



HAL
open science

Heterogeneous Treatment Effects Estimation: When Machine Learning meets multiple treatments regime

Naoufal Acharki, Antoine Bertonecello, Josselin Garnier, Ramiro Lugo

► **To cite this version:**

Naoufal Acharki, Antoine Bertonecello, Josselin Garnier, Ramiro Lugo. Heterogeneous Treatment Effects Estimation: When Machine Learning meets multiple treatments regime. 53èmes Journées de Statistique, Jun 2022, Lyon, France. hal-03698547

HAL Id: hal-03698547

<https://hal.science/hal-03698547v1>

Submitted on 18 Jun 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

HETEROGENEOUS TREATMENT EFFECTS ESTIMATION: WHEN MACHINE LEARNING MEETS MULTIPLE TREATMENTS REGIME

Naoufal Acharki ^{1,2} & Antoine Bertinello ² & Josselin Garnier ¹ & Ramiro Lugo ²

¹ *Centre de Mathématiques Appliquées CMAP, Ecole Polytechnique - Institut Polytechnique de Paris, France; FirstName.LastName@polytechnique.edu*

² *TotalEnergies One Tech, Palaiseau, France; FirstName.LastName@totalenergies.com*

Résumé. Dans beaucoup de domaines scientifiques et techniques, l’inférence de l’effet d’un traitement et l’exploration de son hétérogénéité sont déterminantes pour l’optimisation et la prise de décision. De nombreux méta-algorithmes ont été développés pour estimer la fonction d’effet moyen conditionnel du traitement (CATE) pour un traitement binaire, leur principal avantage est de ne pas restreindre l’estimation à une méthode d’apprentissage supervisée particulière. Dans ce travail, nous étudions le régime de traitement multiple sous le modèle causal de Rubin et nous nous focalisons sur l’estimation des effets hétérogènes de traitement. Nous généralisons les meta-algorithmes pour l’estimation de la fonction CATE pour chaque niveau possible de traitement. Nous évaluons la qualité de chaque méta-algorithme sur des données observationnelles en utilisant un jeu de données Semi-synthétique et nous soulignons en particulier les performances du X-learner.

Mots-clés. Apprentissage Automatique; Inférence Causale; Traitement multiple; Effets hétérogènes.

Abstract. In many scientific and engineering domains, inferring the effect of treatment and exploring its heterogeneity is crucial for optimization and decision making. Several meta-algorithms have been developed to estimate the Conditional Average Treatment Effect (CATE) function in the binary setting, with the main advantage of not restraining the estimation to a specific supervised learning method. In this work, we investigate the multiple treatment regime under Rubin Causal Model and we focus on estimating heterogeneous treatment effects. We generalize Meta-learning algorithms to estimate the CATE for each treatment value. Using semi-synthetic simulation datasets, we assess the quality of each meta-learner in observational data and we highlight in particular the performances of the X-learner.

Keywords. Machine Learning; Causal Inference; Multiple Treatments; Heterogeneous Effects.

1 Introduction.

With the rapid development of Machine Learning (ML) and its efficiency in predicting outcomes, the question of counterfactual prediction “*what would happen if ?*” arises.

Engineers may want to know how the outcome would be affected when a feature is changed to a specific value, not only on average but also within a smaller scale, to personalize treatments at efficient levels and optimize the outcome.

Based on the Potential Outcomes theory Neyman (1923); Rubin (1974), epidemiologists and statisticians developed a set of statistical tools to make causal inference and estimate the effects of a treatment on the outcome, whether on average among the whole population or inside different sub-groups. They have been successfully applied in many fields, but little is known about their efficiency in industrial applications. Furthermore, most existing methods and studies are limited to the setting of a binary treatment, whereas in many real-world applications, the treatment variable can take multiple values. In addition, the heterogeneity of effects may provide valuable information regarding the effectiveness of this treatment. Finally, the ground truth of treatment effects cannot be observed. This is known as the Holland (1986) fundamental problem of causal inference.

2 Problem setting.

Following the potential outcomes framework and the generalization of the Rubin (1974) Causal Model, we suppose the existence of $Y(t)$, the real-valued counterfactual outcome that would have been observed under treatment level $t \in \mathcal{T} = \{t_1, \dots, t_K\}$. We consider $(\mathbf{X}, T, Y(t)_{t \in \mathcal{T}}) \sim \mathbb{P}$ where $\mathbf{X} = (X^{(1)}, \dots, X^{(d)}) \in \mathbb{R}^d$ denotes a random vector of covariates and T denotes the treatment assignment random variable. We suppose finally that we observe data drawn from independent and identically distributed sample of n units $D_{\text{obs},i} = (\mathbf{X}_i, T_i, Y_{\text{obs},i})$ distributed as $(\mathbf{X}, T, Y_{\text{obs}})$ with $Y_{\text{obs}} = Y(T)$. All assumptions (Consistency, Unconfoundedness and Overlap) of this model remain valid.

We aim to infer the effect of the treatment T on the outcome Y . More precisely, we want to estimate the Conditional Average Treatment Effect (CATE), defined as

$$\tau_t(\mathbf{x}) = \mathbb{E}(Y(t) - Y(t_1) | \mathbf{X} = \mathbf{x}), \quad (1)$$

which can be interpreted as is the expected treatment effect between levels $T = t_1$ (defined as baseline treatment value) and t for an individual with covariates $\mathbf{X} = \mathbf{x}$.

3 Proposed model.

To tackle the problem of estimating CATEs in multiple treatment regime, we generalize the notion of meta-learners as initially developed by Künzel et al. (2019) to derive consistent estimators of the CATE. All considered meta-learners below, except the R-learner, can support any supervised regression ML method (e.g. random forest, gradient boosting methods). These ML methods are called *base-learners* when applied to a meta-learner.

3.1 Direct plug-in meta-learners

Direct plug-in meta-learners, also known as *one-step* learners, estimate the CATE in (1) by targeting directly the observed data $\mathcal{D}_{\text{obs}} = (D_{\text{obs},i})_{i=1}^n$.

The **T-learner** builds a CATE estimator using *two* models $\mu_w(\mathbf{x}) = \mathbb{E}(Y(w)|\mathbf{X} = \mathbf{x})$ on $(D_{\text{obs},i})_{i \in \mathbf{S}_w}$ with $\mathbf{S}_w = \{i, T_i = w\}$ for $w \in \{t, t_1\}$ and compute the CATE as $\widehat{\tau}_t^{(T)}(\mathbf{x}) = \widehat{\mu}_t(\mathbf{x}) - \widehat{\mu}_{t_1}(\mathbf{x})$.

The **S-learner** takes the treatment T as a feature similar to all the other covariates, estimates $\mu(\mathbf{x}, t) = \mathbb{E}(Y_{\text{obs}}|\mathbf{X} = \mathbf{x}, T = t)$ by *single* model and computes CATEs as $\widehat{\tau}_t^{(S)}(\mathbf{x}) = \widehat{\mu}(\mathbf{x}, t) - \widehat{\mu}(\mathbf{x}, t_1)$.

Despite their simplicity, direct plug-in meta-learners may be sensitive to the base learner and suffer from regularization bias. We denote **RegT-learner** the regularized **T-learner** using Importance Sampling.

3.2 Pseudo-outcome meta-learners

An alternative possibility for mitigating this bias is to consider a specific representation of the observed outcome Y_{obs} , called *pseudo-outcome*. These representations incorporate *nuisance components* that include valuable information such as the dependence between covariates \mathbf{X} and T (i.e. the Generalized Propensity Score $r(t, \mathbf{x}) := \mathbb{P}(T = t | \mathbf{X} = \mathbf{x})$).

In the **M-Learner**, we estimate the GPS r then build a CATE estimator $\widehat{\tau}_t^{(M)}$ by regressing the Inverse Propensity Weighting pseudo-outcome Z_t^M such that:

$$Z_t^M = \frac{\mathbf{1}\{T = t\}}{r(t, \mathbf{X})} Y_{\text{obs}} - \frac{\mathbf{1}\{T = t_1\}}{r(t_1, \mathbf{X})} Y_{\text{obs}}. \quad (2)$$

In the **DR-Learner**, we estimate the GPS r and the outcome models μ_t , then we build a CATE estimator $\widehat{\tau}_t^{(DR)}$ by regressing the Doubly-Robust pseudo-outcome Z_t^{DR} such that:

$$Z_{\mu, r, t}^{DR} = \frac{Y_{\text{obs}} - \mu_T(\mathbf{X})}{r(t, \mathbf{X})} \mathbf{1}\{T = t\} - \frac{Y_{\text{obs}} - \mu_T(\mathbf{X})}{r(t_1, \mathbf{X})} \mathbf{1}\{T = t_1\} + \mu_t(\mathbf{X}) - \mu_{t_1}(\mathbf{X}). \quad (3)$$

In the **X-Learner**, we estimate all outcome models μ_t , then we build a CATE estimator $\widehat{\tau}_t^{(X)}$ by regressing the Doubly-Robust pseudo-outcome Z_t^X such that:

$$Z_t^X = \mathbf{1}\{T = t\}(Y_{\text{obs}} - \mu_{t_1}(\mathbf{X})) + \sum_{t' \neq t} \mathbf{1}\{T = t'\} \times (\mu_t(\mathbf{X}) - Y_{\text{obs}}) + \sum_{t' \neq t} \mathbf{1}\{T = t'\}(\mu_{t'}(\mathbf{X}) - \mu_{t_1}(\mathbf{X})). \quad (4)$$

3.3 Neyman-Orthogonality meta-learners

The **R-learner** uses the Robinson (1988) decomposition and the Neyman-Orthogonality propriety to provide flexible CATEs estimators. In multiple treatments regime, the R-Learning method estimates all $K - 1$ models $\{\tau_{t,t_1}\}_{t \neq t_1 \in \mathcal{T}}$ within the space \mathcal{F} of candidate models such that

$$\begin{aligned} \{\hat{\tau}_t^{(R)}\}_{t \neq t_1 \in \mathcal{T}} = \arg \min_{\{\tau_t\}_{t \neq t_1} \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n & \left[(Y_{\text{obs},i} - \hat{m}(\mathbf{X}_i)) - \right. \\ & \left. \sum_{t \neq t_1 \in \mathcal{T}} (\mathbf{1}\{T_i = t\} - \hat{r}(t, \mathbf{X}_i)) \tau_t(\mathbf{X}_i) \right]^2, \end{aligned} \quad (5)$$

where \hat{m} (respectively, \hat{r}) is an estimator of $m = \mathbb{E}(Y_{\text{obs}} | \mathbf{X})$ (respectively, the GPS r) and \mathcal{F} is the space of candidate models $[\{\tau_t\}_{t \neq t_1}]$. When \mathcal{F} is the family of linear regression models, then Problem (5) admits a solution that we denote by **Rlin-learner**.

4 Semi-synthetic dataset for validating causal inference method.

Motivation and description. The difficulty in evaluating a causal model’s performance in real-world applications motivates the need to create a semi-synthetic dataset. In this section, we consider a multistage fracturing Enhanced Geothermal System (EGS), Han et al. (2020). We assume that the heat extraction performance satisfies the following physical model $Q_{\text{well}}(\ell_L) = Q_{\text{fracture}} \times \ell_L/d \times \eta_d$ where Q_{fracture} is the *unknown* heat extraction performance from a single fracture, that can be generated using eight parameters including reservoir characteristics and fracture design, ℓ_L is the Lateral Length of the well, d is the average spacing between two fractures and η_d is the stage efficiency penalizing the individual contribution when fractures are close to each other. This model respects the unconfoundedness assumption, and we can control all its variables in the simulations.

A randomized series of numerical experiments using a numerical emulator has been conducted to simulate the heat performance from a single fracture (i.e. Q_{fracture}), leading to an initial full factorial design of experiments dataset covering all possible scenarios of a fracture in a reservoir. The final dataset is created by extrapolating the heat performance of each case to all well’s lateral lengths, fracture spacing and the efficiency coefficient.

Estimating Heterogeneous Treatment Effects on a non-randomized biased dataset. We consider the Lateral Length as treatment T with $K = 13$ possible values and the covariates $\mathbf{X} \in \mathbb{R}^{11}$ are the remaining variables. Following the *preferential selection*, we sample $n = 10000$ units such that wells with high lateral length are likely to have larger fractures and vice versa. The GPS is estimated using gradient boosting models and the




Figure 1: CATEs estimation on semi-synthetic dataset. Each line represents τ_j for $j = 1, \dots, K$. (a): The ground truth model; (b): A biased estimation of CATEs by regressing on Fracture_length_ft; (c): T-learner estimation; (d): X-learner estimation.

outcome models μ_w are estimated by the T-learning approach for RandomForest and the S-learning approach for XGBoost model. Table 1 resumes $\mathbf{mPEHE} = \frac{1}{K-1} \sum_{t \neq t_1} PEHE(\hat{\tau}_t)$ for different meta-learners where $PEHE(\hat{\tau}_t) = \sqrt{\frac{1}{n} \sum_{i=1}^n (\hat{\tau}_t(\mathbf{X}_i) - \tau_t(\mathbf{X}_i))^2}$ is the Precision in Estimation of Heterogeneous Effect (PEHE) as defined by Shalit et al. (2017). Table 1 shows that the XGBoost model is generally a better choice than Random Forests (except for T-learning); The X-learner, followed by DR-learner, outperforms all other learners. Finally, Figure 1 shows the ground truth model, what would one obtain by regressing only on fracture length (correlation) and T-, X-learner’s estimation. It demonstrates the ability of meta-learners, in particular the X-learner, to rebuild the ground truth.

5 Conclusion & Perspectives.

In this work, we investigated heterogeneous treatment effects estimation under multiple treatment regime. We considered standard plug-in, pseudo-outcome and Neyman-orthogonality-based meta-learners. Through a semi-synthetic industrial dataset, we assessed the performances of each meta-learner, we showed the ability of the X-learner to reconstruct the ground truth model and we highlighted the choice of base-learner on the quality of CATEs estimation. The next step would be to extend this approach to

Table 1: **mPEHE** for XGBoost and RandomForest

Meta-learner	XGBoost	RandomForest
T-learner	0.167	0.154
<i>Reg</i> T-Learner	0.153	0.153
S-learner	0.101	0.216
M-learner	1.05	0.907
DR-learner	0.100	0.162
X-learner	0.095	0.175
RLin-learner	0.336	0.338

continuous treatments and use this dataset to validate other Causal Inference methods.

References

- Songcai Han, Yuanfang Cheng, Qi Gao, Chuanliang Yan, and Jincheng Zhang. Numerical study on heat extraction performance of multistage fracturing enhanced geothermal system. *Renewable Energy*, 149:1214–1226, 2020. ISSN 0960-1481. doi: <https://doi.org/10.1016/j.renene.2019.10.114>.
- Paul W. Holland. Statistics and causal inference. *Journal of the American Statistical Association*, 81(396):945–960, 1986.
- Sören R. Künzel, Jasjeet S. Sekhon, Peter J. Bickel, and Bin Yu. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences*, 116(10):4156–4165, Feb 2019. ISSN 1091-6490. doi: 10.1073/pnas.1804597116.
- Jersey Neyman. Sur les applications de la théorie des probabilités aux expériences agricoles: Essai des principes. *Roczniki Nauk Rolniczych*, 10:1–51, 1923.
- Peter M. Robinson. Root-n-consistent semiparametric regression. *Econometrica*, 56(4): 931–954, 1988.
- Donald Rubin. Estimating causal effects if treatment in randomized and nonrandomized studies. *J. Educ. Psychol.*, 66, 01 1974.
- Uri Shalit, Fredrik D. Johansson, and David Sontag. Estimating individual treatment effect: Generalization bounds and algorithms. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ICML’17, page 3076–3085. JMLR.org, 2017.