

CS 506, HW 4

Prof. Jared Saia, University of New Mexico

1. (Arora) Suppose we have two sets P and N of unit vectors in $\mathbb{R}^m$ with the guarantee that there exists a hyperplane $a \cdot x = 0$ such that every point in P is on one side and every point in N is on the other. Furthermore, the L_2 distance of each point in P and N to this hyperplane is at least ϵ .
Let $n = |P| + |N|$. Use the Johnson-Lindenstrauss lemma to show that for a JL linear mapping to $O((\log n)/\epsilon^2)$ dimensions, the points are still separable by a hyperplane with margin $\epsilon/2$.
2. (Arora) There is a grayscale photo at:
<https://www.cs.princeton.edu/courses/archive/fall15/cos521/image.jpg>.
Interpret it as an n by m matrix and run SVD on it. What is the value of k such that a rank k approximation gives a reasonable approximation (visually) to the image? What value of k gives an approximation that looks high quality to your eyes? Attach both pictures and your code. Hint: Matlab has code to perform the SVD. You will also find the `mat2gray` function helpful.)
Next, apply the JL projection to project the m columns to $\mathbb{R}^k$, where k is as chosen above, then project the columns back into $\mathbb{R}^n$ using the inverse projection (this operation will of course be lossy). Compare this JL compressed image with the image obtained via SVD compression. How do they compare?
3. This problem concerns the Winnow algorithm from lecture. Assume that you have the following data set, where the first two values are attributes, and the third value is the classification.

$-1/2$	$+1/3$	$+1$
$+3/4$	$-1/4$	$+1$
$+1/2$	$-1/3$	-1
$-3/4$	$+1/4$	-1

 - (a) You guess that there is a large margin solution to this problem, x^* , such that for all j , $a_j \cdot x^* \geq .05$. So you set $\epsilon = .05$, and for

simplicity, you also set $M = 1$. Give the values for N , T and η that you would use to run Winnow.

- (b) Now code up your Winnow algorithm in Matlab or your favorite programming language. What is the classifier your algorithm returns for the above data set?
 - (c) What if you want to modify Winnow so that it returns not just a classifier that classifies everything correctly, but a large-margin classifier with margin at least say $\epsilon/2$? Sketch how you would modify the algorithm to do this. Next, modify your program, and run it on the above data. Using binary search, what is the largest margin classifier that you can find for the above data?
4. Write a 3-5 page update on your progress on the class project. Also prepare slides for a 15 minute in-class presentation on your project progress.