COMPSCI 240: Reasoning Under Uncertainty

Andrew Lan and Nic Herndon

University of Massachusetts at Amherst

Spring 2019

Midterm II Solution

Problem 1

Problem 1 (5×4=20pts): Consider the pair of random variables X and Y that are uniformly distributed in the region $E = \{(x, y) : |x| + |y| \leq 1\}$, i.e.,

$$f_{X,Y}(X = x, Y = y) = \begin{cases} c & \text{if } x \in E, \\ 0 & \text{otherwise.} \end{cases}$$

The following figure helps you to visualize the region E.

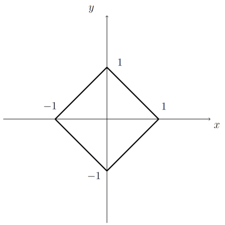

1. What's the value of the constant c?
2. What's the marginal PDF of X?
3. What's the conditional PDF of X given Y, for $0 \leq y \leq 1$?
4. Are X and Y independent? Justify your answer. You will not receive any points if you write only "yes" or "no".

Solution for Problem 1

1. $1 = c \cdot \text{Area of E} = c \cdot \sqrt{2}\sqrt{2} = 2c \Rightarrow c = \frac{1}{2}$
   Alternatively, using the normalization axiom, and the symmetry of PDF:

$$2\int_0^1 \int_{x-1}^{1-x} c dy dx = 2c \int_0^1 y\Big|_{x-1}^{1-x} dx = 2c \int_0^1 [(1-x)-(x-1)]\, dx$$

$$= 4c\left(x - \frac{x^2}{2}\right)\Big|_0^1 = 2c = 1 \Rightarrow c = \frac{1}{2}$$

2.

$$f_X(x) = \begin{cases} \int_{-x-1}^{x+1} \frac{1}{2} dy & , -1 \le x \le 0 \\ \int_{x-1}^{1-x} \frac{1}{2} dy & , 0 \le x \le 1 \\ 0 & , \text{otherwise} \end{cases} = \begin{cases} x+1 & , -1 \le x \le 0 \\ 1-x & , 0 \le x \le 1 \\ 0 & , \text{otherwise} \end{cases}$$

$$= \begin{cases} 1 - |x| & , x \in [-1, 1] \\ 0 & , x \notin [-1, 1] \end{cases}$$

Solution for Problem 1

3. Using a similar calculation we have

$$f_Y(y) = \begin{cases} 1 - |y| & , y \in [-1, 1] \\ 0 & , y \notin [-1, 1] \end{cases}$$

Thus,

$$f_{X|Y}(X|Y = y) = \begin{cases} \dfrac{f_{X,Y}(x, y)}{f_Y(y)} & , x \in [0, 1], y \in [0, 1),\ \text{and } y - 1 \leq x \leq 1 - y \\ 0 & , \text{otherwise} \end{cases}$$

$$= \begin{cases} \dfrac{1/2}{1 - y} & , x \in [0, 1], y \in [0, 1),\ \text{and } y - 1 \leq x \leq 1 - y \\ 0 & , \text{otherwise} \end{cases}$$

We can verify this is a valid PDF:

$$\int_0^1 \int_{y-1}^{1-y} \frac{1}{2(1-y)} dxdy = \int_0^1 \frac{x}{2(1-y)}\Bigg|_{y-1}^{1-y} dy = \int_0^1 \frac{(1-y) - (y-1)}{2(1-y)} dy$$

$$= \int_0^1 dy = y\Big|_0^1 = 1 - 0 = 1$$

Solution for Problem 1

4. X and Y are independent if $f_{X,Y}(x,y) = f_X(x) \cdot f_Y(y)$. However, when $x, y \in [0,1]$, $f_{X,Y}(x,y) \neq f_X(x) \cdot f_Y(y)$, i.e.:

$$\frac{1}{2} \neq (1-x)(1-y)$$

   Therefore, X and Y are not independent.

Problem 2

Problem 2 (10pts): Let X and Y be two random variables, with $var(X) = 4$ and $var(Y) = 9$. If we know that the two random variables $Z = 2X - Y$ and $W = X + Y$ are independent, find $\rho(X, Y)$, i.e., the correlation between X and Y.
Solution: Since Z and W are independent, we have $cov(Z, W) = 0$. Thus,

$$\begin{aligned}
cov(Z, W) &= cov(2X - Y, X + Y) \\
&= E[(2X - Y)(X + Y)] - E[2X - Y]E[X + Y] \\
&= E[2X^2 + XY - Y^2] - (2E[X] - E[Y])(E[X] + E[Y]) \\
&= 2E[X^2] + E[XY] - E[Y^2] - 2E[X]^2 - E[X]E[Y] + E[Y]^2 \\
&= 2(E[X^2] - E[X]^2) + (E[XY] - E[X]E[Y]) - (E[Y^2] - E[Y]^2) \\
&= 2var(X) + cov(X, Y) - var(Y) \\
&= 2 \cdot 4 + cov(X, Y) - 9 \\
&= cov(X, Y) - 1 \\
&= 0 \Rightarrow cov(X, Y) = 1
\end{aligned}$$

With this, we have all the components needed to calculate the correlation.

$$\rho(X, Y) = \frac{cov(X, Y)}{\sqrt{var(X) \cdot var(Y)}} = \frac{1}{\sqrt{4 \cdot 9}} = \frac{1}{6}$$

Problem 3

Problem 3 (5+5=10pts): A coin is biased so that its probability of landing on heads is 20%. Suppose you flip it 20 times.

1. Using Markov's bound, find a bound for the probability it lands on heads at least 16 times.
2. Since we know that the number of times the coin lands on its heads is a Binomial random variable, we can calculate the exact probability of the aforementioned event is 1.38×10^{-8}. Therefore, you can see that the bound we obtained is a loose one. Now, using Chebyshev's bound, find a tighter bound for this probability.

Solution for Problem 3

1. Let X denote the number of times the coin lands on its head. This is a binomial random variable with $n = 20$ and $p = 0.2$.
   Thus, $E[X] = np = 4$. Using this in the Markov inequality we get

$$p(X \geq 16) \leq \frac{E[X]}{16} = \frac{4}{16} = 0.25$$

2. $var(X) = np(1-p) = 20 \cdot 0.2 \cdot 0.8 = 3.2$
   Using this in the Chebyshev inequality we have

$$P(|X - E[X]| \geq b) \leq \frac{var(X)}{b^2}$$
$$P(|X - 4| \geq b) \leq \frac{3.2}{b^2}$$

$$\begin{aligned} P(|X-4| \geq b) &= P(X - 4 \geq b) + P(X - 4 \leq -b) \\ &= P(X \geq 4 + b) + P(X \leq 4 - b) \end{aligned}$$

Solution for Problem 3

Since we want to approximate $P(X \geq 16)$, i.e., $4 + b = 16$, let's set $b = 12$.

$$
\begin{aligned}
P(X \geq 4 + b) + P(X \leq 4 - b) &\leq \frac{3.2}{b^2} \\
P(X \geq 16) + P(X \leq -8) &\leq \frac{3.2}{12^2} \\
P(X \geq 16) + 0 &\leq \frac{3.2}{144} \\
P(X \geq 16) + &\leq 0.022
\end{aligned}
$$

This bound is much tighter than the Markov bound.

Problem 4

Problem 4 (10pts): Let today's high temperature be T. For this time of the year in Amherst, let's assume that T is a normal random variable with mean $\mu = 50$ and variance $\sigma^2 = 25$. Let's say Andrew feels comfortable if today's high temperature is between two integers A and B, i.e., $A \leq T \leq B$. He hasn't been here for long and is a little unsure about what to expect. So his lower temperature threshold A is a discrete random variable and takes two equally-likely values: 40 and 45. Similarly, his high temperature threshold B is also a discrete random variable and takes two equally-likely values: 55 and 60. Further assume that A and B are independent.

What is the probability that Andrew feels comfortable today?

Solution for Problem 4

Since A and B are independent,

$$P(A, B) = P(A) \cdot P(B) = \frac{1}{2} \cdot \frac{1}{2} = \frac{1}{4} \text{ for } (A, B) \in \{(40, 55), (40, 60), (45, 55), (45, 60)\}$$

Let Y denote the binary random variable that indicates whether Andrew feels comfortable.

$$\begin{aligned} P(Y = 1 | A = 40, B = 55) &= P(A \leq T \leq B) = P(40 \leq T \leq 55) \\ &= P\left(\frac{40 - 50}{\sqrt{25}} \leq T' \leq \frac{55 - 50}{\sqrt{25}}\right) \\ &= \Phi(1) - \Phi(-2) = \Phi(1) - [1 - \Phi(2)] \\ &= 0.84134 - (1 - 0.97725) \\ &= 0.81859 \end{aligned}$$

where T' is the standardized random variable.

Solution for Problem 4

Similarly, we can calculate

$$
\begin{aligned}
P(Y=1|A=40, B=60) &= \Phi(2) - \Phi(-2) = 0.95450 \\
P(Y=1|A=45, B=55) &= \Phi(1) - \Phi(-1) = 0.68268 \\
P(Y=1|A=45, B=60) &= \Phi(2) - \Phi(-1) = 0.81859
\end{aligned}
$$

Using the total probability theorem, we have

$$
\begin{aligned}
P(Y=1) = &\; P(A=40, B=55) \cdot P(Y=1|A=40, B=55) \\
&+ P(A=40, B=60) \cdot P(Y=1|A=40, B=60) \\
&+ P(A=45, B=55) \cdot P(Y=1|A=45, B=55) \\
&+ P(A=45, B=60) \cdot P(Y=1|A=45, B=60) \\
&= 3.27436/4 = 0.81859
\end{aligned}
$$

Alternatively, we could have used $\Phi(\cdot)$ in our calculations, to get to the same result:

$$
P(Y=1) = \Phi(1) + \Phi(2) - 1 = 0.81859
$$

Problem EC

Problem EC (10pts): Following the setup of Problem 4, let today's high temperature be T. Assume that T is a normal random variable with mean μ and variance $\sigma^2 = 1$. Let's say you don't know today's date, so your belief about μ follows a normal distribution with mean m and variance $\delta^2 = 1$, i.e., $P(\mu) = \mathcal{N}(m, 1)$. Now, at the end of the day, you observe that today's high temperature is actually t. Given this information, what is your new belief of μ, i.e., what is $P(\mu | T = t)$?

Hint: It has the same functional form as your initial belief, just different moments.

Solution for Problem EC

Quick reminder of PDF for normal random variables:

$$f_X(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

Using Bayes' rule, we know that $P(\mu | T = t) = \dfrac{P(\mu, T = t)}{P(T = t)} = \dfrac{P(T = t|\mu) \cdot P(\mu)}{P(T = t)}$.
The terms in this equation that we know are

$$P(\mu) = \frac{1}{\sqrt{2\pi}} e^{-\frac{(\mu-m)^2}{2}} \text{ and } P(T = t|\mu) = \frac{1}{\sqrt{2\pi}} e^{-\frac{(t-\mu)^2}{2}}$$

Solution for Problem EC

So, we need to calculate $P(T = t)$. We can calculate this by marginalizing over μ:

$$P(T = t) = \int_{-\infty}^{\infty} P(T = t, \mu) d\mu = \int_{-\infty}^{\infty} P(T = t|\mu) P(\mu) d\mu$$

$$= \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-\frac{(t-\mu)^2}{2}} \frac{1}{\sqrt{2\pi}} e^{-\frac{(\mu-m)^2}{2}} d\mu = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-\frac{(t-\mu)^2+(\mu-m)^2}{2}} d\mu$$

$$= \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-\frac{t^2-2\mu t+\mu^2+\mu^2-2\mu m+m^2}{2}} d\mu = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-\frac{2\mu^2-2\mu(t+m)+t^2+m^2}{2}} d\mu$$

$$= \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-\left[\mu^2-\mu(t+m)+\frac{t^2+m^2}{2}\right]} d\mu = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-\left[\mu^2-\mu(t+m)+(\frac{m+t}{2})^2-(\frac{m+t}{2})^2+\frac{t^2+m^2}{2}\right]} d\mu$$

$$= \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-\left(\mu-\frac{m+t}{2}\right)^2} \cdot e^{\left[-(\frac{m+t}{2})^2+\frac{t^2+m^2}{2}\right]} d\mu$$

$$= \frac{1}{2\pi} e^{\left[-(\frac{m+t}{2})^2+\frac{t^2+m^2}{2}\right]} \int_{-\infty}^{\infty} e^{-\left(\mu-\frac{m+t}{2}\right)^2} d\mu$$

$$= \frac{\sqrt{1/2}}{\sqrt{2\pi}} e^{\left[-(\frac{m+t}{2})^2+\frac{t^2+m^2}{2}\right]} \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi \cdot \frac{1}{2}}} e^{-\frac{\left(\mu-\frac{m+t}{2}\right)^2}{2\cdot\frac{1}{2}}} d\mu$$

$$= \frac{\sqrt{1/2}}{\sqrt{2\pi}} e^{\left[-(\frac{m+t}{2})^2+\frac{t^2+m^2}{2}\right]} \cdot 1$$

since the integral is over a normal PDF with mean $\frac{m+t}{2}$ and variance $\frac{1}{2}$.

Solution for Problem EC

Thus, we have

$$P(\mu|T=t) = \frac{\frac{1}{2\pi} e^{-\frac{(\mu-m)^2+(t-\mu)^2}{2}}}{\frac{\sqrt{1/2}}{\sqrt{2\pi}} e^{\left[-(\frac{m+t}{2})^2+\frac{t^2+m^2}{2}\right]}} = \frac{1}{\sqrt{2\pi \cdot \frac{1}{2}}} e^{-\frac{\left(\mu-\frac{m+t}{2}\right)^2}{2\cdot\frac{1}{2}}} \sim \mathcal{N}\left(\frac{m+t}{2}, \frac{1}{2}\right)$$

So, it's nothing more than a different normal distribution with $\mu = \frac{m+t}{2}$ and $\sigma^2 = \frac{1}{2}$.
This means that after a new observation, your new belief gets dragged in the direction of your observation, $T=t$, with increased confidence (variance decreases from 1 to 1/2).
The key here is that your belief maintains the same functional form; which allows you to update it again after another observation $T=t'$, and then another one, and so on. One can show that after n observations, the belief is still normal with

$$\mu_n = \frac{m+t_1+t_2+\cdots+t_n}{n+1} \text{ and } \sigma_n^2 = \frac{1}{n+1}$$

which means that i) your measurement is going to be super accurate ($\sigma_n^2 \to 0$ as $n \to \infty$), and ii) your point estimate is going to be the sample mean (plus counting the mean of your prior belief, m, as one additional sample).