Lecture 2

Lecturer: Nathan Kallus *Scribe: Aman Agarwal*

1 The Potential Outcome Framework

The Potential Outcome Framework is also known as the Neyman-Rubin Potential Outcomes, or the Rubin Causal Model. In this model, a decision, **treatment** or intervention is denoted as $z \in \mathcal{Z}$. Often, we have $\mathcal{Z} = \{0, 1\}$ where 0 is thought of as the control and 1 the treatment. The **potential outcome** $Y(z) \in \mathcal{R}$ for each $z \in \mathcal{Z}$ is the outcome we would see if we were to apply treatment z. Each individual **unit** (such as a patient in the medical setting) has a collection $\{Y(z) : z \in \mathcal{Z}\}$ of fixed outcomes. Each individual unit is also characterized by a collection of observed covariates $x \in \mathcal{X}$ (such as age, blood pressure and medical history in the patient example). Taking the frequentist view of probability, we may consider X and $Y(z)$ as random variables over the population of individual units. Thus, the joint distribution $(X, Y(z) : z \in \mathcal{Z})$ over the population describes a generic unit.

1.1 Some quantities of interest

1. Average Treatment Effect: $ATE = \mathbb{E}[Y(1) - Y(0)]$.

2. Conditional Average Treatment Effect: $CATE(x) = \mathbb{E}[Y(1) - Y(0)|\mathcal{X} = x]$.

3. Average Policy Outcome/Risk: Given poilcy $\pi : \mathcal{X} \mapsto \mathcal{Z}$, $APO(\pi) = \mathbb{E}[Y(\pi(x))]$.
   If $\mathcal{Z} = \{0, 1\}$, we define Average Policy Effect,

$$\begin{aligned} APE(\pi) &= \mathbb{E}[Y(\pi(x)) - Y(1 - \pi(x))] \\ &= \mathbb{E}[\mathbb{E}[(2\pi(x) - 1)(Y(1) - Y(0))|X = x]] \\ &= \mathbb{E}[(2\pi(x) - 1)CATE(x)] \\ &= APO(\pi) - \frac{1}{2}(APO(\pi_1) + APO(\pi_0)) \end{aligned}$$

   where π_1 and π_0 are policies that apply treatment 1 and 0 on all units respectively.

4. Best policy: $\pi^\star(x) = \operatorname{argmin}_{z \in \mathcal{Z}} \mathbb{E}[Y(z)|X = x]$

5. Best-in-class policy: $\pi^\star_\Pi = \operatorname{argmin}_{\pi \in \Pi} APO(\pi)$

More to come: LATE, CLATE, SATE, ...

1.2 Causal Inference Setting

Suppose there is a dataset of n i.i.d. draws of units $X_i, Y_i(z)$. The Central Problem of Causal Inference is that we only observe the outcomes of treatments administered and never any of the other outcomes. Let Z be the treatment administered. We only observe $Y = Y(Z)$. So, $Y(z), z \neq Z$ is not known. We may view Causal Inference as a missing data problem: $X_i, Z_i, Y_i(0), Y_i(1)$ are all the data, but we only observe $X_i, Z_i, Y_i = Y_i(z_i)$ with $Y_i(1 - z_i)$ dropped or missing.

1.3 Assumptions

The following assumptions are implicit in the above formulation.

1. Consistency: Outcome observed actually corresponds to the hypothetical potential outcome of applying Z, i.e. indeed $Y = Y(Z)$.
2. No interference: Outcomes only depend on the treatment applied to the unit, and not on treatments applied to other units.
3. Single Version of Treatment: Each z corresponds to a single version of treatment. (For instance, if the drug treatment is Aspirin, then all treatments are of the same brand, dosage etc.)

Assumptions 2 and 3 are together known as the Stable Unit Treatment Value Assumption, or SUTVA.

2 Learning to Decide with Fully Observed Counterfactuals

We first tackle the easy case where we observe $Y_i(z) \forall z \in Z$. Some examples where this happens:

1. Supervised classification: We observe label L_i and $Y_i(z) = \mathbb{1}(z \neq L_i)$.
2. Supervised regression: We observe value L_i and $Y_i(z)$ is some appropriate loss function such as $(z - L_i)^2$ or $|z - L_i|$.
3. Inventory management: We observe demand D_i and $Y_i(z) = \beta z - \alpha min(D_i, z)$.

The task is to find the best-in-class policy $\pi^\star$ in the class $\Pi \subset [\mathcal{X} \mapsto \mathcal{Z}]$. The true policy risk of a policy π is $R(\pi) = APO(\pi) = \mathbb{E}[Y(\pi(x))]$. So, given the data, we get an empirical policy risk $\hat{R_n}(\pi) = \frac{1}{n}\sum_{i=1}^n Y_i(\pi(x_i))$. How well does $\hat{R_n}(\pi)$ estimate $R(\pi)$?

2.1 Naive Approach

Fix π and look at $|\hat{R_n}(\pi) - R(\pi)|$.

Lemma 1 (Hoeffding)**.** *If $\mathbb{E}V = 0$ for $V \in [a, b]$ then $\mathbb{E}e^{tV} \leq \exp(\frac{1}{8}t^2(b-a)^2)$.*

Theorem 2 (Hoeffding Inequality)**.** *If $V_i \in [a_i, b_i]$ are independent random variables for $i \in [n]$, then*

$$P\big(|\frac{1}{n}\sum_i V_i - \frac{1}{n}\mathbb{E}\sum_i V_i| \geq \epsilon\big) \leq 2\exp\big(\frac{-2n\epsilon^2}{\frac{1}{n}\sum_{i=1}^n (b_i - a_i)^2}\big).$$

Proof. Let $S_n = \sum_i V_i$. Then $\forall t \geq 0$,

$$\begin{aligned} P(S_n - \mathbb{E}S_n \geq \epsilon) &= P(e^{t(S_n - \mathbb{E}S_n)} \geq e^{t\epsilon}) \\ &\leq e^{-t\epsilon}\mathbb{E}e^{t(S_n - \mathbb{E}S_n)} \\ &= e^{-t\epsilon}\prod_{i=1}^n \mathbb{E}e^{t(V_i - \mathbb{E}V_i)} \\ &\leq e^{-t\epsilon}\prod_{i=1}^n \exp(\frac{1}{8}t^2(b_i - a_i)^2) \end{aligned}$$

The first inequality above follows from Markov's inequality, and the second from Hoeffding's Lemma. Now setting $t = \frac{4\epsilon}{\sum_{i=1}^{n}(b_i - a_i)^2}$, gives

$$P(S_n - \mathbb{E}S_n \geq \epsilon) \leq \exp\big(\frac{-2\epsilon^2}{\sum_{i=1}^{n}(b_i - a_i)^2}\big).$$

Finally, repeating the above steps for the other side of deviation, taking Union Bound and then plugging in $n\epsilon$ as ϵ gives the desired result. $\square$

2.2 Better Approach

How well does $\hat{R_n}(\pi)$ estimate $R(\pi)$ over all $\pi \in \Pi$?