- (h) In lecture, we covered  $k$  Nearest Neighbors for classification problems. We decided that the class of a test point would be the plurality of the classes of the  $k$  nearest training points. That algorithm makes sense when the outputs are discrete, so we can vote. Here, the outputs are continuous. **How would you adapt the  $k$  Nearest Neighbors algorithm for a regression problem?**
- (i) **Which countries are the 7 nearest neighbors of the USA (in order)?**
- (j) The most important meta-parameter of  $k$  nearest neighbors is  $k$  itself. **Plot the RMSE of kNN regression versus  $k$ , where  $k$  is the number of neighbors. What is the best value of  $k$ ? What is the RMSE?**

- (k) **Explain your plot in (j) in terms of bias and variance.** This is tricky, so take some time to think about it. Think about the spirit of bias and variance more than their precise definitions.
- (l) We do not need to give every neighbor an equal weight. Maybe closer neighbors are more relevant. For the sake of this problem, let's weight each neighbor by the inverse of its distance to the test point. **Plot the RMSE of kNN regression with distance weighting versus  $k$ , where  $k$  is the number of features. What is the best value of  $k$ ? What is the RMSE?**
- (m) One of the challenges of  $k$  Nearest Neighbors is that it is very sensitive to the scale of the features. For example, if one feature takes on values 0 or 0.1 and another takes on values 0 or 10, then neighbors will almost certainly agree in the second feature. **Which countries are the 7 nearest neighbors of the USA after scaling (in order)? Compare your result to (i).**
- (n) **Add scaling to your  $k$  nearest neighbors pipeline (continue to use distance weighting). Plot RMSE versus  $k$ . What is the best value for  $k$ ? What is the RMSE?**
- (o) (BONUS): **Rather than scaling each feature to have unit variance, explore ways of scaling the features non-uniformly. How much does this help, if at all?**
- (p) You have been given a set of test features: countries where the responses to the world values survey are given but the HDI is not known. **Using the best model developed so far, predict the HDI values of the countries in the test set. Submit your predictions on Gradescope.**
- (q) So far we have dealt with the regression problem. Let's take a brief look at classification. A naive classifier is a classifier which disregards the features and just classifies everything as belonging to a single class. **In any classification problem with  $k$  classes, what accuracy are we guaranteed to get with the best naive classifier?** (Hint: there are  $k$  possible naive classifiers. Use the pigeonhole principle).
- (r) We will split countries into two groups: high HDI (more than 0.7) and low HDI (less than 0.7). **Plot the countries by their first two PCA dimensions again, but now color them by class.**
- (s) Examine the graph generated in (r). **How well do you think a linear SVM would do in classification?**
- (t) We will use an SVM classifier to predict whether a country's HDI is "high" or "low" based on the responses of their citizens to the World Values Survey. **Use the code provided to train and cross-validate an SVM classifier using a linear kernel. What is the accuracy of the classifier?**
- (u) We are going to modify the classifier from (t). **Add a PCA step and Scaling step to the SVM pipeline. Your hyper-parameter search should now be over all possible dimensions for the PCA reduction. Does the accuracy improve?**
- (v) Change the kernel in (t) from linear to "radial basis function" (rbf). For this part, do not use PCA or Scaling. **What is the accuracy?**

- (w) Now we are going to use  $k$  Nearest Neighbors for the same task. That is, we would like to predict whether a country's HDI is "high" or "low" based on the responses of their citizens to the World Values Survey. **Train and cross-validate a  $k$  Nearest Neighbors classifier using distance weighting. What is its accuracy? Does scaling help?**
- (x) (Bonus): Towards the end of the week, we will post the "Berkeley F2017 Values Survey." **If Berkeley were an independent country, what do you predict its HDI would be?**
- (y) (Bonus): **Describe how you would use kNN to revisit the sensor location problem from previous homework. How well do you think it will work?**
- (z) (Bonus): **What did you learn from this problem? Do you have any useful feedback for the problem author?**

## 4 Your Own Question

**Write your own question, and provide a thorough solution.**

Writing your own problems is a very important way to really learn material. The famous "Bloom's Taxonomy" that lists the levels of learning is: Remember, Understand, Apply, Analyze, Evaluate, and Create. Using what you know to create is the top-level. We rarely ask you any HW questions about the lowest level of straight-up remembering, expecting you to be able to do that yourself. (e.g. make yourself flashcards) But we don't want the same to be true about the highest level.

As a practical matter, having some practice at trying to create problems helps you study for exams much better than simply counting on solving existing practice problems. This is because thinking about how to create an interesting problem forces you to really look at the material from the perspective of those who are going to create the exams.

Besides, this is fun. If you want to make a boring problem, go ahead. That is your prerogative. But it is more fun to really engage with the material, discover something interesting, and then come up with a problem that walks others down a journey that lets them share your discovery. You don't have to achieve this every week. But unless you try every week, it probably won't happen ever.