Learning Optimal Group-structured Individualized Treatment Rules with Many Treatments

Haixu Ma*, Donglin Zeng, Yufeng Liu

- Department of Statistics and Operations Research
  University of North Carolina at Chapel Hill

- *International Conference on Machine Learning (ICML)*, Hawaii, 2023
  Journal-to-Conference Track (*Journal of Machine Learning Research*)

- Example in Personalized Medicine
  - *Individualized cancer treatment*: tailoring therapies based on patients' genomic biomarkers to **optimize** future health status

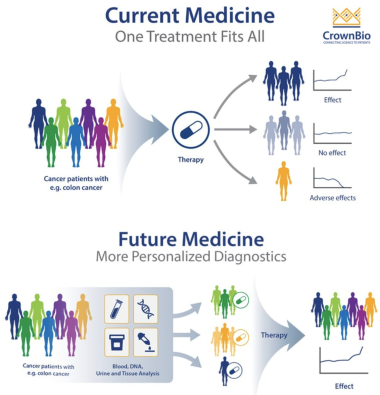

Figure 1: Transition from "one size fits all" to personalized medicine

Setup

- Data $(Z, A, Y) \in \mathcal{Z} \times \mathcal{A} \times \mathbb{R}$
  1. Features $Z \in \mathcal{Z} \subseteq \mathbb{R}^p$:
  2. Assigned treatment $A \in \mathcal{A} = \{1, 2, \ldots, M\}$, where M can be large
  3. Reward $Y \in \mathbb{R}$:
- Propensity score $p(a|\boldsymbol{z}) := \mathbb{P}(A = a|Z = \boldsymbol{z})$ for $a \in \mathcal{A}$ and $\boldsymbol{z} \in \mathbb{R}^p$
- ★ Individualized Treatment Rule (ITR) $D : \mathcal{Z} \to \mathcal{A}$
- Under SUTVA assumptions [Rubin, 1974], value function [Zhao et al., 2012] of an ITR D is

$$\mathcal{V}(D) = \mathbb{E}\left[\frac{\mathbb{I}[D(Z) = A]}{p(A|Z)} Y\right] \Leftarrow \textbf{Inverse Probability Weighting (IPW)}$$

- Goal: Learn optimal ITR $D^* \in \mathcal{D}$ that maximizes the value function

$$D^* \in \arg\max_{D \in \mathcal{D}} \mathcal{V}(D),$$

where for any $\boldsymbol{z} \in \mathcal{Z}$,

$$D^*(\boldsymbol{z}) \in \arg\max_{a \in \mathcal{A}} \underbrace{\mathbb{E}[Y|Z = \boldsymbol{z}, A = a]}_{\star \text{ Heterogeneous Treatment Effect (HTE)}}$$

Setup

- Data $(Z, A, Y) \in \mathcal{Z} \times \mathcal{A} \times \mathbb{R}$
  1. Features $Z \in \mathcal{Z} \subseteq \mathbb{R}^p$:
  2. Assigned treatment $A \in \mathcal{A} = \{1, 2, \ldots, M\}$, where M can be large
  3. Reward $Y \in \mathbb{R}$:
- Propensity score $p(a|\boldsymbol{z}) := \mathbb{P}(A = a|Z = \boldsymbol{z})$ for $a \in \mathcal{A}$ and $\boldsymbol{z} \in \mathbb{R}^p$
- ★ Individualized Treatment Rule (ITR) $D : \mathcal{Z} \to \mathcal{A}$
- Under SUTVA assumptions [Rubin, 1974], value function [Zhao et al., 2012] of an ITR D is

$$\mathcal{V}(D) = \mathbb{E}\left[\frac{\mathbb{I}[D(Z) = A]}{p(A|Z)} Y\right] \Leftarrow$$ **Inverse Probability Weighting (IPW)**

- Goal: Learn optimal ITR $D^* \in \mathcal{D}$ that maximizes the value function

$$D^* \in \arg\max_{D \in \mathcal{D}} \mathcal{V}(D),$$

where for any $\boldsymbol{z} \in \mathcal{Z}$,

$$D^*(\boldsymbol{z}) \in \arg\max_{a \in \mathcal{A}} \underbrace{\mathbb{E}[Y|Z = \boldsymbol{z}, A = a]}_{\bigstar \text{ Heterogeneous Treatment Effect (HTE)}}$$

THE UNIVERSITY of NORTH CAROLINA at CHAPEL HILL

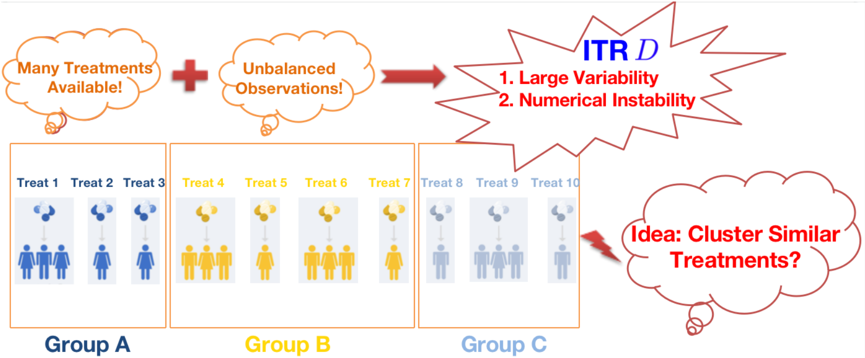

Figure 2: Learning optimal ITRs with many treatments.

1. Many treatments but limited observations for some specific treatments:
   - Patient-Derived Xenograft study: more than 20 treatments
   - Unbalanced treatment assignment
   - Current (direct/indirect) methods suffer from **large variability** $+$ **numerical instability**

     ★ How to learn the optimal ITR for **many** treatments?

2. Treatments in large treatment space may work similarly for patients
   - Depression study: many treatment options are combined into *SSRI/non-SSRI* groups
   - Few methods deal with **clustering** treatments

     ★ How to **cluster** the treatments with similar treatment effects?

1. Many treatments but limited observations for some specific treatments:
   - Patient-Derived Xenograft study: more than 20 treatments
   - Unbalanced treatment assignment
   - Current (direct/indirect) methods suffer from **large variability** $+$ **numerical instability**

     ★ How to learn the optimal ITR for **many** treatments?
2. Treatments in large treatment space may work similarly for patients
   - Depression study: many treatment options are combined into *SSRI/non-SSRI* groups
   - Few methods deal with **clustering** treatments

     ★ How to **cluster** the treatments with similar treatment effects?

THE UNIVERSITY of NORTH CAROLINA at CHAPEL HILL

- ★ Idea: Estimate optimal partition on $\mathcal{A}$ to cluster similar treatments
- Aim to partition $|\mathcal{A}| = M_n$ (large) treatments into K_n treatment groups
- **Supervised clustering**: learn optimal ITR (*supervised learning*), while *at the same time* clustering treatments (*unsupervised learning*)

- Define group-structured ITR class $\mathcal{D} = \bigcup_\delta \mathcal{D}_\delta$:
  - For a fixed δ, a group-structured ITR $\in \mathcal{D}_\delta$ is obtained from a random policy π_δ given as

$$\pi_\delta(a|\boldsymbol{z}) = \underbrace{\mathbb{I}[\delta(a) = D_g(\boldsymbol{z})]}_{\textbf{Deterministic}} \underbrace{\frac{p(a|\boldsymbol{z})}{p(\delta(a)|\boldsymbol{z})}}_{\textbf{Random}}$$

  - D_g: $\mathcal{Z} \to [K_n]$, *group-based* decision rule
  - $p(\delta(a)|\boldsymbol{z})$: propensity score of $\delta(a)$-th group under δ

THE UNIVERSITY *of* NORTH CAROLINA *at* CHAPEL HILL

- Value of group-structured ITR $\mathcal{V}_1(\delta, D_g)$:

$$\mathcal{V}_1(\delta, D_g) = \mathbb{E}\left[\frac{\mathbb{I}[D_g(Z) = \delta(A)]}{p(\delta(A)|Z)} Y\right]$$

- For any δ, let $D_g^\delta \in \arg\max_{D_g} \mathcal{V}_1(\delta, D_g)$ be optimal *group-based* decision rule
- $\mathcal{V}_1^*(\delta) := \mathcal{V}_1(\delta, D_g^\delta)$ is corresponding optimal value for δ
- ★ Optimal partition $\delta^* \in \arg\max_\delta \mathcal{V}_1^*(\delta) := \Delta^*$
- Key observation:

$\mathcal{V}^* = \mathbb{E}_Z\left[\max_{a \in [M_n]} \mathbb{E}[Y|A = a, Z]\right] \Leftarrow$ **Individual Treatment Domain**

$\mathcal{V}_1^*(\delta) = \mathbb{E}_Z\left[\max_{k \in [K_n]} \mathbb{E}[Y|A \in G_k^\delta, Z]\right] \Leftarrow$ **Group Treatment Domain**

- ★ Interpretation: δ^* optimizes expected *group-based* heterogeneous treatment effects

- Value of group-structured ITR $\mathcal{V}_1(\delta, D_g)$:

$$\mathcal{V}_1(\delta, D_g) = \mathbb{E}\left[\frac{\mathbb{I}[D_g(Z) = \delta(A)]}{p(\delta(A)|Z)} Y\right]$$

- For any δ, let $D_g^\delta \in \arg\max_{D_g} \mathcal{V}_1(\delta, D_g)$ be optimal *group-based* decision rule
- $\mathcal{V}_1^*(\delta) := \mathcal{V}_1(\delta, D_g^\delta)$ is corresponding optimal value for δ
- ★ Optimal partition $\delta^* \in \arg\max_\delta \mathcal{V}_1^*(\delta) := \Delta^*$
- Key observation:

$$\mathcal{V}^* = \mathbb{E}_Z\left[\max_{a \in [M_n]} \mathbb{E}[Y|A = a, Z]\right] \Leftarrow \textbf{Individual Treatment Domain}$$

$$\mathcal{V}_1^*(\delta) = \mathbb{E}_Z\left[\max_{k \in [K_n]} \mathbb{E}[Y|A \in G_k^\delta, Z]\right] \Leftarrow \textbf{Group Treatment Domain}$$

- ★ Interpretation: δ^* optimizes expected *group-based* heterogeneous treatment effects

THE UNIVERSITY *of* NORTH CAROLINA *at* CHAPEL HILL

GRoup Outcome Weighted Learning (GROWL)

- Goal: Estimate optimal partition δ^* and *group-based* decision rule D_g
- Maximizing value function $\mathcal{V}_1$ $\Leftrightarrow$ minimizing risk function $\widetilde{\mathcal{R}}$ [Zhao et al., 2012]

  - $$\max_{\delta, D_g} \mathcal{V}_1(\delta, D_g) \Leftrightarrow \min_{\delta, D_g} \left\{ \widetilde{\mathcal{R}}(\delta, D_g) := \mathbb{E}\left[\frac{\mathbb{I}[D_g(Z) \neq \delta(A)]}{p(\delta(A)|Z)} Y\right] - \underbrace{\mathbb{E}\left[\frac{Y}{p(\delta(A)|Z)}\right]}_{\text{free of } \delta} \right\}$$

- Two-step implementation:
  1. For each δ, estimate D_g^δ: minimizing risk $\mathcal{R}$ $\Leftrightarrow$ Weighted Classification

  $$\widehat{D}_g^\delta \in \arg\min_{D_g} \mathbb{E}_n \Big[\underbrace{\frac{Y}{p(\delta(A)|Z)}}_{\text{Weighted}} \underbrace{\mathbb{I}[D_g(Z) \neq \delta(A)]}_{\text{Classification}} \Big]$$

  2. Plug $(\delta, \widehat{f}^\delta)$ back to $\widetilde{\mathcal{R}}_\phi$ and solve integer programming problem for δ

THE UNIVERSITY of NORTH CAROLINA at CHAPEL HILL

- Goal: Estimate optimal partition δ^* and *group-based* decision rule D_g
- Maximizing value function $\mathcal{V}_1$ $\Leftrightarrow$ minimizing risk function $\widetilde{\mathcal{R}}$ [Zhao et al., 2012]

$$\blacktriangleright \max_{\delta, D_g} \mathcal{V}_1(\delta, D_g) \Leftrightarrow \min_{\delta, D_g} \left\{ \widetilde{\mathcal{R}}(\delta, D_g) := \mathbb{E}\left[\frac{\mathbb{I}[D_g(Z) \neq \delta(A)]}{p(\delta(A)|Z)} Y\right] - \underbrace{\mathbb{E}\left[\frac{Y}{p(\delta(A)|Z)}\right]}_{\textbf{free of } \delta} \right\}$$

- Two-step implementation:
  1. For each δ, estimate D_g^δ: minimizing risk $\mathcal{R}$ $\Leftrightarrow$ Weighted Classification

$$\widehat{D}_g^\delta \in \arg\min_{D_g} \mathbb{E}_n \Big[\underbrace{\frac{Y}{p(\delta(A)|Z)}}_{\textbf{Weighted}} \underbrace{\mathbb{I}[D_g(Z) \neq \delta(A)]}_{\textbf{Classification}} \Big]$$

  2. Plug $(\delta, \widehat{\boldsymbol{f}}^\delta)$ back to $\widetilde{\mathcal{R}}_\phi$ and solve integer programming problem for δ

THE UNIVERSITY of NORTH CAROLINA at CHAPEL HILL

Step 1:

- $\underbrace{\mathbb{I}[D_g(\boldsymbol{z}) \neq \delta(a)]}_{\textbf{0-1 loss}} \Longrightarrow \underbrace{L_\phi(\delta(a), \boldsymbol{f}(\boldsymbol{z}))}_{\textbf{Reinforced Angle-based Multicategory SVM loss (RAMSVM)}}$
- Decision function $\boldsymbol{f}$: $\mathcal{Z} \to \mathbb{R}^{K_n - 1}$ for multicatergory classification
- Convex combination of two loss functions by $\gamma \in [0, 1]$:
- ★ *Group-based* decision rule: Maximizing $\langle \cdot, \cdot \rangle \Leftrightarrow$ minimizing angle:

$$D_g(\boldsymbol{z}) = \arg\max_{k \in [K_n]} \langle \mathbf{W}_k, \boldsymbol{f}(\boldsymbol{z}) \rangle$$

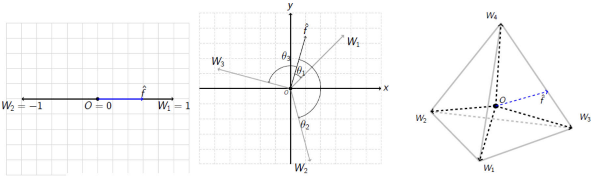

$k = 2$ $k = 3$ $k = 4$

Figure 3: Angle-based multicategory classification.

THE UNIVERSITY *of* NORTH CAROLINA *at* CHAPEL HILL

Step 1:

- $\underbrace{\mathbb{I}[D_g(\boldsymbol{z}) \neq \delta(a)]}_{\text{0-1 loss}} \Longrightarrow \underbrace{L_\phi(\delta(a), \boldsymbol{f}(\boldsymbol{z}))}_{\text{Reinforced Angle-based Multicategory SVM loss (RAMSVM)}}$
- Decision function $\boldsymbol{f}$: $\mathcal{Z} \to \mathbb{R}^{K_n - 1}$ for multicatergory classification
- Convex combination of two loss functions by $\gamma \in [0, 1]$:
- ★ *Group-based* decision rule: Maximizing $\langle \cdot, \cdot \rangle \Leftrightarrow$ minimizing angle:

$$D_g(\boldsymbol{z}) = \arg\max_{k \in [K_n]} \langle \mathbf{W}_k, \boldsymbol{f}(\boldsymbol{z}) \rangle$$

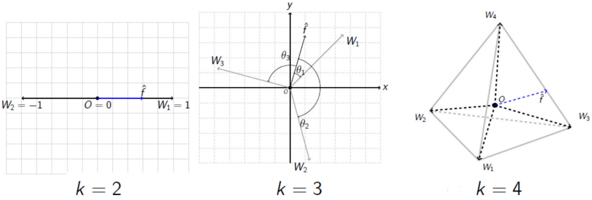

Figure 3: Angle-based multicategory classification.

- Treatment effects have homogeneous grouping structure:

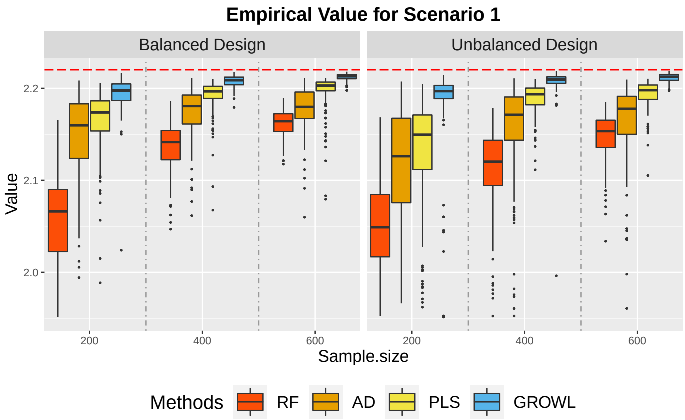

Figure 4: Boxplots of value under *homogeneous* settings and different designs. Red dashed lines demonstrate the oracle value.

★ Trade-off between *bias* and *variance* for value

- As distance between treatments ↑: group structure tends to lose; bias ↑
- Variance of GROWL is small since we consider the group structure

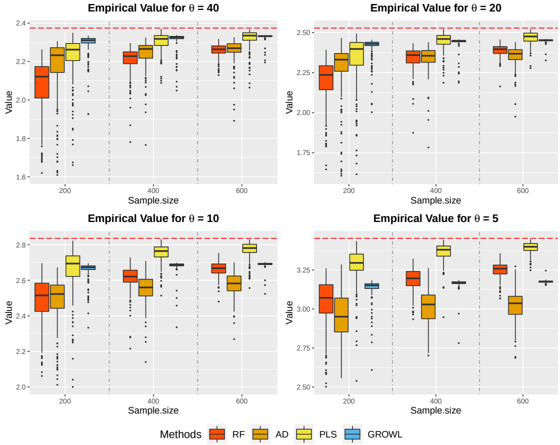

Figure 5: Boxplots of value under *nonhomogeneous* settings and unbalanced design.

We also

- Solved weighted classification problem with RAMSVM effectively
- Proposed coordinate descent type of greedy algorithm to adjust partition δ
- Provided extensive theoretical guarantee for
  - Generalized Fisher consistency
  - Generalized bound for excess risk
  - Convergence rate for value function
- Conducted both simulation studies and real data analysis on depression study

☺ **Thanks for your listening!**

☼ **Welcome to join our poster session:**

Poster Session 2: 2-3:30 pm, July 25th (Tuesday), Exhibit Hall 1, #131

THE UNIVERSITY of NORTH CAROLINA at CHAPEL HILL

References I

Rubin, D. B. (1974).
Estimating causal effects of treatments in randomized and nonrandomized studies.
Journal of educational Psychology, 66(5):688.

Zhang, C., Liu, Y., Wang, J., and Zhu, H. (2016).
Reinforced angle-based multicategory support vector machines.
Journal of Computational and Graphical Statistics, 25(3):806–825.

Zhao, Y., Zeng, D., Rush, A. J., and Kosorok, M. R. (2012).
Estimating individualized treatment rules using outcome weighted learning.
Journal of the American Statistical Association, 107(499):1106–1118.

THE UNIVERSITY *of* NORTH CAROLINA *at* CHAPEL HILL