

eAppendix for “Selecting optimal subgroups for treatment using many covariates”

A. Parameter Definitions

Notation

Let A denote a binary treatment of interest, Y an outcome and C a set of measured baseline covariates. Let Y_1 and Y_0 denote the potential outcomes for each individual under treatment levels 1 and 0 respectively. Let the population of individuals be denoted by Ω . We first assume treatment A is randomized and then consider treatment that may arise from an observational study. For simplicity in the next several section, in order to give intuitive proofs, we will assume a finite population of individuals with no ties at the cut-off for the optimal treatment rule. See Luedtke and van der Laan (2015, 2016ab) for further discussion of these cases without these conditions.

Context 1: Treatment Subgroup Selection Under Limited Resources

Suppose that due to limited resources we can only treat $100q\%$ of the population. We will assume that treatment is beneficial for at least $100q\%$ of the population i.e. $P(Y_1 - Y_0 > 0) > q$. Otherwise, if we know this in advance, the problem reduces to Context 2 described below. We desire to partition Ω into sets S and T such that $P(\omega \in T) = q$ so as to maximize the average outcome if all units in T were treated and all units in S were untreated. In other words, we wish to identify S and T to solve:

$$\arg \max_{S,T: S \cup T = \Omega, S \cap T = \emptyset, P(\omega \in T) = q} [qE(Y_1 | \omega \in T) + (1 - q)E(Y_0 | \omega \in S)].$$

In fact, choosing S and T to maximize the average outcome if all units in T were treated and all units in S were untreated is equivalent to choosing S and T to maximize the treatment effect heterogeneity with $100q\%$ in one group, as stated in the following proposition.

Proposition 1. *The solution to*

$$\arg \max_{S,T: S \cup T = \Omega, S \cap T = \emptyset, P(\omega \in T) = q} [qE(Y_1 | \omega \in T) + (1 - q)E(Y_0 | \omega \in S)]$$

is equivalent to the solution to

$$\arg \max_{S,T: S \cup T = \Omega, S \cap T = \emptyset, P(\omega \in T) = q} [E(Y_1 - Y_0 | \omega \in T) - E(Y_1 - Y_0 | \omega \in S)].$$

Proof. We have that

$$\begin{aligned} & \arg \max_{S,T: S \cup T = \Omega, S \cap T = \emptyset, P(\omega \in T) = q} [E(Y_1 - Y_0 | \omega \in T) - E(Y_1 - Y_0 | \omega \in S)] \\ = & \arg \max_{S,T: S \cup T = \Omega, S \cap T = \emptyset, P(\omega \in T) = q} [E(Y_1 | \omega \in T) - E(Y_1 | \omega \in S) - E(Y_0 | \omega \in T) + E(Y_0 | \omega \in S) \\ & - \frac{1}{1 - q}E(Y_1) + \frac{1}{q}E(Y_0)] \\ = & \arg \max_{S,T: S \cup T = \Omega, S \cap T = \emptyset, P(\omega \in T) = q} [E(Y_1 | \omega \in T) - E(Y_1 | \omega \in S) - E(Y_0 | \omega \in T) + E(Y_0 | \omega \in S) \\ & - \frac{q}{1 - q}E(Y_1 | \omega \in T) - \frac{1 - q}{1 - q}E(Y_1 | \omega \in S) + \frac{q}{q}E(Y_0 | \omega \in T) + \frac{(1 - q)}{q}E(Y_0 | \omega \in S)] \\ = & \arg \max_{S,T: S \cup T = \Omega, S \cap T = \emptyset, P(\omega \in T) = q} [\frac{1 - 2q}{1 - q}E(Y_1 | \omega \in T) + \frac{1 - 2q}{q}E(Y_0 | \omega \in S)] \\ = & \arg \max_{S,T: S \cup T = \Omega, S \cap T = \emptyset, P(\omega \in T) = q} [qE(Y_1 | \omega \in T) + (1 - q)E(Y_0 | \omega \in S)]. \end{aligned}$$

■

The solution to this maximization problem in fact takes a very simple form as stated in the next proposition.

Proposition 2. *The solution to*

$$\arg \max_{S,T: S \cup T = \Omega, S \cap T = \emptyset, P(\omega \in T) = q} [E(Y_1 - Y_0 | \omega \in T) - E(Y_1 - Y_0 | \omega \in S)]$$

almost surely takes the form, for some κ , of $T = \{\omega \in \Omega : Y_1(\omega) - Y_0(\omega) > \kappa\}$ and $S = \{\omega \in \Omega : Y_1(\omega) - Y_0(\omega) \leq \kappa\}$.

Proof. We prove the result by contradiction. Suppose S and T were not of this form. Then there must exist (possibly non-unique) disjoint sets $\Omega' \subseteq T$ and $\Omega^* \subseteq S$ of equal and positive probability such that $Y_1(\omega') - Y_0(\omega') < Y_1(\omega^*) - Y_0(\omega^*)$ for all $\omega' \in \Omega'$ and $\omega^* \in \Omega^*$. Let $T' = (T \cup \Omega^*) \setminus \Omega'$ and $S' = (S \cup \Omega') \setminus \Omega^*$. Then $E(Y_1 - Y_0 | \omega \in T') - E(Y_1 - Y_0 | \omega \in S') > E(Y_1 - Y_0 | \omega \in T) - E(Y_1 - Y_0 | \omega \in S)$ and thus S and T would not be the solution to $\arg \max_{S,T: S \cup T = \Omega, S \cap T = \emptyset, P(\omega \in T) = q} [E(Y_1 - Y_0 | \omega \in T) - E(Y_1 - Y_0 | \omega \in S)]$. ■

If both counterfactual outcomes were known for all individuals, then κ could be obtained as the solution to $P(Y_1 - Y_0 > \kappa) = q$ and the average outcome for the population under the optimal treatment rule of giving treatment if $Y_1 - Y_0 > \kappa$ would be $qE(Y_1 | Y_1 - Y_0 > \kappa) + (1 - q)E(Y_0 | Y_1 - Y_0 \leq \kappa)$. It is of course not possible to partition individuals in this way without complete knowledge of the counterfactual outcomes, which will in general not be available.

However, it is still possible to partition the covariate space to carry out a similar maximization. If we let Γ denote the support of C the task becomes

$$\arg \max_{S,T: S \cup T = \Gamma, S \cap T = \emptyset, P(C \in T) = q} [E(Y_1 - Y_0 | \omega \in T) - E(Y_1 - Y_0 | \omega \in S)].$$

By arguments similar to those presented above, the solution to this takes the form of $T = \{c \in \Gamma : E(Y | A = 1, c) - E(Y | A = 0, c) > k\}$ i.e. the treatment rule is then simply give treatment to those for whom $E(Y | A = 1, c) - E(Y | A = 0, c) > k$ with k given as the solution to

$$\int 1[E(Y_1 - Y_0 | C = c) > k] dP(c) = q.$$

In practice one must model $E(Y | A = 1, C = c) - E(Y | A = 0, C = c)$ and estimate k .

The average outcome under this treatment rule is then given by $q \int E(Y | A = 1, c) dP(c | C \in T) + (1 - q) \int E(Y | A = 0, c) dP(c | C \notin T)$. Note that this average outcome will not in general be as high as the average outcome under the optimal decision rule if both counterfactual outcomes were themselves known for all individuals, in which case the average outcome would be given as above, $qE(Y_1 | Y_1 - Y_0 > \kappa) + (1 - q)E(Y_0 | Y_1 - Y_0 \leq \kappa)$ where κ is defined as the solution to $P(Y_1 - Y_0 > \kappa) = q$. The extent to which the average outcome under the optimal treatment rule using the measured covariates C comes close to that which could be obtained under complete knowledge of the counterfactual outcomes will depend on the extent to which the covariates C are predictive of the outcome itself. Note that it is also thus of course the case that the average outcome under the optimal treatment rule using data on the measured covariates C will always be relative to C .

It may be of interest to compare the average outcome under this optimal treatment decision rule using the measured covariates C to the average outcome with no one treated, $E(Y | A = 0)$, the outcome with 100% treated but selected randomly $qE(Y | A = 1) + (1 - q)E(Y | A = 0)$, and the average outcome with everyone treated, $E(Y | A = 1)$. It might also be of interest to compare the treatment effect for those treated under the optimal rule, $\int \{E(Y | A = 1, c) - E(Y | A = 0, c)\} dP(c | C \in T)$, to those left untreated by the rule $\int \{E(Y | A = 1, c) - E(Y | A = 0, c)\} dP(c | C \notin T)$, and also to the average treatment effect for the population $E(Y | A = 1) - E(Y | A = 0)$. Another relevant metric may be taken as the difference

between the treatment effects comparing those who are assigned treatment by the rule versus those who are not:

$$\begin{aligned} & \int \{E(Y|A = 1, c) - E(Y|A = 0, c)\}dP(c|C \in T) \\ & - \int \{E(Y|A = 1, c) - E(Y|A = 0, c)\}dP(c|C \notin T). \end{aligned}$$

This could be taken as a measure of effect heterogeneity introduced by the optimal treatment rule for the 100q% to receive treatment. The analogous metric under treatment of a random 100q% of the population selected for treatment would simply be 0.

Context 2: Unconstrained Treatment Subgroup Selection

Now suppose that resources are unconstrained and all could be treated who benefit from treatment. The optimal subgroup for treatment would then be

$$T = \{\omega : Y_1(\omega) - Y_0(\omega) > 0\}.$$

Once again, we cannot determine this subgroup as in general we will not have information on both potential outcomes for all individuals. Instead, with covariate data on C , we would select the subgroup to treat as those with covariates C such that the expected value of the treatment effect conditional on C was positive. Thus we would treat those with covariate values that lie in the set

$$T = \{c \in \Gamma : E(Y|A = 1, c) - E(Y|A = 0, c) > 0\}.$$

In practice with high dimensional C , the expectation $E(Y|a, c)$ would have to be modeled. If we were to follow this treatment rule, the average outcome under this rule would be

$$\int_{c \in T} E(Y|A = 1, c)dP(c) + \int_{c \notin T} E(Y|A = 0, c)dP(c).$$

Note that this average outcome will not in general be as high as the average outcome under the optimal decision rule if both counterfactual outcomes were themselves known in which case the average outcome would be $E(Y_1|Y_1 - Y_0 > 0)P(Y_1 - Y_0 > 0) + E(Y_0|Y_1 - Y_0 \leq 0)P(Y_1 - Y_0 \leq 0)$. The extent to which the average outcome under the optimal treatment rule using the measured covariates C comes close to that which could be obtained under complete knowledge of the counterfactual outcomes will depend on the extent to which the covariates C are predictive of the outcome itself. It is also thus of course the case that the average outcome under the optimal treatment rule using data on the measured covariates C will always be relative to C .

It may be of interest to compare this average outcome under the optimal treatment rule using the measured covariates C to the average outcome with no one treated, $E(Y|A = 0)$, and to the average outcome with everyone treated, $E(Y|A = 1)$. It might also be of interest to compare the treatment effect for those treated under the optimal rule $\int \{E(Y|A = 1, c) - E(Y|A = 0, c)\}dP(c|C \in T)$, to those left untreated by the rule $\int \{E(Y|A = 1, c) - E(Y|A = 0, c)\}dP(c|C \notin T)$, and also to the average treatment effect for the population $E(Y|A = 1) - E(Y|A = 0)$. Another relevant metric may be taken as the difference between the treatment effects comparing those who are assigned treatment by the rule versus those who are not:

$$\int \{E(Y|A = 1, c) - E(Y|A = 0, c)\}dP(c|C \in T) - \int \{E(Y|A = 1, c) - E(Y|A = 0, c)\}dP(c|C \notin T).$$

This could be taken as a measure of effect heterogeneity.

Context 3: Unconstrained Treatment Subgroup Selection Under Costs or Side Effects

Another generalization of Setting 1 that might be considered is a setting in which there is a cost constraint. Typically this will require the investigator to only treat those with treatment effect greater than some threshold $\delta(c)$ that relies on the covariate value c . If all counterfactuals were known, we would use individual-level treatment effect $Y_1(\omega) - Y_0(\omega)$, and we will show that this can be replaced by $E(Y|A = 1, c) - E(Y|A = 0, c)$ in the realistic setting where only one counterfactual is observed for each individual.

To formalize our discussion, we wish to estimate the solution to

$$\begin{aligned} & \text{Maximize } \int_{c \in T} E(Y|A = 1, c) dP(c) + \int_{c \notin T} E(Y|A = 0, c) dP(c) \\ & \text{subject to } \int_{c \in T} \text{Cost}(c) dP(c) \leq \text{Cost Constraint}, \end{aligned}$$

where $\text{Cost}(\cdot)$ is pre-defined positive function giving the cost of treating someone in covariate strata and Cost Constraint is the pre-defined constraint. The results in this section easily generalize to the case where $\text{Cost}(c)$ can equal zero, but we omit this case for simplicity. If $b = E[\text{Cost}(C)]$ is less than Cost Constraint then the constraint is not active and we revert to Context 2. Otherwise the T maximizing the above objective takes the following form.

Proposition 3. *If b is finite and greater than Cost Constraint, then the T maximizing the objective function above takes the form, $\{\omega \in \Omega : E(Y|A = 1, c) - E(Y|A = 0, c) > k \text{Cost}(c)\}$, where k is the solution to*

$$\int 1 [E(Y_1 - Y_0|C = c) > k \text{Cost}(c)] \text{Cost}(c) dP(c) = \text{Cost Constraint}.$$

Proof. Let P_c denote the probability measure with density $\frac{dP_c}{dP}(c) = \text{Cost}(c)/b$. Let $\tilde{Y} = bY/\text{Cost}(C)$, where we note that \tilde{Y} is a deterministic function of Y conditional on $C = c$ so that $bE_P(Y|A = a, c)/\text{Cost}(c)$ is equal to $E_P(\tilde{Y}|A = a, c)$. We can write the mean outcome under the optimal rule as

$$\int_{c \in T} E(\tilde{Y}|A = 1, c) dP_c(c) + \int_{c \notin T} E(\tilde{Y}|A = 0, c) dP_c(c).$$

The cost constraint rewrites as $P_c(C \in T) \leq (\text{Cost Constraint})/b$. We are now in Context 1, so maximizing T takes the form $\{\omega \in \Omega : E(\tilde{Y}|A = 1, c) - E(\tilde{Y}|A = 0, c) > \tilde{k}\}$, where \tilde{k} is the solution to

$$\int 1 [E(\tilde{Y}|A = 1, c) - E(\tilde{Y}|A = 0, c) > \tilde{k}] dP_c(c) = (\text{Cost Constraint})/b$$

Multiplying both sides by b , using the definition of \tilde{Y} , and letting $k = \tilde{k}/b$ gives the result. ■

It may be of interest to compare this average outcome under the optimal treatment rule to the average outcome with no one treated, $E(Y|A = 0)$, and to the average outcome with everyone treated, $E(Y|A = 1)$. It might also be of interest to compare the treatment effect for those treated under the optimal rule, $\int \{E(Y|A = 1, c) - E(Y|A = 0, c)\} dP(c|C \in T)$, to those left untreated by the rule $\int \{E(Y|A = 1, c) - E(Y|A = 0, c)\} dP(c|C \notin T)$, and also to the average treatment effect for the population $E(Y|A = 1) - E(Y|A = 0)$.

Context 4: Treatment Subgroup Selection to Maximize Treatment Effect Heterogeneity

In examining subgroup analyses reported in the literature one is sometimes under the impression that a central goal is to find a subgroup division that maximizes effect heterogeneity across the subgroups.

While it is not clear that this goal is of principal policy importance, it can be carried out using a set of measured covariates C . The task could then be stated as finding a partition of individuals into sets T and S to maximize treatment effect heterogeneity when comparing the treatment effects among those in T versus S . The problem could thus formally be stated as

$$\arg \max_{S,T: S \cup T = \Omega, S \cap T = \emptyset} [E(Y_1 - Y_0 | \omega \in T) - E(Y_1 - Y_0 | \omega \in S)].$$

The solution to this maximization problems takes a relatively simple form as in the next Proposition.

Proposition 4. *The solution to*

$$\arg \max_{S,T: S \cup T = \Omega, S \cap T = \emptyset} [E(Y_1 - Y_0 | \omega \in T) - E(Y_1 - Y_0 | \omega \in S)]$$

takes the form, for some κ , of $T = \{\omega \in \Omega : Y_1(\omega) - Y_0(\omega) > \kappa\}$ and $S = \{\omega \in \Omega : Y_1(\omega) - Y_0(\omega) \leq \kappa\}$.

Proof. The proof is similar to that of Proposition 2 above. ■

If we let $V = Y_1 - Y_0$ and let $p(v)$ denote the density of V . To determine κ , we wish to choose κ to maximize $E(Y_1 - Y_0 | Y_1 - Y_0 > \kappa) - E(Y_1 - Y_0 | Y_1 - Y_0 \leq \kappa)$ or equivalently,

$$\int_{\kappa}^{\infty} vp(v)dv / \int_{\kappa}^{\infty} p(v)dv - \int_{-\infty}^{\kappa} vp(v)dv / \int_{-\infty}^{\kappa} p(v)dv.$$

Differentiating with respect to κ we obtain:

$$\frac{\int_{\kappa}^{\infty} p(v)dv [vp(v)]_{v=\kappa}^{v=\infty} - \int_{\kappa}^{\infty} vp(v)dv [p(v)]_{v=\kappa}^{v=\infty}}{[\int_{\kappa}^{\infty} p(v)dv]^2} - \frac{\int_{-\infty}^{\kappa} p(v)dv [vp(v)]_{v=-\infty}^{v=\kappa} - \int_{-\infty}^{\kappa} vp(v)dv [p(v)]_{v=-\infty}^{v=\kappa}}{[\int_{-\infty}^{\kappa} p(v)dv]^2}$$

Setting this equal to 0 and solving for κ gives

$$\begin{aligned} & [\int_{-\infty}^{\kappa} p(v)dv]^2 \{ \int_{\kappa}^{\infty} p(v)dv [vp(v)]_{v=\kappa}^{v=\infty} - \int_{\kappa}^{\infty} vp(v)dv [p(v)]_{v=\kappa}^{v=\infty} \} - [\int_{\kappa}^{\infty} p(v)dv]^2 \{ \int_{-\infty}^{\kappa} p(v)dv [vp(v)]_{v=-\infty}^{v=\kappa} \\ & \quad - \int_{-\infty}^{\kappa} vp(v)dv [p(v)]_{v=-\infty}^{v=\kappa} \} = 0 \\ & P(V \leq \kappa)^2 \{ -\kappa p(\kappa) P(V > \kappa) + p(\kappa) \int_{\kappa}^{\infty} vp(v)dv \} + P(V > \kappa)^2 \{ -\kappa p(\kappa) P(V \leq \kappa) + p(\kappa) \int_{-\infty}^{\kappa} vp(v)dv \} = 0 \\ & P(V \leq \kappa)^2 \{ -\kappa P(V > \kappa) + \int_{\kappa}^{\infty} vp(v)dv \} + P(V > \kappa)^2 \{ -\kappa P(V \leq \kappa) + \int_{-\infty}^{\kappa} vp(v)dv \} = 0. \end{aligned}$$

The solution to the equation gives κ which could be obtained numerically.

The treatment rule that maximizes effect heterogeneity is then either $T = \{\omega \in \Omega : Y_1(\omega) - Y_0(\omega) \geq \kappa\}$ or $T = \{\omega \in \Omega : Y_1(\omega) - Y_0(\omega) > \kappa\}$. The maximum treatment effect heterogeneity that can thus be obtained comparing two subgroups that partition all individuals is thus $E(Y_1 - Y_0 | Y_1 - Y_0 > \kappa) - E(Y_1 - Y_0 | Y_1 - Y_0 \leq \kappa)$.

It is of course not possible to partition individuals in this manner without complete knowledge of the counterfactual outcomes. However, it is still possible to carry out a similar partitioning using measured covariates C . If we let Γ denote the support of C , the task becomes

$$\arg \max_{S,T \subseteq \Gamma: S \cup T = \Gamma, S \cap T = \emptyset} [E(Y_1 - Y_0 | C \in T) - E(Y_1 - Y_0 | C \in S)].$$

By arguments similar to those presented above, the sets S and T are then given by $T = \{c \in \Gamma : E(Y | A = 1, c) - E(Y | A = 0, c) > k\}$ and $S = \{c \in \Gamma : E(Y | A = 1, c) - E(Y | A = 0, c) \leq k\}$ for some k , and once again k could be solved for numerically.

The effect heterogeneity under this treatment rule for maximizing effect heterogeneity with measured covariates C is then given by:

$$\int \{E(Y|A = 1, c) - E(Y|A = 0, c)\}dP(c|C \in T) - \int \{E(Y|A = 1, c) - E(Y|A = 0, c)\}dP(c|C \notin T).$$

Note that this measure of effect heterogeneity will not in general be as high as the maximum effect heterogeneity if both counterfactual outcomes for all individuals were known, which was given above as $E(Y_1 - Y_0|Y_1 - Y_0 \geq \kappa) - E(Y_1 - Y_0|Y_1 - Y_0 < \kappa)$. The extent to which the maximum effect heterogeneity under the treatment rule using the measured covariates C comes close to the maximum which can be obtained under complete knowledge of the counterfactual outcomes will depend on the extent to which the covariates C are predictive of the outcome itself. Note that it is also thus of course the case that the maximum effect heterogeneity under the treatment rule using data on the measured covariates C will always be relative to C . It might also be of interest to compare the treatment effect for those treated under the maximizing rule, $\int \{E(Y|A = 1, c) - E(Y|A = 0, c)\}dP(c|C \in T)$, to those left untreated by the rule $\int \{E(Y|A = 1, c) - E(Y|A = 0, c)\}dP(c|C \notin T)$, and also to the average treatment effect for the population $E(Y|A = 1) - E(Y|A = 0)$.

If using the treatment rule that maximizes effect heterogeneity, one can also estimate the average outcome under this rule as

$$\int_{c \in T} E(Y|A = 1, c)dP(c) + \int_{c \notin T} E(Y|A = 0, c)dP(c).$$

One could compare this to the average outcome with no one treated, $E(Y|A = 0)$, and the average outcome with everyone treated, $E(Y|A = 1)$. Note, however, this average outcome under the treatment rule that maximizes effect heterogeneity will in general be lower than the average outcome under the treatment rule that maximizes the average outcome itself as discussed in Context 2. It is thus not clear that this treatment rule that maximizes effect heterogeneity is of particular use in policy-making, unlike contexts 1, 2 and 3 above.

Further Comments on Observational Studies

The above approaches and results for randomized treatment A apply also to observational studies in which the covariates C suffice to control for confounding of the effect of A on Y but, when reference is made to the average outcome when no one is treated, $E(Y|A = 0)$ must be replaced by $\int E(Y|A = 0, c)dP(c)$; when reference is made to the average outcome when everyone is treated, $E(Y|A = 1)$ must be replaced by $\int E(Y|A = 1, c)dP(c)$; and when reference is made to the average treatment effect for the population $E(Y|A = 1) - E(Y|A = 0)$ must be replaced by $\int E(Y|A = 1, c)dP(c) - \int E(Y|A = 0, c)dP(c)$.

B. Statistical Estimation

Super-Learner for the Conditional Average Treatment Effect

The super-learner algorithm for this estimation problem is given below. For simplicity we use ten-fold cross-validation and assume that the sample size n is a multiple of ten. We suppose that the user has m candidate regression algorithms.

1. Define pseudo-observations $\tilde{Y}_1, \dots, \tilde{Y}_n$.
2. For v in $\{1, \dots, 10\}$:

- (a) For ℓ in $\{1, \dots, m\}$:
 - Fit candidate algorithm ℓ on observations $\left((C_i, \tilde{Y}_i) : i \notin \left\{\frac{(v-1)n}{10} + 1, \dots, \frac{vn}{10}\right\}\right)$ to generate the regression function estimate that takes as input c and outputs $\hat{E}_v^\ell[\tilde{Y}|c]$.
- (b) For i in $\left\{\frac{(v-1)n}{10} + 1, \dots, \frac{vn}{10}\right\}$, define X_i^{CV} to be the m -length column vector $\left(\hat{E}_v^\ell[\tilde{Y}|c_i] : \ell = 1, \dots, m\right)$.
3. Choose α_n be the convex m -length row vector minimizing $\frac{1}{n} \sum_{i=1}^n \left(\tilde{Y}_i - \alpha_n X_i^{\text{CV}}\right)^2$.
4. For ℓ in $\{1, \dots, m\}$:
 - Fit the candidates ℓ on all of the observations, yielding a function which takes as input c and outputs $\hat{E}^\ell[\tilde{Y}|c]$.
5. Return the function \hat{b} defined by $\hat{b}(c) = \alpha_n X(c)$, where $X(c)$ is the column vector $\left(\hat{E}^\ell[\tilde{Y}|c] : \ell = 1, \dots, m\right)$.

Estimator for Contexts 1 and 2

We now present a cross-validated targeted minimum loss-based estimator (CV-TMLE) for Contexts 1 and 2. We omit the derivation of this estimator from this work for brevity, and instead refer the reader to references given in the main text for the derivation of CV-TMLEs for a nearly identical parameters as in Contexts 1 and 2. We assume that the outcome Y is bounded, and without loss of generality we assume that the lower bound is 0 and the upper bound is 1.

For simplicity, we use ten-fold cross-validation and assume that the sample size n is a multiple of ten. For each $v = 1, \dots, 10$, we let T_v denote the observations (C_i, A_i, Y_i) with indices $i \notin \left\{\frac{(v-1)n}{10} + 1, \dots, \frac{vn}{10}\right\}$, i.e. the observations included in the training sample for a given cross-validation split. We also let $v(i)$ denote the v such that $i \notin \left\{\frac{(v-1)n}{10} + 1, \dots, \frac{vn}{10}\right\}$, i.e. the cross-validation index for which i is not in the training sample.

1. For v in $\{1, \dots, 10\}$:
 - (a) Obtain an estimate $\hat{E}_v[Y|\cdot]$ of the function mapping from (a, c) to $E[Y|a, c]$ using observations in T_v .
 - (b) Obtain an estimate \hat{b}_v of the CATE function using observations in T_v .
Note: $\hat{b}_v(c)$ need not equal $\hat{E}[Y|A = 1, c] - \hat{E}[Y|A = 0, c]$.
2. Let $\delta_n = 0$ in Context 2, and in Context 1 let δ_n be the positive part of the smallest solution in δ to

$$\frac{1}{n} \sum_{i=1}^n 1(\hat{b}_{v(i)}(c) > \delta) \leq q.$$

3. Let ϵ_n be the slope estimate in an intercept-free logistic regression with outcome $(Y_i : i = 1, \dots, n)$, covariate $\left(\frac{2A_i - 1}{P(A_i|C_i)} : i = 1, \dots, n\right)$, offset $\left(\text{logit } \hat{E}_{v(i)}[Y|A_i, C_i] : i = 1, \dots, n\right)$, and observation weights $\left(1 \left\{\hat{b}_{v(i)}(C_i) > \delta_n\right\} : i = 1, \dots, n\right)$.
Note: The `glm` function in R will run a logistic regression on any outcome bounded in $[0, 1]$.
4. For each v in $\{1, \dots, 10\}$ and any (a, c) , let $\bar{Q}_v^*(a, c) = \text{logit}^{-1} \left[\text{logit} \left(\hat{E}_v[Y|a, c] \right) + \epsilon_n \frac{2a-1}{P(a|c)} \right]$ represent the fluctuated estimate of $E[Y|a, c]$.
5. Let $\hat{\psi} = \frac{1}{n} \sum_{i=1}^n 1(\hat{b}_{v(i)}(C_i) > \delta_n) \left[\bar{Q}_{v(i)}^*(1, C_i) - \bar{Q}_{v(i)}^*(0, C_i) \right]$.

6. Let $D_i = \frac{2A_i-1}{P(A_i|C_i)} \left[Y_i - \bar{Q}_{v(i)}^*(A_i, C_i) \right] + \bar{Q}_{v(i)}^*(1, C_i) - \bar{Q}_{v(i)}^*(0, C_i) - \hat{\psi}$.
7. Let

$$\hat{\sigma}^2 = \frac{1}{n} \sum_{i=1}^n \left[1(\hat{b}_{v(i)}(C_i) > \delta_n) [D_i - \delta_n] + \delta_n q \right]^2,$$

where we take $q = 1$ in Context 2.

8. Return the point estimate $\hat{\psi}$ and 95% confidence interval $\hat{\psi} \pm 1.96 \frac{\hat{\sigma}}{\sqrt{n}}$.

C. Software Demonstration

In this section we present software for estimating the optimal treatment rule and for obtaining inference for the impact of implementing this rule in the population. This appendix uses functions that appear in version 0.31 of the `sg` package, which is available at <https://github.com/alexluetke12/sg@v0.31>. The most recent version can be found at <https://github.com/alexluetke12/sg>.

We first install the `sg` package from GitHub using the `devtools` package, and subsequently load the package.

```
library(devtools)
devtools::install_github("alexluetke12/sg@v0.31")
```

```
library(sg)
```

The remainder of this demonstration is structured as follows. First, we show how to use the `sg` package to estimate the CATE function. Then, we show how to evaluate the impact of treating an optimal treatment rule, both in the presence and in the absence of a resource constraint.

Estimating the CATE Function

This demonstration will use the following simple simulated data set.

```
# set seed for random number generator
set.seed(1)
# function for simulating data
sim.data = function(nn){
  # covariates
  W = data.frame(W1=rnorm(nn),W2=rnorm(nn),W3=rnorm(nn))
  # treatment indicator
  A = rbinom(nn,1,1/2)
  # binary outcome
  Y = rbinom(nn,1,Qbar(A,W))
  return(list(W=W,A=A,Y=Y))
}

# SuperLearner library
# Note: a restricted library is used in this example to reduce
#       runtime. In practice we recommend using a larger library.
#       Use command 'listWrappers()' to see other prediction algorithm
#       wrappers.
SL.library = c('SL.mean','SL.glm','SL.gam')

# sample size
```



```

n = 1000
# E[Y|A,W]
Qbar = function(a,w){plogis(-1 - a*w$W1^2 + 2*a*w$W2)}

# data to use when running the methods
dat = sim.data(n)
W = dat$W
A = dat$A
Y = dat$Y

# data to use when evaluating true parameter values
n.mc = 2e4
dat = sim.data(n.mc)
W.mc = dat$W
A.mc = dat$A
Y.mc = dat$Y
# conditional average treatment effect function
cate.mc = Qbar(1,W.mc)-Qbar(0,W.mc)

```

We note that $A = 1$ and $A = 0$ can denote two different active treatments, or treatment and control.

The CATE function can be estimated using the `sg.SL` wrapper for the ensemble `SuperLearner` algorithm [Polley et al., 2018].

```

sl.out = sg.SL(W,A,Y,SL.library=SL.library,family=binomial())

# compare estimated CATE to true CATE in data set
plot(Qbar(1,W)-Qbar(0,W),sl.out$est,xlab="True CATE",ylab="Estimated CATE")

```

The output of the above plot is displayed in Figure 1.

The impact of different variables on the CATE function can be evaluated graphically. Here, we set W_2 and W_3 to their sample medians and plot the estimated CATE against W_1 .

```

# W1 values at which to evaluate the CATE
W1.vals = quantile(W$W1,seq(0.001,0.95,by=0.05))
# estimated CATE with W2 and W3 set to their medians
cate.W1 = predict(
  sl.out$SL,
  newdata=data.frame(
    W1=W1.vals,
    W2=median(W$W2),
    W3=median(W$W3)))$pred
# plot estimated CATE against W2
plot(W1.vals,cate.W1,xlab="W1",ylab="CATE",type='l')

```

The output of the above plot is displayed in Figure 2.

Evaluating the Impact of Implementing the Optimal Treatment Rule

We now evaluate the impact of treating the optimal rule in the presence and the absence of resource constraints. Here we use the same simulated data set as was used to estimate the CATE function.

We start with a setting where at most 25% of the population can be treated. This function gives an additive contrast between the mean outcome under the estimated optimal treatment rule and the mean outcome under a stochastic baseline rule that treats all individuals in the population at random, regardless of the covariates. The probability of receiving (treatment $A = 0$, treatment $A = 1$) according




Figure 1: Estimated versus true CATE in the software demonstration.

to this stochastic baseline rule is dictated by the `baseline.probs` input argument. Larger values of `num.SL.rep` and `num.est.rep` reduce stochasticity of output due to cross-validation splits but increase runtime.

```
sg.cvtmle(W,A,Y,
  baseline.probs=c(0.75,0.25), # baseline assigns A=1 with probability 25%
  SL.library=SL.library,
  family=binomial(),
  kappa=0.25, # the optimal rule may treat at most 25% of the population
  num.SL.rep=2,
  num.est.rep=2,
  lib.ests=TRUE)

# R Console Output
# $est
# SuperLearner  SL.mean_All  SL.glm_All  SL.gam_All
# 0.107210896  0.008358105  0.098576837  0.107210896
#
# $ci
#               lb          ub
# SuperLearner  0.083688690  0.13073310
# SL.mean_All  -0.004555034  0.02127124
# SL.glm_All    0.075308929  0.12184474
# SL.gam_All    0.083688690  0.13073310
#
# $est.mat
#           SuperLearner SL.mean_All SL.glm_All SL.gam_All
```




Figure 2: CATE vs. W_1 in the software demonstration, where W_2 and W_3 are set to the sample medians.

```
# Repetition 1    0.1048418 0.008381272 0.09828244 0.1048418
# Repetition 2    0.1095800 0.008334938 0.09887123 0.1095800
```

The output of `sg.cvtmle` is a three-argument list. The `est` argument displays the estimated impact of implementing the treatment strategy estimated by SuperLearner, and also the impact of implementing treatment rules estimated by the candidate algorithms included in the SuperLearner library. The `ci` argument contains the corresponding 95% level confidence intervals. The `est.mat` argument shows the separate estimates obtained during the `num.est.rep` repetitions of the estimation procedure. The final estimate is equal to the mean of these repetition-specific estimates.

The estimates presented above are close to the true improvement under the resource-constrained optimal rule, which can be evaluated in our simulation because we know the true data generating distribution.

```
# truth
mean(
  Qbar(1,W.mc)*((cate.mc>=max(quantile(cate.mc,1-0.25),0))-0.25) +
  Qbar(0,W.mc)*((cate.mc<max(quantile(cate.mc,1-0.25),0))-0.75))
# R Console Output
# [1] 0.105269
```

Note that the `sg.cvtmle` function does not currently support individual-level treatment costs, i.e. it only allows estimating optimal treatment effects when the cost of assigning treatment is fixed across all individuals in the population. In these cases, when $A = 1$ is more expensive than $A = 0$ a cost constraint is equivalent to constraining the proportion of individuals in the population who can be treated.

Below we remove the resource constraint and contrast against a baseline treating individuals with $A = 1$ with probability $1/2$.

```
sg.cvtmle(W,A,Y,
  baseline.probs=c(0.50,0.50), # baseline assigns A=1 with probability 50%
```

```

      SL.library=SL.library,
      family=binomial(),
      kappa=1, # no resource constraint
      num.SL.rep=2,
      num.est.rep=2,
      lib.ests=TRUE)
# R Console Output
# $est
# SuperLearner  SL.mean_All  SL.glm_All  SL.gam_All
# 0.11078319    0.01672297   0.10478232   0.11078319
#
# $ci
#                lb                ub
# SuperLearner  0.085855224 0.13571116
# SL.mean_All  -0.009112482 0.04255842
# SL.glm_All    0.079753790 0.12981086
# SL.gam_All    0.085855224 0.13571116
#
# $est.mat
#           SuperLearner SL.mean_All SL.glm_All SL.gam_All
# Repetition 1    0.1104028 0.01710950 0.1036115 0.1104028
# Repetition 2    0.1111636 0.01633644 0.1059532 0.1111636

```

To contrast against a baseline rule that treats everyone with $A = 0$, let `baseline.probs=c(1,0)`. To estimate the marginal mean outcome under the estimated optimal treatment rule without contrasting against a baseline treatment rule, let `baseline.probs=c(0,0)`.

The above estimate is close to the true improvement under the unconstrained optimal rule.

```

# truth
mean(Qbar(1,W.mc)*((cate.mc>=0)-0.50) + Qbar(0,W.mc)*((cate.mc<0)-0.50))
# R Console Output
# [1] 0.1162848

```

References

Eric Polley, Erin LeDell, Chris Kennedy, and Mark van der Laan. *SuperLearner: Super Learner Prediction*, 2018. URL <https://CRAN.R-project.org/package=SuperLearner>. R package version 2.0-24.