Experimental Evaluation of Individualized Treatment Rules

Kosuke Imai

Department of Government   Department of Statistics
Harvard University

Adaptive Experimentation Workshop, Facebook
February 13, 2020

Joint work with Michael Lingzhi Li (MIT)

Overview

- Individualized treatment rules (ITRs)
  - personalized medicine
  - micro-targeting in business and politics
- Existing literature:
  - development of optimal ITRs
  - estimation of heterogeneous treatment effects
- We propose to use a randomized experiment to evaluate ITRs
  1. Neyman's repeated sampling framework
     - randomized treatment assignment, random sampling
     - no modeling assumption or asymptotic approximation
  2. Cross-validation
     - same experimental data used to estimate and evaluate ITRs
     - additional uncertainty due to the estimation of ITRs
  3. Evaluation measures
     - incorporating a budget constraint
     - Area under the prescriptive effect curve (AUPEC)

Evaluation without a Budget Constraint

- Setup
  - Binary treatment: $T_i \in \{0, 1\}$
  - Pre-treatment covariates: $\mathbf{X} \in \mathcal{X}$
  - No interference:

$$Y_i(T_1 = t_1, T_2 = t_2, \ldots, T_n = t_n) = Y_i(T_i = t_i)$$

  - Random sampling of units:

$$(Y_i(1), Y_i(0), \mathbf{X}_i) \overset{\text{i.i.d.}}{\sim} \mathcal{P}$$

  - Completely randomized treatment assignment:

$$\Pr(T_i = 1 \mid Y_i(1), Y_i(0), \mathbf{X}_i) = \frac{n_1}{n} \quad \text{where} \quad n_1 = \sum_{i=1}^{n} T_i$$

- (Fixed for now) ITR:

$$f : \mathcal{X} \longrightarrow \{0, 1\}$$

Inference for the Standard Metric

- Standard metric (Population Average "Value" or PAV):

$$\lambda_f = \mathbb{E}\{Y_i(f(\mathbf{X}_i))\}$$

- A natural estimator:

$$\hat{\lambda}_f(\mathcal{Z}) = \frac{1}{n_1}\sum_{i=1}^{n} Y_i T_i f(\mathbf{X}_i) + \frac{1}{n_0}\sum_{i=1}^{n} Y_i(1-T_i)(1-f(\mathbf{X}_i)),$$

  where $\mathcal{Z} = \{\mathbf{X}_i, T_i, Y_i\}_{i=1}^n$
- Unbiasedness: $\mathbb{E}\{\hat{\lambda}_f(\mathcal{Z})\} = \lambda_f$
- Variance:

$$\mathbb{V}\{\hat{\lambda}_f(\mathcal{Z})\} = \frac{\mathbb{E}(S_{f1}^2)}{n_1} + \frac{\mathbb{E}(S_{f0}^2)}{n_0},$$

  where $S_{ft}^2 = \sum_{i=1}^n (Y_{fi}(t) - \overline{Y_f(t)})^2/(n-1)$, $Y_{fi}(t) = \mathbf{1}\{f(\mathbf{X}_i) = t\} Y_i(t)$, and $\overline{Y_f(t)} = \sum_{i=1}^n Y_{fi}(t)/n$ for $t = \{0, 1\}$.

Accounting for the Proportion of the Treated Units

- If the treatment is not harmful, then treating everyone is optimal
- Baseline: random (non-individualized) treatment rule
- Called "lift" in applied fields

- The Population Average Prescription Effect

$$\tau_f = \mathbb{E}\{Y_i(f(\mathbf{X}_i)) - p_f Y_i(1) - (1 - p_f) Y_i(0)\}$$

  where $p_f = \Pr(f(\mathbf{X}_i) = 1)$
- We propose an unbiased estimator of τ_f, derive its variance, and propose its consistent estimator

- Not invariant to additive transformation: $Y_i + c$
- Solution: centering $\mathbb{E}(Y_i(1) + Y_i(0)) = 0 \rightsquigarrow$ minimum variance

Estimating and Evaluating ITRs

- We may estimate and evaluate an ITR using the same experimental data
- How should we account for the estimation uncertainty as well as the evaluation uncertainty under the Neyman's framework?
- Setup:
  - Learning algorithm
  $$F : \mathcal{Z} \longrightarrow \mathcal{F}.$$
  - K-fold cross-validation
  $$\hat{f}_{-k} = F(\mathcal{Z}_{-k})$$
  - Evaluation metric estimators:
  $$\hat{\lambda}_F = \frac{1}{K}\sum_{k=1}^{K}\hat{\lambda}_{\hat{f}_{-k}}(\mathcal{Z}_k), \quad \hat{\tau}_F = \frac{1}{K}\sum_{k=1}^{K}\hat{\tau}_{\hat{f}_{-k}}(\mathcal{Z}_k)$$
- What are we estimating? What about uncertainty?

Causal Estimands

- Population Average Value (PAV)
  - Treatment assignment proportion given $\mathbf{X}_i = \mathbf{x}$

$$\bar{f}_F(\mathbf{x}) \; = \; \mathbb{E}\{\hat{f}_{\mathcal{Z}^{tr}}(\mathbf{x}) \mid \mathbf{X}_i = \mathbf{x}\} \; = \; \Pr(\hat{f}_{\mathcal{Z}^{tr}}(\mathbf{x}) = 1 \mid \mathbf{X}_i = \mathbf{x})$$

    averaging over the random sampling of training data $\mathcal{Z}^{tr}$
  - Estimand

$$\lambda_F \; = \; \mathbb{E}\left\{\bar{f}_F(\mathbf{X}_i)\, Y_i(1) + (1 - \bar{f}_F(\mathbf{X}_i))\, Y_i(0)\right\}$$

- Population Average Prescriptive Effect (PAPE)
  - Proportion treated

$$p_F \; = \; \mathbb{E}\{\bar{f}_F(\mathbf{X}_i)\}.$$

  - Estimand

$$\tau_F \; = \; \mathbb{E}\{\lambda_F - p_F\, Y_i(1) - (1 - p_F)\, Y_i(0)\}.$$

Inference under Cross-Validation

- Under Neyman's framework, the cross-validation estimators are unbiased, i.e., $\mathbb{E}(\hat{\lambda}_F) = \lambda_F$ and $\mathbb{E}(\hat{\tau}_F) = \tau_F$
- The variance of the PAV estimator

$$\mathbb{V}(\hat{\lambda}_F) = \frac{\mathbb{E}(S^2_{\hat{f}1})}{m_1} + \frac{\mathbb{E}(S^2_{\hat{f}0})}{m_0} + \underbrace{\mathbb{E}\left\{\mathrm{Cov}(\hat{f}_{\mathcal{Z}^{tr}}(\mathbf{X}_i), \hat{f}_{\mathcal{Z}^{tr}}(\mathbf{X}_j) \mid \mathbf{X}_i, \mathbf{X}_j)\tau_i\tau_j\right\}}_{estimation\ uncertainty\ of\ ITR} - \underbrace{\frac{K-1}{K}\mathbb{E}(S^2_F)}_{\substack{efficiency\ gain\ due \\ to\ cross-validation}}$$

for $i \neq j$ where m_t is the size of the training set with $T_i = t$, $\tau_i = Y_i(1) - Y_i(0)$, $S^2_F = \sum_{k=1}^K \left\{\hat{\lambda}_{\hat{f}_{-k}}(\mathcal{Z}_k) - \overline{\hat{\lambda}_{\hat{f}_{-k}}(\mathcal{Z}_k)}\right\}^2 /(K-1)$

- Estimation of the variance requires care for small K
- Analogous results for the PAPE

Evaluation with a Budget Constraint

- Policy makers often face a binding budget constraint $p < p_f$
- Scoring rule:

$$s : \mathcal{X} \longrightarrow \mathcal{S} \quad \text{where} \quad \mathcal{S} \subset \mathbb{R}$$

- (Fixed) ITR with a budget constraint:

$$f(\mathbf{X}_i, c) \;=\; \mathbf{1}\{s(\mathbf{X}_i) > c\},$$

where $c_p(f) \;=\; \inf\{c \in \mathbb{R} : \Pr(f(\mathbf{X}_i, c) = 1) \leq p\}$

- Prominent example: $s(\mathbf{x}) = \mathbb{E}(Y_i(1) - Y_i(0) \mid \mathbf{X}_i = \mathbf{x})$
- PAPE under a budget constraint

$$\tau_{fp} \;=\; \mathbb{E}\{Y_i(f(\mathbf{X}_i, c_p(f))) - pY_i(1) - (1-p)Y_i(0)\}.$$

- We derive the bias (and its finite sample bound) and variance under the Neyman's framework
- Extensions: cross-validation, diff. in PAPE between two ITRs

The Area Under Prescriptive Effect Curve

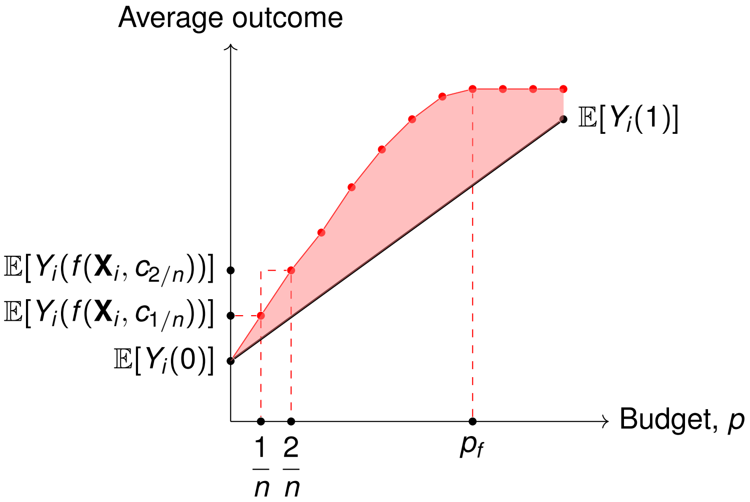

- Measure of performance across different budget constraints
- We show how to do inference with and without cross-validation

Application to the STAR Experiment

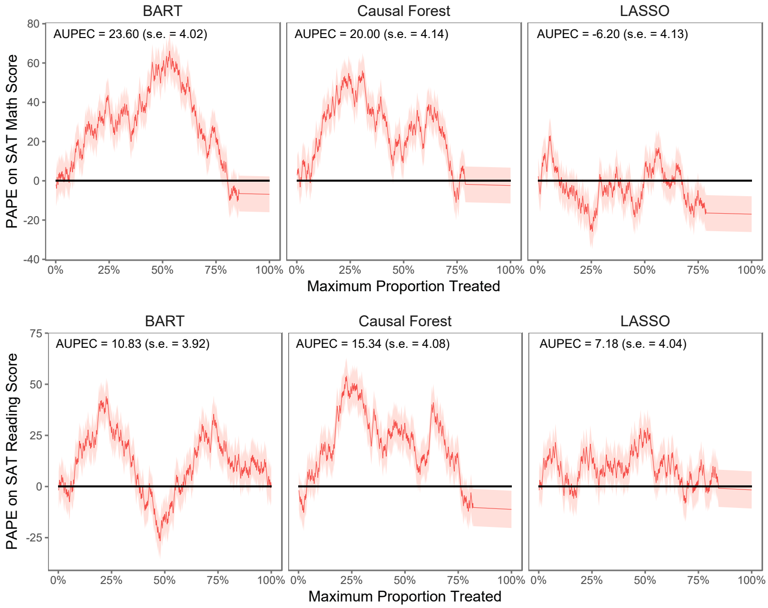

Concluding Remarks

- Individualized treatment rules (ITRs) are used in many fields
- Inference about ITRs has been largely model-based
  - We show how to experimentally evaluate ITRs
  - We incorporate budget constraints
  - No modeling assumption or asymptotic approximation is required
  - Complex machine learning algorithms can be used
  - Applicable to cross-validation estimators
  - Simulations: good small sample performance
- Paper: `https://arxiv.org/abs/1905.05389`
- Software: `https://github.com/MichaelLLi/evalITR`
- Extensions to dynamic ITRs, adaptive experiments?