

Modelling the Hemodynamic Response Function (HRF) by Double Machine Learning

Presented by: Hongbiao Chen

Supervisor: Dr. Michel Besserve

Second Reader: Prof. Dr. Nikos Logothetis

Overview

The **global LFP** would affect the **local LFP** and the **local BOLD** signal, which acts as a high-dimensional **confounder** for the **local HRF** estimation.

Double ML can remove the effect of **global LFP** for the **unbiased** estimation, compared to the **naive regression** method.

- **Introduction**

- From LFP to BOLD
- Modelling HRF

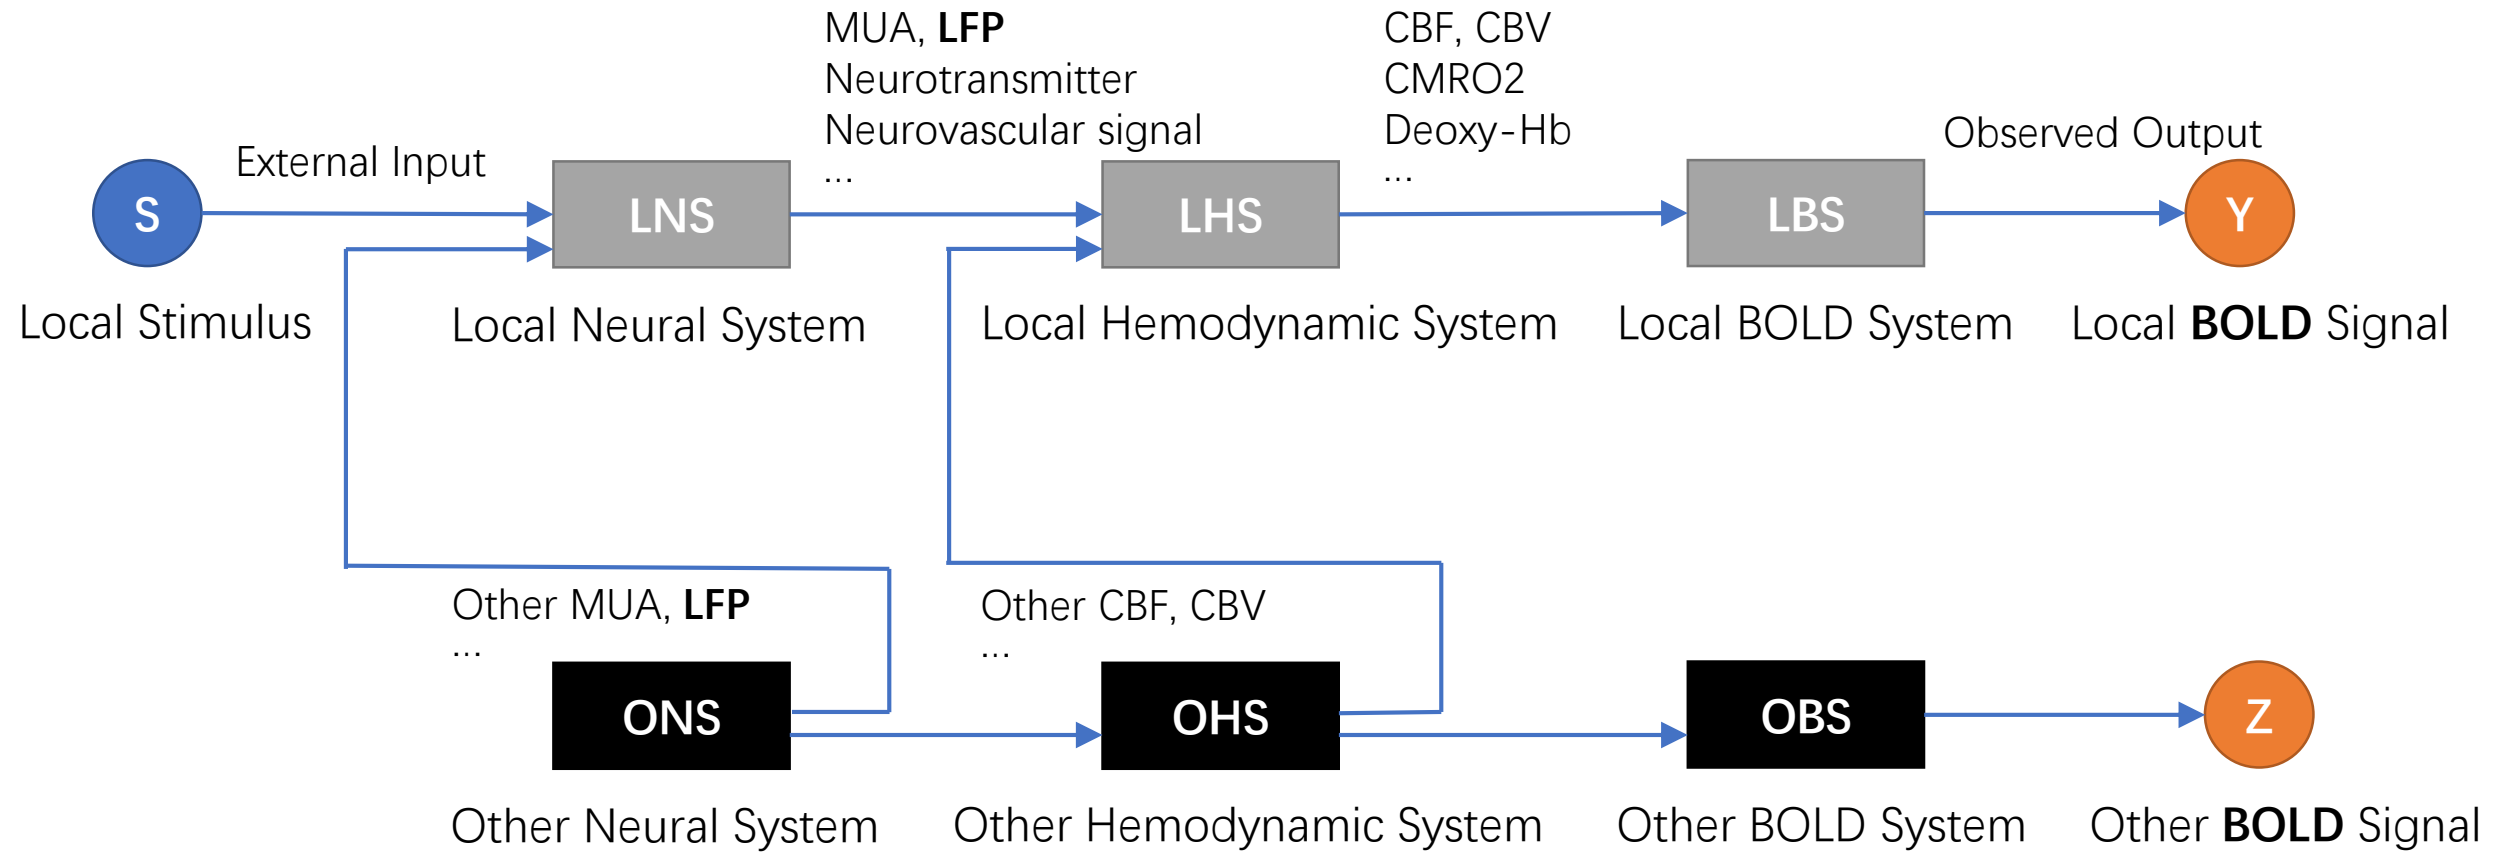
- **Method**

- Double ML
- EconML toolbox


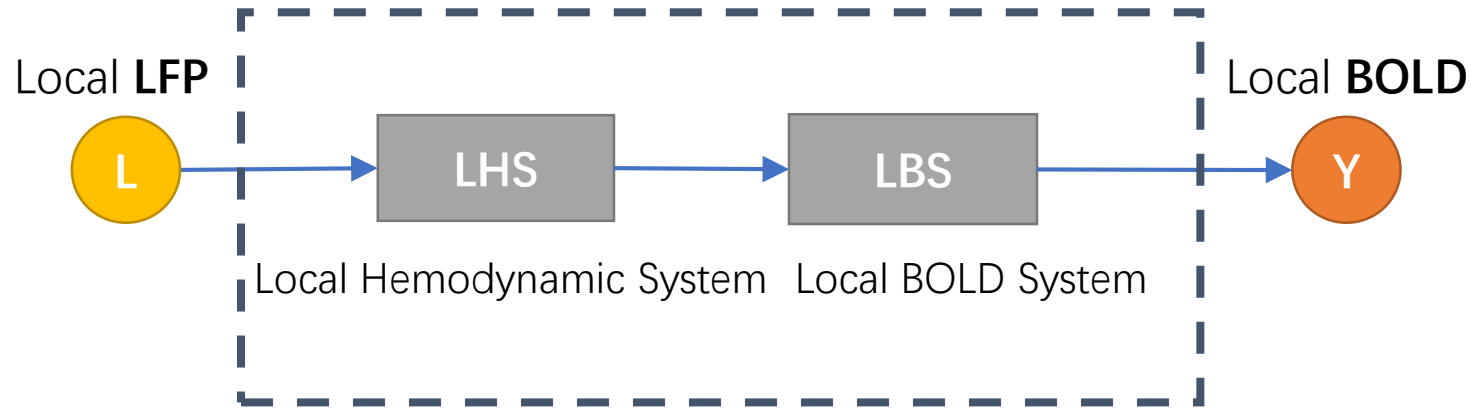
- ~~**Result**~~

- ~~**Summary**~~

Introduction: From LFP to BOLD



Introduction: Modelling HRF



$$y(t) = (s * h)(t).$$

$$h(t) = \sum_{i=1}^B \beta_i g_i(t)$$

$y(t)$: local BOLD at time t

$s(t)$: local LFP at time t

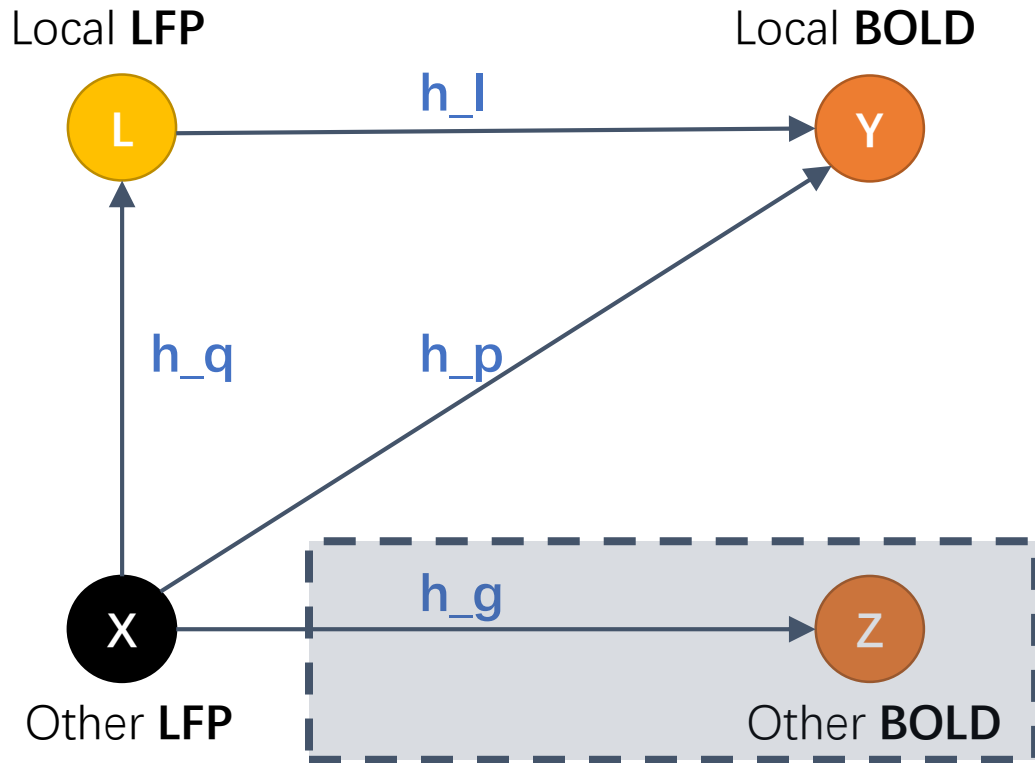
$h(t)$: hemodynamic response at time t

β_i : regression coefficients

g_i : basis functions

Assumed B basis function in total.

Introduction: In the Present of Confounder



$$Y(\mathbf{t}) = \mathbf{h}_p * (X(\mathbf{t}) \cdot \mathbf{a}^T) + \mathbf{h}_l * (\mathbf{L}(\mathbf{t})) + \text{noise}_1$$
$$\mathbf{L}(\mathbf{t}) = \mathbf{h}_q * (X(\mathbf{t}) \cdot \mathbf{b}^T) + \text{noise}_2$$


\mathbf{a}^T : coefficients for h_p link

\mathbf{b}^T : coefficients for h_q link

noise_1: additional noise of local BOLD

noise_2: additional noise of local LFP

Introduction: Partially Linear Regression (PLR)



$$Y = \theta_0(L) + g_0(X) + U$$
$$L = m_0(X) + V$$

$$E[U | X, L] = 0$$
$$E[V | X] = 0$$

Y: outcome variable

L: treatment variable (low-dimension)

X: confounder (high-dimension)

U: noise for Y

V: noise for L

θ_0 : interested parameters

g_0 and **m_0** : nuisance parameters

Some hidden functions, could be constant or nonlinear

Method: Naive Regression

Regress Y by using L and X , and get the predicted Y :

$$\hat{\theta}_0(L) + \hat{g}_0(X)$$

For example, using alternating minimization:

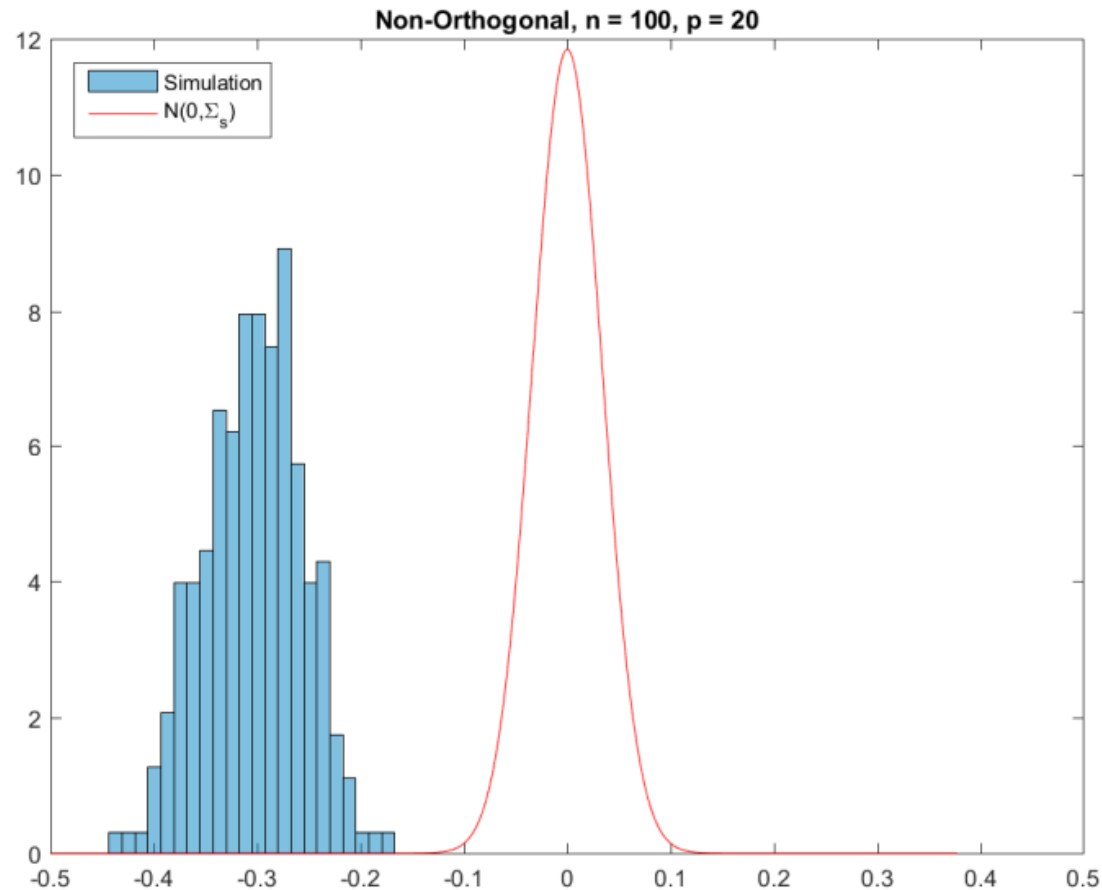
- (1) Given an initial guess of $\hat{\theta}_0()$
- (2) Regress $Y - \hat{\theta}_0(L)$ on X to fit $\hat{g}_0()$ by random forest
- (3) Regress $Y - \text{first_fitted_}\hat{g}_0(X)$ on L to fit updated $\hat{\theta}_0()$
- (4) Repeat (2) and (3) until convergence to get final $\hat{\theta}_0()$ and $\hat{g}_0()$

Didn't use the information from $L = m_0(X) + V$;

Based on the moment condition: $E[(Y - \theta_0(L) - g_0(X))L] = 0$.

Method: Naive Regression is Bad

Good **Prediction Performance**, but Bad Causal **Parameter θ_0 Estimation**.
The distribution of $\hat{\theta}_0 - \theta_0$ looks like this:



Method: Two Sources of Bias from ML

Bias - Variance Tradeoff

$$\text{Error}(x) = \left(E[\hat{f}(x)] - f(x) \right)^2 + E \left[\hat{f}(x) - E[\hat{f}(x)] \right]^2 + \sigma_e^2$$

BY CHRIS ALBON

Bias²
How much predicted values differ from true values.

Variance
How predictions made on the same value vary on different realizations of the model

irreducible error

Machine Learning Flashcard by Chris Albon

- (1) Bias from Regularization
- (2) Bias from Overfitting

Method: Double ML

Regress **Y** and **L** by using **X**, and get the predicted Y and L:

$$\hat{l}_0(X) = \hat{E}[Y|X]$$
$$\hat{m}_0(X) = \hat{E}[L|X]$$

Frisch-Waugh-Lovell (1930s) style

(1) Get $\hat{E}[Y|X]$ and $\hat{E}[L|X]$ by using random forest

(2) Compute the residuals: $\hat{W} = Y - \hat{E}[Y|X]$ and $\hat{V} = L - \hat{E}[L|X]$


(3) Regress \hat{W} on \hat{V} to fit the $\hat{\theta}_0()$

Use the information from $L = \mathbf{m}_0(\mathbf{X}) + V$;

Based on the moment condition: $E[((Y - E[Y|X]) - \theta_0(L - E[L|X])) (L - E[L|X])] = 0$

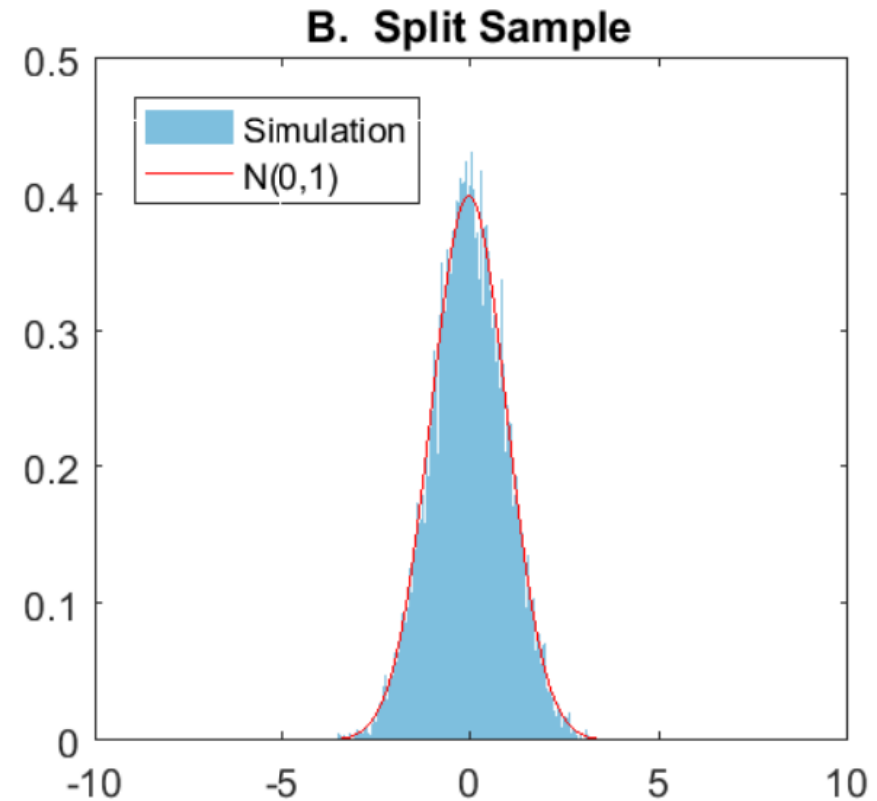

Method: Double ML is Good

The distribution of $\hat{\theta}_0 - \theta_0$ looks like this now:



Method: Deal with Two Sources of Bias from ML

- (1) Correct **Regularization** bias by **Neyman Orthogonal Structure**.
- (2) Correct **Overfitting** bias by **Sample Splitting** with **cross-fitting** for efficiency.



Method: Sample Splitting with cross-fitting

Fit nuisance functions:

$$\hat{g}_0, \hat{m}_0, \hat{l}_0$$

Fit interested function:

$$\hat{\theta}_0$$

Sample splitting: randomly partition the data into two subsets

Fit interested function:

$$\hat{\theta}_0$$

Fit nuisance functions:

$$\hat{g}_0, \hat{m}_0, \hat{l}_0$$

Cross-fitting:

average two estimated $\hat{\theta}_0$
for the final estimated $\hat{\theta}_0$

Method: EconML toolbox

EconML

<https://www.microsoft.com/en-us/research/project/econml/>



[Overview](#) [People](#) [Publications](#) [News & features](#) [Get Started](#) [How To](#) [Use Cases](#) [FAQ](#)

[EconML on GitHub >](#)

[Documentation on EconML >](#)

Overview

EconML is a Python package that applies the power of machine learning techniques to estimate individualized causal responses from observational or experimental data. The suite of estimation methods provided in EconML represents the latest advances in causal machine learning. By incorporating individual machine learning steps into interpretable causal models, these methods improve the reliability of what-if predictions and make causal analysis quicker and easier for a broad set of users.

EconML is open source software developed by the [ALICE team](#) at Microsoft Research.

Method: Conditional Average Treatment Effect (CATE)

We assume we have data that are generated from some collection policy. In particular, we assume that we have data of the form: $\{Y_i(T_i), T_i, X_i, W_i, Z_i\}$, where $Y_i(T_i)$ is the observed outcome for the chosen treatment, T_i is the treatment, X_i are the co-variates used for heterogeneity, W_i are other observable co-variates that we believe are affecting the potential outcome $Y_i(T_i)$ and potentially also the treatment T_i ; and Z_i are variables that affect the treatment T_i but do not directly affect the potential outcome. We will refer to variables W_i as *controls* and variables Z_i as *instruments*. The variables X_i can also be thought of as *control variables*, but they are special in the sense that they are a subset of the controls with respect to which we want to measure treatment effect heterogeneity. We will refer to them as *features*.

Method: Conditional Average Treatment Effect (CATE)

We can equivalently describe the data and the quantities of interest via the means of structural equations. In particular, suppose that we observe i.i.d. samples $\{Y_i, T_i, X_i, W_i, Z_i\}$ from some joint distribution and we assume the following structural equation model of the world:

$$\begin{aligned} Y &= g(T, X, W, \epsilon) \\ T &= f(X, W, Z, \eta) \end{aligned}$$

where ϵ and η are noise random variables that are independent of X, Z, T, W but could be potentially correlated with each other. The target quantity that we want to estimate can then be expressed as:

$$\begin{aligned} \tau(\mathbf{t}_0, \mathbf{t}_1, \mathbf{x}) &= \mathbb{E}[g(\mathbf{t}_1, X, W, \epsilon) - g(\mathbf{t}_0, X, W, \epsilon) | X = \mathbf{x}] && \text{(CATE)} \\ \partial\tau(\mathbf{t}, \mathbf{x}) &= \mathbb{E}[\nabla_{\mathbf{t}} g(\mathbf{t}, X, W, \epsilon) | X = \mathbf{x}] && \text{(marginal CATE)} \end{aligned}$$

where in these expectations, the random variables W, ϵ are taken from the same distribution as the one that generated the data. In other words, there is a one-to-one correspondence between the potential outcomes formulation and the structural equations formulation in that the random variable $Y(\mathbf{t})$ is equal to the random variable $g(\mathbf{t}, X, W, \epsilon)$, where X, W, ϵ is drawn from the distribution that generated each sample in the data set.

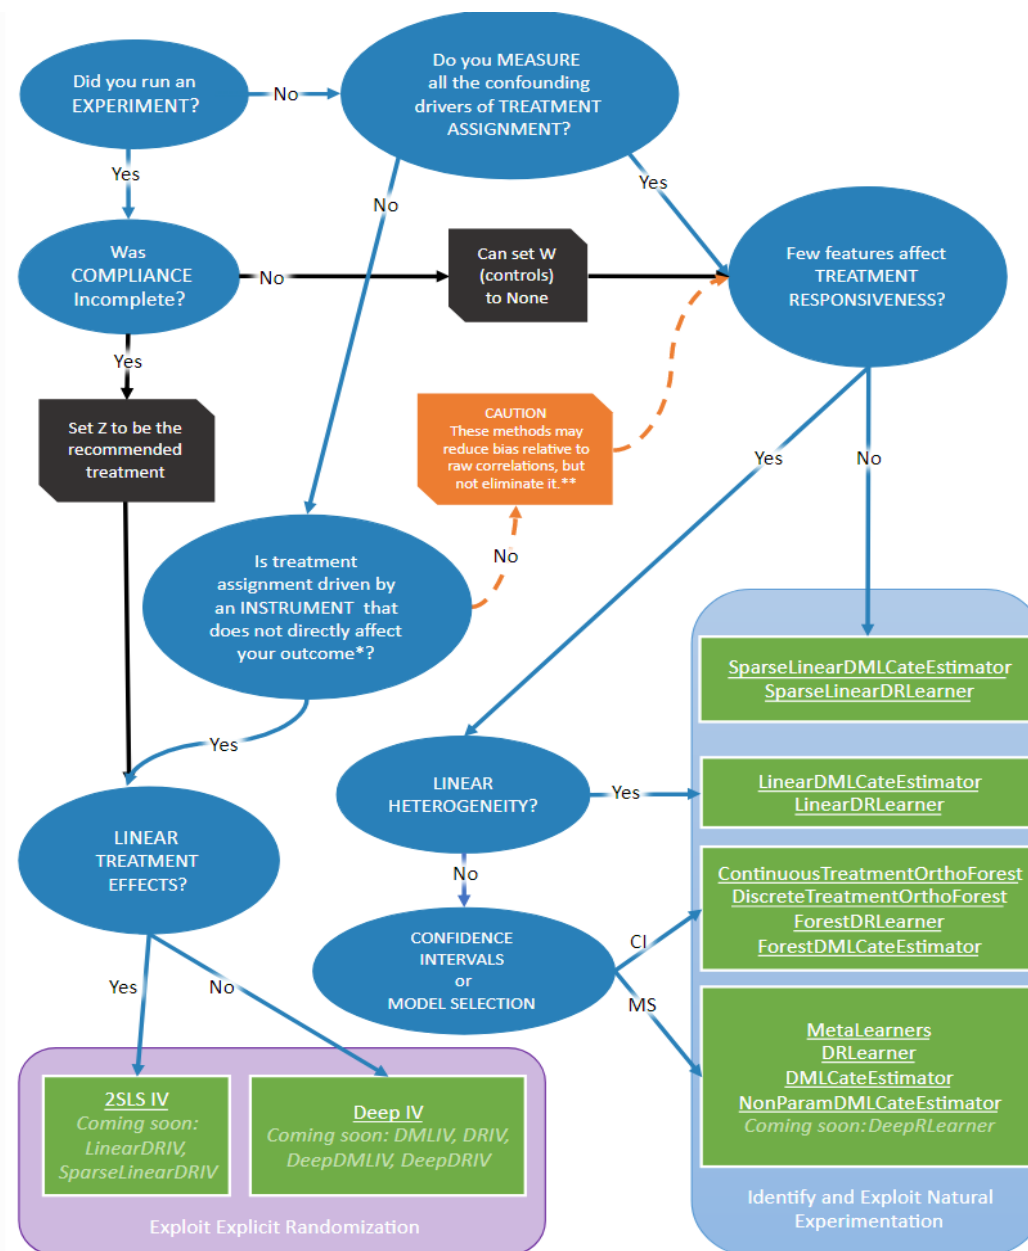
Method: Estimation Methods

Estimation:

- Double Machine Learning (DML)
- Doubly Robust Learning
- Forest Based Estimators
- Meta-Learners

Inference:

- Bootstrap Inference
- OLS Inference
- Debiased Lasso Inference
- Subsampled Honest Forest Inference



Method: Double Machine Learning (DML)

The model makes the following structural equation assumptions on the data generating process.

$$\begin{aligned} Y &= \theta(X) \cdot T + g(X, W) + \epsilon & \mathbb{E}[\epsilon | X, W] &= 0 \\ T &= f(X, W) + \eta & \mathbb{E}[\eta | X, W] &= 0 \\ & & \mathbb{E}[\eta \cdot \epsilon | X, W] &= 0 \end{aligned}$$

What is particularly attractive about DML is that it makes no further structural assumptions on g and f and estimates them non-parametrically using arbitrary non-parametric Machine Learning methods. Our goal is to estimate the constant marginal CATE $\theta(X)$.

Method: DML Example

EXP 1. Single Continuous Treatment

$$T = \langle W, \beta \rangle + \eta$$
$$Y = T \cdot \theta(X) + \langle W, \gamma \rangle + \epsilon$$

noisy error source $\left\{ \begin{array}{l} \eta \sim \text{Uniform}(-1, 1) \\ \epsilon \sim \text{Uniform}(-1, 1) \end{array} \right.$

$\left\{ \begin{array}{l} W \sim \text{Normal}(0, I_{nw}) \\ X \sim \text{Uniform}(0, 1)^{nx} \end{array} \right.$

$\theta(X) = \exp(2 \cdot x_1)$

$\left\{ \begin{array}{l} Y: \text{outcomes for each sample} \\ T: \text{treatments for each sample} \\ X: \text{features for each sample} \\ \theta(X): \text{treatment effects for each sample} \\ W: \text{controllers / confounders for each sample} \\ \quad (\text{high-dimension}) \end{array} \right.$

```
graph TD; W((W)) -- beta --> T((T)); W --> Y((Y)); X((X)) --> T; X --> Y; T --> Y; X --> theta((theta(X))); theta --> Y; epsilon((epsilon)) --> Y;
```

Method: DML Example

$$T = W \cdot \beta + \eta$$

Dimensions: 2000×1 (T), 2000×5 (W), 5×1 (β), 2000×1 (η)

$$Y = T \cdot TE + W \cdot y + \epsilon$$

Dimensions: 2000×1 (Y), 2000×1 (T), 1×1 (TE), 2000×5 (W), 5×1 (y), 2000×1 (ϵ)

Parameters:

- $n_{\text{-sample}} = 2000$
- $n_{\text{-w}} = 30$
- $n_{\text{-x}} = 1$
- $\text{support-size} = 5$

Method: DML Example

