

Section #4

Problems by Chris

- 1. 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 algorithm’s 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 $P(D = 1)$?
- (b) (4 points) What is $P(G = 1|D = 1)$?
- (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?
- (d) (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?
- (e) (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 given that the true outcome is positive ($G = 1|T = 1$) is the same regardless of demographics. Does this algorithm satisfy equality of odds?

2. **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 Gaussian with mean 98 and 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.
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.
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?
 - (b) (4 points) If we charge \$50 per baby that shows up, what is our expected revenue?
 - (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).
 - (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.
 - (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.
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 outcomes? 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 mutual exclusion?
 3. **Random Variables** What is the difference between a random variable and a standard variable? 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?

- 4. Inference** You want to mix Bayes' theorem and random variables to answer a question of the form: What is the probability that $X = 4$ given that $Y = 2$. How could you solve this problem? What would have to happen if Y or X were continuous?