Active Adaptive Experimental Design for Treatment Effect Estimation with Covariate Choices

Oral 6F
Experimental Design and Simulation
Thu 25 Jul 4:30 p.m. - 4:45 p.m.

Masahiro Kato
Akihiro Oga
Wataru Komatsubara
Ryo Inokuchi

MIZUHO
Mizuho-DL Financial Technology

Prologue

- This study began with discussions with data scientists who have an interest in causal inference under distribution shifts.

DS

Hi, I have a quick question. When the covariates in the experimental data differ from those in the population we want to evaluate, does the estimator of the treatment effect become inefficient?

Hi, whether the asymptotic variance increases when the distribution shifts depends on the situation. It might be possible to collect experimental data in a way that reduces the asymptotic variance of the treatment effect estimator for the population of interest.

→ From this discussion we started our research.

1. Contribution

Goal: Estimation of the average treatment effect (ATE) using data gathered from an experiment.

Our task: Design of an experiment that returns an ATE estimator with the smallest asymptotic variance.

Contribution: **Active adaptive experiment**.

1. New framework for experimental design:
   Optimize both propensity score and covariate density.
2. Efficient adaptive experiment:
   Minimize the asymptotic variance of an ATE estimator.
3. Asymptotic analysis of our estimator:
   Asymptotic normality and efficiency.

Technique: Experiment for ATE estimation (Hahn et al., 2011; Kato et al., 2020) + Covariate shift adaptation (Shimoraida, 2000).

2. Potential Outcomes and ATE

- Binary treatment $a \in \{1, 0\}$.
- Potential outcome $Y(a) \in \mathbb{R}$, where $(Y(1), Y(0)) \mid X \sim \zeta(y(1), y(0) \mid X)$.
- Each unit has covariates $X \in \mathcal{X}$.
- Covariate distribution of interest: $q^*(x)$.
- The average treatment effect (ATE) over q^*:

$$\theta_0 := \int \big(y(1) - y(0)\big)\zeta(y(1), y(0)|x)q^*(x)dy(1)dy(0)dx.$$

3. Active Adaptive Experiment

- **Active adaptive experiment:** In each $t \in [T] := \{1,2,\dots,T\}$:
  - Sample a unit from $p_t(x)$ as $X_t \sim p_t(x)$.
  - Assign treatment $A_t \in \{1,0\}$ with probability $\pi_t(A_t \mid X_t)$.
  - Observe outcome $Y_t = \sum_{a\in\{1,0\}} 1[A_t = a]Y(a)$.

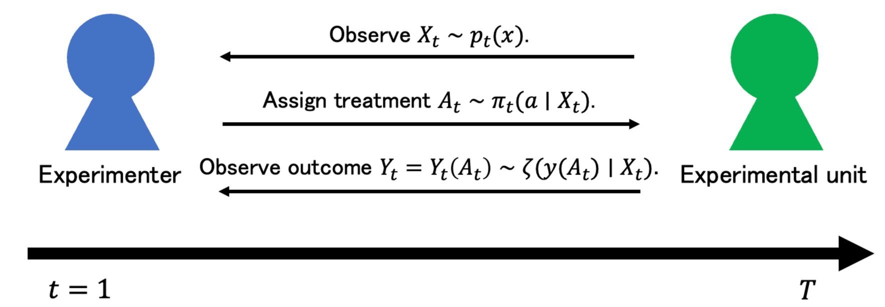

4. Semiparametric Efficiency Bound

- Lower bound of the asymptotic variance of ATE estimators.
- Data-generating process with some fixed $p(x)$:

$$(Y, A, X) \sim p(y, a, x) := \prod_{d\in\{1,0\}} \big(\zeta(y(d) \mid x)\pi(d \mid x)\big)^{1[d=a]} p(x).$$

- Assume that we can obtain infinite samples from $q^*(x)$.
- **Semiparametric efficiency bound** for θ_0 (also see Uehara et al., 2021):

$$V(\pi, p) := \int \left(\frac{\mathrm{Var}(Y(1)|x)}{\pi(1|x)} + \frac{\mathrm{Var}(Y(0)|x)}{\pi(0|x)}\right)\frac{q^*(x)}{p(x)} q^*(x)\,\zeta(y(1), y(0)|x)dy(1)dy(0)dx.$$

5. Propensity Scores and Covariate Density

- **Efficient propensity score and covariate density:**

$$(\pi^*, p^*) := \arg\min_{\pi, p} V(\pi, p).$$

- We obtain the close forms of π^* and p^* as follows:

$$\pi^*(a|x) = \frac{\sqrt{\mathrm{Var}(Y(a)|x)}}{\sqrt{\mathrm{Var}(Y(1)|x)} + \sqrt{\mathrm{Var}(Y(0)|x)}},$$

$$p^*(x) = \frac{\sqrt{\mathrm{Var}(Y(1)|x)} + \sqrt{\mathrm{Var}(Y(0)|x)}}{\int \left(\sqrt{\mathrm{Var}(Y(1)|x)} + \sqrt{\mathrm{Var}(Y(0)|x)}\right) q^*(x)dx}.$$

- π^* is known as the Neyman allocation (Neyman, 1934).
- Kato et al. (2020) proposes optimizing only π^*.
- Sample more units with higher variance.
  Ex. Sugiyama (2000).

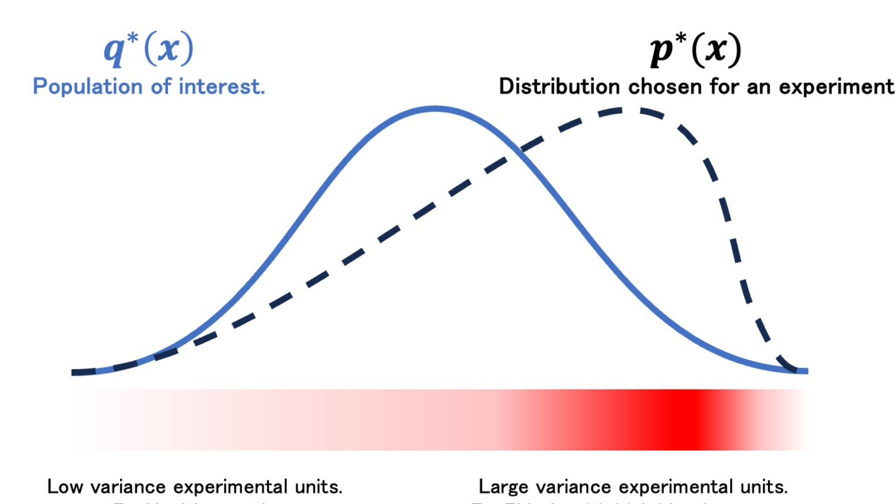

6. Efficient Active Adaptive Experiment

- **Our designed adaptive experiment**: In each round $t \in [T]$:
  - Estimate $\mathrm{Var}(Y(a)|X)$ using $\{(Y_s, A_s, X_s)\}_{s=1}^{t-1}$.
    We denote the estimator by $\widehat{\mathrm{Var}}_t(Y(a)|X)$.
  - Sample $X_t \sim p_t(x) = \dfrac{\sqrt{\widehat{\mathrm{Var}}_t(Y(1)|x)} + \sqrt{\widehat{\mathrm{Var}}_t(Y(0)|x)}}{\int\left(\sqrt{\widehat{\mathrm{Var}}_t(Y(1)|x')} + \sqrt{\widehat{\mathrm{Var}}_t(Y(0)|x')}\right)q^*(x')dx'}$.
  - Assign treatment $A_t \sim \pi_t(a \mid X_t) = \dfrac{\sqrt{\widehat{\mathrm{Var}}_t(Y(a)|X_t)}}{\sqrt{\widehat{\mathrm{Var}}_t(Y(1)|X_t)} + \sqrt{\widehat{\mathrm{Var}}_t(Y(0)|X_t)}}$.
- At the end, we estimate the ATE by using $\{(Y_t, A_t, X_t)\}_{t=1}^{T}$ as

$$\hat{\theta}_T := \frac{1}{T}\sum_{t\in[T]}\left(\frac{1[A_t = 1]\big(Y_t - \hat{\mu}_t(1 \mid X_t)\big)}{\pi_t(1 \mid X_t)} - \frac{1[A_t = 0]\big(Y_t - \hat{\mu}_t(0 \mid X_t)\big)}{\pi_t(0 \mid X_t)}\right)\frac{q^*(X_t)}{p_t(X_t)} + \int \big(\hat{\mu}_t(1 \mid x) - \hat{\mu}_t(0 \mid x)\big)q^*(x)dx,$$

where $\hat{\mu}_t(a \mid X_t)$ is an estimator of $\mathbb{E}[Y(a) \mid X_t]$ only using $\{(X_s, A_s, Y_s)\}_{s=1}^{t-1}$.

- Assume $\pi_t(a \mid x) - \pi^*(a \mid x) \to 0$ and $p_t(x) \to p^*(x)$ $\forall x$ as $t \to \infty$ a.s.

  Then, it holds that $\sqrt{T}\big(\hat{\theta}_T - \theta_0\big) \xrightarrow{d} \mathcal{N}\big(0, V(\pi^*, p^*)\big)$ $(T \to \infty)$.

7. Simulation Studies

- Conditional variances of two treatments →
- Comparison of $q^*(x)$ and $p^*(x)$ in two cases:
- Gaussian and Uniform dist. for $q^*(x)$.

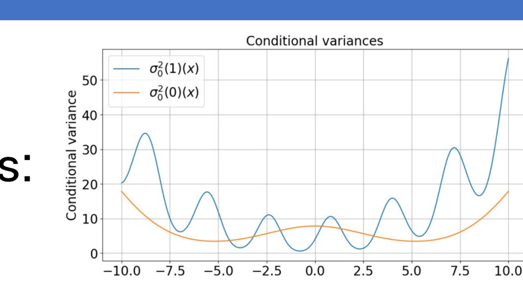

Figure 3: Heterogeneous variances.

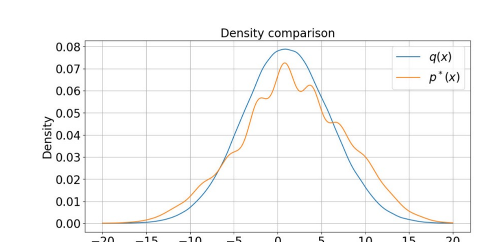

Figure 4: Comparison between $q(x)$ and $p^*(x)$ in the experiment with the covariates following the Gaussian distribution.

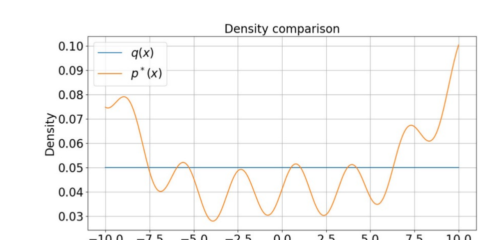

Figure 5: Comparison between $q(x)$ and $p^*(x)$ in the experiment with the covariates following the Uniform distribution.

- Empirical MSEs between ATE estimators and the true ATE.

AS-AIPW: optimize only treatment-assignment probabilities.

AAS-AIPW: optimize treatment-assignment probability and covariate density.

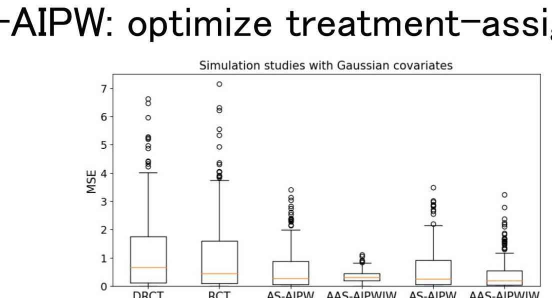

Figure 6: Results of simulation studies with covariates following with a Gaussian distribution $\mathcal{N}(1, 25)$.

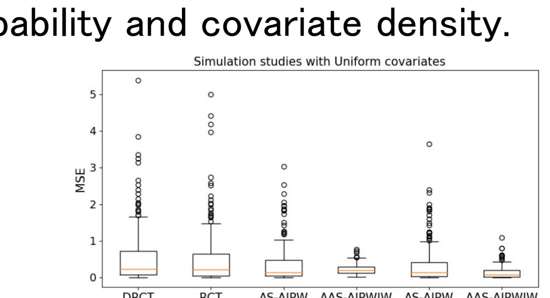

Figure 7: Results of simulation studies with covariates following with a Uniform distribution $\mathcal{U}(-10, 10)$.

Reference

- Hahn, J., Hirano, K., and Karlan, D. Adaptive experimental design using the propensity score. *Journal of Business and Economic Statistics*, 2011.
- Kato, M., Ishihara, T., Honda, J., and Narita, Y. Adaptive experimental design for efficient treatment effect estimation: Randomized allocation via contextual bandit algorithm. arXiv:2002.05308, 2020.
- Shimodaira, H., "Improving predictive inference under covariate shift by weighting the log-likelihood function," *Journal of statistical planning and inference*, 2010.
- Sugiyama, M. Active learning in approximately linear regression based on conditional expectation of generalization error. *Journal of Machine Learning Research*, 2006.
- Uehara, M., Kato, M., and Yasui, S., Off-Policy Evaluation and Learning for External Validity under a Covariate Shift. In *NeurIPS* 2020