15780: GRADUATE AI (SPRING 2018)
Homework 3: Deep Learning and Probabilistic Modeling

Release: March 19, 2018,
Last Updated: March 30, 2018, 7:30 pm
Due: April 2, 2018, 11:59pm

1 Maximum Likelihood Estimation [20 points]

1.1 MLE for Poisson Distribution [5 points]

Given a collection of observed (independent) data points $X = \{x^{(1)}, ..., x^{(m)}\}$ from a Poisson distribution $Poisson(\lambda)$ (for $\lambda > 0$)[1] , where

$$p(x; \lambda) = \frac{\lambda^x e^{-\lambda}}{x!} \quad \text{for } x \in \mathbb{N},$$

derive the maximum likelihood estimator of λ, which maximizes the probability of observing X.

1.2 MAP for Poisson Distribution [5 points]

Given a collection of observed (independent) data points $X = \{x^{(1)}, ..., x^{(m)}\}$ from a Poisson distribution $Poisson(\lambda)$, where λ follows the prior distribution $Gamma(\alpha, \beta)$ (for $\alpha > 1, \beta > 0$)[2] , i.e.

$$p(\lambda; \alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} \lambda^{\alpha-1} e^{-\beta\lambda} \quad \text{for } \lambda > 0,$$

derive the estimator of λ that maximizes the posteriori probability $p(\lambda|X)$.
Hint: Use $p(\lambda|X) \propto p(X|\lambda)p(\lambda)$, and maximize $\log p(\lambda|X)$.

1.3 MLE for Multivariate Gaussian Distribution [10 points]

Given a collection of observed (independent) data points $X = \{\boldsymbol{x}^{(1)}, ..., \boldsymbol{x}^{(m)}\}$ from a multivariate Gaussian distribution $\mathcal{N}(\boldsymbol{x}; \boldsymbol{\mu}, \Sigma)$, where

$$p(\boldsymbol{x}; \boldsymbol{\mu}, \Sigma) = \frac{1}{|2\pi\Sigma|^{1/2}} e^{-\frac{1}{2}(\boldsymbol{x}-\boldsymbol{\mu})^T \Sigma^{-1}(\boldsymbol{x}-\boldsymbol{\mu})},$$

derive the maximum likelihood estimators of $\boldsymbol{\mu}$ and Σ, which maximize the probability of observing X.

---

[1]Added that $\lambda > 0$.
[2]Added that $\alpha > 1, \beta > 0$.

Hint: To find the MLE for Σ, take the derivative of the log likelihood with respect to Σ^{-1}. You can then use the following properties of the matrix trace $\text{tr}(\cdot)$ and matrix calculus equations without proof (but you must be clear about how you are using these):

1. The trace of a scalar is a scalar, i.e., $\text{tr}(c) = c$ for $c \in \mathbb{R}$.

2. Matrices can be cycled within a trace, i.e., $\text{tr}(ABC) = \text{tr}(CAB) = \text{tr}(BCA)$ for matrices A, B, C.

3. $\frac{\partial \ln |A|}{\partial A} = A^{-T}$ and $\frac{\partial tr(AB)}{\partial A} = B^T$ for matrices A, B.

4. The determinant of products of matrices is the product of the determinants of the matrices, i.e., $|AB| = |A| \cdot |B|$ for matrices A, B.

2 Playing Bayes-Ball with LDA [45 points]

In this section, we will understand conditional independencies between variables in a Bayesian network. Specifically, we will consider the Bayesian network corresponding to the Latent Dirichlet Allocation (LDA) generative model for documents.

The LDA Model. We want to model M documents where each document (indexed as $i = 1, \ldots M$) is represented as a sequence of N words $w_{i,j}$ for $j = 1, \ldots, N$. Assume that each word belongs to a dictionary of D words. Under the LDA model, we assume there also exists a set of K topics, where each topic defines a distribution over words. Furthermore, each document corresponds to a distribution over topics. We first 'generate' these topic distributions for each document and the word distributions for each topic, and then 'generate' the words in the document using these distributions. To generate these distributions, we fix some prior distributions $\mathcal{D}_{\text{topics}}$ and $\mathcal{D}_{\text{words}}$ over the spaces $\{(c_1, c_2, \ldots, c_K) | \sum_{k=1}^{K} c_k = 1, \; c_k \in [0, 1] \; \forall k\}$ and $\{(c_1, c_2, \ldots, c_D) | \sum_{d=1}^{D} c_d = 1, \; c_d \in [0, 1] \; \forall d\}$ respectively[3]. Then[4]:

- For each document $i = 1, \ldots M$, we sample $\theta_i \sim \mathcal{D}_{\text{topics}}$, that defines the parameters of the topic distribution for that document.

- For each topic $k = 1, \ldots K$, we sample $\phi_k \sim \mathcal{D}_{\text{words}}$, that defines the parameters of the word distribution for that topic.

Having generated the above two sets of parameters, we generate word $w_{i,j}$ as follows:

- Sample the topic of word $w_{i,j}$ (denoted by $z_{i,j}$) from the categorical distribution over topics $\{1, \ldots K\}$ (defined by the topic distribution of document i), i.e. Categorical(θ_i).

- Sample $w_{i,j}$ from the categorical distribution over the D words in the dictionary (defined by the word distribution of topic $z_{i,j}$), i.e. Categorical($\phi_{z_{i,j}}$).

Bayesian network for the LDA model. Observe that the joint probability distribution for the random variables in the LDA model can be written as:

---

[3]In the actual LDA model, $\mathcal{D}_{\text{topics}}$ and $\mathcal{D}_{\text{words}}$ are both Dirichlet distributions, thus giving the model its name. However, for the purposes of this question, we will keep it a bit more general.

[4]The notations $\mathcal{D}_{\text{words}}$ and $\mathcal{D}_{\text{topics}}$ were swapped below in the earlier model description.

$$p(\theta, z, w, \phi) = \left(\prod_k p(\phi_k)\right)\left(\prod_i p(\theta_i)\right)\left(\prod_{i,j} p(z_{i,j} \mid \theta_i)\right)\left(\prod_{i,j} p(w_{i,j} \mid z_{i,j}, \phi)\right), \qquad (1)$$

where each of the conditional probabilities has been defined by the generative model. Observe that this factorization corresponds to a Bayesian network, say G_{LDA}, with a node corresponding to each random variable in the model. In this problem, we will study the conditional independencies represented by this Bayesian network graph G_{LDA}.

Bayes-Ball Algorithm. The Bayes-Ball algorithm is a technique to infer conditional independencies between the variables in a Bayesian network. Consider a Bayesian network G over a set of variables $X_1, X_2, \ldots, X_T$. Assume we have observed a subset of variables, $\mathcal{K} = \{X_i : X_i \text{ is observed}\}$. Let $\mathcal{J}$ be a subset of the unobserved variables. The Bayes-Ball algorithm finds the set $\mathcal{I}$ of all variables that are conditionally independent of $\mathcal{J}$ according to the Bayesian network graph G[5], given that the variables in $\mathcal{K}$ are observed i.e., the set $\mathcal{I} = \{X_i : X_i \perp X_J | \mathcal{K}\}$. The algorithm places a ball at each of the nodes[6] in $\mathcal{J}$ and sends these around to other nodes in the graph according to some rules. The unobserved variables in the graph not in $\mathcal{J}$ that never received the ball during the Bayes-Ball game are added to $\mathcal{I}$.
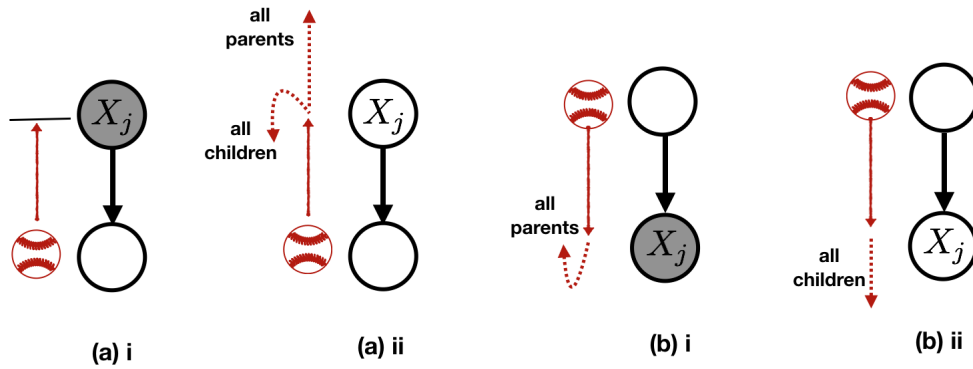

Figure 1: Illustration of the Bayes-Ball game: Observed nodes are shaded in grey.

The algorithm is as follows (illustrated in Figure 1):

1. Every node that belongs to $\mathcal{J}$ receives a ball (as if the ball was sent by one of its children[7]).

2. When a node X_j receives a ball:

   (a) If the ball was sent by a child of X_j:

      i. If $X_j \in \mathcal{K}$, do nothing (i.e., the ball never moves from there!).

---

[5]Note that a Bayesian network is essentially a representation for a family of joint distributions over the T variables, that can be factorized in a specific way. The structure of the network allows us to infer certain conditional independencies that every distribution in the family must satisfy. However, each distribution might have certain independencies specific to it that the Bayesian network may not represent. The Bayes-Ball algorithm only recovers the independencies encoded by the network.

[6]For simplicity, we will use the term 'node' to refer to the variable corresponding to the node and vice versa.

[7]This means you must apply step 2a at all these nodes.

ii. If $X_j \notin \mathcal{K}$, send a copy of the ball to all the parents and all the children of X_j (including the node that sent this ball!).

(b) If the ball was sent by a parent of X_j:

    i. If $X_j \in \mathcal{K}$, then send a copy of the ball to each of X_j's parents (including the parent who sent it!).

    ii. If $X_j \notin \mathcal{K}$, then send a copy of the ball to each of X_j's children.

3. Finally, $\mathcal{I}$ is the set of all unobserved variables in the graph that never received a ball. [8]

Conditional Independences in LDA. We will study some of the conditional independencies in the Bayesian network G_{LDA} of the LDA model by: (1) applying the Bayes-Ball algorithm to G_{LDA}, and (2) deriving them from first principles using the joint distribution [9] . For the rest of this homework, you can assume that we have just 2 documents, 2 words per document and 2 topics, i.e. $M = N = K = 2$ (although the results hold in general). To be able to run the Bayes-Ball algorithm easily, draw out the Bayesian network G_{LDA} for these parameters. For simplicity, you might also want to write out equation (1) for the special case of $M = N = K = 2$.

1. (5+8 points) Let i be a specific document and j be a specific word position. Suppose we observe $z_{i,j}$ (the topic of word $w_{i,j}$). Is θ_i conditionally independent of $w_{i,j}$ in G_{LDA}?

   (a) Run the Bayes-Ball algorithm to show that indeed $\theta_i \perp w_{i,j} | z_{i,j}$. You must place the ball in the node of either θ_i or $w_{i,j}$ and argue why it will never be received by the other node.

   (b) Next, prove this result starting from the factorized joint distribution given in equation 1 and using first principles from probability theory.

2. (5+8 points) Let i be a specific document, and j_1, j_2 be two distinct word positions. Is z_{i,j_1} independent of z_{i,j_2} in the Bayesian network when nothing is observed i.e., is $z_{i,j_1} \perp z_{i,j_2}$?

   (a) Run the Bayes-Ball algorithm to show that z_{i,j_1} and z_{i,j_2} are actually dependent in G_{LDA}. You must place the ball at one of these nodes and state the steps of the given algorithm by which the ball will be received by the other node.

   (b) Verify the result obtained above by constructing a probability distribution that can be represented by the LDA model and show that z_{i,j_1} and z_{i,j_2} [10] are not independent for this distribution.

   Hints for 2.(b):

       i. You can design a $\mathcal{D}_{\text{topics}}$ such that for some $k_1 \neq k_2$ and for some k_3, $p(z_{i,j_2} = k_3 | z_{i,j_1} = k_1) \neq p(z_{i,j_2} = k_3 | z_{i,j_1} = k_2)$ under the given generative process.

       ii. Consider $\mathcal{D}_{\text{topics}}$ to be equally distributed over exactly two points in its support: one point where the k_1th co-ordinate is 1 and all other co-ordinates are zero, and another point where the k_2th co-ordinate is 1 and all other co-ordinates are zero.

---

[8] It may not be clear when the given algorithm terminates. You may assume that when a node receives the ball from a specific direction for the second time, it does nothing.

[9] Specified the point allocation for the 13-points questions as $5 + 8$

[10] A typo has been corrected.

4

    iii. Intuitively, by observing $z_{i,j_1} = k_1$, what can you say about z_{i,j_2}? Now make this argument formal.

3. (5+8 points) Let i be a specific document, j be a specific word position, and k be a specific topic. Is $z_{i,j}$ independent of ϕ_k in G_{LDA} when $w_{i,j}$ is observed, i.e. is $z_{i,j} \perp \phi_k | w_{i,j}$?

    (a) Run the Bayes-Ball algorithm to show that $z_{i,j}$ and ϕ_k are actually dependent in G_{LDA} when $w_{i,j}$ is observed.

    (b) Verify the result obtained above by constructing a probability distribution that can be represented by the LDA model and show that $z_{i,j}$ and ϕ_k are not independent for this distribution, when $w_{i,j}$ is observed.

    Hints for 3b:

        i. Consider a distribution $\mathcal{D}_{\text{words}}$ such that every point in its support has exactly one of the D co-ordinates to be 1 and all other co-ordinates to be 0.

4. (6 points) Let i_1, i_2 be two distinct documents. Suppose all the words of document i_1 are observed, and only the first word of document i_2 is observed. In other words, the complete set of observed variables is $\{w_{i_1,1}, w_{i_1,2}, \ldots w_{i_1,N}, w_{i_2,1}\}$. Now, is the topic of the second word of document i_2 (i.e. $z_{i_2,2}$) conditionally independent of the topic distribution of document i_1 (i.e. θ_{i_1}) in G_{LDA}? Use the Bayes-Ball algorithm to answer this query. [You are not required to prove the result from first principles.]

Hint for Bayes-Ball: $X \perp Y|Z$ if and only if $Y \perp X|Z$. So, you can choose the easier of X or Y to release balls from to run the Bayes-Ball algorithm.

Hints for proving using first principles: Let $\mathbb{X}$ denote the set of all variables. And, suppose we want to prove (or disprove) that $X \perp Y|Z$.

1. For independence, it must be true that $p(X|Y, Z) = p(X|Z)$.

2. By definition of conditional probability $p(X|Y, Z) = \frac{p(X,Y,Z)}{p(Y,Z)}$.

3. To obtain expressions for both the numerator and denominator, we can marginalize the joint distribution. Mathematically, $p(X, Y, Z) = \sum_{\mathbb{X}\setminus\{X,Y,Z\}} p(\mathbb{X})$, and $p(Y, Z) = \sum_{\mathbb{X}\setminus\{Y,Z\}} p(\mathbb{X})$. [Recall that we have access to the expression for the joint distribution $p(\mathbb{X})$ from equation (1).]

4. Each summation can be substantially simplified by moving some of the summations inside the factored $p(\mathbb{X})$, and observing that $\sum_U p(U|V) = 1$ for any V, as long as U does not appear in any other parts of the whole expression.

5. Some terms can be taken completely out of the summation, and canceled if they are common to both the numerator and denominator.

6. This can help you completely eliminate Y from the expression for $p(X|Y, Z)$, to prove independence.

3 Metropolis-Hastings Implementation [10 points]

This portion is not graded on Autolab. Instead, we will run your code and check your outputs below.

In this section, we will implement the Metropolis-Hastings algorithms to generate samples from the unnormalized distribution

$$\tilde{p}(x) = e^{-x^2} + 1.3e^{-(x-2)^2}$$

with a Gaussian sampling function centered at the previous point with some standard deviation σ.

1. (5 points) Run your implementation for every combination of $x_0 \in \{-2, 2\}$ and $\sigma \in \{0.1, 1.0, 10\}$. Include plots (1. samples over time and 2. a histogram) for all of these in your writeup.

2. (3 points) How and why does σ impact the samples? For all three values, you should state what happens and why it happens.

3. (2 points) What is the impact of x_0 on the samples?

3.1 Getting started

Download the handout source code and edit the `mh.py` source file. This file contains a `main` method that will run your implementation and generate sampling plots and histograms. In this file, finish the Metropolis-Hastings implementation in the `mh` function from the description in Slides 17 (Probabilistic Modeling III: MCMC).

4 Programming a Deep Neural Network [25 points]

In Homework 2, you developed a linear classification model to classify MNIST digits. In this homework, you will implement a deep neural network for the same classification task. As a reminder, MNIST digit classification is a k-class classification task (where in this case $k = 10$). Our training set is of the form $(x^{(i)}, y^{(i)})$, $i = 1, \ldots, m$, with $y^{(i)} \in \{0, 1\}^k$, where $y_j^{(i)} = 1$ when j is the target class, and 0 otherwise (i.e., the output is one-hot encoded). Under the model, our hypothesis function $\hat{y} = h_\theta(x)$ outputs *vectors* in $\mathbb{R}^k$, where the relative size of the $\hat{y}_j$ corresponds roughly to how likely we believe that the output is really in class j. If we want the model to predict a single class label for the output, we simply predict class class j for which $\hat{y}_j$ takes on the largest value.

Our loss function $\ell : \mathbb{R}^k \times \{0, 1\}^k \to \mathbb{R}_+$ is the softmax loss, given by

$$\ell(\hat{y}, y) = \log \left(\sum_{j=1}^{k} e^{\hat{y}_j} \right) - \hat{y}^T y. \tag{2}$$

This loss function has the gradient

$$\nabla_{\hat{y}} \ell(\hat{y}, y) = \frac{e^{\hat{y}}}{\sum_{j=1}^{k} e^{\hat{y}_j}} - y \tag{3}$$

where the exponent $e^{\hat{y}}$ is taken elementwise.

To repeat a note from last time: In practice, you would probably want to implement regularized loss minimization, but for the sake of this problem set, we'll just consider minimizing loss without any regularization (at the expense of overfitting a little bit).

4.1 Implementation task [20 points]

You will implement a deep neural network to classify MNIST digits, trained by stochastic gradient descent. Specifically, you must implement the following three functions in `cls_nn.py`, with the inputs and outputs as described in the comments[11].

---

[11] An earlier version of the document incorrectly suggested that you can change the conventions for these functions as you like. However please make sure to stick to the specified conventions when you submit to autolab.

```python
def nn(x, W, b, f):
    """
    Compute output of a neural network.

    Input:
    x: numpy array of input
    W: list of numpy arrays for W parameters
    b: list of numpy arrays for b parameters
    f: list of activation functions for each layer

    Output:
    z: list of activationsn, where each element in the list is a tuple:
    (z_{i}, z'_{i}),
    where z_{i}=f_{i-1}(W_{i-1}z_{i-1}+b_{i-1}) (for i>=2),
    z'_{i}=f'_{i-1}(W_{i-1}z_{i-1}+b_{i-1}) (for i>=2),
    z_{1}=x and z'_{1}=None.

    """


def nn_loss(x, y, W, b, f):
    """
    Compute loss of a neural net prediction, plus gradients of parameters

    Input:
    x: numpy array of input
    y: numpy array of output
    W: list of numpy arrays for W parameters
    b: list of numpy arrays for b parameters
    f: list of activation functions for each layer

    Output tuple: (L, dW, db)
    L: softmax loss on this example
    dW: list of numpy arrays for gradients of W parameters
    db: list of numpy arrays for gradients of b parameters
    """


def nn_sgd(X, y, Xt, yt, W, b, f, epochs=10, alpha = 0.005):
    """
    Run stochastic gradient descent to solve linear softmax regression.

    Inputs:
    X: numpy array of training inputs
    y: numpy array of training outputs
    Xt: numpy array of testing inputs
    yt: numpy array of testing outputs
    W: list of W parameters (with initial values)
    b: list of b parameters (with initial values)
    f: list of activation functions
    epochs: number of passes to make over the whole training set
    alpha: step size

    Output: None (you can directly update the W and b inputs in place)
    """
```

The size of the inputs `W,b, f` determine the effective size of the neural network. We have included starter code that constructs the network in such a format, so you just need to implement the above function. In particular, a deep network for the MNIST problem, initialized with random weights, with rectified linear units in all layers except just a linear unit in the last layer, is constructed by the code:

```
np.random.seed(0)
layer_sizes = [784, 200, 100, 10]
W = [0.1*np.random.randn(n,m) for m,n in zip(layer_sizes[:-1], layer_sizes[1:])]
b = [0.1*np.random.randn(n) for n in layer_sizes[1:]]
f = [f_relu]*(len(layer_sizes)-2) + [f_lin]
```

In particular, this creates a deep network with 4 total layers (2 hidden layers), of sizes 784 (the size of the input), 200, 100, and 10 (the size of the output) respectively. This means, for instance, that there will be 3 W terms: $W_1 \in \mathbb{R}^{200 \times 784}$, $W_2 \in \mathbb{R}^{100 \times 200}$, and $W_3 \in \mathbb{R}^{10 \times 100}$.

You'll use virtually the same SGD procedure as in Homework 2, so the main challenge will be just to compute the network output and the gradients in the `nn` and `nn_loss` functions, using the backpropagation algorithm as specified in the class slides. We have included the various activation functions, which return both the non-linear function f and its elementsize subgradient f'.

Note that training a neural network with pure stochastic gradient descent (i.e., no minibatches) can be fairly slow, so we recommend you test your system on a small subset of (say) the first 1000 training and testing examples, and only run it on the final data set after you have debugged its performance on a smaller data set. As a reference point, our (quite unoptimized) implementation takes about a minute per epoch on the full training set, with the following performance over the first few epochs for the default $\alpha = 0.01$:[12]

```
Test Err |Train Err |Test Loss |Train Loss
   0.9033|    0.9044|    2.4827|    2.4800
   0.0514|    0.0430|    0.1578|    0.1351
   0.0349|    0.0265|    0.1218|    0.0829
   0.0291|    0.0173|    0.1024|    0.0527
   0.0339|    0.0180|    0.1187|    0.0550
```

4.2 Written portion [5 points]

In addition to the code that you'll submit (which will be evaluated on simpler toy examples to check for correctness), also include with your submission

1. A figure showing the training error and testing error versus epoch for the two linear softmax regression algorithms from Homework 2 (using either your own code or the solution code in `cls_nn.py`) plus the neural network, as measured on the full MNIST digit classification data set.

2. A figure showing the average training and testing loss versus epoch for all three algorithms, again as measured on the MNIST problem.

---

[12]This was added to provide more details.

For the above, you can use 10 epochs and the default learning rates given in the starter code (that is, $\alpha = 0.5$ for softmax GD, and $\alpha = 0.01$ for softmax SGD and the neural network).[13]

5   Submitting to Autolab

Create a tar file containing your writeup and the completed `mh.py` and `cls_nn.py` modules for the programming problems. Make sure that your tar has these files at the root and not in a subdirectory. Use the following commands from a directory with your files to create a `handin.tgz` file for submission.

```
$ ls
cls_nn.py  mh.py  writeup.pdf  [...]
$ tar cvzf handin.tgz writeup.pdf mh.py cls_nn.py
a writeup.pdf
a mh.py
a cls_nn.py
$ ls
handin.tgz  cls_nn.py  mh.py  writeup.pdf  [...]
```

---

[13]This sentence was added to provide more details.