Homework #0

CSE 546: Machine Learning
Prof. Kevin Jamieson
Due: 10/4/18 11:59 PM

1 Analysis

1. *[1 points]* A set $A \subseteq \mathbb{R}^n$ is *convex* if $\lambda x + (1-\lambda)y \in A$ for all $x, y \in A$ and $\lambda \in [0,1]$. A *norm* $\|\cdot\|$ over $\mathbb{R}^n$ is defined by the properties: i) non-negative: $\|x\| \geq 0$ for all $x \in \mathbb{R}^n$ with equality if and only if $x = 0$, ii) absolute scalability: $\|a\, x\| = |a|\, \|x\|$ for all $a \in \mathbb{R}$ and $x \in \mathbb{R}^n$, iii) triangle inequality: $\|x+y\| \leq \|x\| + \|y\|$ for all $x, y \in \mathbb{R}^n$.

   a. Using just the definitions above, show that the set $\{x \in \mathbb{R}^n : \|x\| \leq 1\}$ is convex for any norm $\|\cdot\|$.

   b. Show that $\left(\sum_{i=1}^n |x_i|^{1/2}\right)^2$ is or is not a norm.

2. *[1 points]* For any $x \in \mathbb{R}^n$, define the following norms: $\|x\|_1 = \sum_{i=1}^n |x_i|$, $\|x\|_2 = \sqrt{\sum_{i=1}^n |x_i|^2}$, $\|x\|_\infty = \max_{i=1,\dots,n} |x_i|$. Show that $\|x\|_\infty \leq \|x\|_2 \leq \|x\|_1$.

3. *[1 points]* For possibly non-symmetric $\boldsymbol{A}, \boldsymbol{B} \in \mathbb{R}^{n \times n}$ and $c \in \mathbb{R}$, let $f(x,y) = x^T \boldsymbol{A} x + y^T \boldsymbol{B} x + c$. Define $\nabla_z f(x,y) = \begin{bmatrix} \frac{\partial f(x,y)}{\partial z_1} & \frac{\partial f(x,y)}{\partial z_2} & \dots & \frac{\partial f(x,y)}{\partial z_n} \end{bmatrix}^T$. What is $\nabla_x f(x,y)$ and $\nabla_y f(x,y)$?

4. *[1 points]* Let $\boldsymbol{A}$ and $\boldsymbol{B}$ be two $\mathbb{R}^{n \times n}$ symmetric matrices. Suppose $\boldsymbol{A}$ and $\boldsymbol{B}$ have the exact same set of eigenvectors $\boldsymbol{u}_1, \boldsymbol{u}_2, \cdots, \boldsymbol{u}_n$ with the corresponding eigenvalues $\alpha_1, \alpha_2, \cdots, \alpha_n$ for $\boldsymbol{A}$, and $\beta_1, \beta_2, \cdots, \beta_n$ for $\boldsymbol{B}$. Please write down the eigenvectors and their corresponding eigenvalues for the following matrices:

   a. $\boldsymbol{C} = \boldsymbol{A} + \boldsymbol{B}$

   b. $\boldsymbol{D} = \boldsymbol{A} - \boldsymbol{B}$

   c. $\boldsymbol{E} = \boldsymbol{A}\boldsymbol{B}$

   d. $\boldsymbol{F} = \boldsymbol{A}^{-1}\boldsymbol{B}$ (assume $\boldsymbol{A}$ is invertible)

5. *[1 points]* A symmetric matrix $\boldsymbol{A} \in \mathbb{R}^{n \times n}$ is *positive-semidefinite (PSD)* if $x^T \boldsymbol{A} x \geq 0$ for all $x \in \mathbb{R}^n$.

   a. For any $y \in \mathbb{R}^n$, show that yy^T is PSD.

   b. Let X be a random vector in $\mathbb{R}^n$ with covariance matrix $\boldsymbol{\Sigma} = \mathbb{E}[(X - \mathbb{E}[X])(X - \mathbb{E}[X])^{\mathrm{T}}]$. Show that $\boldsymbol{\Sigma}$ is PSD.

   c. Assume $\boldsymbol{A}$ is a symmetric matrix so that $\boldsymbol{A} = \boldsymbol{U}\mathrm{diag}(\alpha)\boldsymbol{U}^T$ where $\mathrm{diag}(\alpha)$ is an all zeros matrix with the entries of α on the diagonal and $\boldsymbol{U}^T\boldsymbol{U} = I$. Show that $\boldsymbol{A}$ is PSD if and only if $\min_i \alpha_i \geq 0$. (Hint: compute $x^T \boldsymbol{A} x$ and consider values of x proportional to the columns of $\boldsymbol{U}$, i.e., the orthonormal eigenvectors).

6. *[1 points]* Let X and Y be real independent random variables with PDFs given by f and g, respectively. Let h be the PDF of the random variable $Z = X + Y$.

   a. Derive a general expression for h in terms of f and g

   b. If X and Y are both independent and uniformly distributed on $[0,1]$ (i.e. $f(x) = g(x) = 1$ for $x \in [0,1]$ and 0 otherwise) what is h, the PDF of $Z = X + Y$?

   c. For these given explicit distributions, what is $\mathbb{P}(X \leq 1/2 | X + Y \geq 5/4)$?

7. *[1 points]* A random variable $X \sim \mathcal{N}(\mu, \sigma^2)$ is Gaussian distributed with mean μ and variance σ^2. Given that for any $a, b \in \mathbb{R}$, we have that $Y = aX + b$ is also Gaussian, find a, b such that $Y \sim \mathcal{N}(0, 1)$.

8. *[1 points]* If $f(x)$ is a PDF, we define the cumulative distribution function (CDF) as $F(x) = \int_{-\infty}^{x} f(y)dy$. For any function $g : \mathbb{R} \to \mathbb{R}$ and random variable X with PDF $f(x)$, define the expected value of $g(X)$ as $\mathbb{E}[g(X)] = \int_{-\infty}^{\infty} g(y)f(y)dy$. For a boolean event A, define $\mathbf{1}\{A\}$ as 1 if A is true, and 0 otherwise. Thus, $\mathbf{1}\{x \leq a\}$ is 1 whenever $x \leq a$ and 0 whenever $x > a$. Note that $F(x) = \mathbb{E}[\mathbf{1}\{X \leq x\}]$. Let $X_1, \ldots, X_n$ be *independent and identically distributed* random variables with CDF $F(x)$. Define $\widehat{F}_n(x) = \frac{1}{n}\sum_{i=1}^{n} \mathbf{1}\{X_i \leq x\}$.

a. For any x, what is $\mathbb{E}[\widehat{F}_n(x)]$?

b. For any x, show that $\mathbb{E}[(\widehat{F}_n(x) - F(x))^2] = \frac{F(x)(1-F(x))}{n}$

c. Using part b., show that $\sup_{x \in \mathbb{R}} \mathbb{E}[(\widehat{F}_n(x) - F(x))^2] \leq \frac{1}{4n}$.

2 Programming

9. *[2 points]* Two random variables X and Y have equal distributions if their CDFs, F_X and F_Y, respectively, are equal: $\sup_x |F_X(x) - F_Y(x)| = 0$. The central limit theorem says that the sum of k independent, zero-mean, variance-$1/k$ random variables converges to a Gaussian distribution as k goes off to infinity. We will study this phenomenon empirically (you will use the Python packages Numpy and Matplotlib). Define $Y^{(k)} = \frac{1}{\sqrt{k}}\sum_{i=1}^{k} B_i$ where each B_i is equal to -1 and 1 with equal probability. It is easy to verify (you should) that $\frac{1}{\sqrt{k}} B_i$ is zero-mean and has variance $1/k$.

a. For $i = 1, \ldots, n$ let $Z_i \sim \mathcal{N}(0, 1)$. If $F(x)$ is the true CDF from which each Z_i is drawn (i.e., Gaussian) and $\widehat{F}_n(x) = \frac{1}{n}\sum_{i=1}^{n} \mathbf{1}\{Z_i \leq x)$, use the homework problem above to choose n large enough such that $\sup_x \sqrt{\mathbb{E}[(\widehat{F}_n(x) - F(x))^2]} \leq 0.0025$, and plot $\widehat{F}_n(x)$ from -3 to 3. (Hint: use `Z=numpy.random.randn(n)` to generate the random variables, and `import matplotlib.pyplot as plt`; `plt.step(sorted(Z), np.arange(1,n+1)/float(n))` to plot).

b. For each $k \in \{1, 8, 64, 512\}$ generate n independent copies $Y^{(k)}$ and plot their empirical CDF on the same plot as part a. (Hint: you can use `np.sum(np.sign(np.random.randn(n, k))*np.sqrt(1./k), axis=1)` to generate n of the $Y^{(k)}$ random variables.)

Be sure to always label your axes. Your plot should look something like the following (Tip: checkout `seaborn` for instantly better looking plots.)