

Homework 0: Probability Review and Survey

CS 585, UMass Amherst, Fall 2017

Note

Wikipedia is a useful resource for basic probability.

Make a PDF file of your answers, and upload it to Gradescope by the end of Friday. We will only accept PDF format. Make sure to clearly mark the question number of your answers, and to select the locations of your answers with Gradescope's interface.

1 Domain of a joint distribution

1.1

A and B are discrete random variables. A could take on one of 4 possible values. B could take on one of 3 possible values. (In other words, the size of $\text{domain}(A)$ is 4, and the size of $\text{domain}(B)$ is 3.) How many possible outcomes does the joint distribution $P(A, B)$ define probabilities for?

1.2

Say we have a sequence of n binary random variables A_1, A_2, \dots, A_n . How many possible outcomes does the joint distribution $P(A_1, A_2, \dots, A_n)$ define probabilities for?

2 Independence versus Basic Definitions

Say we have three random variables A and B and C . Note that we're using standard probability theory notation where $P(A, B) = P(B, A)$, which simply means the joint probability of both A and B occurring.

2.1

Which of the following statements is always true?

1. $P(A|B) = P(B|A)$
2. $P(A, B) = P(A|B)P(B)$
3. $P(A, B) = P(A)P(B)$
4. $P(A|B) = P(A)$
5. $P(A, B, C) = P(A)P(C)$

6. $P(A, B, C) = P(A)P(B)P(C)$
7. $P(A, B, C) = P(A)P(B|A)P(C|A, B)$
8. $P(A) = \sum_{b \in \text{domain}(B)} P(A, B = b)$
9. $P(A) = \sum_{b \in \text{domain}(B)} P(A|B = b)P(B = b)$

2.2

Now assume that A , B , and C are all independent of each other. Which of these statements is true?

3 Logarithms

3.1 Log-probs

Let p be a probability, so it is bounded to $[0, 1]$ (between 0 and 1, inclusive). What is the range of possible values for $\log(p)$? Please be specific about open versus closed intervals.

3.2 Prob ratios

Let p and q both be probabilities. What is the range of possible values for p/q ?

3.3 Log prob ratios

What is the range of possible values for $\log(p/q)$?

4 Deriving Bayes Rule

The definition of conditional probability can be written as $P(A, B) = P(A|B)P(B)$ or alternatively as $P(A|B) = P(A, B)/P(B)$. Starting from this, derive Bayes Rule, in this form:

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}$$

It should take only a few lines to prove/derive this. Note that you can apply the definition of conditional probability not just for $P(A, B)$ but also for $P(B, A)$. Here we are using standard notation where comma indicates conjunction/intersection so order of variables doesn't matter when defining a joint event.