Augmented inverse propensity weighting for randomized experiments

Steve Howard
The Voleon Group

March 14, 2023

Some common analyses of randomized experiments

- Estimate average treatment effect / absolute lift
- Estimate relative lift
- Estimate quantile treatment effects
- Estimate a model for heterogeneous treatment effects
- Estimate the value of a personalized policy (and perhaps optimize such a policy)

We want to do all these things *efficiently*, making use of covariates to reducing variance.

Some common analyses of randomized experiments

- Estimate average treatment effect / absolute lift
- Estimate relative lift
- Estimate quantile treatment effects
- Estimate a model for heterogeneous treatment effects
- Estimate the value of a personalized policy (and perhaps optimize such a policy)

We want to do all these things *efficiently*, making use of covariates to reducing variance.

Augmented inverse propensity weighting (AIPW) provides a unified approach to solving all of these.

Setup: two-armed experiment with potential outcomes

- We run an experiment with two arms, treatment and control.
- We sample i.i.d. from a population of units $(X, W, Y(0), Y(1))$ *(e.g., users)*
  - X: covariates measured before treatment *(e.g., pre-experiment weekly sessions)*
  - $W \in 0, 1$: treatment indicator *(e.g., send new notifications)*
  - $Y(0), Y(1)$: potential outcomes under control and treatment, respectively *(e.g., weekly sessions if assigned to control or treatment)*
- We observe either $Y(0)$ or $Y(1)$, not both: our *observed* sample consists of $(X_i, W_i, Y_i(W_i))$
  - We make the usual SUTVA / no-interference assumption.
- We assign W_i by i.i.d. coin flips with $P(W_i = 1) = p$.

Example potential outcomes table

Full table (unobserved):

	X_i	W_i	$Y_i(0)$	$Y_i(1)$
Unit 1	6	1	5	8
Unit 2	2	0	2	0
Unit 3	3	1	4	4
⋮	⋮	⋮	⋮	⋮

Observed table:

	X_i	W_i	$Y_i(0)$	$Y_i(1)$
Unit 1	6	1	?	8
Unit 2	2	0	2	?
Unit 3	3	1	?	4
⋮	⋮	⋮	⋮	⋮

The estimands and some (inefficient) estimators

Estimand	Definition	Naive estimator
ATE / Absolute lift	$\mathbb{E}\left[Y(1) - Y(0)\right]$	$\mathbb{E}_n\left[Y_i(1) \mid W_i = 1\right] - \mathbb{E}_n\left[Y_i(0) \mid W_i = 0\right]$

The estimands and some (inefficient) estimators

Estimand	Definition	Naive estimator
ATE / Absolute lift	$\mathbb{E}\left[Y(1) - Y(0)\right]$	$\mathbb{E}_n\left[Y_i(1) \mid W_i = 1\right] - \mathbb{E}_n\left[Y_i(0) \mid W_i = 0\right]$
Relative lift	$\frac{\mathbb{E}[Y(1)]}{\mathbb{E}[Y(0)]}$	$\frac{\mathbb{E}_n[Y_i(1) \mid W_i=1]}{\mathbb{E}_n[Y_i(0) \mid W_i=0]}$

The estimands and some (inefficient) estimators

Estimand	Definition	Naive estimator
ATE / Absolute lift	$\mathbb{E}\left[Y(1) - Y(0)\right]$	$\mathbb{E}_n\left[Y_i(1) \mid W_i = 1\right] - \mathbb{E}_n\left[Y_i(0) \mid W_i = 0\right]$
Relative lift	$\frac{\mathbb{E}[Y(1)]}{\mathbb{E}[Y(0)]}$	$\frac{\mathbb{E}_n[Y_i(1) \mid W_i=1]}{\mathbb{E}_n[Y_i(0) \mid W_i=0]}$
Policy value	$V(\pi) = \mathbb{E}\left[Y(\pi(X))\right]$	$\widehat{V}(\pi) = \mathbb{E}_n\left[\frac{1_{A_i=\pi(X_i)} Y_i(\pi(X_i))}{\mathbb{P}(A_i=\pi(X_i))}\right]$

The estimands and some (inefficient) estimators

Estimand	Definition	Naive estimator
ATE / Absolute lift	$\mathbb{E}\left[Y(1) - Y(0)\right]$	$\mathbb{E}_n\left[Y_i(1) \mid W_i = 1\right] - \mathbb{E}_n\left[Y_i(0) \mid W_i = 0\right]$
Relative lift	$\frac{\mathbb{E}[Y(1)]}{\mathbb{E}[Y(0)]}$	$\frac{\mathbb{E}_n[Y_i(1) \mid W_i=1]}{\mathbb{E}_n[Y_i(0) \mid W_i=0]}$
Policy value	$V(\pi) = \mathbb{E}\left[Y(\pi(X))\right]$	$\widehat{V}(\pi) = \mathbb{E}_n\left[\frac{1_{A_i=\pi(X_i)} Y_i(\pi(X_i))}{\mathbb{P}(A_i=\pi(X_i))}\right]$
CATE(x)	$\mathbb{E}\left[Y(1) - Y(0) \mid X = x\right]$	$Y_i(1) \sim \mu_1(X_i) \mid W_i = 1$ $Y_i(0) \sim \mu_0(X_i) \mid W_i = 0$ $\widehat{\text{CATE}}(x) = \hat{\mu}_1(x) - \hat{\mu}_0(x)$

The AIPW "imputation principle"

1. Regress $Y_i(1)$ on X_i among units with $W_i = 1$ to get $\hat{\mu}_1(x)$.
2. Regress $Y_i(0)$ on X_i among units with $W_i = 0$ to get $\hat{\mu}_0(x)$.
3. For each unit i compute "pseudo-outcomes"

$$\Gamma_i(0) = \hat{\mu}_0(X_i) + \frac{1 - W_i}{1 - p} \left(Y_i(0) - \hat{\mu}_0(0)\right)$$

$$\Gamma_i(1) = \hat{\mu}_1(X_i) + \frac{W_i}{p} \left(Y_i(1) - \hat{\mu}_1(X_i)\right)$$

(In practice, need to use cross-fitting. Also a good idea to use empirical propensity.)

The AIPW "imputation principle"

1. Regress $Y_i(1)$ on X_i among units with $W_i = 1$ to get $\hat{\mu}_1(x)$.
2. Regress $Y_i(0)$ on X_i among units with $W_i = 0$ to get $\hat{\mu}_0(x)$.
3. For each unit i compute "pseudo-outcomes"

$$\Gamma_i(0) = \hat{\mu}_0(X_i) + \frac{1 - W_i}{1 - p} \left(Y_i(0) - \hat{\mu}_0(0)\right)$$
$$\Gamma_i(1) = \hat{\mu}_1(X_i) + \frac{W_i}{p} \left(Y_i(1) - \hat{\mu}_1(X_i)\right)$$

(In practice, need to use cross-fitting. Also a good idea to use empirical propensity.)

Then treat $\Gamma_i(0), \Gamma_i(1)$ as if they were observed potential outcomes.

Imputed potential outcomes table

Full table (unobserved):

	X_i	W_i	$Y_i(0)$	$Y_i(1)$
Unit 1	6	1	5	8
Unit 2	2	0	2	0
Unit 3	3	1	4	4
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$

Imputed table:

	X_i	W_i	$\Gamma_i(0)$	$\Gamma_i(1)$
Unit 1	6	1	6	10
Unit 2	2	0	2	2
Unit 3	3	1	3	5
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$

Efficient AIPW estimators

Estimand	Definition	AIPW estimator
ATE / Absolute lift	$\mathbb{E}[Y(1) - Y(0)]$	$\mathbb{E}_n[\Gamma_i(1) - \Gamma_i(0)]$

Efficient AIPW estimators

Estimand	Definition	AIPW estimator
ATE / Absolute lift	$\mathbb{E}\left[Y(1) - Y(0)\right]$	$\mathbb{E}_n\left[\Gamma_i(1) - \Gamma_i(0)\right]$
Relative lift	$\frac{\mathbb{E}[Y(1)]}{\mathbb{E}[Y(0)]}$	$\frac{\mathbb{E}_n[\Gamma_i(1)]}{\mathbb{E}_n[\Gamma_i(0)]}$

Efficient AIPW estimators

Estimand	Definition	AIPW estimator
ATE / Absolute lift	$\mathbb{E}\left[Y(1) - Y(0)\right]$	$\mathbb{E}_n\left[\Gamma_i(1) - \Gamma_i(0)\right]$
Relative lift	$\frac{\mathbb{E}[Y(1)]}{\mathbb{E}[Y(0)]}$	$\frac{\mathbb{E}_n[\Gamma_i(1)]}{\mathbb{E}_n[\Gamma_i(0)]}$
Policy value	$V(\pi) = \mathbb{E}\left[Y(\pi(X))\right]$	$\widehat{V}(\pi) = \mathbb{E}_n\left[\Gamma_i(\pi(X_i))\right]$

Efficient AIPW estimators

Estimand	Definition	AIPW estimator
ATE / Absolute lift	$\mathbb{E}[Y(1) - Y(0)]$	$\mathbb{E}_n[\Gamma_i(1) - \Gamma_i(0)]$
Relative lift	$\frac{\mathbb{E}[Y(1)]}{\mathbb{E}[Y(0)]}$	$\frac{\mathbb{E}_n[\Gamma_i(1)]}{\mathbb{E}_n[\Gamma_i(0)]}$
Policy value	$V(\pi) = \mathbb{E}[Y(\pi(X))]$	$\widehat{V}(\pi) = \mathbb{E}_n[\Gamma_i(\pi(X_i))]$
CATE(x)	$\mathbb{E}[Y(1) - Y(0) \mid X = x]$	Regress $\Gamma_i(1) - \Gamma_i(0)$ on X_i

Attractive properties of AIPW

- Unbiased no matter what outcome model you use

Attractive properties of AIPW

- Unbiased no matter what outcome model you use
- Lowest possible variance if the outcome model is good

Attractive properties of AIPW

- Unbiased no matter what outcome model you use
- Lowest possible variance if the outcome model is good
- Can estimate variance ignoring randomness of the outcome model. For example:

$$\widehat{\text{ATE}} = \frac{1}{n}\sum_{i=1}^{n}[\Gamma_i(1) - \Gamma_i(0)]$$

$$\widehat{\text{Var}}(\widehat{\text{ATE}}) = \frac{1}{n^2}\sum_{i=1}^{n}\left[\Gamma_i(1) - \Gamma_i(0) - \widehat{\text{ATE}}\right]^2.$$

(The story is more complicated for CATE estimation.)

Attractive properties of AIPW

- Unbiased no matter what outcome model you use
- Lowest possible variance if the outcome model is good
- Can estimate variance ignoring randomness of the outcome model. For example:

$$\widehat{\text{ATE}} = \frac{1}{n}\sum_{i=1}^{n}\left[\Gamma_i(1) - \Gamma_i(0)\right]$$

$$\widehat{\text{Var}}(\widehat{\text{ATE}}) = \frac{1}{n^2}\sum_{i=1}^{n}\left[\Gamma_i(1) - \Gamma_i(0) - \widehat{\text{ATE}}\right]^2.$$

(The story is more complicated for CATE estimation.)

Generalizes to non-constant propensity score: replace p with $e(x) = P\left(W_i = 1 \mid X_i = x\right)$.

AIPW for policy optimization

Recall AIPW estimate of the value of a personalized policy π:

$$\widehat{V}(\pi) = \mathbb{E}_n \left[\Gamma_i(\pi(X_i))\right].$$

Can treat this as an "empirical welfare" objective and optimize over π.

Reduces to cost-sensitive multiclass classification.

In the binary treatment case, this is weighted binary classification:

$$\begin{aligned} \text{Label:} &\quad 1_{\Gamma_i(1) > \Gamma_i(0)} \\ \text{Weight:} &\quad |\Gamma_i(1) - \Gamma_i(0)| \end{aligned}$$

References

Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., & Robins, J (2018). Double/Debiased Machine Learning for Treatment and Structural Parameters. The Econometrics Journal 21(1), C1–C68.

Jin, Y., & Ba, S. (2022). Toward Optimal Variance Reduction in Online Controlled Experiments. Technometrics, 1-12.

Dudík, M., Langford, J., & Li, L. (2011). Doubly robust policy evaluation and learning. arXiv preprint arXiv:1103.4601.

Kennedy, E. H. (2020). Towards optimal doubly robust estimation of heterogeneous causal effects. arXiv preprint arXiv:2004.14497.

Angelopoulos, A. N., Bates, S., Fannjiang, C., Jordan, M. I., & Zrnic, T. (2023). Prediction-Powered Inference. arXiv preprint arXiv:2301.09633.

Questions?

steve@stevehoward.org

More details in the blog post at stevehoward.org

Bonus: quantiles

Suppose we want to estimate the π-quantile of treatment outcomes $Y(1)$.

Bonus: quantiles

Suppose we want to estimate the π-quantile of treatment outcomes $Y(1)$.

This is defined by the population estimating equation

$$\theta^\star = \underset{\theta}{\text{solve}}\, \mathbb{E}\left[1_{Y(1)\leq\theta} - \pi\right],$$

and ordinarily our (M-)estimator would be the sample quantile

$$\hat{\theta} = \underset{\theta}{\text{solve}}\, \frac{1}{n}\sum_{i=1}^{n} 1_{Y_i(1)\leq\theta} - \pi,$$

but $Y_i(1)$ is missing for units with $A_i = 0$. So **apply AIPW to the estimating equation** itself:

$$\hat{\theta}_{\text{AIPW}} = \underset{\theta}{\text{solve}}\, \frac{1}{n}\sum_{i=1}^{n}\left[1_{\hat{\mu}_1(X_i)\leq\theta} + \frac{A_i}{p}\left(1_{Y_i(1)\leq\theta} - 1_{\hat{\mu}_1(X_i)\leq\theta}\right)\right].$$