Lec 23-24

11/15/17 11:39

Recap: Individual-level effects & personalization

$$CATE(x) = \mathbb{E}[TE \mid X=x]$$
$$= \mathbb{E}[Y(1) \mid X=x] - \mathbb{E}[Y(0) \mid X=x]$$
$$= f_1(x) - f_0(x)$$

Ways to estimate:

- $\hat{f}_1 - \hat{f}_0$ diffs of two regressions

  $$\hat{f}_z(x) = \mathbb{E}[Y \mid X=x, Z=z]$$

- $\tilde{Y}_i = \left(\frac{T_i}{\hat{e}(X_i)} - \frac{1-T_i}{1-\hat{e}(X_i)}\right) Y_i$ regressed $\tilde{Y}_i$ on X_i

  works b/c

  $$\mathbb{E}[TE \mid X=x] = \mathbb{E}\left[\left(\frac{T_i}{e(X_i)} - \frac{1-T_i}{1-e(X_i)}\right) Y_i \mid X=x\right]$$

- Recursive partitioning approach

Policy learning → recall learning w/ fully observed counterfactuals

Want to find $\pi : \mathcal{X} \to [m]$ with high reward

$$\mathcal{R}(\pi) = \mathbb{E}[Y(\pi(X))]$$

from data $X_{1:n}, Z_{1:n}, Y_{1:n}$

Looking at $\pi \in \Pi = [\mathcal{X} \to [m]]$ then
estimating $R(\pi)$ for all $\pi \in \Pi$
$\Leftrightarrow$ estimating $CATE(x)$ for all x

Other extreme $\Pi = \{x \mapsto z : z \in [m]\}$
estimating $\mathcal{R}(\pi)$ for all $\pi \in \Pi$
$\Leftrightarrow$ estimating ATEs

Middle ground: Π that's rich but not everything
→ Π w/ finite VC-dim / Vanishing Rademacher complexity

Two important settings:

General m: $\qquad$ $m=2$ $(Z = \{\pm 1\})$

$$R(\pi) = \mathbb{E}[f_{\pi(X)}(X)]$$

$$\bar{R}(\pi) = R(\pi) - \frac{R(x \mapsto 1) + R(x \mapsto -1)}{2}$$
$$= \frac{1}{2}\mathbb{E}[\pi(X)\, CATE(X)]$$
$$= \frac{1}{2}\mathbb{E}[Y(\pi(X)) - Y(-\pi(X))]$$

How to learn $\pi \in \Pi$?

Maximize the estimated policy reward

① Direct method (regression adjustment for ATEs)

$$\hat{R}_n^{DM}(\pi) = \frac{1}{n}\sum_i \hat{f}_{\pi(X_i)}(X_i)$$

for $m=2$

$$\hat{R}_n^{DM}(\pi) = \frac{1}{n}\sum_i \pi(X_i) \cdot \frac{\widehat{CATE}(X_i)}{2}$$

② IPW

Let $e(z|x) = \mathbb{P}(Z=z|X=x)$

$e(Z_i|X_i)$ = Generalized propensity score

$$\hat{R}_n^{IPW}(\pi) = \frac{1}{n}\sum_i Y_i \frac{\mathbb{I}[Z_i = \pi(X_i)]}{\hat{e}(Z_i|X_i)}$$

$\updownarrow$ diff by a const

$\mathbb{E}[B|X] = CATE(x)$

$$m=2 \quad \hat{R}_n^{IPW}(\pi) = \frac{1}{n}\sum_i \frac{1}{2}\underbrace{Y_i \frac{1}{\hat{e}(Z_i|X_i)} \cdot Z_i}_{\text{Just like } \tilde{Y}_i} \pi(X_i)$$

③ AIPW (same as DR for ATEs)

$$\hat{R}_n^{AIPW}(\pi) = \frac{1}{n}\sum_i \left(f_{\pi(X_i)}(X_i) + \frac{\mathbb{I}[\pi(X_i) = Z_i]}{\hat{e}(Z_i|X_i)}\left(Y_i - \hat{f}_{Z_i}(X_i)\right)\right)$$

for $m=2$

$$\hat{R}_n^{AIPW}(\pi) = \frac{1}{n}\sum_i \frac{1}{2}\left(\hat{f}_+(X_i) - \hat{f}_-(X_i) + \frac{Z_i}{\hat{e}(Z_i|X_i)}\left(Y_i - \hat{f}_{Z_i}(X_i)\right)\right)\pi(X_i)$$

ERM-Style

Policy learning: $\max_{\pi\in\Pi} \hat{R}_n(\pi)$

for $m=2$ all of the above had the form

$$\hat{R}_n(\pi) = \frac{1}{n}\sum_i \Gamma_i \, \pi(X_i)$$

$$= \frac{1}{n}\sum_i |\Gamma_i| \underbrace{\operatorname{sign}(\Gamma_i)\pi(X_i)}_{= \begin{cases} +1 & \pi(X_i) = \operatorname{sign}(\Gamma_i) \\ -1 & \pi(X_i) \neq \operatorname{sign}(\Gamma_i)\end{cases}}$$

= weighted classification

provides a way to implement this

for $m>2$: weighted multi class classification
& different weights
for different errors

Target is low regret in Π

$$\text{Regret}_\Pi(\hat{\pi}) = \max_{\pi\in\Pi} R(\pi) - R(\hat{\pi})$$

Generalization bounds:

Consider $\hat{\pi}_n \in \operatorname{argmax}_{\pi\in\Pi} \hat{R}_n^{IPW}(\pi)$

$m=2$ $\qquad = \operatorname*{argmax}_{\pi\in\Pi} \hat{R}_n^{IPW}(\pi)$

$$= \operatorname*{argmax}_{\pi\in\Pi} \frac{1}{2}\frac{1}{n}\sum_i \frac{Y_i Z_i \pi(X_i)}{\hat{e}(Z_i|X_i)}$$

Thm: Suppose $\hat{e}(1|X) - e(1|X) = O_p(r(n))$, $\quad |Y_i(z)| \le B$

$$0 < \eta \le e(1|X),\ \hat{e}(1|X) \le 1-\eta < 1$$

$$D_n(\Pi) = \sup_{\pi\in\Pi} \left|\hat{R}_n^{IPW}(\pi) - R(\pi)\right| = O_p\left(\frac{1}{\sqrt{n}} + r(n) + \mathcal{R}_n(\Pi)\right)$$

Recall: $\bar{R}(\hat{\pi}_n) \ge \hat{R}_n^{IPW}(\hat{\pi}_n) - D_n(\Pi)$ $\qquad \pi^*_\Pi = \operatorname*{argmax}_{\pi\in\Pi} R(\pi)$

$$\ge \hat{R}_n^{IPW}(\pi^*_\Pi) - D_n(\Pi)$$

$$\ge \bar{R}(\pi^*_\Pi) - 2D_n(\Pi)$$

$\Rightarrow$ add $\frac{R(x\mapsto 1)+R(x\mapsto -1)}{2}$ to each side

$$R(\hat{\pi}_n) \ge R(\pi^*_\Pi) - 2D_n(\Pi)$$

$$\text{Regret}_\Pi(\hat{\pi}_n) \le 2D_n(\Pi)$$

Pf: Let $\bar{R}_n(\pi) = \frac{1}{n}\cdot\frac{1}{2}\sum_i \frac{1}{e(Z_i|X_i)} Y_i Z_i \pi(X_i)$

$$\sup_{\pi\in\Pi}\left|\hat{R}_n(\pi) - \bar{R}(\pi)\right| \le \sup_{\pi\in\Pi}\left|\hat{R}_n(\pi) - \bar{R}_n(\pi)\right| \quad Ⓐ$$

$$+ \sup_{\pi\in\Pi}\left|\bar{R}_n(\pi) - \bar{R}(\pi)\right| \quad Ⓑ$$

$$Ⓐ = \sup_{\pi\in\Pi}\left|\frac{1}{2n}\sum_i Y_i Z_i \pi(X_i)\left(\frac{1}{e(Z_i|X_i)} - \frac{1}{\hat{e}(Z_i|X_i)}\right)\right|$$

$$\le \frac{1}{2n}\sum_i B \left|\frac{1}{e(Z_i|X_i)} - \frac{1}{\hat{e}(Z_i|X_i)}\right|$$

$$\le \frac{B}{2n}\sum_i\left(\left|\frac{1}{e(1|X_i)} - \frac{1}{\hat{e}(1|X_i)}\right| + \left|\frac{1}{1-\underbrace{e(1|X_i)}_{e_i}} - \frac{1}{1-\underbrace{\hat{e}(1|X_i)}_{\hat{e}_i}}\right|\right)$$

$$= \frac{B}{2n}\sum_i\left(\frac{|\hat{e}_i - e_i|}{e_i\hat{e}_i} + \frac{|\hat{e}_i - e_i|}{(1-e_i)(1-\hat{e}_i)}\right)$$

$$\le \frac{B}{\eta^2 n}\sum_i |\hat{e}_i - e_i| = O_p(r(n))$$

$$Ⓑ = \sup_{\pi\in\Pi}\left|\frac{1}{2}\frac{1}{n}\sum_i \frac{Y_i Z_i \pi(X_i)}{e(Z_i|X_i)} - \overbrace{\frac{1}{2}\mathbb{E}\left[\frac{YZ\pi(X)}{e(Z|X)}\right]}^{\bar{R}(\pi)}\right|$$

$$= \frac{1}{2}\sup_{\pi\in\Pi}\left|\frac{1}{n}\sum_i\left(\frac{Y_i Z_i \pi(X_i)}{e(Z_i|X_i)} - \mathbb{E}\left[\frac{Y_i Z_i \pi(X_i)}{\underbrace{e(Z_i|X_i)}}\right]\right)\right|$$

ghost sample

$$\mathbb{E}\,\textcircled{B} = \mathbb{E}\frac{1}{2}\sup_{\pi\in\Pi}\left|\frac{1}{n}\sum_i \frac{Y_i Z_i \pi(X_i)}{e} - \mathbb{E}\left[\frac{Y_i' Z_i' \pi(X_i')}{e'}\right]\right|$$

only on X_i's, Z_i's, Y_i's

Rademacher RVs

$$\le \mathbb{E}\frac{1}{2}\sup_{\pi\in\Pi}\left|\frac{1}{n}\sum_i \sigma_i\left(\frac{Y_i Z_i \pi(X_i)}{e} - \frac{Y_i' Z_i' \pi(X_i')}{e'}\right)\right|$$

also on ghost sample & σ_i's

$$\le \mathbb{E}\sup_{\pi\in\Pi}\left|\frac{1}{n}\sum_i \sigma_i Y_i Z_i \pi(X_i)/e(Z_i|X_i)\right| \qquad \left(\left|\frac{Y_i Z_i}{e(Z_i|X_i)}\right| \le B/\eta\right)$$

$$\le B/\eta\ \mathbb{E}\sup_{\pi\in\Pi}\left|\frac{1}{n}\sum_i \sigma_i \pi(X_i)\right|$$

$$= B/\eta \cdot \mathcal{R}_n(\Pi)$$

$\bar{R}_n(\Pi)$ as a fn of $((X_1, Y_1, Z_1), \dots)$

satisfies bdd diffs condition

with $c = 2B/\eta \cdot \frac{1}{n}$

$\Rightarrow$ By McDiarmid's inequality

$$\textcircled{B} = O_p\left(\mathcal{R}_n(\Pi) + 1/\sqrt{n}\right)$$

(w/ another McDiarmid step we can get $= O_p\left(\hat{\mathcal{R}}_n(\Pi) + 1/\sqrt{n}\right)$), D)

Conclusion: If $r(n) = O_p(1/\sqrt{n})$

Π has finite VC dim then

$$D_n(\Pi) = O_p\left(\sqrt{\log n / n}\right)$$

& combine Massart's lemma + Sauer's lemma

So: can learn best-in-class policies

w/o fully observed counterfactuals

w/o specifying or learning outcome model

Role of self-normalization in policy learning

IPW: $\frac{1}{n}\sum_i \frac{Y_i}{\hat{e}(Z_i|X_i)}\mathbb{I}[\pi(X_i) = Z_i]$

π can overfit by simply putting less weight on the data

SNIPW: $$\frac{\sum_i Y_i/\hat{e}(Z_i|X_i)\ \mathbb{I}[\pi(X_i) = Z_i]}{\sum_i 1/\hat{e}(Z_i|X_i)\ \mathbb{I}[\pi(X_i) = Z_i]} \leftarrow \mathbb{E}[\cdot] = 1$$

Can be more stable

putting all π's on a somewhat more level playing field

is not of the form of weighted classification

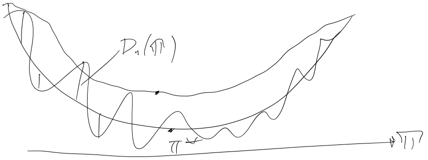

Combining DR, self-normalization, balancing for policy learning.

Problems: - Plug in estimated nuisance in denominator
- clipping — just heuristic for bias/variance
- For $m>2$: high rejection rate
- DR/SN helps but not a solution by themselves.

Take a balancing approach as before.

First consider more generally randomized policies

$$\pi: \mathcal{X} \to \Delta^m = \{\lambda \geq 0 : \textstyle\sum_i \lambda_i = 1\}$$

$\pi_z(x) =$ prob of choosing z when $X=x$

Target if we had fully observed counterfactuals or even if we know $f_z(x)$, we'd just maximize the sample reward

$$R_n(\pi) = \frac{1}{n} \sum_i \sum_z \pi_z(X_i) f_z(X_i)$$

Consider any weighted average $\hat{\tau}_W = \frac{1}{n} \sum_i W_i Y_i$
w/ $W = W(X_{1:n}, Z_{1:n})$

We'd want to minimize

$$\mathrm{CMSE}(\hat{\tau}, \pi) = \mathbb{E}\left[(\hat{\tau} - R_n(\pi))^2 \mid X_{1:n}, Z_{1:n}\right]$$

Prop Let $\varepsilon_i = Y_i - f_{Z_i}(X_i)$ $\sigma_i^2 = \mathbb{E}[\varepsilon_i^2 \mid X_i, Z_i]$ $\Sigma = \mathrm{diag}(\sigma^2)$

Under ignorability,

$$\mathrm{CMSE}(\hat{\tau}_W, \pi) = B^2(W, \pi; f_{1:m}) + \frac{1}{n^2} W^T \Sigma W$$

$$B(W, \pi; f_{1:m}) = \sum_z B_z(W, \pi_z; f_z)$$

$$B_z(W, \pi_z; f_z) = \frac{1}{n} \sum_i \left(W_i \mathbb{I}[Z_i = z] - \pi_z(X_i)\right) f_z(X_i)$$

Pf: Using $Y_i = f_{Z_i}(X_i) + \varepsilon_i = \sum_z \mathbb{I}[Z_i = z] f_z(X_i) + \varepsilon_i$

rewrite $\hat{\tau}_W = \frac{1}{n} \sum_z \sum_i W_i \mathbb{I}[Z_i = z] f_z(X_i) + \frac{1}{n} \sum_i W_i \varepsilon_i$

$$R_n(\pi) = \frac{1}{n}\sum_z \sum_i \pi_z(x_i) f_z(x_i)$$

$$\hat{\tau}_W - R_n(\pi) = \sum_z \frac{1}{n}\sum_i \left(W_i \mathbb{I}[Z_i = z] - \pi_z(X_i)\right) f_z(X_i) + \frac{1}{n}\sum_i W_i \epsilon_i$$

By igno, $\mathbb{E}[W_i \epsilon_i \mid X_{1:n}, Z_{1:n}] = W_i \mathbb{E}[\epsilon_i \mid X_i, Z_i] = 0$

$$\Rightarrow \mathbb{E}[\hat{\tau}_W \mid X_{1:n}, Z_{1:n}] = R_n(\pi) + B(W, \pi; f_{1:m})$$

$$\mathrm{CMSE}(\hat{\tau}_W, \pi) = \left(\mathbb{E}[\hat{\tau}_W \mid X_{1:n}, Z_{1:n}] - R_n(\pi)\right)^2 + \mathrm{Var}(\hat{\tau}_W \mid X_{1:n}, Z_{1:n})$$

$$\hat{\tau}_W - \mathbb{E}[\hat{\tau}_W \mid X_{1:n}, Z_{1:n}] = \frac{1}{n}\sum_i W_i \epsilon_i$$

$$\begin{aligned}\mathrm{Var}(\hat{\tau}_W \mid X_{1:n}, Z_{1:n}) &= \mathbb{E}[(\cdot)^2 \mid X_{1:n}, Z_{1:n}] \\ &= \frac{1}{n^2}\sum_{ij} W_i W_j \mathbb{E}[\epsilon_i \epsilon_j \mid X_{1:n}, Z_{1:n}] \\ &= \frac{1}{n^2}\sum_i W_i^2 \mathbb{E}[\epsilon_i^2 \mid X_i, Z_i] \\ &= \frac{1}{n^2} W^T \Sigma W. \qquad \square\end{aligned}$$

Can also consider a DR-style estimator:

$$\hat{\tau}_{W, \hat{f}_{1:m}} = \frac{1}{n}\sum_i \sum_z \pi_z(X_i) \hat{f}_z(X_i) + \frac{1}{n}\sum_i W_i (Y_i - \hat{f}_{Z_i}(X_i))$$

Prop: If $\hat{f}_{1:m} \perp\!\!\!\perp Y_{1:n} \mid X_{1:n}, Z_{1:n}$

$$\mathrm{CMSE}(\hat{\tau}_{W, \hat{f}_{1:m}}, \pi) = B^2(W, \pi; f_{1:m} - \hat{f}_{1:m}) + \frac{1}{n^2} W^T \Sigma W$$

Balanced policy evaluation:

Use $\hat{\tau}_W$ or $\hat{\tau}_{W, \hat{f}_{1:m}}$ w/

$$W = \underset{\underset{\uparrow \text{ self normalization}}{W \geq 0 : \sum_i W_i = n}}{\operatorname{argmin}} \underbrace{\sup_{\|f_{1:m}\| \leq 1} B^2(W, \pi; f_{1:m})}_{\mathbb{B}(W, \pi; \|\cdot\|)} + \frac{1}{n^2} W^T \Lambda W$$

Focus on product norms

$$\|f_{1:m}\|^2 = \sum_z \|f_z\|_z^2 / \gamma_z^2$$

Then

$$\mathbb{B}^2(W, \pi; \|\cdot\|) = \sup_{\sum_z \|f_z\|_z^2 / \gamma_z^2 \leq 1} \left(\sum_z B_z(W, \pi; f_z)\right)^2$$

$$= \sup_{\substack{\|f_z\|_z \leq \gamma_z \beta_z \\ \sum_z \beta_z^2 \leq 1}} \quad - // -$$

$$= \sup_{\sum_z \beta_z^2 \leq 1} \Big(\sum_z \gamma_z \beta_z \underbrace{\sup_{\|f_z\|_z \leq 1} B_z(W, \pi; f_z)}_{\mathbb{B}_z(W, \pi; \|\cdot\|_z)} \Big)^2$$

$$= \sup_{\|\beta\| \leq 1} \Big(\sum_z \beta_z \gamma_z \mathbb{B}_z(W, \pi; \|\cdot\|_z) \Big)^2$$

$$= \sum_z \gamma_z^2 \mathbb{B}_z^2(W, \pi; \|\cdot\|_z)$$

Next focus on RKHS norm
$\|f_z\|_z$ given by kernel K_z

$$\mathbb{B}_z^2(W, \pi; \|\cdot\|_{K_z}) = \| B_z(W, \pi; \cdot) \|_{K_z}^2$$

$$= \Big\| \frac{1}{n} \sum_i \big(W_i \mathbb{I}[Z_i = z] - \pi_z(X_i)\big) K_z(X_i, \cdot) \Big\|^2$$

$$= \langle \overset{b}{\circ}, \overset{b}{\circ} \rangle$$

$$= \frac{1}{n^2} \sum_{ij} \big(W_i \mathbb{I}[Z_i = z] - \pi_z(X_i)\big)\big(W_j \mathbb{I}[Z_j = z] - \pi_z(X_j)\big) K_z(X_i, X_j)$$

$\Rightarrow \mathbb{B}^2(W, \pi; \|\cdot\|) =$ convex quadratic in W

Also: If $f_z \sim GP(c, K_z)$ then

$$\text{CMSE}(\hat{\tau}_{W}, \pi) = \mathbb{B}^2(W, \pi; \|\cdot\|) + \frac{1}{n^2} W^T \Sigma W$$

If $f_z \sim GP(\hat{f}_z, K_z)$ then

$$\text{CMSE}(\hat{\tau}_{W, \hat{f}_{1:m}}) = \mathbb{B}^2(W, \pi; \|\cdot\|) + \frac{1}{n^2} W^T \Sigma W$$

Balanced policy learning

$$\max_{\pi \in \Pi} \hat{\tau}_{W^*(\pi)}, \leftarrow \text{potentially DR } \hat{f}_{1:m}$$

$$W^*(\pi) \in \operatorname*{argmin}_{W \geq 0 : \sum_i W_i = n} \mathbb{B}^2(W, \pi; \|\cdot\|) + \frac{1}{n^2} W^T \Lambda W$$

Thm: Suppose ⓐ $0 < \underline{\lambda} I \preceq \Lambda_n \preceq \bar{\lambda} I < \infty$

ⓑ $|\varepsilon_i| \leq B$ a.s.

ⓒ $\sqrt{K_z(x, x)} \leq \Gamma \quad \forall z$

ⓓ Strong overlap $\varphi(z|X) \geq 1/\alpha \quad \forall z$

ⓔ Ignorability

Then whp:

If $\|f_{1:m}\| < \infty$

$$\sup_{\pi\in\Pi} \left| \hat{\tau}_{W^*(\pi)} - R_n(\pi) \right| \le O_p\left(\mathfrak{R}_n(\Pi) + 1/\sqrt{n}\right)$$

If K_z C_0-universal

$$= o_p(1)$$

Conclusion would be: $\text{Regret}_{\Pi}(\hat{\pi}_n) = O_p\left(\mathfrak{R}_n(\Pi) + 1/\sqrt{n}\right)$

Pf: Let $Q(\pi) = \frac{1}{n}\sum_i \pi_{Z_i}(X_i) / \varphi(Z_i \mid X_i)$

$$\tilde{W}_i(\pi) = \frac{1}{Q(\pi)} \cdot \pi_{Z_i}(X_i) / \varphi(Z_i \mid X_i)$$

Note $\tilde{W}$ is feasible

Also note that

$$\sup_{\pi\in\Pi} \mathfrak{B}_z(\tilde{W}, \pi_z; \|\cdot\|_{K_z}) = \sup_{\substack{\pi\in\Pi \\ \|f\|\le 1}} \frac{1}{n}\sum_i \left(\frac{1}{Q(\pi)} \cdot \frac{\overset{\text{repl w/ } z}{\pi_{Z_i}}(X_i)}{\underset{\text{repl w/ } z}{\varphi}(Z_i \mid X_i)} \mathbb{I}[Z_i = z] - \pi_z(X_i) \right) f_z(X_i)$$

$$= \sup_{\substack{\pi\in\Pi \\ \|f\|\le 1}} \frac{1}{Q(\pi)} \frac{1}{n} \sum_i \left(\frac{\mathbb{I}[Z_i = z]}{\varphi(Z_i \mid X_i)} - Q(\pi) \right) \pi_z(X_i) f_z(X_i)$$

$$\le \sup_{\substack{\pi\in\Pi \\ \|f\|\le 1}} \frac{1}{n} \frac{1}{Q(\pi)} \sum_i \left(\frac{\mathbb{I}[Z_i = z]}{\varphi(Z_i \mid X_i)} - 1 \right) \pi_z(X_i) f_z(X_i)$$

$$+ \sup_{\substack{\pi\in\Pi \\ \|f\|\le 1}} \frac{1}{n} \left(Q^{-1}(\pi) - 1 \right) \cdot \sum_i \pi_z(X_i) f_z(X_i)$$

$$\le \sup_{\pi\in\Pi} \left(Q^{-1}(\pi) \right) \cdot \sup_{\substack{\pi_z\in\Pi_z \\ \|f_z\|\le 1}} \frac{1}{n} \sum_i \left(\frac{\mathbb{I}[Z_i = z]}{\varphi(Z_i \mid X_i)} - 1 \right) \pi_z(X_i) f_z(X_i)$$

$$+ \sup_{\pi\in\Pi} \left| Q^{-1}(\pi) - 1 \right| \cdot \overbrace{\frac{1}{n} \sum_i \underbrace{\sup_{\|f\|\le 1} f_z(X_i)}_{\|K_z(X_i, \cdot)\| = \sqrt{K_z(X_i, X_i)} \le \Gamma} \cdot 1}^{\le \Gamma}$$

Define $\Xi_z(X_{1:n}, Z_{1:n}) = \sup_{\substack{\pi_z\in\Pi_z \\ \|f_z\|\le 1}} \frac{1}{n} \sum_i \left(\frac{\mathbb{I}[Z_i = z]}{\varphi(Z_i \mid X_i)} - 1 \right) \pi_z(X_i) f_z(X_i)$

by $\sup_{\|f\|} f_z(X_i) \le \Gamma$ $\quad 1/\varphi(Z_i \mid X_i) \le \alpha$

$\Rightarrow$ Ξ_z sat bdd diffs condition w/ $c_i = \frac{2}{n} \alpha \Gamma$

$\Rightarrow$ By McD's ineq

$$\mathbb{P}\Big(\Xi_z \geq \underbrace{\mathbb{E}\,\Xi_z}_{\text{want to bound}} + \eta\Big) \leq e^{-n\eta^2\alpha^{-2}\Gamma^{-2}/2}$$

$$\xi_i(\pi_z, f_z) = \underbrace{\left(\frac{\mathbb{I}[z_i = z]}{q(z|x_i)} - 1\right)}_{\leq \alpha} \pi_z(x_i)\, f_z(x_i)$$

Note that $\forall \pi_z, f_z$, $\mathbb{E}\,\xi_i(\pi_z, f_z) = 0$

& ξ_i's are iid

Let ξ_i' be a ghost sample

ρ_i be iid Rademacher RVs

$$\mathbb{E}\,\Xi_z = \mathbb{E} \sup_{\substack{\pi_z \in \Pi_z \\ \|f_z\| \leq 1}} \frac{1}{n}\sum_i \left(\xi_i(\pi_z, f_z) - \mathbb{E}\,\xi_i'(\pi_z, f_z)\right)$$

$$\leq 2\,\mathbb{E} \sup_{\substack{\pi_z \in \Pi_z \\ \|f_z\| \leq 1}} \frac{1}{n}\sum_i \rho_i\, \xi_i(\pi_z, f_z)$$

By the Lipschitz comparison lemma

$$\leq 2\alpha\, \underbrace{\mathbb{E} \sup_{\pi_z, f_z} \frac{1}{n}\sum_i \rho_i\, \pi_z(x_i) f_z(x_i)}_{\substack{\text{Rademacher complexity of} \\ \{x \mapsto \pi_z(x) f_z(x) :\ \pi_z \in \Pi_z,\ \|f_z\| \leq 1\}}}$$

$$(a+b)^2 = a^2 + b^2 + 2ab$$

$$ab = \frac{1}{2}\left((a+b)^2 - a^2 - b^2\right)$$

$$\leq \alpha\, \mathbb{E} \sup_{\pi_z, f_z} \frac{1}{n}\sum_i \rho_i \left(\pi_z(x_i) + f_z(x_i)\right)^2$$

$$+ \alpha\, \mathbb{E} \sup_{\pi_z, \cancel{f_z}} \frac{1}{n}\sum_i \rho_i\, \pi_z^2(x_i)$$

$$\sup_{\|f_z\| \leq 1} f_z(x_i) = \Gamma$$

$$+ \alpha\, \mathbb{E} \sup_{\cancel{\pi_z}, f_z} \frac{1}{n}\sum_i \rho_i\, f_z^2(x_i)$$

$$\leq 4\Gamma\alpha\, \mathbb{E} \sup_{\pi_z} \frac{1}{n}\sum_i \rho_i\, \pi_z(x_i)$$

$$+ 4\Gamma\alpha\, \mathbb{E} \sup_{f_z} \frac{1}{n}\sum_i \rho_i\, f_z(x_i)$$

$$= 4\Gamma\alpha\,\mathcal{R}_n(\Pi_z) + 4\alpha\Gamma^2/\sqrt{n}$$

. . . .

Review of class

Decisions → Causal effect/consequences

⇒ If we want to make a decision or if we want to design a system that makes decisions then we need to understand causal effects.

Key framework: potential outcomes

$Y(z)$ = counterfactual of what would happen if z is the decision

What we looked at:

- Designing policies w/ fully observed counterfactuals.

don't get to observe counterfactuals →

- Randomized controlled trials
  - optimal design

don't get to control exposure →

- Noncompliance & IVs
- RDDs

don't get to interact at all →

- General observational studies
  - optimal design
- Heterogeneous effects
- Policy learning from observational