

1 Getting Started

Read through this page carefully. You may typeset your homework in latex or submit neatly handwritten/scanned solutions. Please start each question on a new page. Deliverables:

1. Submit a PDF of your writeup, **with an appendix for your code**, to assignment on Gradescope, “HW14 Write-Up”. If there are graphs, include those graphs in the correct sections. Do not simply reference your appendix.
2. If there is code, submit all code needed to reproduce your results, “HW14 Code”.
3. If there is a test set, submit your test set evaluation results, “HW14 Test Set”.

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

- (a) Who else did you work with on this homework? 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. We just want to make it *extra* clear so that no one inadvertently cheats.

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.

This homework is due **Thursday, May 3rd at 10pm.**

2 K-SVD

As you have seen in earlier homework problems, sparse representations are powerful. When we know the right dictionary of features within which our input is likely to have a sparse representation (possibly with some additional small non-sparse noise), we can use techniques like the LASSO or OMP (or even just rounding/thresholding) to recover the coefficients. As we have seen, this has a fundamentally better bias/variance trade-off.

But what if we don't know the dictionary? How can the dictionary itself be found from the training data? This is known as *dictionary learning*. In this problem, we will introduce the K-SVD algorithm (Aharon, M., Elad, M. and Bruckstein, A., 2006. “*k*-SVD: An algorithm for designing overcomplete dictionaries for sparse representation.” IEEE Transactions on signal processing, 54(11), pp.4311-4322.) for unsupervised dictionary learning. (The naive perspective on directly supervised dictionary learning would reveal the dictionary to us, so that is not very interesting. However, it should be clear that the dictionary learning algorithm here can be adapted to deal with the case where we already know some dictionary elements to start with and want to learn more. It should also be clear that there are EM-based approaches to do dictionary learning in a more fully Bayesian perspective, as well as that various deep neural-network architectures are in effect trying to do dictionary learning at early layers to find the features that combine to best represent the signals being presented.) The K-SVD setting is where we would like to find a dictionary that enables good sparse fits to our data in the standard least-squares sense.

Mathematically, given an input matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$, where n is the number of data points and d is the dimension of each data point, we want to solve the following optimization problem:

$$\begin{aligned} & \underset{\mathbf{D} \in \mathbb{R}^{K \times d}, \mathbf{Z} \in \mathbb{R}^{n \times K}}{\text{minimize}} && \|\mathbf{X} - \mathbf{Z}\mathbf{D}\|_F^2 \\ & \text{subject to} && \|\mathbf{z}_i^\top\|_0 \leq s \text{ for all } i. \end{aligned} \tag{1}$$

We explain the notation a bit more here:

- The matrix $\mathbf{D} \in \mathbb{R}^{K \times d}$ is called a *dictionary* or a *codebook*. Each row \mathbf{D}_k represents a dictionary vector/element/codeword/atom.
- Each data sample \mathbf{x}_i is presumed to be a combination of the dictionary vectors and the coefficients are given by \mathbf{z}_i^\top , the i th row of \mathbf{Z} . See equation (2).
- We denote by $\|\mathbf{v}\|_0$ the “zeroth-norm” of vector \mathbf{v} , which is defined to be the number of nonzero entries in the vector \mathbf{v} .

To summarize, we have $\mathbf{X} \in \mathbb{R}^{n \times d}$, $\mathbf{Z} \in \mathbb{R}^{n \times K}$, and $\mathbf{D} \in \mathbb{R}^{K \times d}$ and that

$$\mathbf{X} = \begin{bmatrix} \mathbf{x}_1^\top \\ \mathbf{x}_2^\top \\ \vdots \\ \mathbf{x}_n^\top \end{bmatrix}, \quad \mathbf{D} = \begin{bmatrix} \mathbf{D}_1^\top \\ \mathbf{D}_2^\top \\ \vdots \\ \mathbf{D}_K^\top \end{bmatrix} \quad \text{and} \quad \mathbf{Z} = \begin{bmatrix} \mathbf{z}_1^\top \\ \mathbf{z}_2^\top \\ \vdots \\ \mathbf{z}_n^\top \end{bmatrix} = [\mathbf{Z}_1, \mathbf{Z}_2, \dots, \mathbf{Z}_K]$$

with $\mathbf{D}_k \in \mathbb{R}^d$, $\mathbf{z}_i \in \mathbb{R}^K$ and $\mathbf{Z}_k \in \mathbb{R}^n$.

And to summarize the goal: Our goal is to find a dictionary \mathbf{D} such that we can approximate the data samples \mathbf{x}_i by a sparse linear combination of atoms in \mathbf{D} :

$$\mathbf{x}_i^\top \approx \sum_{k=1}^K z_{ik} \mathbf{D}_k^\top, \quad (2)$$

where the coefficients of linear combination are given in \mathbf{Z} (which also needs to be determined) and should have at most s nonzero entries for each sample.

A greedy iterative algorithm called K-SVD is proposed for the above optimization problem. The algorithm optimizes the objective over \mathbf{D} and \mathbf{Z} in an alternating fashion. The pseudo code of the algorithm is given in Algorithm 1. It alternates between two stages: Sparse coding (Algorithm 2) and Update Codebook (Algorithm 4). The Sparse Coding procedure finds a sparse representation for each data sample with Algorithm 3 based on a given dictionary. The procedure for updating the dictionary updates each row of the dictionary with Algorithm 5 in a for loop, while modifying the effected representation coefficients correspondingly.

We have two goals in this problem: (1) to establish a relationship of K-SVD to the K-means algorithm you already know, and (2) to show that this K-SVD algorithm converges to a local minimum. (For this reason in practice, K-SVD is run with multiple random initial conditions to make sure that a good minimum is found. Just like we have seen for K-means in earlier homeworks.)

Algorithm 1: K-SVD

Input: Data matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$; Number of atoms K ; Sparsity constraint s

Output: A dictionary $\mathbf{D} \in \mathbb{R}^{K \times d}$, and the coefficient matrix $\mathbf{Z} \in \mathbb{R}^{n \times K}$

function K-SVD (\mathbf{X}, K, s):

 Initialize $\mathbf{D} \leftarrow$ randomly chosen K rows of the dataset \mathbf{X} (without replacement)

 Initialize $\mathbf{Z} \leftarrow$ Sparse-coding ($\mathbf{D}, \mathbf{0}, \mathbf{X}, s$)

while not converged **do**

$\mathbf{Z} \leftarrow$ Sparse-coding ($\mathbf{D}, \mathbf{Z}, \mathbf{X}, s$)

 Update-dictionary ($\mathbf{D}, \mathbf{Z}, \mathbf{X}$)

end

return \mathbf{D}, \mathbf{Z}

end function

- (a) (Relationship to K-means) Recall the set-up of K-means: Given a data matrix \mathbf{X} , K-means algorithm *partitions* the data into K disjoint groups $\pi = \{\pi_1, \pi_2, \dots, \pi_K\}$. Each sample \mathbf{x}_i is contained in *exactly one partition* π_j . Recall that the K-means algorithm tries to solve the following optimization problem:

$$\underset{\pi, \mu \in \mathbb{R}^{K \times d}}{\text{minimize}} \sum_{k=1}^K \sum_{i \in \pi_k} \|\mathbf{x}_i - \boldsymbol{\mu}_k\|_2^2. \quad (3)$$

Algorithm 2: Sparse-coding

Input: Dictionary \mathbf{D} ; Coefficient matrix $\mathbf{Z} \in \mathbb{R}^{n \times K}$; Data matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$; Sparsity constraint s

Output: The coefficient matrix $\mathbf{Z} \in \mathbb{R}^{n \times K}$

function Sparse-coding($\mathbf{D}, \mathbf{Z}, \mathbf{X}, s$):

```
  for  $i = 1, \dots, n$  do
     $\mathbf{z}'_i = \text{Sparse-coding-single}(\mathbf{D}, \mathbf{x}_i, s)$ 
    if  $\|\mathbf{x}_i^\top - \mathbf{z}_i^\top \mathbf{D}\|_2 > \|\mathbf{x}_i^\top - (\mathbf{z}'_i)^\top \mathbf{D}\|_2$  then
       $\mathbf{z}_i \leftarrow \mathbf{z}'_i$ 
    end
```

```
  end
```

```
  return  $\mathbf{Z} = \begin{bmatrix} \mathbf{z}_1^\top \\ \vdots \\ \mathbf{z}_n^\top \end{bmatrix}$ 
```

```
end function
```

Algorithm 3: Sparse-coding-single

Input: Dictionary \mathbf{D} ; Data sample $\mathbf{x} \in \mathbb{R}^d$; Sparsity constraint s

Output: A coefficient vector $\mathbf{z} \in \mathbb{R}^K$

function Sparse-coding-single($\mathbf{D}, \mathbf{x}, s$):

```
  Initialize  $\mathbf{z} \leftarrow \mathbf{0} \in \mathbb{R}^k$ 
```

```
  Initialize the basis  $\mathcal{B}_0 = \{\}$ 
```

```
  for  $j = 1, \dots, s$  do
```

```
    Find the index of the best dictionary vector  $\mathbf{D}_{k^{(j)}} \in \mathbb{R}^d$  by solving:
```

$$\hat{\beta}^{(j)}, k^{(j)} = \arg \min_{\beta \in \mathbb{R}^j} \min_{k \in [K]} \|\mathbf{x} - \sum_{\mathbf{D}_\ell \in \mathcal{B}_{j-1} \cup \{\mathbf{D}_k\}} \beta_\ell \mathbf{D}_\ell\|_2^2$$

```
    Update  $\mathcal{B}_j \leftarrow \mathcal{B}_{j-1} \cup \{\mathbf{D}_{k^{(j)}}\}$ 
```

```
  end
```

```
  for  $\mathbf{D}_\ell \in \mathcal{B}_s$  do
```

```
     $\mathbf{z}_\ell \leftarrow \hat{\beta}_\ell^{(s)}$ 
```

```
  end
```

```
  return  $\mathbf{z}$ 
```

```
end function
```

Algorithm 4: Update-dictionary

Input: Dictionary \mathbf{D} ; Coefficient matrix $\mathbf{Z} \in \mathbb{R}^{n \times K}$; Data matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$

Output: Dictionary $\mathbf{D} \in \mathbb{R}^{K \times d}$

function Update-dictionary ($\mathbf{D}, \mathbf{Z}, \mathbf{X}$):

for $k = 1, \dots, K$ **do**

 | Update k -th row $\mathbf{D}_k^\top \leftarrow$ Update-dictionary-single ($\mathbf{D}, \mathbf{Z}, \mathbf{X}, k$)

end

return $\mathbf{D} = \begin{bmatrix} \mathbf{D}_1^\top \\ \vdots \\ \mathbf{D}_K^\top \end{bmatrix}$.

end function

Let us connect the K-means optimization problem (3) with that given by equation (1). Let $\boldsymbol{\mu}$ denote a $K \times d$ matrix and $\tilde{\mathbf{Z}}$ denote an $n \times K$ matrix which are given by

$$\boldsymbol{\mu} = \begin{bmatrix} \boldsymbol{\mu}_1^\top \\ \vdots \\ \boldsymbol{\mu}_K^\top \end{bmatrix} \text{ and } \tilde{\mathbf{Z}} = \begin{bmatrix} \tilde{\mathbf{z}}_1^\top \\ \vdots \\ \tilde{\mathbf{z}}_n^\top \end{bmatrix}.$$

Show that the objective function given in the problem (3) is equivalent to

$$\begin{aligned} & \underset{\tilde{\mathbf{Z}} \in \mathbb{R}^{n \times K}, \boldsymbol{\mu} \in \mathbb{R}^{K \times d}}{\text{minimize}} && \|\mathbf{X} - \tilde{\mathbf{Z}}\boldsymbol{\mu}\|_F^2 \\ & \text{subject to} && \tilde{\mathbf{z}}_i \in \{0, 1\}^K \text{ and } \|\tilde{\mathbf{z}}_i\|_0 = 1 \text{ for all } i = 1, \dots, n. \end{aligned} \tag{4}$$

Conclude that K-means is equivalent to K-SVD with $s = 1$ with an additional constraint of forcing the coefficient entries in the matrix $\tilde{\mathbf{Z}}$ to take values in the set $\{0, 1\}$.

- (b) Show that at each step $j = 1, 2, \dots, s$ in Algorithm 3 for finding a sparse solution, the min value of the objective ($\min_{\beta \in \mathbb{R}^j} \min_{k \in [K]} \|\mathbf{x} - \sum_{\mathbf{D}_\ell \in \mathcal{B}_{j-1} \cup \{\mathbf{D}_k\}} \beta_\ell \mathbf{D}_\ell\|_2^2$) is non-increasing.

The following hint may or may not be useful for you but is worth thinking through. *Hint: Can you identify this algorithm as OMP?*

- (c) (Sparse coding decreases the objective) Conclude from the above part that Algorithm 2 cannot increase the value of the objective function (1). (In practice, people use straight matching pursuit — Algorithm 6 — to replace Algorithm 3, which is much more efficient, but gets comparable performance in most practical settings where there isn't a huge dynamic range of coefficients and the target sparsity s is low. But showing that matching pursuit usually works is omitted here.)
- (d) Show that the single-atom dictionary update given by Algorithm 5 does not increase the objective function while preserving the sparsity constraint. (You may or may not find the following hint useful.)

Hint: Recall the Eckart-Young theorem.

Algorithm 5: Update-dictionary-single

Input: Dictionary \mathbf{D} ; Coefficient matrix $\mathbf{Z} \in \mathbb{R}^{n \times K}$; Data matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$; Index k

Output: Return a row vector $\tilde{\mathbf{d}} \in \mathbb{R}^{1 \times d}$

function Update-dictionary-single ($\mathbf{D}, \mathbf{Z}, \mathbf{X}, k$):

for $k = 1, \dots, K$ **do**

 Compute the error matrix $\mathbf{E}_k \in \mathbb{R}^{n \times d}$ that represent how well things are represented without the dictionary vector \mathbf{D}_k^\top :

$$\mathbf{E}_k = \mathbf{X} - \sum_{j \neq k} \mathbf{Z}_j \mathbf{D}_j^\top$$

 Find the indices of data samples whose sparse representation uses the dictionary vector \mathbf{D}_k^\top . In other words, find the indices corresponding to non-zero entries in the coefficient vector \mathbf{z}_k

$$\omega_k = \{i | i \in [n], \mathbf{z}_{k,i} \neq 0\}.$$

 and let us denote $\omega_k = \{i_1, \dots, i_{|\omega_k|}\}$.

 Find the matrix $\mathbf{E}_k^{\omega_k} \in \mathbb{R}^{|\omega_k| \times d}$ by choosing those rows of \mathbf{E}_k whose index belongs to the set ω_k , i.e., matrix $\mathbf{E}_k^{\omega_k}$ is computed by picking the rows with indices $i_1, \dots, i_{|\omega_k|}$.

 Compute the SVD of $\mathbf{E}_k^{\omega_k}$:

$$\mathbf{E}_k^{\omega_k} = \mathbf{U} \mathbf{\Lambda} \mathbf{V}^\top$$

 where $\mathbf{U} \in \mathbb{R}^{|\omega_k| \times |\omega_k|}$, $\mathbf{\Lambda} \in \mathbb{R}^{|\omega_k| \times d}$ and $\mathbf{V} \in \mathbb{R}^{d \times d}$. We assume $\mathbf{\Lambda}$ is a diagonal matrix with non-increasing diagonal entries (Λ_{11} is its largest entry). Let $U_1 \in \mathbb{R}^{|\omega_k|}$ and $V_1 \in \mathbb{R}^d$ denote the first column of the matrices \mathbf{U} and \mathbf{V} respectively.

 Update $\mathbf{D}_k \leftarrow V_1$

 Update only the non-zero entries of the vector \mathbf{z}_k^\top :

for $j = 1, \dots, |\omega_k|$ **do**

 | $\mathbf{z}_{k,i_j} = U_{1,i_j} \Lambda_{11}$

end

end

return \mathbf{D}_k^\top .

end function

Algorithm 6: Matching-pursuit

Input: Dictionary \mathbf{D} ; Data sample $\mathbf{x} \in \mathbb{R}^d$; Sparsity constraint s

Output: A coefficient vector $\mathbf{z} \in \mathbb{R}^K$.

function Sparse-coding-single-in-practice ($\mathbf{D}, \mathbf{x}, s$):

 Initialize $\mathbf{z} \leftarrow \mathbf{0} \in \mathbb{R}^k$

 Initialize the basis $\mathcal{B}_0 = \{\}$

 Initialize the residue $\mathbf{r}_0 = \mathbf{x}$

for $j = 1, \dots, s$ **do**

 Find the *best* dictionary vector $\mathbf{D}_{k^{(j)}}^\top \in \mathbb{R}^d$ by solving:

$$k^{(j)} = \arg \max_{k \in [K]} \frac{|\mathbf{r}_{j-1}^\top \mathbf{D}_k|}{\|\mathbf{D}_k\|_2} = \arg \min_{k \in [K]} \min_{\beta_k \in \mathbb{R}} \|\mathbf{r}_{j-1} - \beta_k \mathbf{D}_k\|_2^2$$

 Update the coefficient $\mathbf{z}_{k^{(j)}} \leftarrow \frac{|\mathbf{r}_{j-1}^\top \mathbf{D}_k|}{\|\mathbf{D}_k\|_2^2}$

 Update the residue $\mathbf{r}_j \leftarrow \mathbf{r}_{j-1} - \mathbf{z}_{k^{(j)}} \mathbf{D}_{k^{(j)}}$

end

return \mathbf{z} .

end function

(e) (Updating dictionary decreases the objective) **Conclude that the iterations of the K-SVD algorithm put together cannot increase the objective function and hence the objective function must converge.**

(f) (Implementation, BONUS) Alas! Here we have a bad news. Due to time and resource constraints, we were unable to provide a good and debugged implementation part to this problem. You are free to wander around and play with the starter code and construct simulated datasets (or find real world data sets) where K-SVD recovers the underlying dictionary and/or provides a good reconstruction of the data using sparse linear combinations. No submission is required but you are encouraged to share your empirical findings on the piazza post associated with this problem. Feel free to share the datasets/results that you find. Extra credit will be given to those who help put together a nice exploration part of this problem. Given that the original K-SVD paper has over six thousand citations, we are confident that it should be possible to do this, and fun too.

HINT: For those who have taken EE16A in a semester where Massive-IoT-CDMA was used to motivate OMP, can you come up with an example where an eavesdropper listens in without knowledge of the wireless device codes but after lurking for a while and hearing small numbers of devices talk simultaneously can now figure out what all the individual devices' codes were? Alternatively, how about imaging lots of very sparse unknown images but not knowing what the imaging masks were. Can you learn the masks from lots of unknown sparse images being observed using the unknown random masks? Can you replicate the seminal work of Olshausen and Field from 1996 that showed that dictionary learning can help explain mammalian vision? See the Proceedings of the IEEE survey article on dictionary learning 10.1109/JPROC.2010.2040551 for more potential references.

3 Dropout in Neural Network and Linear Regression

Dropout is an algorithm to “prevent over-fitting” that is used while training lot of modern neural network implementations. This question will introduce how dropout works and why we can understand dropout as method to do regularization. Consider the simplest neural network – linear regression, which tries to find a vector \mathbf{w} to minimize

$$\ell(\mathbf{w}) = \|\mathbf{y} - \mathbf{X}\mathbf{w}\|_2^2 = \sum_{i=1}^n \left(y_i - \mathbf{x}_i^\top \mathbf{w} \right)^2. \quad (5)$$

Dropout is usually used with stochastic gradient descent. Instead of doing a stochastic gradient (with mini-batch-size 1) step using $\mathbf{w}_{t+1} = \mathbf{w}_t - \alpha \nabla_{\mathbf{w}} (y_i - \mathbf{x}_i^\top \mathbf{w})^2$, we instead apply the formula

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \alpha \nabla_{\mathbf{w}} \left[\left(y_i - \frac{1}{p} (\mathbf{x}_i \circ \mathbf{m})^\top \mathbf{w} \right)^2 \right]. \quad (6)$$

Here α is the learning rate, p is a user-specified dropout retention rate from the $(0, 1)$ interval, $i \sim U(1, n)$ is an index drawn from the uniform distribution, \circ is the element-wise multiplication operator, and \mathbf{m} is a random vector where $m_i \sim \text{Bernoulli}(p)$ iid. In other words, we will randomly mask each feature of a sample to be zero with a probability $1 - p$ during SGD steps.

While this problem only deals with dropout on linear regression, SGD steps with dropout seen here can easily generalize to any layers of neural networks that contain trainable parameters. We can randomly mask each element of the output vector from the previous layer to be zero with a probability $1 - p$ to regularize the training of the current layer. Dropout is commonly used with fully-connected layers and recurrent neural networks as they often contain a large number of parameters, making them prone to overfitting by nature.

(a) For simplicity, let us say that you are only solving a linear regression problem with only two features, i.e.,

$$\mathbf{X} = [\mathbf{x}^{(1)} \quad \mathbf{x}^{(2)}].$$

And for concreteness, suppose that we are using $p = \frac{2}{3}$ as the dropout retention rate.

Find a dataset \mathbf{X}' and \mathbf{y}' so that the normal SGD steps on the dataset \mathbf{X}' and \mathbf{y}' is effectively the same as the dropout SGD steps on the original \mathbf{X} and \mathbf{y} . Hint: You can augment the dataset and make it effectively be a linear regression problem by making multiple copies of data points. You can define

$$\mathbf{X}' = \begin{bmatrix} \alpha_1 \mathbf{x}^{(1)} & \alpha_1 \mathbf{x}^{(2)} \\ \alpha_2 \mathbf{x}^{(1)} & \mathbf{0} \\ \mathbf{0} & \alpha_3 \mathbf{x}^{(2)} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \quad (7)$$

and

$$\mathbf{y}' = \begin{bmatrix} \beta_1 \mathbf{y} \\ \beta_2 \mathbf{y} \\ \beta_3 \mathbf{y} \\ \beta_4 \mathbf{y} \end{bmatrix} \quad (8)$$

and find the expression for $\alpha_1, \alpha_2, \alpha_3, \beta_1, \beta_2, \beta_3$, and β_4 , along with how many copies you want to include of each “row” above in your augmented data set. Verify that the SGD steps are indeed the same by showing that the probability of making each kind of update to the weights is the same.

- (b) Leveraging what you have learned from the previous part, **explain why** the SGD steps with dropout for linear regression is the same as doing standard SGD without dropout for the following modified loss function

$$\ell'(\mathbf{w}) = \mathbf{E}_{\mathfrak{M}} \left[\left\| \mathbf{y} - \frac{1}{p} (\mathbf{X} \circ \mathfrak{M}) \mathbf{w} \right\|_2^2 \right], \quad (9)$$

where \circ is an element-wise matrix multiplication operator and random matrix $\mathfrak{M}_{ij} \sim \text{Bernoulli}(p)$ iid.

Hint: You can “unroll” the expectation here so that

$$\ell'(\mathbf{w}) = \frac{1}{n} \sum_{i=1}^n \sum_{\mathfrak{m}} f(\mathbf{x}_i \circ \mathfrak{m}, y_i; \mathbf{w}) \quad (10)$$

and then show that SGD step above is as same as taking a gradient step gradient $f(\cdot)$ with a random sample. Computing $\mathbf{x}_i \circ \mathfrak{m}_i$ can be viewed as augmenting the data like we did in the previous part.

- (c) To see the fact that dropout is a way of doing regularization, **prove that**

$$\arg \min_{\mathbf{w}} \ell'(\mathbf{w}) = \arg \min_{\mathbf{w}} \|\mathbf{y} - \mathbf{X}\mathbf{w}\|_2^2 + \frac{1-p}{p} \|\mathbf{\Gamma}\mathbf{w}\|_2^2, \quad (11)$$

where $\mathbf{\Gamma} = \text{diag}(\mathbf{X}^T \mathbf{X})^{1/2}$ — i.e. $\mathbf{\Gamma}$ is a diagonal matrix with just the Euclidean norms of the feature vectors along the diagonal. This is saying that applying dropout *on the linear regression problem* is equivalent to doing the ridge regression with a special diagonal Tikhonov regularization matrix. For the effect of dropout on more complex models such as logistic regression or convolution neural networks, it is hard to obtain a closed-form solution. But the idea dropout is equivalent to some kind of regularization still applies.

Hint: If random variables $X, Y \sim \text{Bernoulli}(p)$, we have $\mathbf{E}[X^2] = p$ and $\mathbf{E}[XY] = p^2$.

- (d) If you normalize the columns of \mathbf{X} (i.e. scale the features) so that $\|\mathbf{x}^{(i)}\|_2^2 = 1$ for each column $\mathbf{x}^{(i)}$, **what is the equivalent Tikhonov regularization matrix if you use dropout? Explain why (or in what context) dropout’s induced implicit Tikhonov regularization might be better than the identity matrix that vanilla ridge regression uses when we are working with unnormalized columns of \mathbf{X} .**
- (e) Run the code in Part A. **Explain what the code is doing and the meaning of the plot.**
- (f) The code in Part B compares dropout with bagged linear regression. Read the code and **explain the differences and similarities between bagged linear regression and linear regression with dropout (implemented using ridge regression in the code).**

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 the 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.