Sufficiency

Outline

1) Sufficiency
2) Factorization Theorem
3) Examples
4) Minimal sufficiency

Three models for coin flipping

Model 3 $\quad X_{i,j} \overset{ind}{\sim} \text{Bernoulli}(\theta_{i,t}) \quad \begin{array}{l} i=1,\ldots,48 \\ j=1,\ldots,n_i \end{array} \quad \theta_{i,j} \searrow \text{ in } j$

Model 2 $\quad X_{i+} \overset{ind.}{\sim} \text{Binom}(n_i, \theta_i) \qquad X_{i+} = \sum_{j=1}^{n_i} X_{i,j}$

Model 1 $\quad X_{++} \overset{ind.}{\sim} \text{Binom}(n, \theta) \qquad X_{++} = \sum_{i=1}^{48} \sum_{j=1}^{n_i} X_{i,j}$

most assumptions $\qquad$ fewest assumptions

These models are nested: $\mathcal{P}_1 \subseteq \mathcal{P}_2 \subseteq \mathcal{P}_3$

Data keeps getting compressed too...
are we losing anything by doing this?

Answer No. X_{++} is a sufficient statistic for $\mathcal{P}_1$, and $(X_{1+}, \ldots, X_{48+})$ is also sufficient for $\mathcal{P}_2$

Def A statistic $T(X)$ is any function of data X

Sufficiency

Def A statistic $T(x)$ is *sufficient* for model $\mathcal{P}$ if the conditional distribution of $X \mid T(x)$ is the same for all $P \in \mathcal{P}$

Check definition for $T(x) = X_{++}$ in $\mathcal{P}_1$:

$$p_\theta(x) = \prod_{i=1}^{48} \prod_{j=1}^{n_i} \theta^{x_{ij}} (1-\theta)^{1-x_{ij}}$$

$$= \theta^{x_{++}} (1-\theta)^{n - x_{++}} \qquad \text{(why no } \tbinom{n}{x_{++}}\text{?)}$$

$$\mathbb{P}_\theta(X = x \mid X_{++} = t) = \frac{\mathbb{P}_\theta(X = x, X_{++} = t)}{\mathbb{P}_\theta(X_{++} = t)}$$

$$= \frac{1\{x_{++} = t\} \cdot \theta^t (1-\theta)^{n-t}}{\binom{n}{t} \theta^t (1-\theta)^{n-t}}$$

$$= 1\{x_{++} = t\} / \tbinom{n}{t}$$

Intuition Suppose we believe Model 1.
Big/small X_{++} more likely with big/small θ
But once we know $X_{++} = 178{,}079$,
all data sets X with that many same-side flips are *equally likely*, *regardless of θ*

Not true in Models 2 & 3 $\Rightarrow$ X_{++} no longer sufficient

Factorization Theorem

Usually, we can recognize sufficient stats by inspecting the density

Theorem (Factorization Theorem)

Let $\mathcal{P} = \{P_\theta : \theta \in \Theta\}$ be a model with densities $p_\theta(x)$ wrt common measure μ.

$T(X)$ is sufficient iff there exist $g_\theta(t)$, $h(x) \geq 0$ with

$$p_\theta(x) = g_\theta(T(x))\, h(x) \qquad (\text{for } \mu\text{-a.e. } x)$$

Note we could absorb h into μ as density

(define new base measure ν, $\nu(A) = \int_A h(x)\, d\mu(x)$)

$\Rightarrow \mathcal{P}$ has densities $p_\theta(x) = g_\theta(T(x))$ wrt ν

Interp: after changing base measure, density depends on x only through $T(x)$

(Can't absorb $g_\theta(T(x))$ into μ: depends on θ)

Proof (discrete $\mathcal{X}$): Assume wlog $\mu = \#$ on $\mathcal{X}$

$$(\Leftarrow) \quad \mathbb{P}_\theta(X = x \mid T = t) = \frac{\mathbb{P}_\theta(X = x, T(x) = t)}{\mathbb{P}_\theta(T(x) = t)}$$

$$= \frac{\cancel{g_\theta(t)}\, h(x)\, 1\{T(x) = t\}}{\sum_{T(z) = t} \cancel{g_\theta(t)}\, h(z)}$$

$(\Rightarrow)$ Assume $T(x)$ sufficient, let

$$g_\theta(t) = \mathbb{P}_\theta(T(X) = t)$$

$$h(x) = \mathbb{P}(X = x \mid T(X) = T(x))$$

no dep. on θ

$$\Rightarrow \quad g_\theta(T(x))\, h(x) = \mathbb{P}_\theta(T(X) = T(x) \text{ and } X = x)$$

$$= \mathbb{P}_\theta(X = x) = p_\theta(x) \qquad \boxtimes$$

Proof similar for general densities
- careful about conditioning in cts spaces

Examples

Ex. Normal location family

$$X_1, \dots, X_n \overset{iid}{\sim} N(\theta, 1) = \frac{1}{\sqrt{2\pi}} e^{-(x-\theta)^2/2}$$

$$p_\theta(x) = (2\pi)^{-n/2} \prod_{i=1}^{n} e^{-(x_i-\theta)^2/2}$$

$$= e^{\theta \sum_i x_i - n\theta^2/2} \cdot \frac{e^{-\sum x_i^2/2}}{(2\pi)^{n/2}}$$

(collect factors with no dep. on θ)

$\Longrightarrow$ $\sum X_i$ is sufficient

Ex. Poisson family

$$X_1, \dots, X_n \overset{iid}{\sim} \mathrm{Pois}(\theta) = \frac{\theta^x e^{-\theta}}{x!} \quad \text{for } x = 0, 1, 2, \dots$$

$$p_\theta(x) = \prod_{i=1}^{n} \frac{\theta^{x_i} e^{-\theta}}{x_i!} = \theta^{\sum x_i} e^{-n\theta} \cdot \frac{1}{\prod x_i!}$$

$\Longrightarrow$ $\sum X_i$ is sufficient

There two examples have something important in common! (next lecture)

Ex. Uniform location family

$$X_1, \dots, X_n \overset{iid}{\sim} U[\theta, \theta+1] = 1\{\theta \le x \le \theta + 1\}$$

$$p_\theta(x) = \prod_{i=1}^{n} 1\{\theta \le x_i \le \theta + 1\}$$

$$= 1\{\theta \le X_{(1)}\}\, 1\{X_{(n)} \le \theta + 1\}$$

$\Longrightarrow$ $(X_{(1)}, X_{(n)})$ is sufficient.

Interpretations of Sufficiency

X is informative about θ only because its distribution depends on θ.

We can think of the data as being generated in two stages:

1) Generate T : distribution dep. on θ
2) Generate $X \mid T$: does <u>not</u> dep on θ

Sufficiency Principle

If $T(X)$ is sufficient for $\mathcal{P}$ then any statistical procedure should depend on X only through $T(X)$

In fact, we could throw away X and generate a new $\tilde{X} \sim P(X \mid T)$ and it would be just as good as X since $\tilde{X} \sim P_\theta$

no θ

In graphical model form:

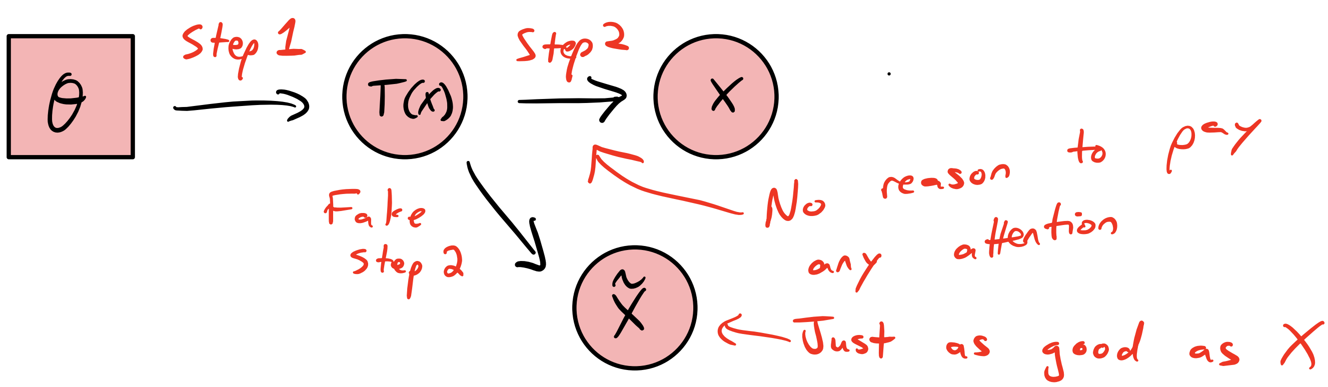

Order Statistics

For $x_1, \dots, x_n \in \mathbb{R}$, define order statistics

$$\min_i x_i = x_{(1)} \le x_{(2)} \le \cdots \le x_{(n)} = \max_i x_i$$

Ex (iid sampling on $\mathbb{R}$) $X_1, \dots, X_n \overset{iid}{\sim} P_\theta$,

any model $\mathcal{P} = \{ P_\theta^n : \theta \in \Theta \}$ on $\mathcal{X} \subseteq \mathbb{R}$

P_θ^n invariant to perm.s of $X = (X_1, \dots, X_n)$

$\Rightarrow$ All permutations of x are equally likely

$\Rightarrow$ Order statistics $S(X) = (X_{(i)})_{i=1}^n$ sufficient

$X \rightsquigarrow S(X)$ forgets orig. ordering of observations

Empirical Distribution

Order statistics depend on total ordering of $\mathcal{X}$

What about more general sample space?

Define Dirac measure $\delta_x(A) = 1\{x \in A\}$

Empirical distribution $\hat{P}_n(\cdot) = \frac{1}{n} \sum_{i=1}^{n} \delta_{X_i}(\cdot)$

random measure on $\mathcal{X}$, determined by sample

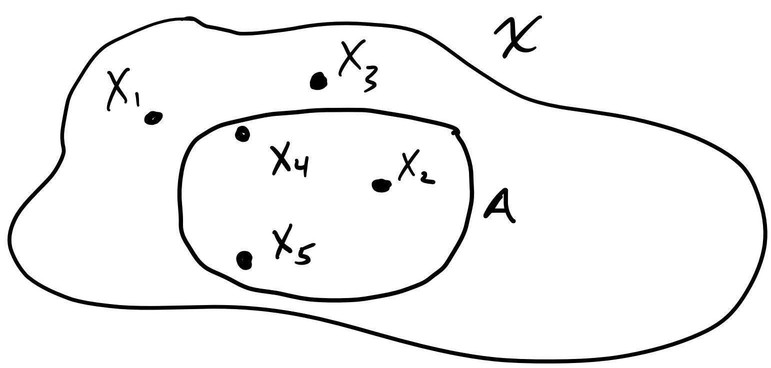

$\hat{P}_n(A) = \frac{3}{5}$

Ex (iid sampling) $X_1, \ldots, X_n \overset{iid}{\sim} P_\theta$

any model $\mathcal{P} = \{P_\theta^n : \theta \in \Theta\}$ on any $\mathcal{X}$

$\hat{P}_n$ is sufficient

$X \rightsquigarrow \hat{P}_n$ records which values observed, how many times

Minimal Sufficiency

Consider $X_1, \dots, X_n \overset{iid}{\sim} N(\theta, 1)$

$T(X) = \sum X_i$ sufficient

$\bar{X} = \frac{1}{n} \sum X_i$ also

$S(X) = (X_{(1)}, \dots, X_{(n)})$ too

$X = (X_1, \dots, X_n)$ too

Which can be recovered from which others?

$$
\begin{tikzcd}
 & X \arrow[d] \arrow[ddl, bend right] \arrow[ddr, bend left] & \\
 & S(X) \arrow[dl] \arrow[dr] & \\
\sum X_i \arrow[rr, leftrightarrow] & & \bar{X}
\end{tikzcd}
$$

these can be compressed further

These are the most compressed. Are they as compressed as possible?

Prop If $T(X)$ is sufficient and $T(x) = f(S(x))$ then $S(X)$ is sufficient

Proof: $$p_\theta(x) = g_\theta(T(x))\,h(x)$$
$$= (g_\theta \circ f)(S(x))\,h(x)$$ $\boxtimes$

Definition: $T(X)$ is <u>minimal sufficient</u> if

1) $T(X)$ is sufficient
2) For any other sufficient $S(X)$,
$$T(X) = f(S(X)) \text{ for some } f$$
(a.s. in $\mathcal{P}$)

So, no matter how many more suff. stats we add to our diagram, they will all have arrows pointing to ΣX_i

Recognizing minimal sufficiency

Assume $\mathcal{P}$ has densities p_θ, sample space $\mathcal{X}$

Define equivalence relation on $\mathcal{X}$:

$$x \equiv_{\mathcal{P}} y \quad \text{if} \quad \frac{p_\theta(x)}{p_\theta(y)} \text{ doesn't depend on } \theta$$

Note any sufficient statistic T can only collapse together equivalent values: if $T(x) = T(y) = t$

$$\frac{p_\theta(x)}{p_\theta(y)} = \frac{\mathbb{P}_\theta(X = x, T(X) = t)}{\mathbb{P}_\theta(X = y, T(X) = t)} = \frac{\mathbb{P}(X = x \mid T(X) = t)}{\mathbb{P}(X = y \mid T(X) = t)}$$

So, for any sufficient stat $T(X)$, $\quad T(x) = T(y) \Rightarrow x \equiv_{\mathcal{P}} y$

For minimal sufficient stats, the reverse implication also holds:

<u>Theorem (Bahadur)</u> $T(X)$ is minimal sufficient if

$$x \equiv_{\mathcal{P}} y \Leftrightarrow T(x) = T(y)$$

<u>Interp</u>: a minimal sufficient stat. collapses the sample space into exactly these equiv. classes.

<u>Proof</u>: First show $T(X)$ sufficient:

For any x with $T(x) = t$, we have

$$\mathbb{P}_\theta(X = x \mid T(X) = t) = \frac{p_\theta(x)}{\sum_{z:T(z)=t} p_\theta(z)} = \frac{1}{\sum_{z:T(z)=t} p_\theta(z)/p_\theta(x)}$$

which doesn't depend on θ because

$T(z) = t = T(x) \Rightarrow p_\theta(z)/p_\theta(x)$ doesn't depend on θ

Next assume $S(X)$ sufficient. If $S(x) = S(y) = s$ then $x \equiv_{\mathcal{P}} y$ so $T(x) = T(y)$. Set $f(s) = T(x)$.

Any other z with $S(z) = s$ has

$$z \equiv_{\mathcal{P}} x \Rightarrow T(z) = T(x) = f(s) = f(S(z)).$$

$\boxtimes$

(log-) Likelihood functions

Definition

Assume $\mathcal{P} = \{P_\theta : \theta \in \Theta\}$ has densities $p_\theta(x)$

The likelihood function is the (random) function

$$\mathrm{Lik}(\theta; X) = p_\theta(x)$$

function of x with parameter θ

function of θ

data X determines which function

The log-likelihood function is its log:

$$\ell(\theta; X) = \log \mathrm{Lik}(\theta; X)$$

Note if $X \equiv_{\mathcal{P}} Y$ is same as saying

$$\ell(\theta; x) - \ell(\theta; y) = \frac{p_\theta(x)}{p_\theta(y)} \quad \text{is } \underline{\text{constant}}$$

Ex Laplace location family

$$X_1, \ldots, X_n \overset{iid}{\sim} p_\theta^{(1)}(x) = \frac{1}{2} e^{-|x-\theta|}$$

$$\ell(\theta; X) = -\sum_{i=1}^{n} |X_i - \theta| - n \log 2$$

Piecewise linear in θ, knots at $X_{(i)}$

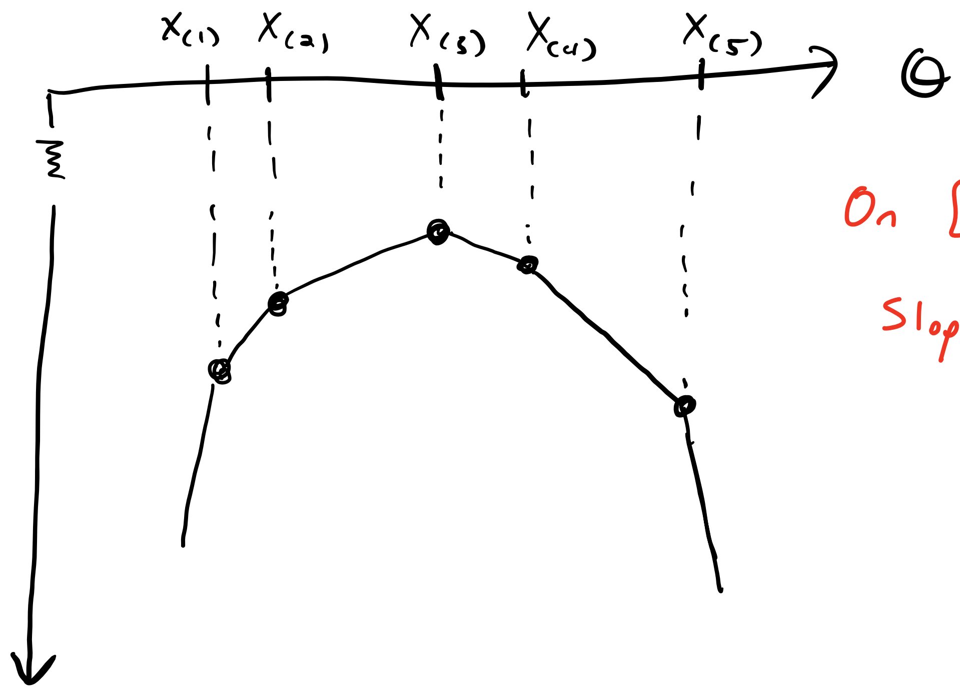

On $[X_{(k)}, X_{(k+1)}]$,
Slope $= n - 2k$

$$\ell(\theta; X) = \ell(\theta; Y) + \text{const} \iff X, Y \text{ same order statistics}$$

$\Rightarrow$ order stats are minimal suff.