STATS 200: Solutions to Homework 2*

1 Monte Carlo integration

(a) The expected value of $\frac{f(X)}{g(X)}$ when $X \sim g$ is

$$\mathbb{E}\left[\frac{f(X)}{g(X)}\right] = \int_a^b \frac{f(x)}{g(x)} g(x)\mathrm{d}x = \int_a^b f(x)\mathrm{d}x = I(f).$$

Then

$$\mathbb{E}\left[\hat{I}_n(f)\right] = \mathbb{E}\left[\frac{1}{n}\sum_{i=1}^n \frac{f(X_i)}{g(X_i)}\right] = \frac{1}{n}\sum_{i=1}^n \mathbb{E}\left[\frac{f(X_i)}{g(X_i)}\right] = I(f).$$

As $f(X_i)/g(X_i)$ are IID, by the Law of Large Numbers,

$$\hat{I}_n(f) = \frac{1}{n}\sum_{i=1}^n \frac{f(X_i)}{g(X_i)} \to \mathbb{E}\left[\frac{f(X)}{g(X)}\right] = I(f)$$

in probability as $n \to \infty$.

(b) For $X \sim g$, set

$$\begin{aligned}\sigma^2 := \mathrm{Var}\left[\frac{f(X)}{g(X)}\right] &= \mathbb{E}\left[\left(\frac{f(X)}{g(X)}\right)^2\right] - \left(\mathbb{E}\left[\frac{f(X)}{g(X)}\right]\right)^2 \\ &= \int_a^b \frac{f(x)^2}{g(x)}\mathrm{d}x - (I(f))^2.\end{aligned}$$

*Edited from the solutions by Zhenpeng Zhou; thanks to Zhenpeng for sharing

Then

$$
\begin{aligned}
\text{Var}\left[\hat{I}_n(f)\right] &= \text{Var}\left[\frac{1}{n}\sum_{i=1}^{n}\frac{f(X_i)}{g(X_i)}\right] \\
&= \frac{1}{n^2}\sum_{i=1}^{n}\text{Var}\left[\frac{f(X_i)}{g(X_i)}\right] \\
&= \frac{1}{n}\text{Var}\left[\frac{f(X)}{g(X)}\right] \\
&= \frac{\sigma^2}{n}
\end{aligned}
$$

Let $Y_i = \frac{f(X_i)}{g(X_i)}$ and $\bar{Y} = \frac{1}{n}\sum_{i=1}^{n} Y_i$. Then by the Central Limit Theorem,

$$
\sqrt{n}\left(\frac{\bar{Y} - \mathbb{E}[Y_i]}{\sigma}\right) \to \mathcal{N}(0,1)
$$

in distribution as $n \to \infty$. Note $\bar{Y} = \hat{I}_n(f)$ and $\mathbb{E}[Y] = I(f)$, so

$$
\frac{\sqrt{n}}{\sigma}\left(\hat{I}_n(f) - I(f)\right) \to \mathcal{N}(0,1).
$$

Therefore $c_n = \frac{\sqrt{n}}{\sigma}$.

(c) For $f(x) = \cos(2\pi x)$ and $g(x) = 1$, we have

$$
\mathbb{E}\left[\frac{f(X)}{g(X)}\right] = \int_0^1 f(x)\mathrm{d}x = 0
$$

and

$$
\begin{aligned}
\sigma^2 = \text{Var}\left[\frac{f(X)}{g(X)}\right] &= \int_0^1 \cos^2(2\pi x)\mathrm{d}x - \left(\int_0^1 \cos(2\pi x)\mathrm{d}x\right)^2 \\
&= \frac{1}{2} - 0 = \frac{1}{2}
\end{aligned}
$$

By part (b), $\hat{I}_n(f) - I(f)$ is approximately distributed as $\mathcal{N}(0, \frac{\sigma^2}{n}) = \mathcal{N}(0, \frac{1}{2000})$ for large n, so

$$
\mathbb{P}[|\hat{I}_n(f) - I(f)| > 0.05] \approx 2 \times \Phi\left(-0.05/\frac{1}{\sqrt{2000}}\right) \approx 0.02535
$$

where $\Phi(x)$ is the standard normal CDF.

(d) There are many possible answers; the intuition is that to reduce $\mathrm{Var}[f(X)/g(X)]$, we would like $g(x)$ to be larger when $f(x)^2$ is larger. One possibility is

$$g(x) = \frac{\pi}{2}|\cos(2\pi x)|.$$

One may verify that

$$\int_0^1 |\cos(2\pi x)|\, dx = \frac{2}{\pi},$$

so $\int_0^1 g(x)dx = 1$ and

$$\mathrm{Var}[f(X)/g(X)] = \int_0^1 \frac{\cos^2(2\pi x)}{\frac{\pi}{2}|\cos(2\pi x)|}dx = \frac{2}{\pi} \times \frac{2}{\pi} = \frac{4}{\pi^2} \approx 0.41.$$

Then $\hat{I}_n(f) - I(f)$ is approximately distributed as $\mathcal{N}(0, \frac{\sigma^2}{n}) = \mathcal{N}(0, 0.00041)$, and

$$\mathbb{P}[|\hat{I}_n(f) - I(f)| > 0.05] \approx 2 \times \Phi\left(-0.05/\sqrt{0.00041}\right) \approx 0.013.$$

2 Continuous mapping

We must show, for any $\varepsilon > 0$,

$$\mathbb{P}[|g(X_n) - g(c)| > \varepsilon] \to 0$$

as $n \to \infty$.

For any $\varepsilon > 0$, as $g : \mathbb{R} \to \mathbb{R}$ is continuous at c, there exists $\delta > 0$ such that if $|x - c| \leq \delta$ then $|g(x) - g(c)| \leq \varepsilon$. Hence if $|g(x) - g(c)| > \varepsilon$, then $|x - c| > \delta$. This implies

$$\mathbb{P}[g(X_n) - g(c)| > \varepsilon] \leq \mathbb{P}[|X_n - c| > \delta].$$

But since $X_n \to c$ in probability,

$$\mathbb{P}[|X_n - c| > \delta] \to 0$$

as $n \to \infty$. So $\mathbb{P}[g(X_n) - g(c)| > \varepsilon] \to 0$ also as $n \to \infty$, as desired.

3 Testing gender ratios

(a) There are many possible answers. We may take T_1 to be the average number of male children per family,

$$T_1 = \bar{X},$$

and perform a two-sided test based on T_1 to check whether roughly half of the children are male. We may take T_2 to be Pearson's chi-squared statistic

$$T_2 = \sum_{k=0}^{12} (O_k - E_k)^2 / E_k$$

where O_k is the number of families with k male children and E_k is the expected number under the hypothesized binomial distribution, i.e. $E_k = 6115 \times \binom{12}{k}(0.5)^{12}$, and perform a one-sided test that rejects for large T_2 to check whether the shape of the observed distribution of $X_1, \ldots, X_{6115}$ matches the shape of the binomial PDF.

(b) R code corresponding to the above T_1 and T_2 is as follows:

```
ks = seq(0,12)
counts = c(7,45,181,478,829,1112,1343,1033,670,286,104,24,3)
expected = 6115*choose(12,ks)*(0.5^12)

T1_obs = sum(ks*counts)/6115
T2_obs = sum((counts-expected)^2/expected)

T1s = numeric(1000)
T2s = numeric(1000)
for (i in 1:1000) {
  X = rbinom(6115, 12, 0.5)
  T1s[i] = mean(X)
  counts = numeric(13)
  for (k in 0:12) {
    counts[k+1] = length(which(X==k))
  }
  T2s[i] = sum((counts-expected)^2/expected)
}
hist(T1s)
hist(T2s)
T1_pvalue = length(which(T1s<T1_obs))/1000 * 2
T2_pvalue = length(which(T2s>T2_obs))/1000
```

Histograms of the null distributions of T_1 and T_2 are below:

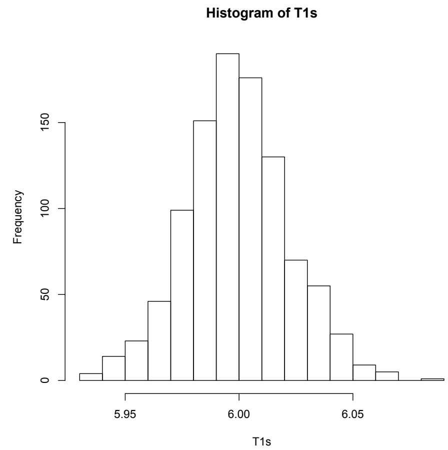

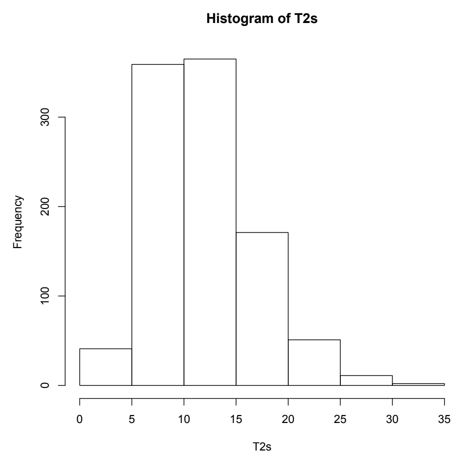

The values of the test statistics for the observed data are $T_1 = 5.77$ and $T_2 = 249$, which are both far outside the range of the simulated null distributions above. The simulated p-values for the two tests are both < 0.001, and there is strong evidence that H_0 is not correct.

(c) There may be both biological and sociological reasons why H_0 is false. Biologically, the human male-to-female sex ratio at birth is not exactly 1:1. The probability p that a child is male might also vary from family to family. The sexes of children within a family might be dependent; in particular, one source of dependence is the presence of identical twins.

Sociologically, there may be a relationship between family size and the sex ratio of children in the family, because the current sex ratio influences parents' decision of whether to have another child. Note that the given data is only for families with 12 children, which is quite large even for that time. There is a noticeable bias towards families with more girls than boys, which may be explained if parents tended to continue having children when their current children were predominantly female.

4 Most-powerful test for the normal variance

(a) The joint PDF under H_0 is

$$f_0(x_1, ..., x_n) = \left(\frac{1}{\sqrt{2\pi\sigma_0^2}}\right)^n \exp\left(-\frac{\sum_{i=1}^n x_i^2}{2\sigma_0^2}\right)$$

The joint PDF under H_1 is

$$f_1(x_1, ..., x_n) = \left(\frac{1}{\sqrt{2\pi\sigma_1^2}}\right)^n \exp\left(-\frac{\sum_{i=1}^n x_i^2}{2\sigma_1^2}\right)$$

So the likelihood ratio statistic is

$$L(X_1, ..., X_n) = \frac{f_0(X_1, ..., X_n)}{f_1(X_1, ..., X_n)} = \left(\frac{\sigma_1}{\sigma_0}\right)^n \exp\left(\frac{\sigma_0^2 - \sigma_1^2}{2\sigma_0^2\sigma_1^2} \sum_{i=1}^n x_i^2\right)$$

Since $\sigma_0^2 < \sigma_1^2$, L is a decreasing function of $T := \sum_{i=1}^n X_i^2$. Then rejecting for small values of L is the same as rejecting for large values of T.

Since under H_0, $\sum_{i=1}^n \left(\frac{X_i}{\sigma_0}\right)^2 \sim \chi_n^2$, we have $\frac{1}{\sigma_0^2}T \sim \chi_n^2$, so $T \sim \sigma_0^2\chi_n^2$. Then the rejection threshold should be $c = \sigma_0^2\chi_n^2(\alpha)$, and the most powerful test rejects H_0 when $T > c$.

(b) Under H_1, $\sum_{i=1}^n \left(\frac{X_i}{\sigma_1}\right)^2 \sim \chi_n^2$, so $T \sim \sigma_1^2\chi_n^2$. Then the probability of type II error is

$$\begin{aligned}\beta = \mathbb{P}_{H_1}[\text{accept } H_0] &= \mathbb{P}_{H_1}[T \le \sigma_0^2\chi_n^2(\alpha)] \\ &= \mathbb{P}_{H_1}\left[\frac{T}{\sigma_1^2} \le \frac{\sigma_0^2}{\sigma_1^2}\chi_n^2(\alpha)\right] = F\left(\frac{\sigma_0^2}{\sigma_1^2}\chi_n^2(\alpha)\right)\end{aligned}$$

where F is the χ_n^2 CDF. The power of the test is then

$$\text{Power} = 1 - \beta = 1 - \text{F}\left(\frac{\sigma_0^2}{\sigma_1^2}\chi_\text{n}^2(\alpha)\right)$$

As $\sigma_1^2 \to \infty$, $\beta \to F(0) = 0$ and the power of the test $\to 1$.

5 Testing a uniform null

The likelihood ratio statistic is

$$L(X) = \frac{f_0(X)}{f_1(X)} = \frac{1}{2X}$$

The condition $L(X) < c$ is then equivalent to $X > \tilde{c}$, where $\tilde{c} = \frac{1}{2c}$.

Under the hypothesis H_0, $X \sim \text{Uniform}(0, 1)$, so the rejection threshold $\tilde{c}$ should be $1 - 0.1 = 0.9$, i.e. the most powerful tests rejects H_0 when $X > 0.9$.

Under the hypothesis H_1, $X \sim f_1(x) = 2x$. Then the type II error probability is

$$\beta = \mathbb{P}_{H_1}[\text{accept } H_0] = \mathbb{P}_{H_1}[X \leq 0.9] = \int_0^{0.9} 2x\, dx = 0.81.$$

Thus the power of the test is

$$\text{Power} = 1 - \beta = 0.19$$

This is the maximum power that can be achieved: According to the Neyman-Pearson lemma, for any other test of H_0 with significance level at most 0.1, its power against H_1 is at most 0.19.