28 2026-04-15 | Week 13 | Lecture 28

28.1 The Coupon Collector Problem

Example 128. Joe is collecting coupons, numbered $1, \ldots, k$. There are n different coupons, and when he gets all n of them, he gets a prize. Each time he gets a coupon, it is chosen uniformly at random (i.e., all coupons are equally likely).

Find a value m such that if Joe collects at least m coupons, then his probability of winning the prize is at least 99%.

Solution: Let m be any positive integer greater than n. Let

$$V = \text{number of coupons Joe collects until he gets the prize}$$

and let

$$A_k = \text{coupon } k \text{ is not collected in the first } m \text{ coupons.}$$

To solve the problem, want to find a value of m sufficiently large that

$$\mathbb{P}[V \leq m] \geq .99.$$

Equivalently, we want

$$\mathbb{P}[V > m] \leq 0.01.$$

$$\begin{aligned}
\mathbb{P}[V > m] &= \mathbb{P}[A_1 \cup A_2 \cup \ldots \cup A_n] \\
&\leq \mathbb{P}[A_1] + \mathbb{P}[A_2] + \ldots + \mathbb{P}[A_n] && \text{(the union bound)} \\
&\leq n\left(1 - \frac{1}{n}\right)^m \\
&= n\left[\left(1 - \frac{1}{n}\right)^n\right]^{\frac{m}{n}} \\
&\approx n\left[e^{-1}\right]^{\frac{m}{n}} \\
&= ne^{-m/n}
\end{aligned}$$

Setting

$$ne^{-m/n} < 0.01$$

we can solve for m:

$$e^{-m/n} < \frac{0.01}{n}$$

$$-\frac{m}{n} < \log\left(\frac{0.01}{n}\right)$$

or

$$m > -n\log\left(\frac{0.01}{n}\right)$$

or equivalently,

$$m > n\log(100n)$$

In other words, if Joe collects at least $n\log(100n)$ coupons, then with probability at least 99%, he will get the prize.

For example, if $n = 15$, then $15\log(100 \cdot 15) = 109.7$ so $m = 110$ coupons is sufficient to ensure that Joe has 99% probability of getting the prize.

End of Example 128. □

28.2 Hypothesis Testing

Textbook section 8.1

The hypothesis testing framework: suppose you observe some random quantity, which is assumed to depend on some unknown numerical parameter θ. You make a *hypothesis*, which can be any claim about the value of some parameter.

We are interested in testing the following two hypotheses:

- The ***null hypothesis:*** whatever effect you observed was due to random chance. This is denoted H_0.
- The ***alternative hypothesis:*** the claim you made about the parameter. This is usually denoted H_a

Based on our observation(s) we will make a decision to either reject H_0, or fail to reject H_0.

Suppose the observed data is X. The ***p-value*** of X is defined as

$$\text{p-value} := \mathbb{P}\left[\text{observing a result at least as extreme as } X \text{ when } H_0 \text{ is true}\right]$$

When the p-value is small, we **reject** H_0.

In particular, we set a ***significance level***, denoted α, as the cutoff such that we reject H_0 is $p < \alpha$ and we fail to reject H_0 if $p \geq \alpha$.

A typical value is $\alpha = .05$ or $\alpha = .01$.

This is the basic story. Let's do an example:

Example 129. Suppose you have a coin but you don't know if it is a fair coin or not. You think maybe the coin is biased in favor of flipping heads. As an experiment, you flip the coin 10 times and you get 8 heads. Is this convincing evidence that the coin is biased?

In this case, the data is an IID sample

$$X_1, \ldots, X_{10}$$

where

$$X_i = \begin{cases} 1 & : \text{ with probability } \theta \\ 0 & : \text{ with probability } 1-\theta \end{cases}$$

The null hypothesis is

$$H_0 : \theta = \frac{1}{2}$$

The alternative hypothesis is that heads is more likely than tails:

$$H_1 : \theta > \frac{1}{2}$$

The p-value is

$$\begin{aligned} p &= \mathbb{P}\left[\text{at least 8 heads when } \theta = \frac{1}{2}\right] \\ &= \binom{10}{8}\theta^8(1-\theta)^2 + \binom{10}{9}\theta^9(1-\theta)^1 + \binom{10}{10}\theta^{10}(1-\theta)^0 \text{ with } \theta = \frac{1}{2} \\ &= \binom{10}{8}\left(\frac{1}{2}\right)^{10} + \binom{10}{9}\left(\frac{1}{2}\right)^{10} + \binom{10}{10}\left(\frac{1}{2}\right)^{10} \\ &= \frac{7}{128} \\ &\approx 0.0546. \end{aligned}$$

In words, if the coin were fair, we would expect it to flip 8 or more heads (out of ten) about 5% of the time. That's not actually super rare, so while this is consistent with the coin being biased, it's also consistent with the null hypothesis. Hence, I wouldn't find this to be convincing evidence that the coin isn't fair, so I wouldn't reject the null hypothesis.

Maybe you would. The cutoff for what is "convincing" is necessarily subjective.

End of Example 129. □