

Section #4 Solutions

Problems by Chris

- 1. Conditional Flu** If a person has the flu, the distribution of their temperature is Gaussian with mean 101 and variance 1. If a person does not have the flu, the distribution of their temperature is 98 with variance 1. All you know about a person is that they have a temperature of 100. What is the probability they have the flu? Historically, 20% of people you analyze have had the flu.

We are going to solve this problem using joint random variables. We are going to define two random variables:

F is an indicator variable which is 1 if the person has the flu.

X is the distribution of the person's temperature.

The question asks, what is $P(F = 1|X = 100)$.

The problem tell us that $F \sim \text{Bern}(p = 0.2)$ and that

$X|F = 1 \sim N(\mu = 101, \sigma^2 = 1)$

$X|F = 0 \sim N(\mu = 98, \sigma^2 = 1)$

We can solve this using the inference version of Bayes, which allows for a mixture of discrete and continuous random variables.

$$P(F = 1|X = 100) = \frac{f(X = 100|F = 1)P(F = 1)}{f(X = 100|F = 1)P(F = 1) + f(X = 100|F = 0)P(F = 0)}$$

The next step is to substitute the PDF of the normal distribution

$$\begin{aligned} P(F = 1|X = 100) &= \frac{\frac{1}{\sqrt{2\pi}}e^{-\frac{1}{2}(100-101)^2}0.2}{\frac{1}{\sqrt{2\pi}}e^{-\frac{1}{2}(100-101)^2}0.2 + \frac{1}{\sqrt{2\pi}}e^{-\frac{1}{2}(100-98)^2}0.8} \\ &= \frac{e^{-\frac{1}{2}}0.2}{e^{-\frac{1}{2}}0.2 + e^{-2}0.8} \\ &\approx .528 \end{aligned}$$

- 2. Algorithmic Fairness** An artificial intelligence algorithm is being used to make a binary prediction (G for guess) for whether a person will repay a microloan. The question has come up: is the algorithm fair with respect to a binary demographic (D for demographic)? To answer this question we are going to analyze the historical predictions of the algorithm and compare the predictions to the true outcome (T for truth). Consider the following joint probability table from the history of the algorithms predictions:

	D = 0		D = 1	
	G = 0	G = 1	G = 0	G = 1
T = 0	0.21	0.32	0.01	0.01
T = 1	0.07	0.28	0.02	0.08

D : is the demographic of an individual (binary).

G : is the “repay” prediction made by the algorithm. 1 means predicted repay.

T : is the true “repay” result. 1 means did repay.

Recall that cell $(D = i, G = j, T = k)$ is the probability $P(D = i, G = j, T = k)$. For all questions, justify your answer. You may leave your answers with terms that could be input into a calculator.

- (a) (4 points) What is $P(D = 1)$?

$$P(D = 1) = 0.01 + 0.01 + 0.02 + 0.08 = 0.12$$

- (b) (4 points) What is $P(G = 1|D = 1)$?

$$P(G = 1|D = 1) = (0.01 + 0.08) / 0.12 = 0.75$$

- (c) (6 points) Fairness definition 1: Parity

An algorithm satisfies parity if the probability that the algorithm makes a positive prediction ($G = 1$) is the same regardless of the demographic variable. Does this algorithm satisfy parity?

We want to see if $P(G = 1|D = 1) = P(G = 1|D = 0)$.

$$P(G = 1|D = 0) = (0.32 + 0.28) / (0.21 + 0.07 + 0.32 + 0.28) = 0.6/0.88 = 0.68.$$

Thus, we see that $P(G = 1|D = 1) > P(G = 1|D = 0)$ and the algorithm does not satisfy parity.

- (6 points) Fairness definition 2: Calibration

An algorithm satisfies calibration if the probability that the algorithm is correct ($G = T$) is the same regardless of demographics. Does this algorithm satisfy calibration?

We essentially want to see if $P(G = 0, T = 0|D = 0) = P(G = 0, T = 0|D = 1)$ and if $P(G = 1, T = 1|D = 0) = P(G = 1, T = 1|D = 1)$.

First we check if $P(G = 0, T = 0|D = 0) = P(G = 0, T = 0|D = 1)$.

$$P(G = 0, T = 0|D = 0) = 0.21 / (0.21 + 0.07 + 0.32 + 0.28) = 0.239$$

$$P(G = 0, T = 0|D = 1) = 0.01 / (0.01 + 0.02 + 0.01 + 0.08) = 0.083$$

So we can see that the algorithm does not satisfy calibration.

(d) (6 points) Fairness definition 3: Equality of odds

An algorithm satisfies equality of odds if the probability that the algorithm predicts a positive outcome given that the true outcome is positive ($G = 1|T = 1$) is the same regardless of demographics. Does this algorithm satisfy equality of odds?

$$P(G = 1|T = 1, D = 0) = 0.28 / (0.28 + 0.07) = 0.8$$

$$P(G = 1|T = 1, D = 1) = 0.08 / (0.08 + 0.02) = 0.8$$

We see that $P(G = 1|T = 1, D = 0) = P(G = 1|T = 1, D = 1)$ and thus, the algorithm does satisfy equality of the odds!

3. **Approximating Normal:** (10 points) Your website has 100 users and each day each user independently has a 20% chance of logging into your website. Use a normal approximation to estimate the probability that more than 21 users log in.

The number of users that log in B is binomial: $B \sim \text{Bin}(n = 100, p = 0.2)$. It can be approximated with a normal that matches the mean and variance. Let C be the normal that approximates B . We have $E[B] = np = 20$ and $\text{Var}(B) = np(1 - p) = 16$, so $C \sim N(\mu = 20, \sigma^2 = 16)$. Note that because we are approximating a discrete value with a continuous random variable, we need to use the continuity correction:

$$\begin{aligned} P(B > 21) &\approx P(C > 21.5) \\ &= P\left(\frac{C - 20}{\sqrt{16}} > \frac{21.5 - 20}{\sqrt{16}}\right) \\ &= P(Z > 0.375) \\ &= 1 - P(Z < 0.375) \\ &= 1 - \phi(0.375) = 1 - 0.6462 = 0.3538 \end{aligned}$$

4. **Daycare.ai** Providing affordable (or better, free) daycare would have a tremendously positive effect on society. California mandates that the ratio of babies to staff must be ≤ 4 . We have a challenge: just because a baby is **enrolled**, doesn't mean they will **show up**. At a particular location, 6 babies are enrolled. We estimate that the probability an enrolled child actually shows up on a given day is $\frac{5}{6}$. Assume that babies show up independent of one another.

(a) (4 points) What is the probability that either 5 or 6 babies show up?

Let X be the number of babies that show up. $X \sim \text{Bin}(n = 6, p = \frac{5}{6})$.

$$\begin{aligned}
 P(X = 5 \text{ or } X = 6) &= P(X = 5) + P(X = 6) \\
 &= \binom{n}{5} p^5 (1-p)^{n-5} + \binom{n}{6} p^6 (1-p)^{n-6} \\
 &= \binom{6}{5} \frac{5^5}{6} \left(1 - \frac{5}{6}\right) + \frac{5^6}{6}
 \end{aligned}$$

- (b) (4 points) If we charge \$50 per baby that shows up, what is our expected revenue?

Let $R = 50X$ be the revenue.

$$E[R] = E[50X] = 50 \cdot E[X] = 50 \cdot 6 \cdot \frac{5}{6}$$

- (c) (6 points) If 0 to 4 babies show up our costs are \$200. If 5 or 6 babies show up our costs are \$500. What are our expected costs? You may express your answer in terms of a , the answer to part (a).

Let C be our cost and X be the number of babies which show up.

$$\begin{aligned}
 E[C] &= P(0 \leq X \leq 4)200 + P(5 \leq X \leq 6)500 \\
 &= (1 - P(5 \leq X \leq 6))200 + P(5 \leq X \leq 6)500
 \end{aligned}$$

We know $P(5 \leq X \leq 6)$ from part (a), so we can just plug in!

- (d) (8 points) What is the lowest value $\$k$ that we can charge per child in order to have an expected profit of \$0? Recall that Profit = Revenue - Cost. You may express your answer in terms of a , b or c , the answers to part (a), (b) and (c) respectively.

Let $P = R - C$ be our revenue.

$$E[P] = E[R - C] = E[R] - E[C]$$

We know $E[R]$ and $E[C]$ from parts (b) and (c) and so we can just plug in!

- (e) (8 points) Each family is unique. With our advanced analytics we were able to estimate a show-up probability for each of the six enrolled babies: p_1, p_2, \dots, p_6 where p_i is the probability that baby i shows up. Write a new expression for the probability that 5 or 6 babies show up. You may still assume that babies show up independent of one another.

Let X be the number of babies who show up and B_i be a Bernoulli variable representing if baby i shows up.

$$P(X = 5 \text{ or } X = 6) = P(X = 5) + P(X = 6)$$

$$= \sum_{n=0}^i (P(B_i = 0) \left(\prod_{j=0, i \neq j}^6 B_j = 1 \right)) + \prod_{i=0}^6 (P(B_i = 1))$$

The term $\prod_{i=0}^6 (P(B_i = 1))$ is for all babies showing up, and the term $\sum_{n=0}^i (P(B_i = 0) \left(\prod_{j=0, i \neq j}^6 B_j = 1 \right))$ is for 5 of the 6 showing up.

5. Midterm Prep Guiding Questions The midterm exam is coming up. Below are a few broad, guiding questions you might use to help solidify your thinking, prepare a study guide, etc.

- 1. Counting** What are event and sample spaces? What's the significance of equally likely events in probability problem-solving? How do we reason differently about distinct vs. indistinct events? What's the difference between combinations and permutations? What are the sum rule, product rule, inclusion-exclusion, and when do we use them?
- 2. Probability Rules** When do we use the definition of conditional probability, the chain rule, the law of total probability, Bayes' theorem, the Complement Rule, DeMorgan's law etc.? What are independence and conditional independence?
- 3. Random Variables** What does randomness mean? What are expectation and variance, generally? What's the difference between continuous and discrete random variables? We've seen lots of random variables - in which situations would each of them be appropriate? Which ones can be used to approximate others, and in which cases? What's the difference between PMF, PDF, and CDF? What are multivariate distributions and conditional distributions, and how do we reason about them in both discrete and continuous situations?