Robustifying Trial-Derived Treatment Rules to a Target Population

Yingqi Zhao

Public Health Sciences Division
Fred Hutchinson Cancer Research Center

Workshop on Perspectives and Analysis Methods for Personalized Medicine,
Singapore, July 11, 2017

Outline

1. Personalized Medicine
2. Minimax Linear Decisions
   - Framework and Motivation
   - Methods
   - Simulation Studies and Data Analysis
3. Discussion

SWOG 0421 study

- 1038 men with metastatic castration-resistant prostate cancer
- Docetaxel administered every 21 days at a dose of 75 mg/m^2 with or without the bone targeted atrasentan for up to 12 cycles
- Co-primary outcomes: progression-free survival and overall survival
- Data from 751 patients for anlaysis; 371 patients are in the docetaxel + atrasentan arm and 380 patients are in the docetaxel + placebo arm
- 10 covariates: age, serum prostate-specific antigen (PSA), indicator of bisphosphonate usage, indicator of metastatic disease beyond the bones, indicator of pain at basline, indicator of performance status, and 4 bone marker levels

SWOG 0421 study

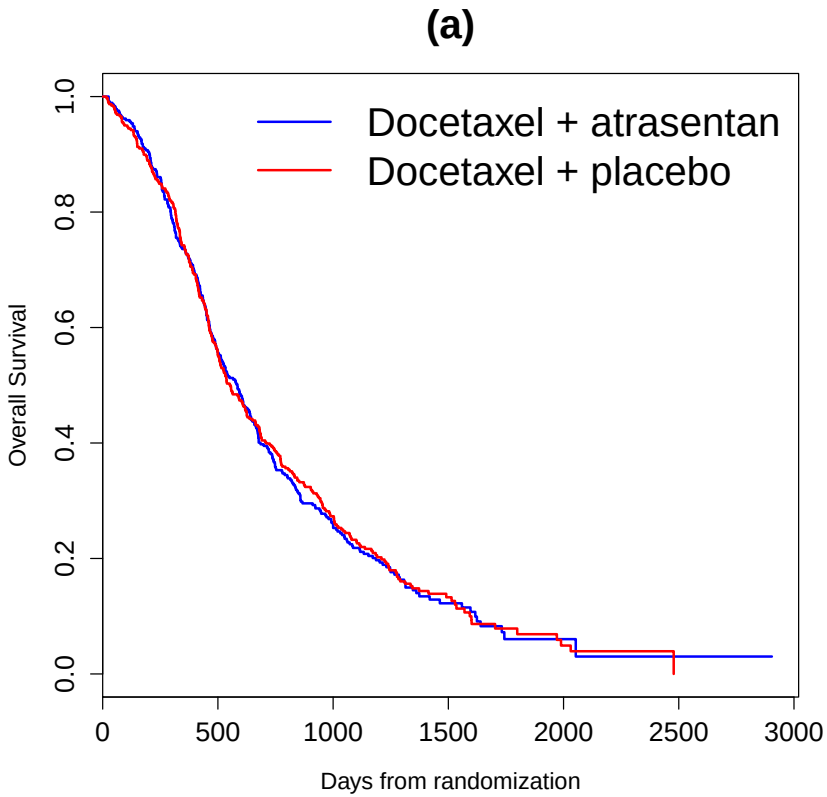

Figure 1: Kaplan-Meier survival curve of overall survival in castration-resistant prostate cancer patients by treatment received.

Tailored therapies

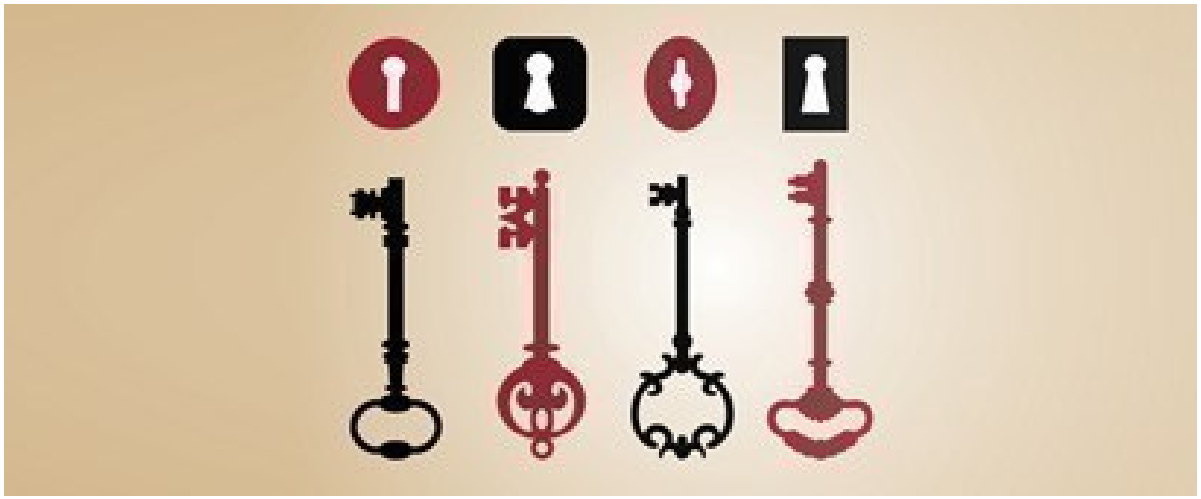

"The right treatment for the right patient (at the right time)"

Single decision: background

- T: outcome of interest
- A: binary treatment options, $A \in \{-1, 1\}$,
- X: baseline variables, $X \in \mathbb{R}^p$,

Single decision: background

- Individualized treatment rule

$$d(X) : \mathbb{R}^p \to \{-1, 1\},$$

patient presenting with $X = x$ recommended $d(x)$.

- Example: if bone marker NTx level $>$ 75 percentile $\Rightarrow$ Docetaxel $+$ atrasentan

- Value: $V(d) = E^d(T)$; the average outcome if all patients are assigned treatment according to d.

- Optimal rule d^* satisfies $d^* \in \text{argmax}_d V(d)$.

Single decision: background

- Under standard causal inference assumptions,

$$d^*(x) = \text{sign}\{f^*(x)\},$$

where $f^*(x) = E(T|X = x, A = 1) - E(T|X = x, A = -1)$.

- Let $p_X(x)$ denote the distribution of X in the trial population, and $q_X(x)$ denote the distribution of X in the target population
  - the support of $q_X(x)$ is contained in the support of $p_X(x)$
  - the potential outcome mean given X is the same between the trial population and the target population

Developing robust and interpretable rules

- A black box decision rule: generalize the rule to a larger target population without introducing any bias
- A parsimonious and interpretable decision rule, e.g., a linear rule: possibly not the optimal linear rule in the target population, unless the optimal rule is linear itself.

Developing robust and interpretable rules

- Optimize a general criterion on assessing the quality of a treatment rule on the target population
- Easily deployed in clinical practice.

- E.g. clinical trials: susceptible to a lack of generalizability.
- Covariate distribution over patients in a trial data may not be representative of a future population targeted by a treatment rule.
- Also likely in other types of data.

The General Criterion

Assuming that we already know the optimal treatment rule d^*, we maximize the benefit function

$$\widetilde{\mathcal{B}}(d) = \widetilde{E}[W(X)I\{d(X) = d^*(X)\}],$$

- $W(X)$: a predefined and non-negative function
- $\widetilde{E}$ indicates that the expectation is taken with respect to the distribution of X in the target population.
- Optimal treatment assigned: $d(x) = d^*(x)$

The General Criterion

$$\begin{aligned}\widetilde{\mathcal{B}}(d) &= \widetilde{E}[W(X)I\{d(X)=1\}|d^*(X)=1]P\{d^*(X)=1\} \\ &\quad +\widetilde{E}[W(X)I\{d(X)=-1\}|d^*(X)=-1]P\{d^*(X)=-1\}.\end{aligned}$$

- Choose linear decision rule that maximizes a lower bound of

$$\widetilde{\mathcal{B}}(d) = \widetilde{E}[W(X)I(\text{optimal treatment assigned}|X)],$$

regardless of X and regardless of which treatment is optimal.

The General Criterion: Choices of $W(X)$

- $W(x) = W_1(x) = E\{T|A = d^*(x), X = x\} - E\{T|A \neq d^*(x), X = x\} = |f^*(x)|$, where

$$f^*(x) = E(T|X = x, A = 1) - E(T|X = x, A = -1)$$

$\rightarrow$ maximize the value function in the target population.

- $W(X) = W_2(X) \equiv 1$

$\rightarrow$ minimize the mis-allocation rate of the optimal treatment in the target population.

Minimax Linear Decisions (MiLD)

- Goal: a high-quality rule for future patients to follow, using a dataset that may be subject to biases.
- Minimax Linear Decisions (MiLD)
  - A linear decision rule has the form of $d(x) = \text{sign}(x^\intercal \beta_1 + \beta_0)$.
  - $\widetilde{E}[W(X)I\{\text{sign}(X^\intercal \beta_1 + \beta_0) = j\} | d^*(X) = j]$ represents the expected benefit that would have obtained if the patients were to receive treatment j, whose optimal treatments would indeed be j in the target population, $j = \pm 1$
  - Guarantee that the expected benefit for either group of patients is not small

Minimax Linear Decisions (MiLD)

Objective:

$$\max_{\alpha,\beta_1,\beta_0} \alpha \ s.t. \quad \inf_{X\sim f_1} \widetilde{E}\{W(X)I(X^\intercal\beta_1+\beta_0\geq 0)|d^*(X)=1\}\geq\alpha, (1)$$

$$\inf_{X\sim f_{-1}} \widetilde{E}\{W(X)I(X^\intercal\beta_1+\beta_0< 0)|d^*(X)=-1\}\geq\alpha,$$

where f_1 and f_{-1} are the density of X in patients whose optimal treatment are 1 or -1 respectively.

Minimax Linear Decisions (MiLD)

Let $\tilde{q}_1(x)$ denote the density of X for patients with $d^*(X) = 1$ in the target population. Then

$$\begin{aligned} & \widetilde{E}\{W(X)I(X^{\mathsf{T}}\beta_1 + \beta_0 \geq 0)|d^*(X) = 1\} \\ = & \int W(x)I(x^{\mathsf{T}}\beta_1 + \beta_0 \geq 0)q_1(x)dx \\ \propto & \widetilde{P}(X^{\dagger\mathsf{T}}\beta_1 + \beta_0 \geq 0), \end{aligned}$$

where the density of $X^\dagger$ is proportional to $\tilde{q}_1(x^\dagger)W(x^\dagger)$.

Minimax Linear Decisions (MiLD)

Hence, (1) is equivalent to

$$\max_{\alpha,\beta_1,\beta_0} \alpha \quad s.t. \quad \inf_{X^\dagger \sim f_1^\dagger} \widetilde{P}(X^{\dagger\mathsf{T}}\beta_1 + \beta_0 \geq 0) \geq \alpha, \qquad (2)$$
$$\inf_{X^\dagger \sim f_{-1}^\dagger} \widetilde{P}(X^{\dagger\mathsf{T}}\beta_1 + \beta_0 < 0) \geq \alpha,$$

where $f_1^\dagger$ and $f_{-1}^\dagger$ are the density of $X^\dagger$ in patients with optimal treatment being 1 or -1 respectively.

- Problem: $f_1^\dagger$ and $f_{-1}^\dagger$ could be very different, and difficult to characterize based on the trial data.

Quantifying differences

Quantify the difference between the target population and the trial population in terms of the first two moment conditions.

Table 1: Moment-based conditions, $j = \pm 1$

	$E\{X^\dagger \mid d^*(X) = j\}$	$Cov\{X^\dagger \mid d^*(X) = j\}$
Trial population	$\mu_j^\dagger$	$\Sigma_j^\dagger$
Target population	$\tilde{\mu}_j^\dagger$	$\widetilde{\Sigma}_j^\dagger$

- Assume that the quantities in the target population belong to

$$\mathcal{U}_j^\dagger = \{(\tilde{\mu}_j^\dagger, \widetilde{\Sigma}_j^\dagger) : (\tilde{\mu}_j^\dagger - \mu_j^\dagger)^\intercal \widetilde{\Sigma}_j^\dagger (\tilde{\mu}_j^\dagger - \mu_j^\dagger) \leq \nu^2, \|\widetilde{\Sigma}_j^\dagger - \Sigma_j^\dagger\|_F \leq \rho_j\},$$

$j = \pm 1$, where $\nu \geq 0$ and $\rho_j \geq 0$ are known constants, and $\|\cdot\|_F$ is the Frobenius norm defined as $\|M\|_F^2 = \mathrm{Tr}(M^\intercal M)$.

Minimax Linear Decisions (MiLD)

(2) is consequently written as

$$\max_{\alpha,\beta_1,\beta_0} \quad \alpha \quad s.t. \quad \inf_{X^\dagger \sim (\tilde{\mu}_1^\dagger, \tilde{\Sigma}_1^\dagger) \in \mathcal{U}_1^\dagger} \widetilde{P}(X^{\dagger\mathsf{T}}\beta_1 + \beta_0 \geq 0) \geq \alpha, \qquad (3)$$

$$\inf_{X^\dagger \sim (\tilde{\mu}_{-1}^\dagger, \tilde{\Sigma}_{-1}^\dagger) \in \mathcal{U}_{-1}^\dagger} \widetilde{P}(X^{\dagger\mathsf{T}}\beta_1 + \beta_0 < 0) \geq \alpha.$$

- A linear decision rule that safeguards against the possible difference of the distribution of X between the trial and the target populations.

Solving Minimax Linear Decisions (MiLD)

- Employ minimax probability machine techniques (Lanckriet et al, 2013)
- Key step: generalized Chebychev inequality

$$\inf_{X^\dagger \sim (\tilde{\mu}_1^\dagger, \widetilde{\Sigma}_1^\dagger)} P(X^{\dagger\mathsf{T}}\beta_1 + \beta_0 \geq 0) \geq \alpha,$$

holds if and only if

$$\beta_0 + \tilde{\mu}_1^{\dagger\mathsf{T}}\beta_1 \geq \kappa(\alpha)\sqrt{\beta_1^\mathsf{T}\widetilde{\Sigma}_1^\dagger\beta_1},$$

where $\kappa(\alpha) = \sqrt{\alpha/(1-\alpha)}$

- $(\tilde{\mu}_j^\dagger, \widetilde{\Sigma}_j^\dagger)$ is unknown, $j = \pm 1$.

Solving Minimax Linear Decisions (MiLD)

A further simplified objective function can be obtained as

$$\min_{\beta_1} \sqrt{\beta_1^\intercal(\Sigma_1^\dagger + \rho_1 I_p)\beta_1} + \sqrt{\beta_1^\intercal(\Sigma_{-1}^\dagger + \rho_{-1} I_p)\beta_1}, \tag{4}$$

such that $\beta_1^\intercal(\mu_1^\dagger - \mu_{-1}^\dagger) = 1$.

- Eliminate the equality constraint
- $\beta_1 = \beta_{10} + Fu$, where $u \in \mathbb{R}^{p-1}$, $\beta_{10} = (\mu_1^\dagger - \mu_{-1}^\dagger)/\|\mu_1^\dagger - \mu_{-1}^\dagger\|_2^2$, and $F \in \mathbb{R}^{p\times(p-1)}$ is an orthogonal matrix whose columns span the subspace of vectors orthogonal to $\mu_1^\dagger - \mu_{-1}^\dagger$.

Implementing MiLD

Step 1. Estimate $d^*(x)$ using a nonparametric method with the trial data, denoted by $\hat{d}(x)$.

Step 2. Estimate $(\tilde{\mu}_j^\dagger, \widetilde{\Sigma}_j^\dagger)$ using the initial estimate $\hat{d}(x)$, denoted by $(\hat{\mu}_j^\dagger, \widehat{\Sigma}_j^\dagger)$.

Step 3. Implement MiLD based on the estimated $(\hat{\mu}_j^\dagger, \widehat{\Sigma}_j^\dagger)$.

Implementing MiLD: Survival Data

- $T = \min(\tau, \tilde{T})$, where $\tilde{T}$ denotes survival time, and τ is the end of the study; C: the censoring time.
- $\{Y_i = T_i \wedge C_i, \Delta_i = I(T_i \leq C_i), X_i, A_i\}, \quad i = 1, ..., n$, where $\Delta = I(T \leq C)$ denotes the censoring indicator.
- Random forest survival tree to estimate $E(T|X, A)$
- Techniques from importance sampling to estimate $(\mu_j^\dagger, \Sigma_j^\dagger)$ with a general weight $W(x)$
- Sensitivity analysis on different combinations of (ν, ρ_j).

Implementing MiLD: Survival Data

- $\mu_j^\dagger$ can be estimated by

$$\hat{\mu}_j^\dagger = \frac{\sum_{i=1}^n X_i W(X_i) I\{\hat{d}(X_i) = j\}}{\sum_{i=1}^n W(X_i) I\{\hat{d}(X_i) = j\}};$$

- $\Sigma_j^\dagger$ can be estimated by

$$\hat{\Sigma}_j^\dagger = \sum_{i=1}^n \left[\frac{W(X_i) I\{\hat{d}(X_i) = j\}}{\sum_{i=1}^n W(X_i) I\{\hat{d}(X_i) = j\}} \right]^2 (X_i - \hat{\mu}_j^\dagger)^\intercal (X_i - \hat{\mu}_j^\dagger).$$

Theoretical Results

- Let β_1^* be the unique solution to (1) so that $\beta_1^* X$ is the optimal linear rule maximizing $\widetilde{\mathcal{B}}(d)$ for the target population. β_1^* lies in the interior of a compact set $\boldsymbol{B}$.
- $\|\hat{f} - f^*\|_2 = O_p(r_n)$, where $\|f\|_2 = E\{f(X)^2\}^{1/2}$.
- Margin condition: there exist $K_1, \gamma > 0$ such that for all $t > 0$

$$\widetilde{P}(|f^*(X)| \leq t) \leq K_1 t^{\gamma}.$$

- Under certain assumptions, it holds that

$$|\hat{\beta}_1 - \beta_1^*| = O_p(r_n^{\frac{2\gamma}{\gamma+2}} + n^{-1/2}).$$

Simulation Setup: Scenario 1

- The first half patients from $X_1, X_2 \sim N(1,1), X_3, \ldots, X_{10} \sim N(0,1)$ and the second half patients with $X_1, \ldots, X_{10}$ from $N(0,1)$.

- $T = \min(\widetilde{T}, \tau)$, where $\tau = 0.5$ and

$$\widetilde{T} = \exp[\exp\{0.6 * X_1 - 0.8 * X_2 + A * c(X)\}]\epsilon.$$

  Here, $c(X) = 1$ for the first half patients and $c(X) = -1$ for the other half, and $\epsilon \sim \exp(1)$.

- Censoring time C is generated from Uniform$[0,1]$.

- The optimal decision boundary is $d^*(X) = \text{sign}(X_1 + X_2 - 1)$.

Simulation Setup: Scenario 2

- $X_1, \ldots, X_{10} \sim N(0, 1)$.
- $T = \min(\widetilde{T}, \tau)$, where $\tau = 4$ and

  $$\lambda_{\widetilde{T}}(t|X, A) = \exp[0.6X_1 + 0.8X_2 - 1 + \{2X_1 + 3(X_2 + 1)^2 - 2\}A$$

- Censoring time C is generated from Uniform$[0, 5]$.
- The optimal decision boundary is $d^*(X) = -\text{sign}\{2X_1 + 3(X_2 + 1)^2 - 2\}$.

Simulation Setup: Scenario 3

- $X_1, \ldots, X_{10} \sim N(0, 1)$.
- $T = \min(\widetilde{T}, \tau)$, where $\tau = 4$ and

$$\log(\widetilde{T}) = X_1 + X_2 + 1 + A(2X_1^3 + 2X_2 + 0.5) + N(0, 1).$$

- Censoring time C is generated from

$$\log(C) \sim X_1 + X_2 + X_3 + N(0, 1).$$

- The optimal decision boundary is $d^*(X) = \text{sign}(2X_1^3 + 2X_2 + 0.5)$.

Simulation Setup: Introduce mismatch between two populations

Scen. 1' Let the proportion of patients with $d^*(x) = 1$ is $1/3$ in the trial population, and $1/2$ in the target population

Scen. 2' $X_1 \sim N(-0.25, 1.5)$ and other covariates following $N(0, 1.5)$ in the trial data

Scen. 3' Selection bias: covariates predictive of participation in trial could be predictive of treatment effects

$$\text{logit}\{P(\text{Participating the trial}|X)\} = -2X_1^3 - 1.$$

Patients with $d^*(x) = 1$ are less likely to participate in the trial.

Simulation setup

- $n = 250, 500$ and 1000.
- Methods for comparison
  - Cox regression
  - Inverse weighed outcome weighted learning: minimize $\mathbb{P}_n \left[Y \min\{1 - Af(X), 0\} / \{P(A|X)\hat{S}_C(Y|A, X)\} \right]$, where $f(X)$ is in a linear form, $\mathbb{P}_n$ denotes the empirical averages, and $\hat{S}_C(t|A, X)$ is the estimated censoring probability conditional on patients characteristics.
  - MiLD-P ($W(X) = 1$) and MiLD-V ($W(X) = |f^*(X)|$) with different set of parameters (ρ, ν_j).
- Different methods are validated on the large testing set of size 10000; 500 replications.

Simulation results: misallocation rates

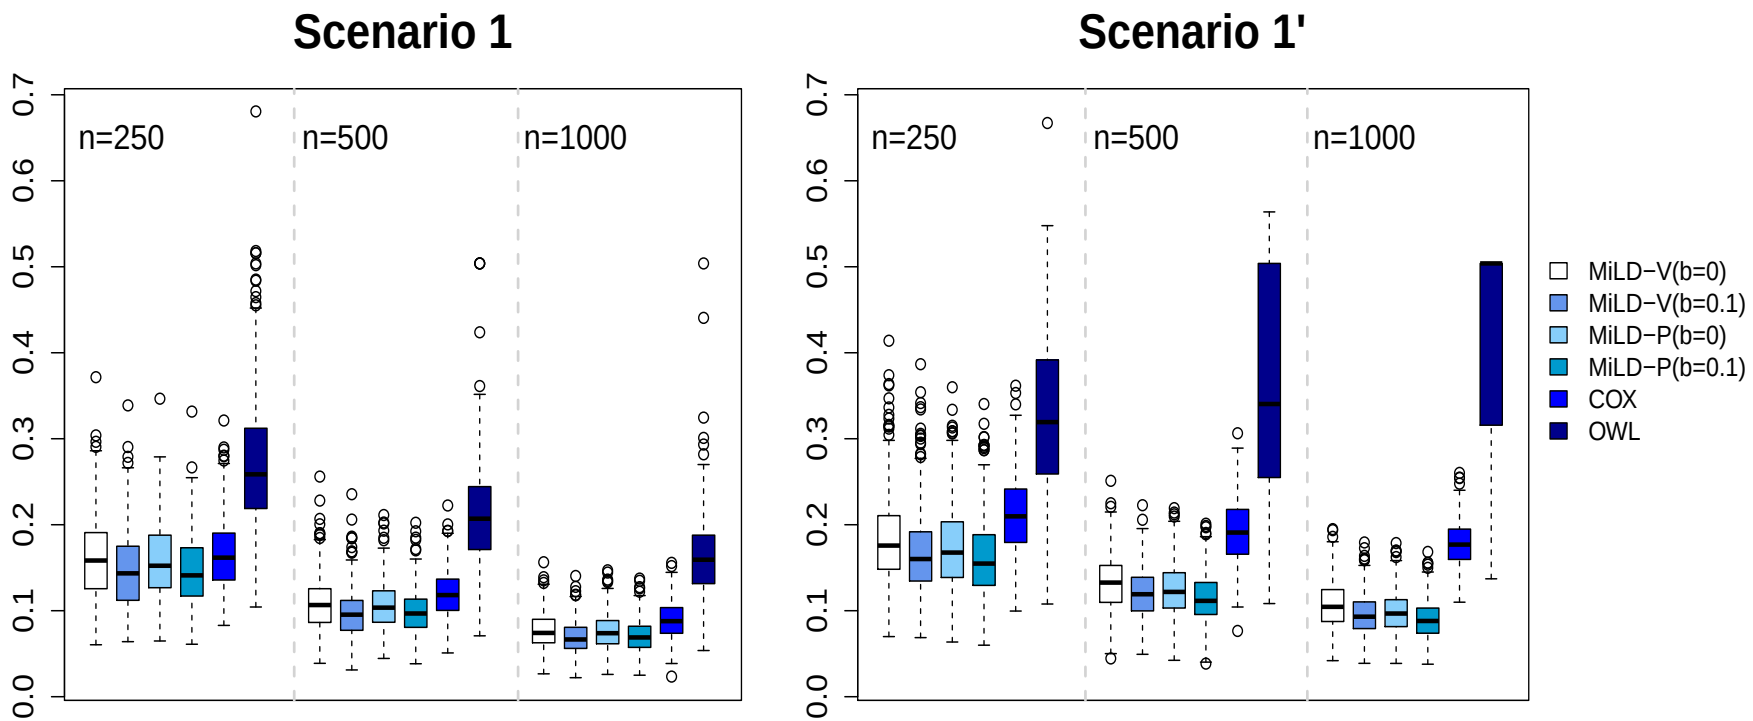

Simulation results: misallocation rates

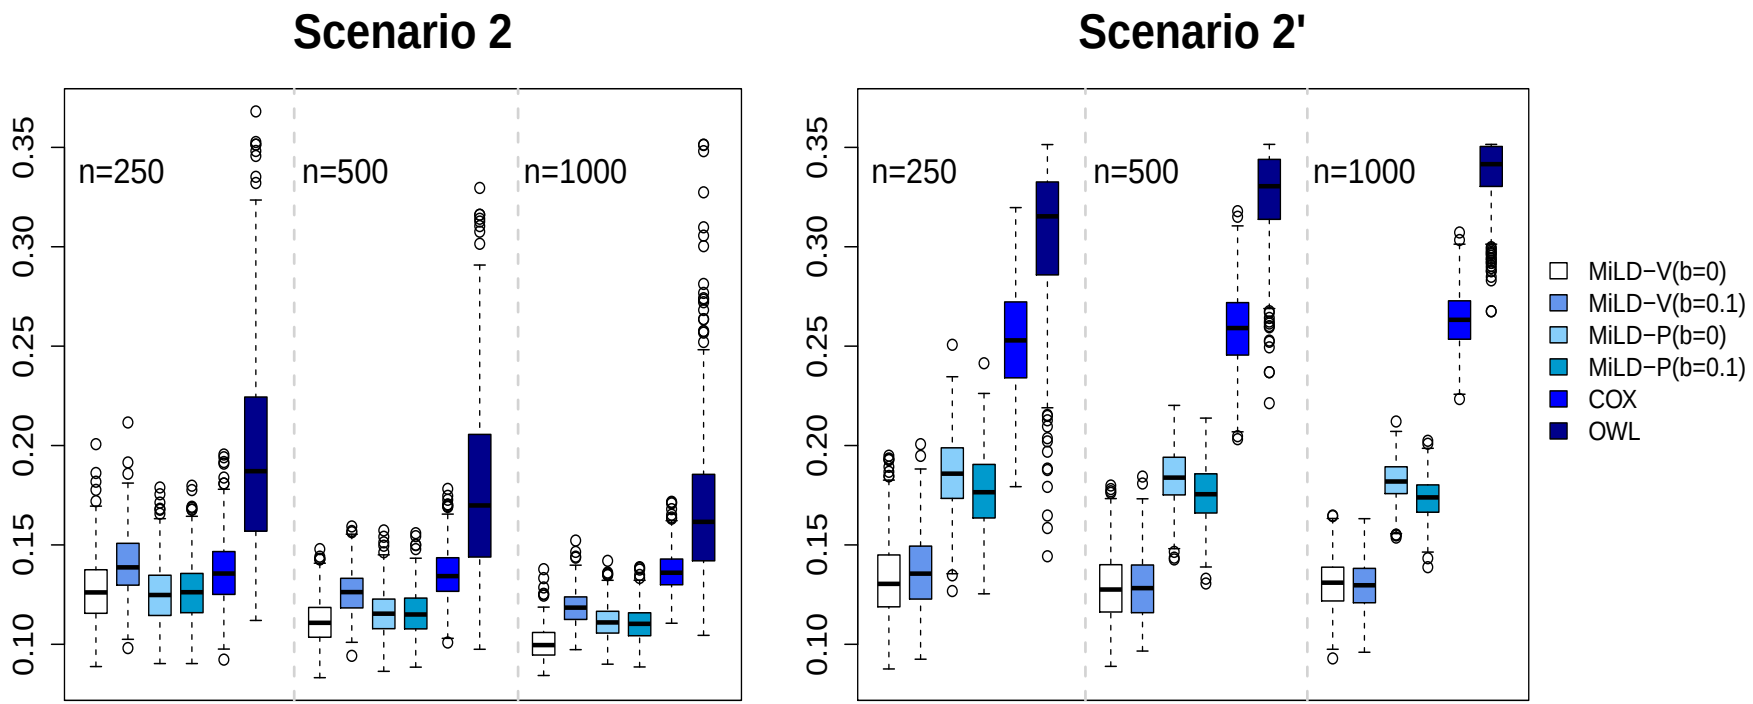

Simulation results: misallocation rates

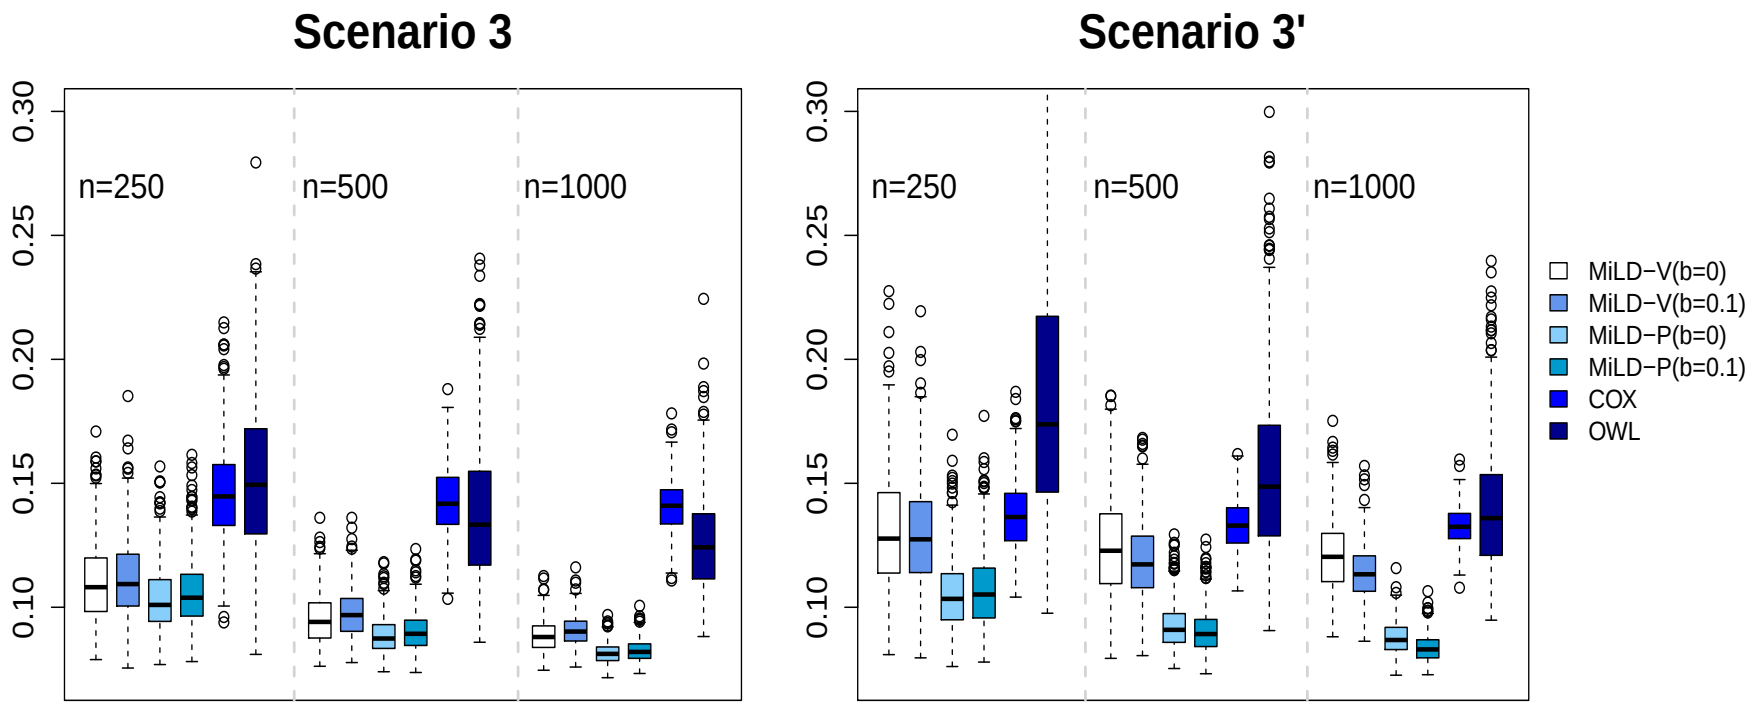

Data analysis: SWOG 0421 study

- 1038 men with metastatic castration-resistant prostate cancer
- Docetaxel administered every 21 days at a dose of 75 mg/m^2 with or without the bone targeted atrasentan for up to 12 cycles
- Co-primary outcome: progression-free survival and overall survival
- Data from 751 patients for anlayis; 371 patients are in the docetaxel + atrasentan arm and 380 patients are in the docetaxel + placebo arm
- 10 covariates: age, serum prostate-specific antigen (PSA), indicator of bisphosphonate usage, indicator of extraskeletalmets, indicator of pain at basline, indicator of performance status, and 4 bone marker levels

Data analysis

Table 2: Mean (s.e.) cross-validated values (days)

COX	OWL	MiLD-V		MiLD-P	
		$b = 0$	$b = 0.1$	$b = 0$	$b = 0.1$
749.2 (69.8)	711.1 (71.2)	764.6 (70.2)	765.0 (68.9)	753.5 (69.4)	753.7 (69.4)

"s.e." denotes standard errors.

Data analysis

Table 3: Coefficients for the estimated linear decision rules by MiLD-V and MiLD-P using the SWOG0421 data

	MiLD-V	MiLD-P
Intercept	0.021	0.037
Age	0.538	0.150
Baseline serum PSA	-0.510	-0.734
Bisphosphonate usage (YES = 1)	0.787	1.235
Metastatic disease beyond the bones (YES = 1)	-0.301	0.008
Pain (YES = 1)	-0.272	0.050
Performance Status ('2-3' = 1)	0.515	0.796
BAP	0.382	0.972
CICP	-0.023	0.302
NTx	0.137	0.657
PYD	-0.339	-0.326

Data analysis

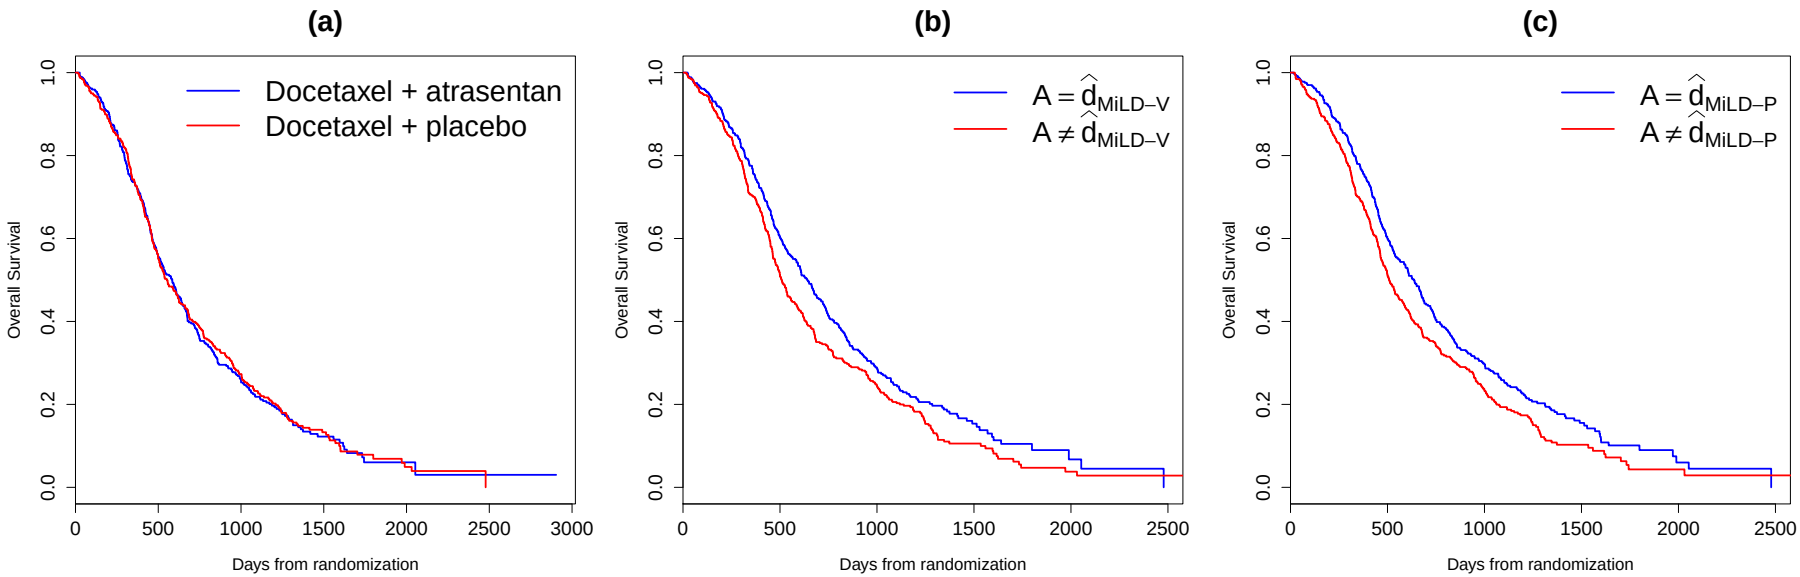

Figure 2: Kaplan-Meier survival curve of overall survival in castration-resistant prostate cancer patients: (a) by treatment received; (b) by accordance between treatment recommended by MiLD-V and treatment received; (c) by accordance between treatment recommended by MiLD-P and treatment received.

Open questions

- Can be applied to other types of data.
- Better tools for higher dimensional data.
- Multi-category treatments.
- Multi-stage treatments.

Acknowledgement

Donglin Zeng, UNC at Chapel Hill
Catherine Tangen, FHCRC
Michael Leblanc, FHCRC

Thank You!