Causal Strategic Learning with Competitive Selection

Kiet Q. H. Vo,
Muneeb Aadil,
Siu Lun Chau,
Krikamol Muandet

(0). TL; DR

Prior work studied predictions (incl. classification) but did not consider the impact of selection.

We show that:

- The optimal selection rule is a trade-off between picking the most capable agents and maximally incentivising everyone,
- Agent improvement requires not only causal inference but also a benevolent regulator to enforce it (see (3). and (6).).

(2). Core assumptions

- Linear & additive structures.
- Agents differ in their base covariates, but improve by the same amount.

(3). Minimising the trade-off

Additive structures result in:

$$\mathcal{Q}(\theta) = \text{cBP}(\theta) + \text{cPI}(\theta)$$

Further linearity leads to:

$$\arg\max_{\theta} \; Q(\theta) = \frac{\alpha + MM^{\top}\theta^*}{\|\alpha + MM^{\top}\theta^*\|}$$

To align the two objectives:

$$\alpha = (k-1)MM^{\top}\theta^*$$

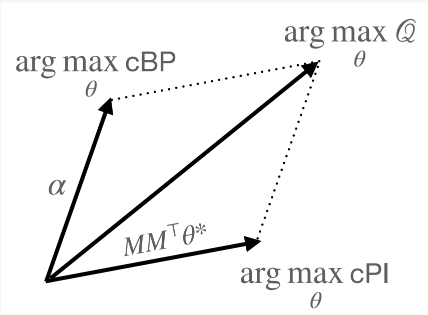

(5). Prior work vs Ours (cont'd)

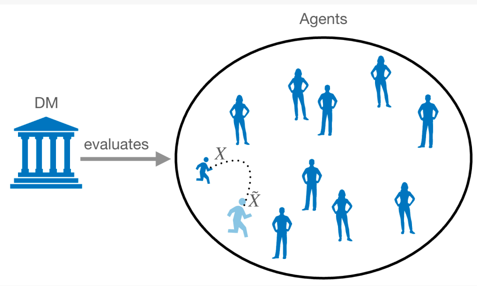

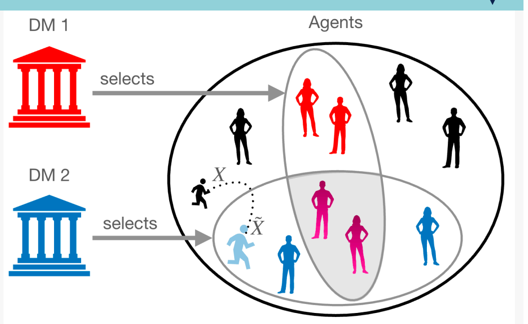

$$\mathbb{E}[\phi_{i,\theta^{all}}(X, Y_i) \mid Z = i; \theta^{all}]$$

(6). Learning under competitive selection

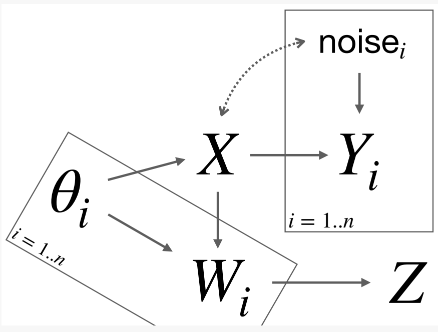

A cooperative protocol for all DMs to obtain partitions of data without correlated noise:

$$\exists\{t, t'\}, \forall i : \theta_{it} = k_i \theta_{it'}$$

(1). Prior work vs. Ours…

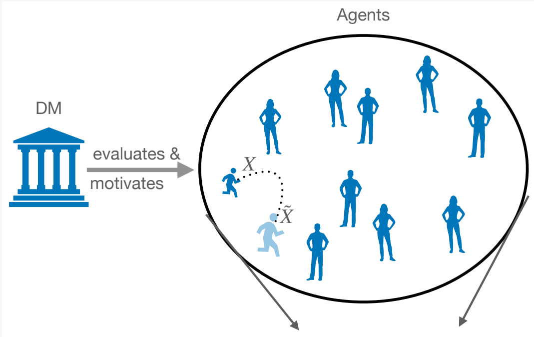

$$\mathcal{Q}(\theta) = \mathbb{E}[\phi_\theta(X, Y)]$$

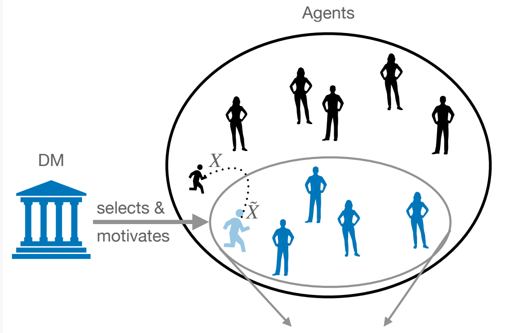

$$\mathcal{Q}(\theta) = \mathbb{E}[\phi_\theta(X, Y) \mid W_\theta = 1]$$

for some objective function $\phi_\theta(X, Y)$ encompassing BP (base performance) and PI (performance improvement).

With additive structures:

$$\begin{aligned} &\mathbb{E}[\phi_\theta(X, Y)] \\ &= \mathbb{E}[BP] + \mathbb{E}[PI_\theta] \\ &= \text{const} + \mathbb{E}[PI_\theta] \end{aligned}$$

$$\begin{aligned} &\mathbb{E}[\phi_\theta(X, Y) \mid W_\theta = 1] \\ &= \mathbb{E}[BP \mid W_\theta = 1] + \mathbb{E}[PI_\theta \mid W_\theta = 1] \\ &= \text{cBP}(\theta) + \text{cPI}(\theta) \end{aligned}$$

(4). Learning θ^* under selection bias

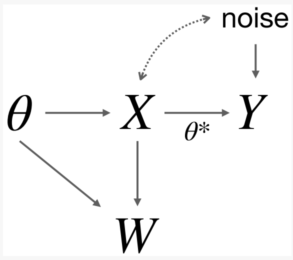

$$\exists\{\theta_\dagger, \theta_\diamond\} :$$

$$P(\text{noise} \mid W = 1; \theta_\dagger) = P(\text{noise} \mid W = 1; \theta_\diamond)$$

$$\Delta\theta \longrightarrow \Delta\mathbb{E}[X \mid W = 1] \xrightarrow[\theta^*]{} \Delta\mathbb{E}[Y \mid W = 1]$$

Certain partitioning schemes can remove the correlated noise.

(7). Experiments

Our MSLR has unbiased estimation of θ^*.

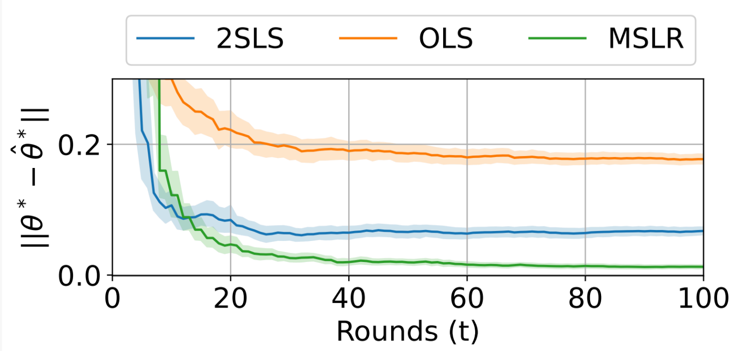

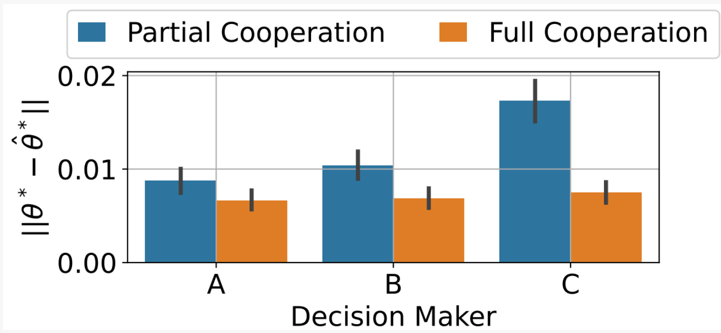

More cooperation,

less estimation error.

CISPA
HELMHOLTZ CENTER FOR INFORMATION SECURITY