

6.867 Machine learning and neural networks

Problem set 3

Deadline: November 2, 11am in recitation

What and how to turn in?

Turn in short written answers to the questions explicitly stated, and when requested to explain or prove. Do **not** turn in answers when requested to “think”, “consider”, “try” or “experiment” (except when specifically instructed). You are encouraged to think about, and turn in, answers to questions marked “*optional*”—they will be read and corrected, but a grade will not be recorded for them.

Turn in all MATLAB code explicitly requested, or that you used to calculate requested values. It should be clear exactly what command was used to get the answer to each question.

To help the graders (including yourself...), please be neat, answer the questions briefly, and in the order they are stated. **Staple each “Problem” separately**, and be sure to write your name on the top of every page.

Problem 1: Boosting

This question is about AdaBoost, the simple boosting algorithm presented in the lectures. We assume that there is a method unknown to us that produces a component (weak) hypothesis $h(\mathbf{x})$ in response to a weighted training set. The resulting weak classifier should do at least a bit better than random guessing on the weighted training set it was trained on. We make no assumptions about how well it does on other possibly differently weighted training sets.

Let $D = \{(\mathbf{x}_1, y_1) \dots, (\mathbf{x}_n, y_n)\}$ be the training set of examples \mathbf{x}_i and binary (± 1) labels y_i . We denote the weights on the training examples at the beginning of the k^{th} boosting iteration as $p_k(1), \dots, p_k(n)$, where $\sum_{i=1}^n p_k(i) = 1$. Unlike in the lecture, the component hypotheses $h(\mathbf{x})$ can produce a real valued output (confidence rated classifiers). The sign of the output indicates the label and the magnitude specifies a measure of “confidence” in the classification decision. So, for example, $h(\mathbf{x}) = 10$ would be interpreted as a relatively confident prediction of label 1.

Now, starting with equal weights $p_1(i) = 1/n$, AdaBoost generates a sequence of hypotheses $h_1(\mathbf{x}), \dots, h_m(\mathbf{x})$, where each $h_k(\mathbf{x})$ was trained with a different set of example weights. After having generated a hypothesis $h_k(\mathbf{x})$ in response to weights $p_k(i)$ at the k^{th} iteration, the example weights are updated according to

$$p_{k+1}(i) = c \cdot p_k(i) \exp(-\hat{\alpha}_k y_i h_k(\mathbf{x}_i)), \quad i = 1, \dots, n \quad (1)$$

where c is the normalization constant ensuring that $\sum_{i=1}^n p_{k+1}(i) = 1$ and $\hat{\alpha}_k$ gives the votes that we assign to the new component classifier $h_k(\mathbf{x})$. We would like to assign these votes so that the new component classifier $h_k(\mathbf{x})$ is at chance performance relative to the new weights $p_{k+1}(i)$. One way to express this is as follows:

$$\sum_{i=1}^n p_{k+1}(i) y_i h_k(\mathbf{x}_i) = 0 \quad (2)$$

1. Show that if $h_k(\mathbf{x}_i)$ generates only binary ± 1 outputs, then the above condition corresponds to requiring that the training error is exactly 0.5 relative to the new weights

Things are a bit different here since the component classifier can generate real valued outputs. We can view the condition instead as a way of *decorrelating* the predictions and the labels relative to the new weights.

2. Show that if we choose $\hat{\alpha}_k$ as the minimizing argument of

$$J(\alpha) = \log \left(\sum_{i=1}^n p_k(i) \exp\{-\alpha y_i h_k(\mathbf{x}_i)\} \right) \quad (3)$$

we will indeed ensure that the decorrelation condition (2) holds for the new weights. (Hint: Do not attempt to compute $\hat{\alpha}_k$. Instead, set the derivative to zero and use the resulting equation, which $\hat{\alpha}_k$ must satisfy, in order to show (2) is satisfied.

This should look a bit suspicious... as if there were an objective function that the boosting algorithm were minimizing at each iteration. This is indeed the case. Let $H_m(\mathbf{x})$ be the combined hypothesis resulting from m boosting iterations:

$$H_m(\mathbf{x}) = \hat{\alpha}_1 h_1(\mathbf{x}) + \dots + \hat{\alpha}_m h_m(\mathbf{x}) \quad (4)$$

We'd like you to show that

$$J(H_m) = \log \left(\sum_{i=1}^n \exp\{-y_i H_m(\mathbf{x}_i)\} \right) \quad (5)$$

serves as an objective function for the boosting algorithm. In other words, we claim that every time we add a new component hypothesis to the combined classifier $H_m(\mathbf{x})$, we decrease the objective.

3. Show that $J(H_m) \geq J(H_{m+1})$, where

$$H_{m+1}(\mathbf{x}) = H_m(\mathbf{x}) + \hat{\alpha}_{m+1}h_{m+1}(\mathbf{x}), \quad (6)$$

$h_{m+1}(\mathbf{x})$ is the new component hypothesis, and $\hat{\alpha}_{m+1}$ is optimized as shown above. (Hint: expand the weights $p_m(i)$ and use the optimization of $\hat{\alpha}_{m+1}$).

Having now put some effort into understanding the boosting algorithm, let's explore a bit how it behaves in practice. We have provided you with MATLAB code that finds and evaluates (confidence rated) decision stumps. These are the hypothesis that our boosting algorithm assumes we can generate. The relevant MATLAB files are "boost_digit.m", "boost.m", "eval_boost.m", "find_stump.m", "eval_stump.m". You'll only have to make minor modifications to "boost.m" and, a bit later, to "eval_boost.m" and "boost_digit.m" to make these work.

4. Complete the weight update in "boost.m" and run "boost_digit" to plot the training and test errors for the combined classifier as well as the corresponding training error of the decision stump, as a function of the number of iterations. Are the errors what you would expect them to be? Why or why not?

We will now investigate the classification margins of training examples. Recall that the classification margin of a training point in the boosting context reflects the "confidence" in which the point was classified correctly. You can view the margin of a training example as the difference between the weighted fraction of votes assigned to the correct label and those assigned to the incorrect one. Note that this is not a geometric notion of "margin" but one based on votes. The margin will be positive for correctly classified training points and negative for others.

5. Modify "eval_boost.m" so that it returns normalized predictions (normalized by the total number of votes). The resulting predictions should be in the range $[-1, 1]$. Fill in the missing computation of the training set margins in "boost_digit.m" (that is, the classification margins for each of the training points). You should also uncomment the plotting script for cumulative margin distributions (what is plotted is, for each $-1 < r < 1$ on the horizontal axis, what fraction of the training points have a margin of at least r). Explain the differences between the cumulative distributions after 4 and 16 boosting iterations.

Problem 2: support vector machines

Lagrange multipliers and optimization problems

We'll present here a very simple tutorial example of using and understanding Lagrange multipliers. Let w be a scalar parameter we wish to estimate and x a fixed scalar. We wish

to solve the following (tiny) SVM like optimization problem:

$$\text{minimize } \frac{1}{2}w^2 \text{ subject to } wx - 1 \geq 0 \tag{7}$$

This is difficult only because of the constraint. We'd rather solve an unconstrained version of the problem but, somehow, we have to take into account the constraint. We can do this by including the constraint itself in the minimization objective as it allows us to twist the solution towards satisfying the constraint. We need to know how much to emphasize the constraint and this is what the Lagrange multiplier is doing. We will denote the Lagrange multiplier by α to be consistent with the SVM problem. So we have now constructed a new minimization problem (still minimizing with respect to w) that includes the constraint as an additional linear term:

$$J(w; \alpha) = \frac{1}{2}w^2 - \alpha(wx - 1) \tag{8}$$

The Lagrange multiplier α appears here as a parameter. You might view this new objective a bit suspiciously since we appear to have lost the information about what type of constraint we had, i.e., whether the constraint was $wx - 1 \geq 0$, $wx - 1 \leq 0$, or $wx - 1 = 0$. How is this information encoded? We can encode this by constraining the values of the Lagrange multipliers:

$$\begin{aligned} wx - 1 \geq 0 &\Rightarrow \alpha \geq 0 \\ wx - 1 \leq 0 &\Rightarrow \alpha \leq 0 \\ wx - 1 = 0 &\Rightarrow \alpha \text{ is unconstrained} \end{aligned}$$

Note, for example, that when the constraint is $wx - 1 \geq 0$, as we have above, large positive values of α will encourage choices of w that result in large positive values for $wx - 1$. This is because in the above objective, $J(w; \alpha)$, we try to minimize $-\alpha(wx - 1)$ in addition to $w^2/2$; minimizing $-\alpha(wx - 1)$ is the same as maximizing $\alpha(wx - 1)$ or $wx - 1$ since α is positive. Figure 1 tries to illustrate this effect. Assuming $x = 1$ we can plot the new objective function as a function of w for different values of α . Larger values of α clearly move the solution (minimizing w) towards satisfying $w - 1 \geq 0$ (we assume here that $x = 1$). Based on the figure we can see that setting $\alpha = 1$ produces just the right solution, i.e., $w^* = 1$, which satisfies the constraint $wx - 1 \geq 0$ (when $x = 1$) with minimal distortion of the original objective. There's no reason to consider negative values for α since they would push the solution away from satisfying our inequality constraint.

Effectively what we are doing here is solving a large number of optimization problems, once for each setting of the Lagrange multiplier α . Indeed, we can express the solution (the minimizing w) as a parametric function of α :

$$\frac{\partial}{\partial w} J(w; \alpha) = w - \alpha x = 0 \tag{9}$$

meaning that $w_\alpha^* = \alpha x$. We could now find the setting of α such that the constraint $w_\alpha^* x - 1 \geq 0$ is satisfied. There are multiple answers to this since larger values of α would




Figure 1: $J(w; \alpha)$ as a function of w for different values of α . The minimizing w values are indicated with dashed line segments. x was set to 1.

better satisfy the constraint. Finding the smallest (non-negative) α for which the constraint is satisfied would in this case produce the right solution (one corresponding to the minimal change of the original problem).

We can proceed a bit more generally, however, the way we handled the quadratic optimization problem for SVMs. Let's insert our solution w_α^* back into the objective function:

$$J(w_\alpha^*; \alpha) = \frac{1}{2}(w_\alpha^*)^2 - \alpha(w_\alpha^*x - 1) = \frac{1}{2}(\alpha x)^2 - \alpha(\alpha x^2 - 1) = \alpha - \frac{1}{2}(\alpha x)^2 \quad (10)$$

The result, which we denote as $J(\alpha)$, is a function of the Lagrange multiplier α only. Let's understand this function a bit better. In Figure 1, the values of the objective at the dashed lines correspond exactly to $J(w_\alpha^*; \alpha)$ or $J(\alpha)$, evaluated at $\alpha = 0, 1, 2$. Isn't it strange that the right solution ($\alpha = 1$) appears to yield the maximum of $J(\alpha)$? This is a very useful property. Let's verify this by finding the maximum of $J(\alpha)$ a bit more formally:

$$J(\alpha) = \alpha - \frac{1}{2}(\alpha x)^2 \quad (11)$$

$$\frac{\partial}{\partial \alpha} J(\alpha) = 1 - \alpha x^2 = 1 - w_\alpha^* x = 0 \quad (12)$$

where we have used our previous result $w_\alpha^* = \alpha x$. So, the constraint is satisfied with equality at the maximum of $J(\alpha)$. More rigorously, since $\alpha \geq 0$ in our setting, the maximum is obtained either at $\alpha = 0$ or at the point where $1 - w_\alpha^* x = 0$. We can express this more concisely by saying that their product vanishes, i.e., $\alpha(w_\alpha^* x - 1) = 0$ at the optimum. This is generally true, i.e., either the Lagrange multiplier is not used and $\alpha = 0$ (the constraint is satisfied without any modification) or the Lagrange multiplier is positive and the constraint is satisfied with equality.

The remaining question for us here is why

$$\text{maximize } \alpha - \frac{1}{2}(\alpha x)^2 \text{ subject to } \alpha \geq 0 \quad (13)$$

is any better than the problem we started with. The short answer is that the constraints here are very simple non-negativity constraints that are easy to deal with in the optimization. In the SVM context, we have another reason to prefer this formulation.

Examples of SVM training

We provided a skeleton for a SVM training routine `trainsvm.m`. The routine takes as input the training set, a constant C which will be an upper bound on the Lagrange multipliers, and a specification for a kernel function. Your first task is to complete this routine. Two parts of the routine are missing.

First, you should set up and solve the quadratic programming problem:

$$\text{minimize } \frac{1}{2} \sum_{i,j} \alpha_i \alpha_j y_i y_j K(\mathbf{x}_i, \mathbf{x}_j) - \sum_i \alpha_i$$

subject to: $\sum_i \alpha_i y_i = 0$ and for all i , $0 \leq \alpha_i \leq C$. Solve the quadratic programming using the MATLAB routine `quadprog`. (Note that in lecture we saw this as a maximization problem. MATLAB's `quadprog` solves quadratic *minimization* problems, and so we changed the sign of the objective function to get an equivalent minimization problem).

Next, add code to calculate the bias term w_0 . We have already provided code for finding vectors that lie on the margin (note that we account for numeric inaccuracies by ignoring α values very close to zero or to C). Although it is enough to use one of these input vectors for calculating w_0 , in order to increase accuracy, we will calculate w_0 using each one of these input vectors separately, and then take w_0 to be the median of all the calculated values (all of which should theoretically be the same).

Note that the routine also returns a list of the support vectors, and the vectors that lie exactly on the margin.

1. Why do we use only training vectors \mathbf{x}_i with $0 < \alpha_i < C$ for calculating w_0 ?
2. Complete the routine `trainsvm.m`.
3. Write a routine `svmpred.m` that classifies a set of input vectors based on a specified SVM classifier. The routine should take as input a matrix of input vectors, and the specifications of the SVM classifier. You may use the provided skeleton.

In studying SVM training, it will be useful to plot the decision boundary, the margin, and the support vectors. The provided routine `plotsvm2d.m` does just that. The routine takes as input the training set, the lists of support and margin vectors, the name of the prediction function (which will probably be `svmpred`), and the classifier specifications to be passed to the prediction function. (If you deviated from the provided skeleton of `svmpred.m`, you might need to change `plotsvm2d.m` appropriately).

The example data sets used in this question can be found in the file `svmexample.mat`, and can be read using:

```
>> load svmexample
```

4. We start with a simple example, given in `train1x`, `train1y`. Use the routines above to train a SVM using a linear kernel function $K(\mathbf{x}_1, \mathbf{x}_2) = (1 + \mathbf{x}_1^T \mathbf{x}_2)$ (note the provided `polyK.m`), and without bounding α_i from above. Turn in the plot created by `plotsvm2d.m`.
5. Why is it not necessary to bound the values of α_i from above in this case ?
6. What is the classification margin ? (Explain how you calculated it, providing any relevant MATLAB code)

We now move on to a slightly more complex data set given in `train2x`, `train2y`. This time it is necessary to bound the values of α_i by some finite C .

7. In order to understand what happens when we do not impose such a bound, try bounding with increasingly higher constants C and plot the value of the quadratic programming objective function as a function of C (turn in this plot). Why can we not solve the quadratic optimization problem when the values of α_i are unbounded ?
8. Recall that we are using Lagrange multipliers to solve a minimization problem on the weights \mathbf{w} , where we would like to minimize $\|\mathbf{w}\|^2$ subject to $y_i(\mathbf{w}^T \mathbf{x}_i + w_0) \geq 1$ for all training examples i . In this case, no weight vector satisfies all the constraints. How is this addressed by bounding the Lagrange multipliers ?
9. Train a SVM using the a linear kernel, imposing a bound of $C = 5$. Turn in the `plotsvm2d` plot. How many support vectors are there ? How many training points are misclassified ? How many training points do not satisfy the SVM constraints $y_i(w^T \mathbf{x}_i + w_0) \geq 1$?
10. In order to enrich the space of allowable classifiers, we would like to use a kernel function representing a more complex feature space. Try training a SVM, on the same example data, using polynomial kernels of the form $(1 + \mathbf{x}_1^T \mathbf{x}_2)^p$, with various settings of p , and look at the resulting graphs (no need to turn them in). What is the lowest degree under which the training set is separable ?

Also try using the Gaussian kernel $e^{-\frac{|\mathbf{x}_1 - \mathbf{x}_2|^2}{2}}$ (given in `expK.m`).

11. Note that in many of the graphs you produced, the graphical margin, as it appears in the graphs, is not uniform as with the linear kernel: it is narrower in some areas and wider in others. Why does this happen ?

12. Even if the training set is separable, it is sometimes useful to set a bound on the values of α_i . Train a SVM using a fourth degree kernel $(1 + \mathbf{x}_1^T \mathbf{x}_2)^4$, once without a bound, and once with a bound $C = 2$. Turn in both resulting graphs, and the margin implied by the SVM weights (i.e. the classification margins for those training points satisfying the SVM constraints).
13. What is the affect of imposing a low value for C ? In this example, what happened when C was set to two ?
14. Using the training data `train3x`, `train3y`, investigate training a SVM with different kernels and different values of C . Is the training set separable using a Gaussian kernel ? Do you think the classification given by a SVM with a Gaussian kernel would generalize well ?

Problem 3: VC-dimentionality

In this problem, we will investigate the VC-dimension of various sets of classifiers. In this context, a *classifier* is a function from some input space to the binary class labels $+1, -1$. A classifier can also be described as a subset of the input space which gets the label $+1$. For example, a linear classifier in the plane \mathcal{R}^2 , can be described by a half-plane. For this reason, we can discuss the family of linear classifiers as the set of all half-planes (and possibly also the plane itself and the empty set).

We say that a class (i.e. set) \mathcal{H} of classifiers *shatters* a set of points $X = \{x_1, x_2, \dots, x_n\}$ if we can classify the points in X in all possible ways. More precisely, for all 2^n possible labeling vectors $y_1, y_2, \dots, y_n \in \{-1, 1\}^n$, there exists a classifier $h \in \mathcal{H}$ such that $h(x_i) = y_i$ for all i . For any possible labelings of the points, there has to be a classifier in our set that reproduces those labels. Using the set notation for classifiers, this means that for any subset $X' \subseteq X$ (indicating the subset of points labeled $+1$), there exist a classifier $h \in \mathcal{H}$ such that $X \cap h = X'$ (the set of points for which h assigns label $+1$ includes X' but not the rest of X). It is important to understand that shattering is a property of a set of classifiers—not of a single classifier (a single classifier cannot even shatter a single point).

The *VC-dimension* of a set \mathcal{H} of classifiers is the size of the largest set X that can be shattered by \mathcal{H} .

We first analyze the VC-dimension of the class of axis-parallel rectangles in the plane. That is, a classifier is defined in terms of an axis-parallel rectangle, and classifies points inside the rectangle as positive, and points outside of it as negative. More formally (or perhaps just using more notation), a classifier belongs to this class if it is specified by the left, right, top and bottom coordinates and

$$h_{l,r,t,b}(x_1, x_2) = \begin{cases} +1 & \text{If } l \leq x_1 \leq r \text{ and } b \leq x_2 \leq t \\ -1 & \text{Otherwise} \end{cases}$$

Or written as a set:

$$h_{l,r,t,b} = \{(x_1, x_2) | l \leq x_1 \leq r \text{ and } b \leq x_2 \leq t\}$$

And the class of axis-parallel rectangles is

$$\mathcal{H}_{rect} = \{h_{l,r,t,b} | l < r \text{ and } b < t\},$$

which is just the set of all axis-parallel rectangles. (Your answers need not use such formal notation)

1. Find a set of four points that can be shattered by the class \mathcal{H}_{rect} of axis-parallel rectangles. That is, specify four points, and show how for all 16 possible subsets (each in turn specifying the points labeled +1), you can find a rectangle that captures exactly this subset.
2. Prove that *no* set of five points can be shattered by \mathcal{H}_{rect} . The claim here is that for *any* set of five points, there exists a labeling that cannot be obtained with an axis-parallel rectangle. Conclude that the VC-dimension of \mathcal{H}_{rect} is exactly four.

Now consider classification of points on the one-dimensional line \mathcal{R}^d using d -degree polynomials as discriminant functions. These are linear classifiers with basis functions x^k , $k = 1, \dots, d$. Any classifier in this class has the form $h_{\mathbf{w}}(x) = \text{sign}(w_0 + w_1x + w_2x^2 + \dots + w_dx^d)$. The set of all such classifiers is

$$\mathcal{H}_d = \{h_{\mathbf{w}}, \mathbf{w} \in \mathcal{R}^{d+1}\}$$

3. (*Optional*) Prove that the VC-dimension of \mathcal{H}_d is $d + 1$.

As mentioned in class, this results extends to more general types of basis functions. Consider a class \mathcal{H} of classifier that can all be written as $h_{\mathbf{w}}(\mathbf{x}) = \text{sign}(w_0 + w_1\phi_1(\mathbf{x}) + w_2\phi_2(\mathbf{x}) + \dots + w_d\phi_d(\mathbf{x}) + \phi_{\text{fixed}}(\mathbf{x}))$ for some fixed set of basis functions (by *fixed* basis functions we mean that the basis functions are the same for all classifiers, and the classifiers differ only in their weight). Note that we also allow a special term $\phi_{\text{fixed}}(\mathbf{x})$, for which we cannot vary the weight. Regardless of the basis functions, \mathcal{H} has a VC-dimension of at most $d + 1$. In general, if the basis functions are *independent* (i.e. no basis function is a function of other basis functions), and all weight vectors are allowed, the VC-dimension will be exactly $d + 1$.

4. Determine the VC-dimension of each of the following classes, by considering them as linear discriminants with specific basis functions. Briefly explain your reasoning in each case.
 - (a) Circles in the plane. (Hint: use the basis-function formulation only as an upper bound).

- (b) (*Optional*) Ellipsoids in the plane.
- (c) Classifiers representing decision regions of 2-component Gaussian mixture models in the plane. (Hint: rewrite the class posterior as $\frac{1}{1+e^{z(\mathbf{x})}}$ and use $z(\mathbf{x})$ as a discriminant function).
- (d) (*Optional*) Classifiers representing decision regions of 2-component Gaussian mixture models in R^d .

5. In this question we will explore the VC-dimension of decision stumps. Recall that a stump classifier $h_{i,a,b}$ in R^d is specified by an axis-parallel half-space: $h_{i,a,b} = \{(x_1, x_2, \dots, x_n) | ax_i \leq b\}$ (a can always be taken to be $+1$ or -1).

- (a) (*Optional*) Consider a $d = 2^\delta$ dimensional space. Suggest a set of δ point in $\mathcal{R}^d = \mathcal{R}^{(2^\delta)}$ that can be shattered by stumps, and show how it can be shattered. Conclude that the VC-dimension of stumps in R^d is at least $\lfloor \log_2 d \rfloor$ (i.e. $\log d$ rounded down to the nearest integer).
- (b) For any set of n points in \mathcal{R}^d , show that the stumps can only classify the n points in at most $2dn$ different ways. That is, there are at most $2dn$ different labeling on the points which are attainable using stump classifiers. (Hint: count the number of possible classifications using stumps on a specific axis)
- (c) Use the above bound to show that the VC-dimension of stumps is at most $2(\log_2 d + 1)$. (Hint: use the fact that $\log_2 n < n/2$ and so $n - \log_2 n > n/2$)

Combining the results, we see that the VC-dimension of stumps is roughly $\log d$ (more formally, it is $\Theta(\log d)$: for large enough d , it differs from $\log d$ only by a multiplicative constant).

Although the VC-dimension of stumps is fairly low, combining stumps, as we did in AdaBoost, is much more expressive. A linear combination of m stumps is a classifier given by $h(x) = \text{sign}(\alpha_1 h_1 + \alpha_2 h_2 + \dots + \alpha_m h_m)$, where h_j are decision stumps, and $\alpha_m \in \mathcal{R}$ are their corresponding weights. We will focus on linear combinations of stumps on the one-dimensional line \mathcal{R} .

Note that this is not a linear combination of fixed basis functions, since the stumps h_j vary from from classifier to classifier.

- (d) Show how $m/2$ points in \mathcal{R} can be shattered by convex combinations of m stumps. (Hint: for each point use a combination of two stumps that classifies the point correctly, and is ambivalent on all other points). State the resulting bound on the VC-dimension of convex combinations of m stumps in the one-dimensional line.
- (e) (*Optional*) For $m+2$ points in \mathcal{R} , show a specific labeling that cannot be attained by a convex combination of m stumps. State the resulting bound on the VC-dimension of convex combinations of m stumps in the one-dimensional line.

- (f) (*Optional*) Bound the VC-dimension of a convex combination of m stumps in higher dimensions.

Now consider sinusoidal classifiers over the real line \mathcal{R} . A sinusoidal classifier h_ω is specified by:

$$h_\omega(x) = \text{sign}(\sin(\omega x))$$

And the class of sinusoidal classifiers contains all such classifiers:

$$\mathcal{H}_{\text{sin}} = \{h_\omega | \omega > 0\}$$

6. (*Optional*) Consider a set of n points $x_i = 4^{-i}$ for $i = 1, \dots, n$. Show how this set of points can be shattered by \mathcal{H}_{sin} .
7. What is the VC-dimension of \mathcal{H}_{sin} ? (You may use the previous question, even if you did not answer it)