Inference in Targeted Covariate-Adjusted Response-Adaptive Randomized Clinical Trials

Antoine Chambaz

Université Paris-Ouest Nanterre

April 9th, 2014

Kick-off SPADRO

Joint work with W. Zheng and M. van der Laan (UC Berkeley)

RCTs

- **Reinforcement learning/Dynamic resources allocation**

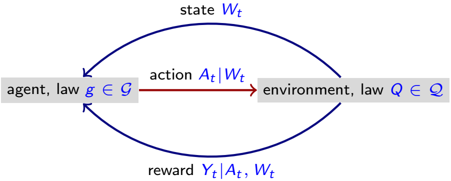

infer $\Psi(Q)$

under **optimal** action law

$$g^\star(Q) = \arg\min_{g \in \mathcal{G}} \int L_Q(g) d(Q, g)$$

  - fields of application: on-line marketing, recommender systems, **randomized clinical trials (RCTs)**

- **RCTs**
  - observed data structure: $O_t = (W_t, A_t, Y_t)$, t-th observation
    - W_t, t-th subject's baseline covariates (possibly high-dimensional)
    - $A_t \in \{0, 1\}$, placebo/treatment (randomly) assigned to t-th subject
    - $Y_t \in [0, 1]$, t-th subject's outcome of disease
  - chosen by the investigator:
    - $\Psi(Q)$, **effect of treatment on disease**
    - L_Q, loss function
    - $\mathcal{G}$, class of randomization schemes

Example of an investigator's choices

- **Excess risk Ψ**

$$\Psi(Q) = E_Q\{E_Q(Y|A=1,W) - E_Q(Y|A=0,W)\}, \quad Q \in \mathcal{Q}$$

  - statistical interpretation:
  - may be interpreted **causally** too

- **Loss function $Q \mapsto L_Q$**
  - objective: **minimizing the asymptotic variance of the estimator** of $\Psi(Q)$...
  - ... drives the characterization of $Q \mapsto L_Q$

- **Class of randomization schemes $\mathcal{G}$**
  - **covariate-adjusted** treatment assignments...
  - ... *choose* $\mathcal{G} = \{g_\theta : \theta \in \Theta \subset \mathbb{R}^d\}$ *a* parametric class

    $g_\theta(W) \equiv$ conditional probability to get $A = 1$ given W

Example of an investigator's choices

- **Excess risk Ψ**

$$\Psi(Q) = E_Q\big\{E_Q(Y|A=1,W) - E_Q(Y|A=0,W)\big\}, \quad Q \in \mathcal{Q}$$

  - statistical interpretation:
    1. for each "context" W, compare the conditional means of Y given $A=1$ versus $A=0$

$$E_Q(Y|A=1,W) - E_Q(Y|A=0,W)$$

  - may be interpreted **causally** too
- **Loss function $Q \mapsto L_Q$**
  - objective: **minimizing the asymptotic variance of the estimator** of $\Psi(Q)$...
  - ... drives the characterization of $Q \mapsto L_Q$
- **Class of randomization schemes $\mathcal{G}$**
  - **covariate-adjusted** treatment assignments...
  - ... *choose* $\mathcal{G} = \{g_\theta : \theta \in \Theta \subset \mathbb{R}^d\}$ *a* parametric class

    $g_\theta(W) \equiv$ conditional probability to get $A=1$ given W

Example of an investigator's choices

- **Excess risk Ψ**

$$\Psi(Q) = E_Q\big\{E_Q(Y|A=1,W) - E_Q(Y|A=0,W)\big\}, \quad Q \in \mathcal{Q}$$

  - statistical interpretation:
    1. for each "context" W, compare the conditional means of Y given $A=1$ versus $A=0$
    2. average out the context

$$E_Q\big\{E_Q(Y|A=1,W) - E_Q(Y|A=0,W)\big\}$$

  - may be interpreted **causally** too

- **Loss function $Q \mapsto L_Q$**
  - objective: **minimizing the asymptotic variance of the estimator** of $\Psi(Q)$...
  - ... drives the characterization of $Q \mapsto L_Q$

- **Class of randomization schemes $\mathcal{G}$**
  - **covariate-adjusted** treatment assignments...
  - ... *choose* $\mathcal{G} = \{g_\theta : \theta \in \Theta \subset \mathbb{R}^d\}$ *a* parametric class

    $g_\theta(W) \equiv$ conditional probability to get $A=1$ given W

Bibliography (non exhaustive!)

- **Sequential designs**
  - Thompson (1933), Robbins (1952)
  - specifically in the context of medical trials
    - Anscombe (1963), Colton (1963)
    - **response-adaptive designs**: Cornfield et al. (1969), Zelen (1969) and many more since then
- **Covariate-adjusted Response-Adaptive (CARA) designs**
  - Rosenberger et al. (2001), Bandyopadhyay and Biswas (2001), Zhang et al. (2007), Zhang and Hu (2009), Shao et al (2010)... *typically* study
    - **convergence of design**
    - in **correctly specified** parametric model $\mathcal{Q}$ for Q
  - (Chambaz and van der Laan, 2013) concerns
    - convergence of design *and* **asymptotic behavior** of estimator of $\Psi(Q)$
    - **without assuming** correctly specified parametric model $\mathcal{Q}$ for Q
    - ⋆ **using** (mis-specified) parametric model $\mathcal{Q}$ for Q
    - ⋆ **choosing** $\mathcal{G} = \{g_\theta : \theta \in \Theta\}$ such that $g_\theta(W) \equiv g_\theta(V)$ where $V \subset W$ only takes finitely-many values

The sampling scheme

- **Reminder**
  - objective is to estimate $\Psi(Q)$ under optimal allocation $g^\star(Q) = \arg\min_{g \in \mathcal{G}} \int L_{\bar{Q}}(g) d(Q, g)$
    $\bar{Q}(A, W) = E_Q(Y|A, W)$
  - $\mathcal{G} = \{g_\theta : \theta \in \Theta \subset \mathbb{R}^d\}$ chosen by the investigator
    covariate-adjustment: $g_\theta(W)$ (**not** $g_\theta(V)$, $V \subset W$)
- **Description of sampling scheme** (recursion)
  - starting from i.i.d. sampling from $(Q, 50\%)$
  - sample $O_1, \ldots, O_n \sim (Q, \vec{g}_n)$
    $\vec{g}_n = (g_1, \ldots, g_n) \in \mathcal{G}^n$ sequence of (known) allocation probabilities
  - conditional on $(O_1, \ldots, O_n)$
    1. estimate $\bar{Q}(A, W) = E_Q(Y|A, W)$ with $\bar{Q}_n(A, W)$ based on **lasso-regression** through choice of $\mathcal{Q}_n$, model for Q, and dealing with dependency
    2. define $g_{n+1} = \arg\min_{g \in \mathcal{G}} \int L_{\bar{Q}_n}(g) d(\mathbb{Q}_n, g)$
       $\mathbb{Q}_n$ tilted empirical measure
    3. sample $O_{n+1} \sim (Q, g_{n+1})$

Targeted inference: why?

- **Initial substitution estimator** of $\Psi(Q)$

$$\psi_n^0 = \frac{1}{n}\sum_{i=1}^{n} \bar{Q}_n(1, W_i) - \bar{Q}_n(0, W_i)$$

  - is **biased** if lasso-regression is mis-specified
  - may **fail to be $\sqrt{n}$-consistent** even if lasso-regression were correctly specified...

- **Targeted minimum loss estimation (TMLE)** of $\Psi(Q)$
  - see van der Laan and Rubin (2006), van der Laan and Rose (2012)...
  - based on semiparametrics theory, strong links with estimating function methodology
    see Bickel et al (1998), van der Vaart (1998, chapter 25), van der Laan and Robins (2003)

Targeted inference: modus operandi

- Logistic loss function: $\ell(\bar{Q})(O) = -Y \log(\bar{Q}(A,W)) - (1-Y)\log(1-\bar{Q}(A,W))$
- **Fluctuating the initial** $\bar{Q}_n$
  - characterize

$$\bar{Q}_n(\varepsilon)(A,W) = \operatorname{expit}\left\{\operatorname{logit}\left(\bar{Q}(A,W)\right) + \varepsilon\left(\frac{A}{g_n(W)} - \frac{1-A}{1-g_n(W)}\right)\right\}, \qquad \varepsilon \in \mathbb{R}$$

  such that $\bar{Q}_n(0) = \bar{Q}_n$ and $\frac{\partial}{\partial\varepsilon}\ell(\bar{Q}_n(\varepsilon))\big|_{\varepsilon=0} =$ proper direction in $L^2(Q, 50\%)$
  - define optimal fluctuation parameter

$$\varepsilon_n = \operatorname*{arg\,min}_{\varepsilon\in\mathbb{R}} \int \ell(\bar{Q}_n(\varepsilon))d\mathbb{Q}_n$$

- **TMLE** of $\Psi(Q)$

$$\psi_n^1 = \frac{1}{n}\sum_{i=1}^n \bar{Q}_n(\varepsilon_n)(1,W_i) - \bar{Q}_n(\varepsilon_n)(0,W_i)$$

Targeted inference: main results

Convergence of design

There exist Q_∞ and g_∞ such that $Q_n \to Q_\infty$ in $L^2(Q, 50\%)$ and $g_n \to g_\infty$ in $L^2(Q_W)$.

- Q_∞ is a "lasso"-projection of true Q
- g_∞ may differ from projection of Neyman allocation on $\mathcal{G}$

Asymptotics of ψ_n^1

The TMLE ψ_n^1 consistently estimates the truth $\Psi(Q)$.

Moreover, $\sqrt{n}(\psi_n^1 - \Psi(Q)) \rightsquigarrow N(0, \sigma^2)$, and we know a conservative estimator σ_n^2 of σ^2.

- **Keys**
  - Ψ is pathwise differentiable (smooth) with derivative $\nabla\Psi$
  - "robustness property of $\nabla\Psi$":
    - if $\int \nabla\Psi(Q', g) d(Q, g) = 0$ then $\Psi(Q') = \Psi(Q)$ even if $Q' \neq Q$
    - $\int \nabla\Psi(\bar{Q}_n(\varepsilon_n), g_n) d(\mathbb{Q}_n, g_n) = 0$!
  - concentration result for martingales (van Handel, 2010)

8 different models $\mathcal{G}$:

working model	parametric form	dimension	optimal variance
$\mathcal{G}_{11}$	θ_0	1	18.50
$\mathcal{G}_{12}$	$\sum_{v=1}^{3} \theta_v \mathbf{1}\{V = v\}$	3	18.18
$\mathcal{G}_{13}$	$\theta_0 + \theta_1 U$	2	18.37
$\mathcal{G}_{14}$	$\sum_{v=1}^{3} \theta_v \mathbf{1}\{V = v\} + \theta_4 U$	4	18.05
$\mathcal{G}_{15}$	$\theta_0 + \sum_{v=1}^{3} \theta_v \mathbf{1}\{V = v\} U$	4	18.12
$\mathcal{G}_{16}$	$\sum_{v=1}^{3} \theta_v \mathbf{1}\{V = v\} + \theta_4 U + \sum_{v=2}^{3} \theta_{3+v} \mathbf{1}\{V = v\} U$	6	18.01
$\mathcal{G}_{17}$	$\theta_0 + \sum_{v=1}^{3} \theta_v \mathbf{1}\{V = v\} U + \sum_{v=1}^{3} \theta_{4+v} \mathbf{1}\{V = v\} U^2$	7	18.36
$\mathcal{G}_{18}$	$\sum_{v=1}^{3} \theta_v \mathbf{1}\{V = v\} + \theta_4 U + \theta_5 U^2 + \sum_{v=2}^{3} \theta_{4+v} \mathbf{1}\{V = v\} U$ $+ \sum_{v=2}^{3} \theta_{6+v} \mathbf{1}\{V = v\} U^2$	9	18.03

4 different parametric models and 4 different "lasso" procedures $\mathcal{Q}$:

	working model	parametric form	dimension
parametric	$\mathcal{Q}_{11}$	$\sum_{v=1}^{3} \theta_v \mathbf{1}\{V = v\} + \theta_4 U + \theta_5 A$	5
parametric	$\mathcal{Q}_{12}$	$\theta_0 + A\left(\theta_1 U + \sum_{v=2}^{3} \theta_v \mathbf{1}\{V = v\}\right) +(1 - A)\left(\theta_4 U + \sum_{v=2}^{3} \theta_{3+v} \mathbf{1}\{V = v\}\right)$	7
parametric	$\mathcal{Q}_{13}$	$A\left(\sum_{v=1}^{3} \theta_v \mathbf{1}\{V = v\} + \theta_4 U\right) +(1 - A)\left(\sum_{v=1}^{3} \theta_{4+v} \mathbf{1}\{V = v\} + \theta_8 U\right)$	8
parametric	$\mathcal{Q}_{14}$	$A\left(\sum_{v=1}^{3} \theta_v \mathbf{1}\{V = v\} + \theta_4 U + \theta_5 U^2\right) +(1 - A)\left(\sum_{v=1}^{3} \theta_{5+v} \mathbf{1}\{V = v\} + \theta_9 U + \theta_{10} U^2\right)$	10
lasso	$\mathcal{Q}_{15}$	$A\left(\sum_{v=1}^{3} \theta_v \mathbf{1}\{V = v\} + \theta_4 U + \theta_5 U^2\right) +(1 - A)\left(\sum_{v=1}^{3} \theta_{5+v} \mathbf{1}\{V = v\} + \theta_9 U + \theta_{10} U^2\right)$	10
lasso	$\mathcal{Q}_{16}$	$A\left(\sum_{v=1}^{3} \theta_v \mathbf{1}\{V = v\} + \sum_{l=1}^{5} \theta_{3+l} U^l\right) +(1 - A)\left(\sum_{v=1}^{3} \theta_{8+v} \mathbf{1}\{V = v\} + \sum_{l=1}^{5} \theta_{11+l} U^l\right)$	16
lasso	$\mathcal{Q}_{17}$	$A\left(\sum_{v=1}^{3} \theta_v \mathbf{1}\{V = v\} + \sum_{l=1}^{10} \theta_{3+l} U^l\right) +(1 - A)\left(\sum_{v=1}^{3} \theta_{13+v} \mathbf{1}\{V = v\} + \sum_{l=1}^{10} \theta_{16+l} U^l\right)$	26
lasso	$\mathcal{Q}_{18}$	$A\left(\sum_{v=1}^{3} \theta_v \mathbf{1}\{V = v\} + \sum_{l=1}^{20} \theta_{3+l} U^l\right) +(1 - A)\left(\sum_{v=1}^{3} \theta_{23+v} \mathbf{1}\{V = v\} + \sum_{l=1}^{20} \theta_{26+l} U^l\right)$	46

- 1000 independent replications for each of the 64 combinations
- Updating and targeting at intermediate sample sizes $250, 500, \ldots, 1750, 2000$
- Summary:
  - Consistency guaranteed
  - 95%-confidence intervals guarantee at least 94%-coverage for all sample sizes
  - Variances of TMLEs nearly coincide with the targeted values

On-going

- Replacing parametric $\mathcal{G} = \{g_\theta : \theta \in \Theta\}$ with nonparametric $\mathcal{G}$

  e.g., using lasso-regression to better target the Neyman allocation
- Deriving finite sample results
- Assessing the sensitivity to protocol violations in real-life RCTs
- Assessing the sensitivity to non-stationarity in real-life RCTs

lasso-regression

- "lasso-regression" for "ℓ^1-restricted least-squares regression"
- Consider
  - $\{\phi_j : j \geq 0\}$ a basis of a given class of functions, $\|\phi_j\|_\infty = 1$ (all $j \geq 0$)
  - $\Phi_\beta(W) = \sum_{j\geq 0} \beta_j \phi_j(W)$ (all $\beta \in \ell^1(\mathbb{N})$)
  - $\{d_n\}_{n\geq 0}$ and $\{b_n\}_{n\geq 0}$, both non-decreasing and unbounded, $d_n = o(n^r)$
  - $B_n = \{\beta \in \ell^1(\mathbb{N}) : j \geq d_n \implies \beta_j = 0, \text{ and } \|\beta\|_1 \leq b_n \wedge M\}$
- Define $\mathcal{Q}_n = \{O \mapsto A\Phi_\beta(W) + (1-A)\Phi_{\beta'}(W) : \beta, \beta' \in B_n\}$