Estimation of Optimal Individualized Treatment Rules for Multistate Disease Processes

Giorgos Bakoyannis

Department of Biostatistics and Health Data Science
Fairbanks School of Public Health and School of Medicine
Indiana University

Seminar at the Department of Statistics
Athens University of Economics and Business
Athens, Greece
19 November 2021

Introduction

- Many chronic diseases evolve through multiple clinical states (**multistate disease processes**)

- There is a variety of methods to evaluate the **overall** effectiveness of treatments on multistate disease processes:
  - (i) Aalen–Johansen estimator[1]
  - (ii) Simultaneous confidence bands by Bluhmki et al. (2018)[2]
  - (iii) Two-sample tests by Bakoyannis (2020)[3]

- No methods for estimating optimal **patient-tailored treatment rules** for multistate disease processes

---

[1] Aalen, O.O. and Johansen, S., 1978. An empirical transition matrix for non-homogeneous Markov chains based on censored observations. *Scandinavian Journal of Statistics* **5**, 141–150

[2] Bluhmki, T., Schmoor, C., Dobler, D., Pauly, M., Finke, J., Schumacher, M. and Beyersmann, J., 2018. A wild bootstrap approach for the Aalen–Johansen estimator. *Biometrics* **74**, 977–985

[3] Bakoyannis, G., 2020. Nonparametric tests for transition probabilities in nonhomogeneous Markov processes. *Journal of Nonparametric Statistics* **32**, 131–156

Examples of multistate processes under treatment

A.

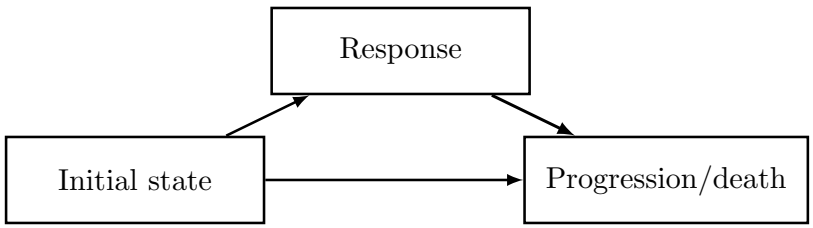

B.

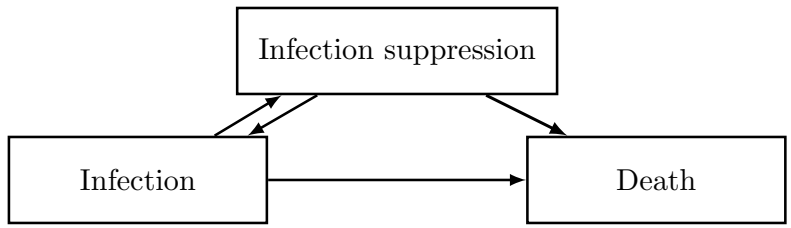

C.

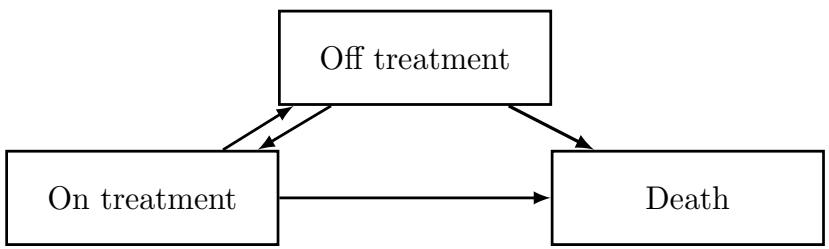

Example A: Probability of being in response

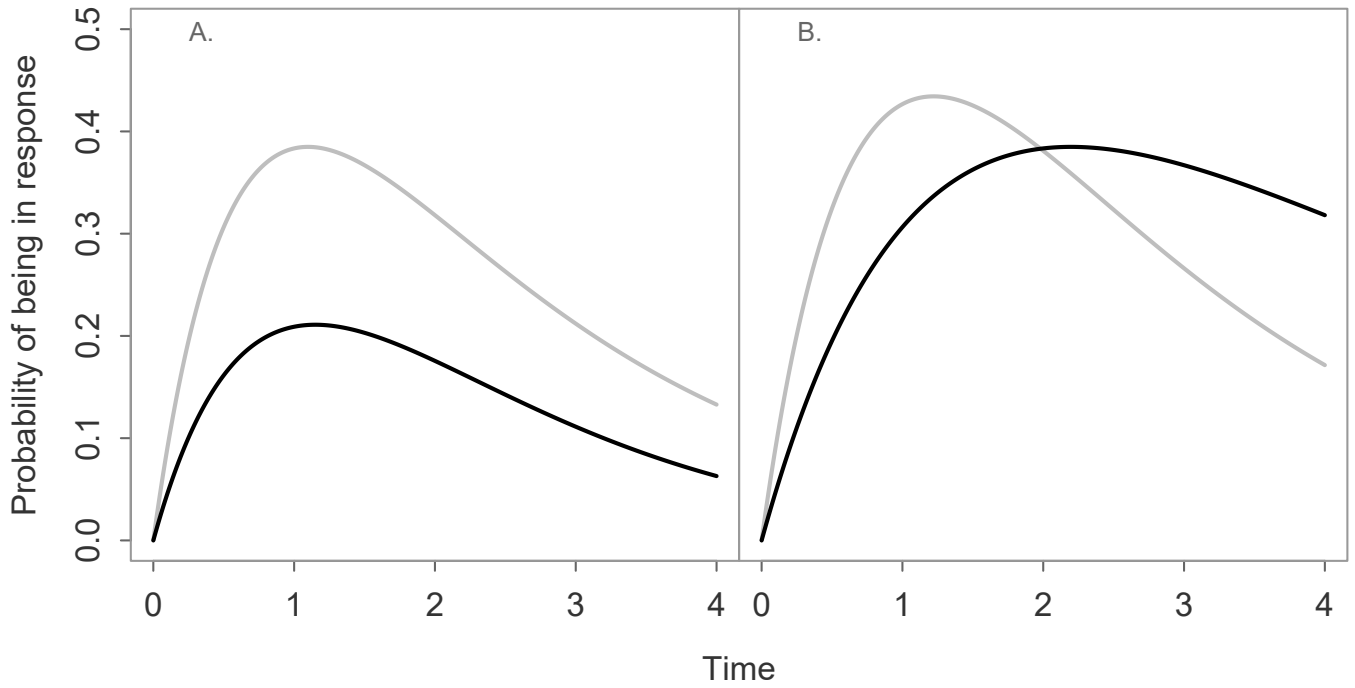

- The probability of being in **response** (**non-monotonic function of time**) provides a more direct insight into treatment effect compared to crude events such as overall survival or progression-free survival
- **Response** is an outcome which is endorsed by the FDA for drug evaluation in cancer trials

Metastatic squamous-cell carcinoma trial

- Randomized controlled trial on metastatic squamous-cell carcinoma of the head and neck
- Clinical trial with two interventions:
  - (i) Chemotherapy alone
  - (ii) Chemotherapy + panitumumab
- Multistate disease process:

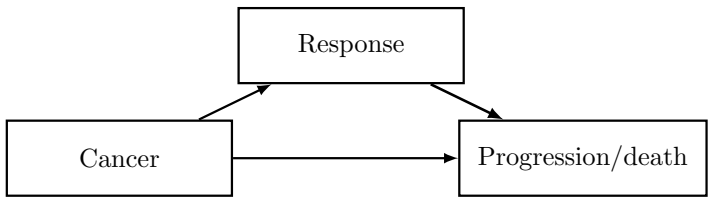

- The main outcome in this analysis is response to treatment
- 243 patients where randomized assigned in the chemotherapy alone group and 236 in the chemotherapy + panitumumab group

Event probabilities and 95% confidence bands

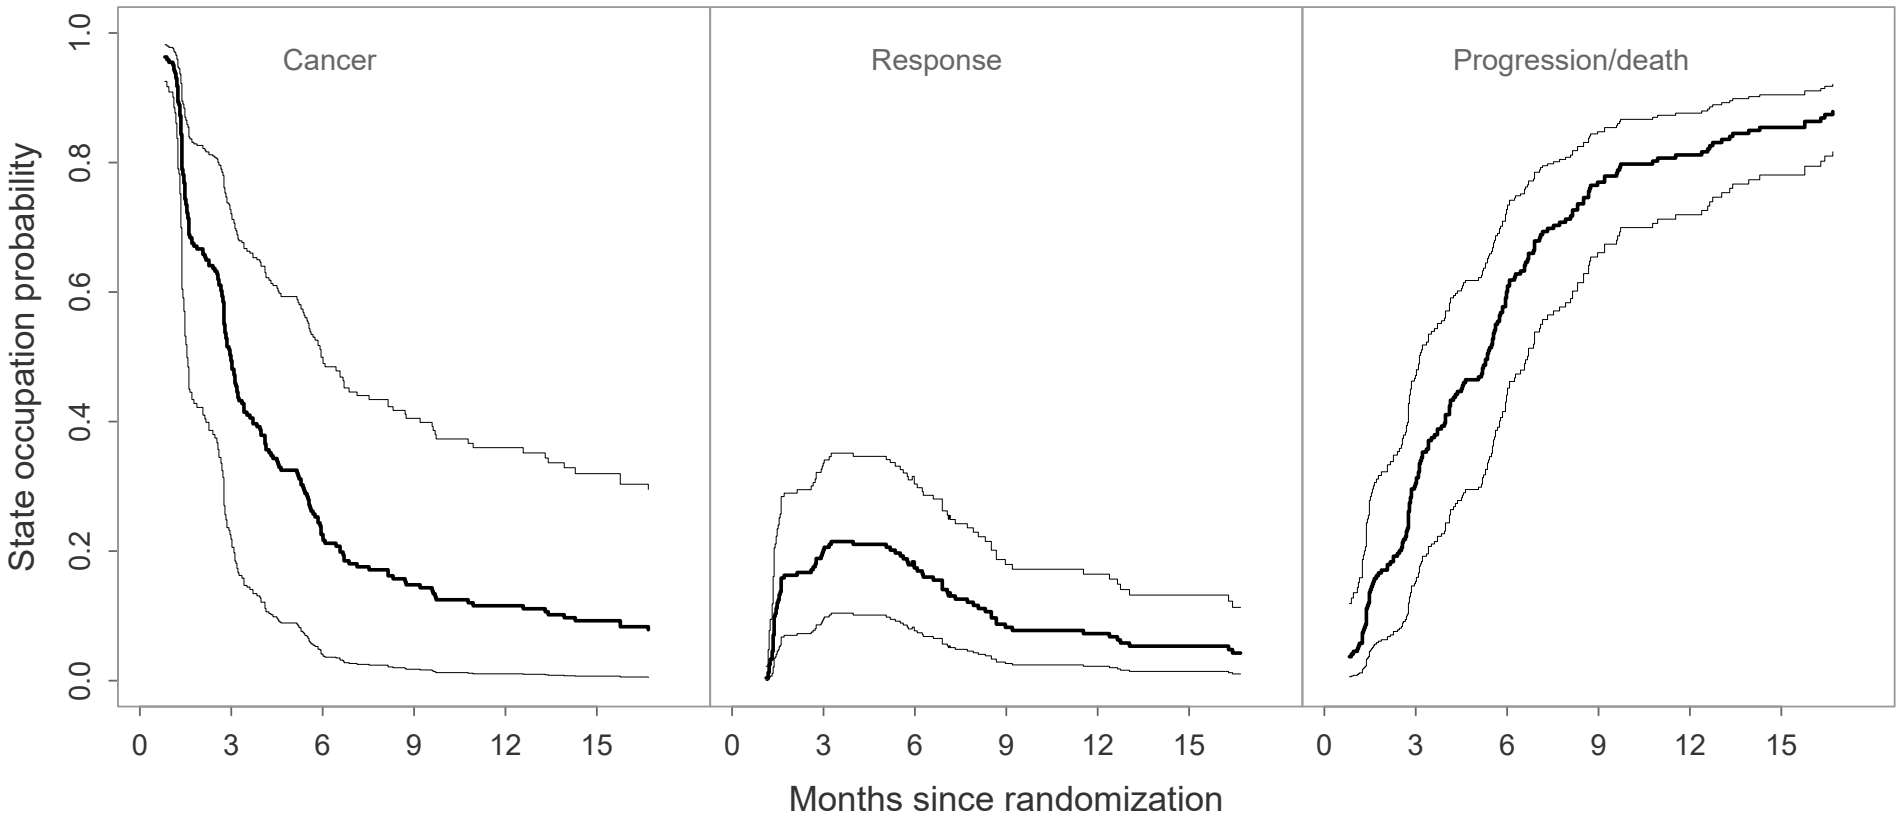

Probability of response by treatment group

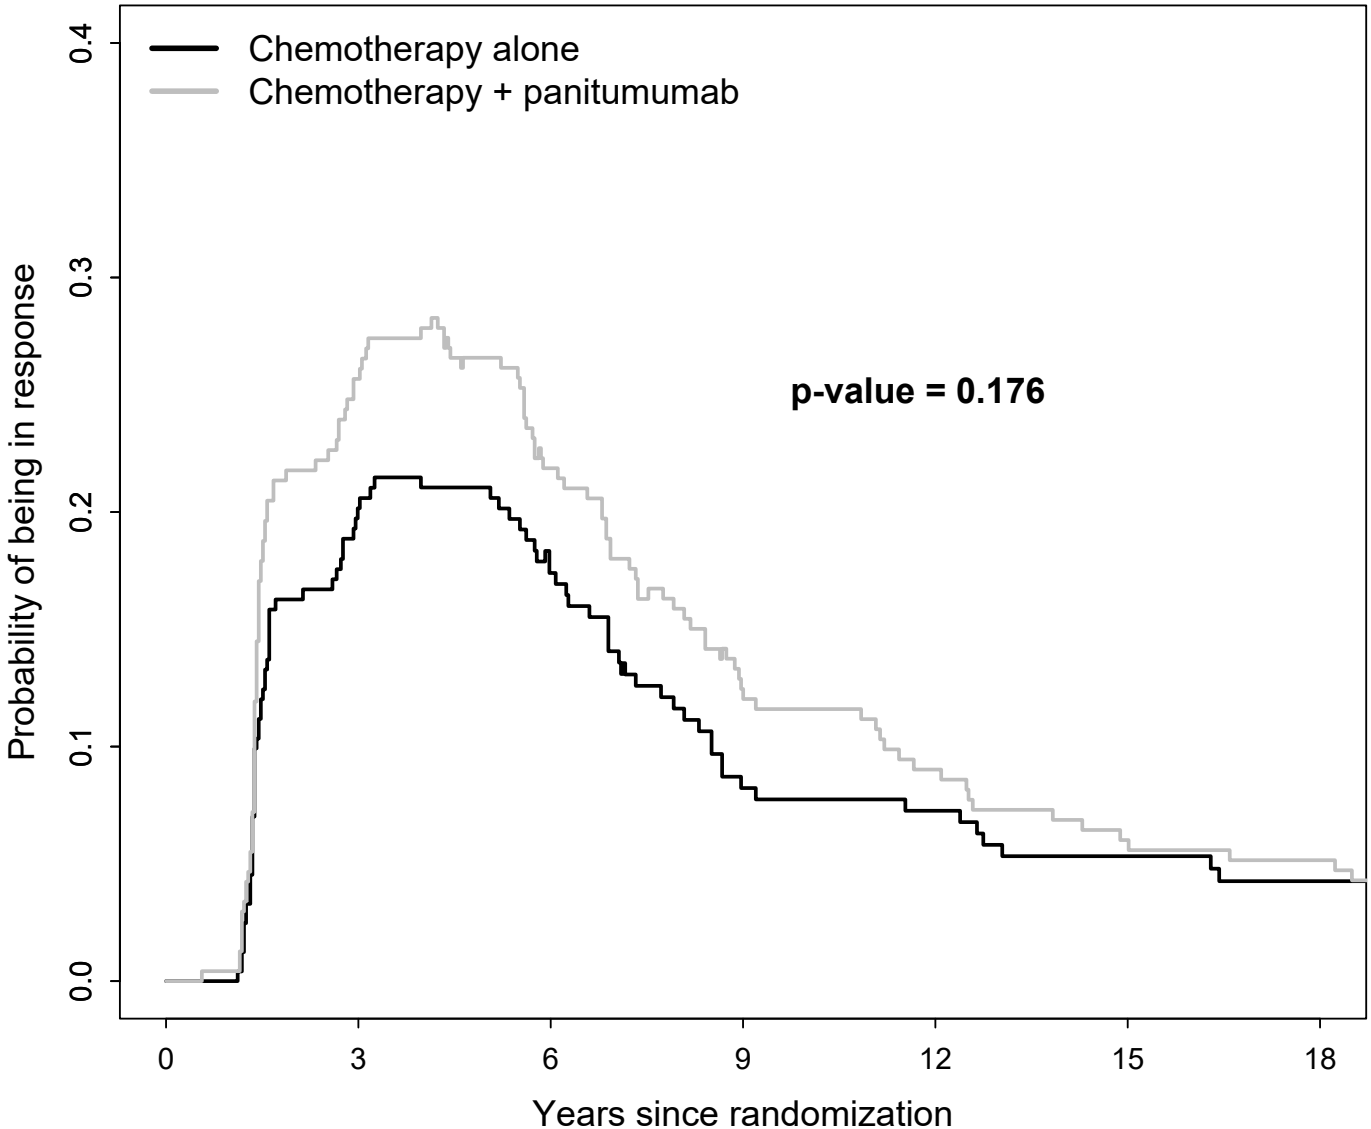

[3] Bakoyannis, G., 2020. Nonparametric tests for transition probabilities in nonhomogeneous Markov processes. *Journal of Nonparametric Statistics* **32**, 131–156

Individualized treatment rules

- Are these results discouraging regarding the potential of chemotherapy + panitumumab?
- Classical two-sample comparisons do not take into account **patient heterogeneity**.
- However, a treatment option that works for one individual **may not work** for another.
- There is a possibility to achieve better health outcomes by providing treatments that are **tailored to the individual patient**.

Traditional paradigm: One size fits all

Everyone receives treatment #1

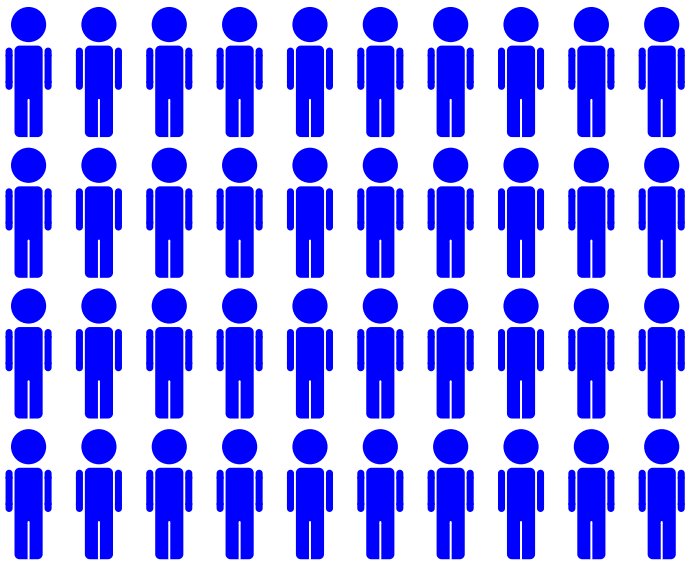

vs

Everyone receives treatment #2

Modern paradigm: Patient-tailored treatments

Patient-tailored treatment

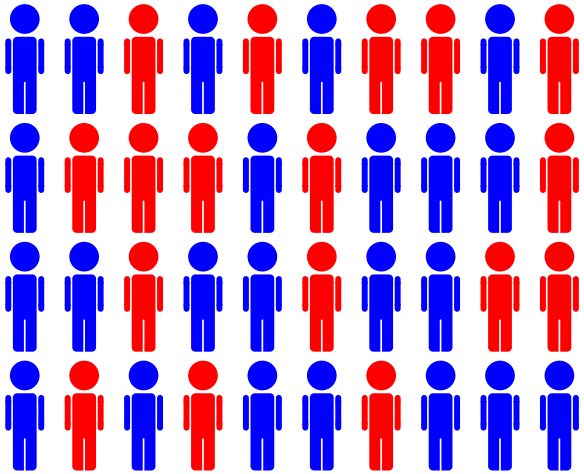

vs

Everyone receives treatment #1

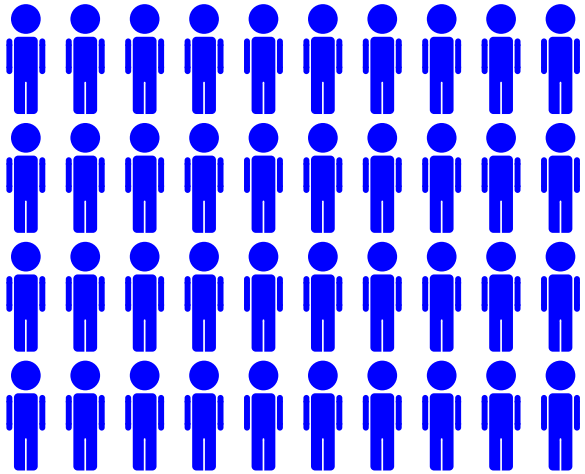

Patient-tailored treatment

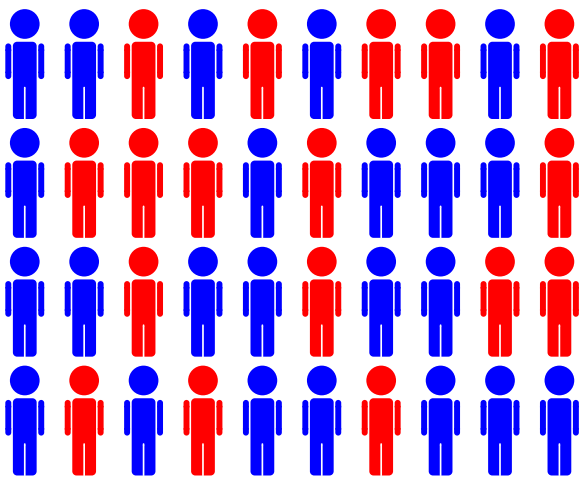

vs

Everyone receives treatment #2

Estimating optimal individualized treatment rules

- Estimating optimal individualized treatment rules (ITRs) is challenging because of:
  - (i) Complex nonlinear associations between different variables and the disease of interest
  - (ii) Complex nonlinear and high order interactions between treatment and other variables

- Modern machine learning methods that tackle the above challenges have been employed for the estimation of optimal ITRs with simple outcomes (e.g. continuous, binary, survival)

- In this work I develop the **first** method for **complex multistate disease processes**

- The method utilizes **support vector machines**

Multistate processes

- Consider a continuous time non-homogeneous multistate process $\{X(t) : t \in [0, \tau]\}$, for $\tau \in (0, \infty)$, with a finite state space $\mathcal{S} = \{1, \ldots, S\}$
- The marginal behavior of the process can be described by the **state occupation probabilities**

$$P_{0,j}(t) \equiv P(X(t) = j), \quad j \in \mathcal{S}, \quad t \in [0, \tau]$$

- Inference about $P_{0,j}(t)$ does not require Markov assumptions[4]
- If the state $j \in \mathcal{S}$ corresponds to the response state, then we define the **response process**

$$Y(t) \equiv I\{X(t) = j\}, \quad t \in [0, \tau]$$

[4] Datta, S. and Satten, G.A., 2001. Validity of the Aalen–Johansen estimators of stage occupation probabilities and Nelson–Aalen estimators of integrated transition hazards for non-Markov models. *Statistics & Probability Letters* **55**, 403–411.

Definitions

- Let
  - $Z \in \mathcal{Z} \subset \mathbb{R}^p$ be a vector of variables that are potentially useful for **tailoring treatment** to the individual patient
  - $A \in \{-1, 1\}$ be the treatment variable
  - T be the time to entering an absorbing state (e.g., death)
  - C be the (random) right censoring time (e.g., loss-to-follow-up time)

- The outcome of interest is **time spent at the response state** by time τ, i.e.

$$\int_0^\tau Y(t) dm(t),$$

where $m(t) = t$ induces the lebesgue measure on the Borel σ-algebra on $[0, \tau]$

Can we just use methods for survival data?

- Let $\tilde{C}$ be the minimum of the random right censoring time and the maximum follow-up time, i.e. $\tilde{C} = C \wedge \tau$, where $a \wedge b = \min(a, b)$
- Survival analysis methods are applicable to situations where the **observed time** is the minimum between the (uncensored) **time of interest** and the the right censoring time, i.e.

$$\tilde{T} = T \wedge \tilde{C}$$

- Under right censoring, the **observed time** spent at the response state by time τ is

$$\int_0^\tau Y(t) I(C \geq T \wedge t) dm(t)$$

- If $C < T \wedge \tau$, $\int_0^C Y(t) dm(t) < C$, and $\int_C^\tau Y(t) dm(t) > 0$, then

$$0 \leq \int_0^\tau Y(t) I(C \geq T \wedge t) dm(t) < \int_0^\tau Y(t) dm(t) \wedge \tilde{C},$$

and thus survival analysis methods are **not applicable** here

Individualized treatment rules

- An ITR d is a map

$$d : \mathcal{Z} \mapsto \{-1, 1\}$$

- Define the sign function

$$\text{sgn}(x) = \begin{cases} 1, & \text{if } x \geq 0 \\ -1, & \text{otherwise} \end{cases}$$

and let WBC stand for white blood cell count

- ITR example #1[5]
If $\text{age} + 8.7 \times \log(\text{WBC}) - 60 \geq 0$ then give treatment 1, otherwise give treatment -1:

$$d(z) = \text{sgn}\{\text{age} + 8.7 \times \log(\text{WBC}) - 60\}$$

[5] Tsiatis, A.A., Davidian, M., Holloway, S.T. and Laber, E.B., 2019. Dynamic Treatment Regimes: Statistical Methods for Precision Medicine. CRC press.

Individualized treatment rules (cont.)

- ITR example #2[5]
  If $\text{age} < 50 \;\; \& \;\; \text{WBC} < 10$ then give treatment 1, otherwise give treatment -1:

$$d(z) = 2 \times I(\text{age} < 50 \;\; \& \;\; \text{WBC} < 10) - 1$$

- Any binary classification rule can be expressed as

$$d(z) = \text{sgn}\{f(z)\}$$

  for some measurable function $f : \mathcal{Z} \mapsto \mathbb{R}$

- The **goal** is to estimate the **optimal** decision function f

- How do we define optimality?
  Find f such that $d(z) = sgn\{f(z)\}$ maximizes some **benefit function** *for any $z \in \mathcal{Z}$*

Definitions (cont.)

- Let $\pi_0 = P(A = 1)$
  - In randomized clinical trials $\pi_0 = 0.5$
- Estimating optimal ITRs is a **causal inference** problem
  - We seek the best ITR in an effort to **cause** the best possible health outcome
- Need to utilize the **potential outcomes** approach
- Let $Y^*(t; a)$ be the response status at time $t \in [0, \tau]$ if the patient received treatment $a \in \{-1, 1\}$ (regardless of the actual treatment received)

Basic causal assumptions

Causal assumptions

A1. Stable unit treatment value assumption:
$Y(\cdot) = Y^*(\cdot; 1)I(A = 1) + Y^*(\cdot; -1)I(A = -1)$

A2. $\{Y^*(\cdot; 1), Y^*(\cdot; -1)\} \perp\!\!\!\perp A$

A3. Positivity assumption: $\pi_0 \in [c_1, c_2]$, with $0 < c_1 < c_2 < 1$

- In a randomized clinical trial, assumptions A2 and A3 and are automatically satisfied

Value functions (i.e. benefit functions)

- For a given ITR d, let

$$Y^*(t;d) = Y^*(t;1)I\{d(Z)=1\} + Y^*(t;-1)I\{d(Z)=-1\}$$

be the potential response status at time $t \in [0,\tau]$ if the patient received treatment according to d (regardless of the actual treatment received)

- Define the **value function** (i.e. benefit function)

$$\mathcal{V}(d) = E\left\{\int_0^\tau Y^*(t;d)dm(t)\right\}$$

  * $\mathcal{V}(d)$ is the potential **expected time spent in the response state** under ITR d by time τ

Optimal ITRs

- An optimal ITR d^* is a maximizer of the value function, i.e.

$$d^* \in \arg\max_d \mathcal{V}(d).$$

- Under assumptions A1–A3 and the independent right censoring assumption, the value function can be expressed in terms of the observable data as

$$\mathcal{V}(d) = E\left[\int_0^\tau \frac{Y_1(t)I(C_1 > T_1 \wedge t)I(A_1 = d(Z_1))}{\exp\{-\Lambda_0(\tilde{T}_1 \wedge t)\}\{A_1\pi_0 + (1-A_1)/2\}} dm(t)\right],$$

where $\Lambda_0(t)$ is the cumulative hazard function of the right censoring variable C at time t and $\tilde{T} = T \wedge C \wedge \tau$

Optimal ITRs (cont.)

- Any ITR of the form $d : \mathcal{Z} \mapsto \{-1, 1\}$ can be expressed as $d(z) = \text{sgn}\{f(z)\}$
- It is easy to see that

$$\begin{aligned}\mathcal{V}(d) &= E\left[\int_0^\tau \frac{Y_1(t)I(C_1 > T_1 \wedge t)}{\exp\{-\Lambda_0(\tilde{T}_1 \wedge t)\}\{A_1\pi_0 + (1-A_1)/2\}} dm(t)\right] \\ &\quad -E\left(\left[\int_0^\tau \frac{Y_1(t)I(C_1 > T_1 \wedge t)}{\exp\{-\Lambda_0(\tilde{T}_1 \wedge t)\}\{A_1\pi_0 + (1-A_1)/2\}} dm(t)\right] I(A_1 \neq d(Z_1))\right)\end{aligned}$$

- Then the optimal ITR is the **minimizer** of the **risk function**

$$\begin{aligned}\mathcal{R}(f) &= E\left(\left[\int_0^\tau \frac{Y_1(t)I(C_1 > T_1 \wedge t)}{\exp\{-\Lambda_0(\tilde{T}_1 \wedge t)\}\{A_1\pi_0 + (1-A_1)/2\}} dm(t)\right] I(A_1 \neq \text{sgn}(f(Z_1)))\right) \\ &= E\left(\left[\int_0^\tau \frac{Y_1(t)I(C_1 > T_1 \wedge t)}{\exp\{-\Lambda_0(\tilde{T}_1 \wedge t)\}\{A_1\pi_0 + (1-A_1)/2\}} dm(t)\right] \times \right. \\ &\quad \left. I(A_1 f(Z_1) < 0)\right)\end{aligned}$$

Discontinuity and nonconvexity of the risk function

- Minimizing the **empirical version** of $\mathcal{R}(f)$ is challenging computationally as it involves a discontinuous and nonconvex function of f

- To alleviate, we follow the paradigm of outcome weighting learning[6] and support vector machines, and utilize the **hinge loss** $\phi(x) = \max(0, 1 - x)$ which leads to the **surrogate risk**

$$\mathcal{R}_\phi(f) = E\left(\left[\int_0^\tau \frac{Y_1(t)I(C_1 > T_1 \wedge t)}{\exp\{-\Lambda_0(\tilde{T}_1 \wedge t)\}\{A_1\pi_0 + (1 - A_1)/2\}} dm(t)\right] \phi(A_1 f(Z_1))\right),$$

[6] 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**, 1106–1118.

Hinge loss vs discontinuous loss*

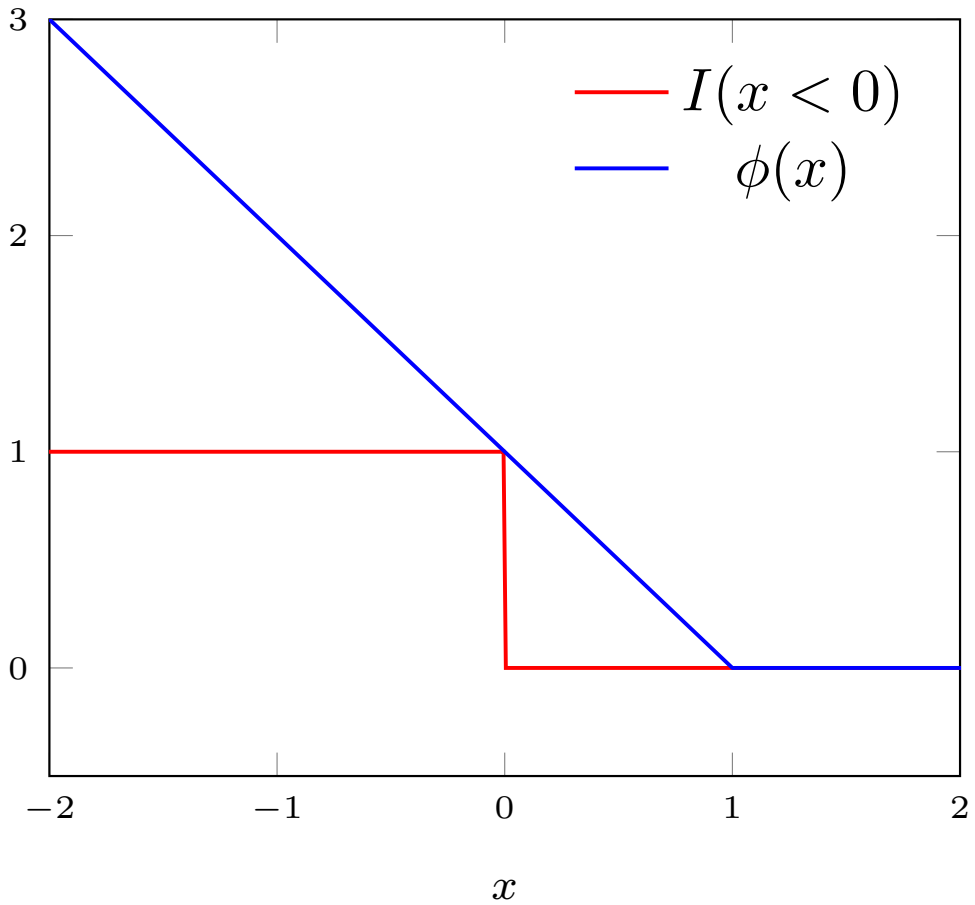

*Note that $I(x < 0) \leq \phi(x)$, $x \in \mathbb{R}$

Estimation of optimal ITR

- The surrogate loss $\mathcal{R}_\phi$ can be estimated as

$$\hat{\mathcal{R}}_\phi(f) = \frac{1}{n}\sum_{i=1}^{n}\left[\int_0^\tau \frac{Y_i(t)I(C_i > T_i \wedge t)}{\exp\{-\hat{\Lambda}_n(\tilde{T}_i \wedge t)\}\{A_i\hat{\pi}_n + (1-A_i)/2\}} dm(t)\right]\phi(A_i f(Z_i))$$

  where

$$\hat{\Lambda}_n(t) = \int_0^t \frac{\sum_{i=1}^n dN_i(u)}{\sum_{i=1}^n Y_i(u)}, \qquad t \in [0,\tau],$$

  with $N_i(t) = (1-\Delta_i)I(\tilde{T}_i \le t)$ and $Y_i(t) = I(\tilde{T}_i \ge t)$, and $\hat{\pi}_n = n^{-1}\sum_{i=1}^n I(A_i = 1)$

- $\hat{\mathcal{R}}_\phi(f)$ is a continuous and convex function of f

Estimation of optimal ITR (cont.)

- The optimal decision function within a class of functions $\mathcal{F}$ can be estimated as

$$\hat{f}_n = \arg\min_{f \in \mathcal{F}} \left\{ \hat{\mathcal{R}}_\phi(f) + \lambda_n \| f \|^2 \right\},$$

  where λ_n is a penalty term depending on n and $\| \cdot \|$ is a norm on $\mathcal{F}$

- Mininimization of over the class of all measurable functions is **infeasible**

- Therefore, we need to use a **restricted class** $\mathcal{F}$

Estimation of optimal ITR (cont.)

- In this work we use either of the following **restricted classes**
  - Class of linear functions $\{f(\cdot) = \beta_0 + \langle \beta, \cdot \rangle : \beta_0 \in \mathbb{R}, \beta \in \mathbb{R}^p\}$, where $\langle \beta, z \rangle = \beta' z$ is the inner product on the Euclidean space
  - Reproducing kernel Hilbert space (RKHS) with kernel k, which is the completion of the space

$$\left\{ f(\cdot) = \sum_{j=1}^{m} \alpha_j k(\cdot, z_j) + \beta_0 : m \in \mathbb{N}, z_j \in \mathcal{Z}, \alpha_j \in \mathbb{R}, \beta_0 \in \mathbb{R} \right\}.$$

    Here we consider the RKHS with the Gaussian kernel, $k(z_1, z_2) = \exp(-\|z_1 - z_2\|^2/\sigma^2)$, $z_1, z_2 \in \mathcal{Z}$.

- Then the estimated optimal ITR is

$$\hat{d}_n(z) = \operatorname{sgn}\{\hat{f}_n(z)\}, \qquad z \in \mathcal{Z}$$

Regularity conditions

C1. The potential right censoring time C is independent of the response process $\{Y(t) : t \in [0, \tau]\}$ and the time T to the absorbing state

C2. The response process has a square-integrable total variation, i.e. $E\{\int_0^\tau |dY(t)|\}^2 < \infty$

C3. The covariate space $\mathcal{Z}$ is a compact subset of $\mathbb{R}^p$

C4. The true state occupation probability of response $EY(t)$ is a continuous function on $[0, \tau]$

C5. The true cumulative baseline hazard $\Lambda_0(t)$ of the right censoring distribution is a continuous function on $[0, \tau]$

Fisher consistency

Theorem 1 (Fisher consistency)

If f^* minimizes $\mathcal{R}_\phi$, then $d^*(z) = sgn\{f^*(z)\}$ for all $z \in \mathcal{Z}$.

Theorem 1 justifies the use of the surrogate risk $\mathcal{R}_\phi$ instead of the original risk $\mathcal{R}$

Consistency of the estimated ITR

Theorem 2

Suppose that assumptions A1–A3 and conditions C1–C5 hold. Then, for $\lambda_n > 0$ *with* $\lambda_n \to 0$ *and* $n\lambda_n \to \infty$,

$$\left| \mathcal{R}_\phi(\hat{f}_n) - \inf_{f \in \mathcal{F}} \mathcal{R}_\phi(f) \right| \xrightarrow{p} 0,$$

as $n \to \infty$, *for any distribution* P *of the data* D. *Moreover, if (i)* $\mathcal{F}$ *is the space of linear functions and* $f^* \in \mathcal{F}$ *or (ii)* $\mathcal{F}$ *is the RKHS with the Gaussian kernel and the marginal distribution* μ *of* Z *is regular, then*

$$\left| \mathcal{V}(\hat{d}_n) - \mathcal{V}(d^*) \right| \xrightarrow{p} 0,$$

as $n \to \infty$.

No free lunch...

- Remarkably, when $\mathcal{F}$ is the RKHS with the Gaussian kernel, the value of estimated ITR $\mathcal{V}(\hat{d}_n)$ converges to the optimal value $\mathcal{V}(d^*)$
- However, the so-called no-free-lunch theorem[7] implies that the corresponding rate of convergence can be extremely slow for at least some distributions of the data D
- This means that an extremely large sample size may be required in real-life settings in order to obtain an ITR $\hat{d}_n$ with a value reasonably close to the optimal value

[7] Steinwart, I. and Christmann, A., 2008. *Support Vector Machines*. Springer Science & Business Media.

Choose space of linear functions

- Due to the no-free-lunch theorem, We will restrict our attention to the case where $\mathcal{F}$ is the space of linear functions for the remainder of the presentation
- If $f^* \notin \mathcal{F}$, $\mathcal{V}(\hat{d}_n)$ converges to a value lower than the optimal value $\mathcal{V}(d^*)$
- Nevertheless, the limit of $\mathcal{V}(\hat{d}_n)$ can be seen as an approximation to the optimal value $\mathcal{V}(d^*)$ because

$$\mathcal{R}(f^*) \leq \mathcal{R}(f) \leq \mathcal{R}_\phi(f) \qquad f \in \mathcal{F}$$

- The performance can be improved by considering an enlarged covariate space $\tilde{\mathcal{Z}}$ that includes polynomial terms and/or two-way interaction terms between the original covariates Z

Asymptotic normality

Theorem 3

Suppose that $\mathcal{F}$ is the space of linear functions. Then, under assumptions A1–A3 and conditions C1–C5, we have

$$\sqrt{n}\left\{\hat{\mathcal{V}}_n(sgn(f)) - \mathcal{V}(sgn(f))\right\} = \frac{1}{\sqrt{n}}\sum_{i=1}^{n}\psi_i(f) + \epsilon(f), \qquad f \in \mathcal{F}$$

with $\sup_{f\in\mathcal{F}}|\epsilon(f)| = o_p(1)$. Moreover, the class of influence functions $\{\psi(f) : f \in \mathcal{F}\}$ is P-Donsker.

Asymptotic normality (cont.)

Theorem 4

Suppose that $\mathcal{F}$ is the space of linear functions and let

$$\tilde{f} = \arg\min_{f \in \mathcal{F}} \mathcal{R}_\phi(f).$$

Then, under assumptions A1–A3 and conditions C1–C5, the additional assumption that $P(\tilde{f}(Z) = 0) = 0$, and for $\lambda_n > 0$ with $\lambda_n \to 0$ and $n\lambda_n \to \infty$, we have

$$\left| \sqrt{n} \left[\hat{\mathcal{V}}_n(sgn(\hat{f}_n)) - \mathcal{V}(sgn(\hat{f}_n)) \right] - \sqrt{n} \left[\hat{\mathcal{V}}_n(sgn(\tilde{f})) - \mathcal{V}(sgn(\tilde{f})) \right] \right| = o_p(1)$$

- Note that Theorem 4 **does not assume** that $f^* \in \mathcal{F}$
- Theorem 4 is very important for conducting rigorous **inference** about the benefit of the estimated optimal ITR $\mathcal{V}(\hat{d}_n)$

Simulation study

- Simulations under an **illness-death model**:

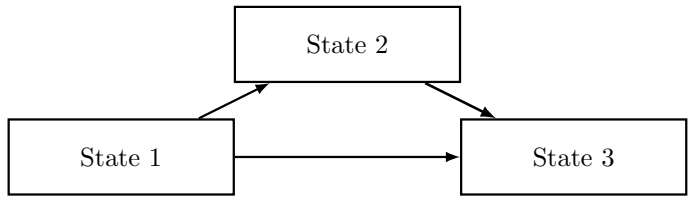

- Two tailoring variables were simulated $Z_1, Z_2 \sim U(-1, 1)$
- Data were simulated under the true optimal ITR $d^*(Z) = \text{sgn}(Z_1 - Z_2)$
- Right censoring times were simulated from Exp(θ), with $\theta \in \{e^{-1.6}, e^{-1}, e^{0.4}\}$
- 1,000 simulated data sets for each scenario

Simulation results regarding $\hat{d}_n$

Censoring	n	$\mathcal{V}(d^*)$	$\mathcal{V}(\hat{d}_n)$	$\hat{\mathcal{V}}_n(\hat{d}_n)$	MCSD	MR
29%	200	1.271	1.214	1.294	0.148	0.191
	400	1.271	1.245	1.285	0.107	0.125
	800	1.271	1.259	1.278	0.077	0.087
45%	200	1.271	1.205	1.310	0.165	0.206
	400	1.271	1.238	1.289	0.117	0.141
	800	1.271	1.255	1.281	0.085	0.098
62%	200	1.271	1.180	1.322	0.198	0.251
	400	1.271	1.219	1.289	0.140	0.179
	800	1.271	1.246	1.282	0.102	0.124

MCSD: Monte Carlo standard deviation

MR: Misclassification rate

Normal Q-Q plots of $\hat{\mathcal{V}}_n(\hat{d}_n)$

Censoring	$n = 200$	$n = 400$	$n = 800$
29%	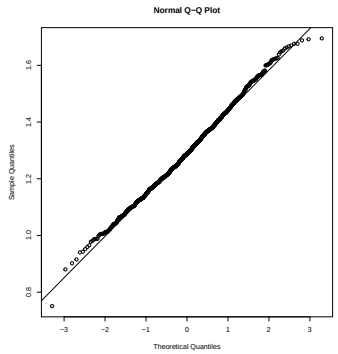	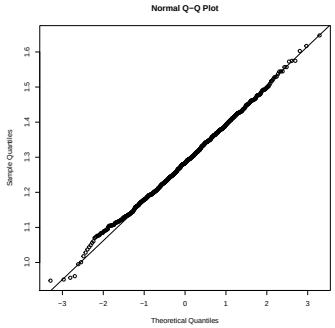	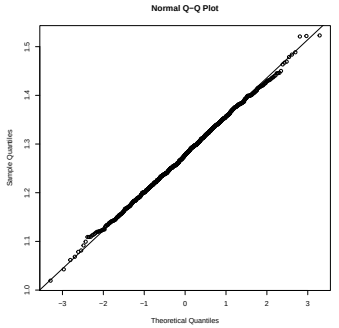
45%	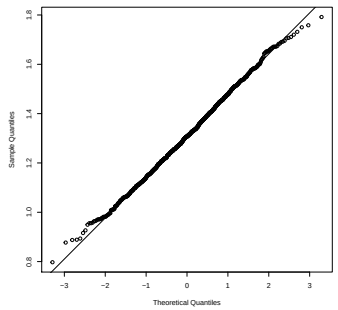	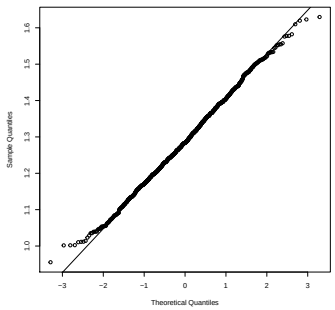	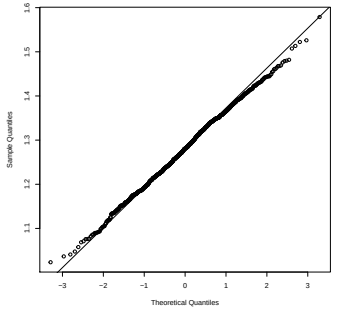
62%	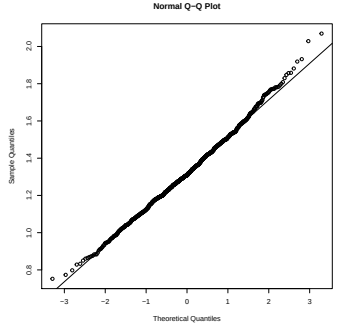	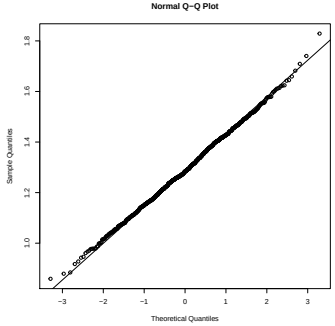	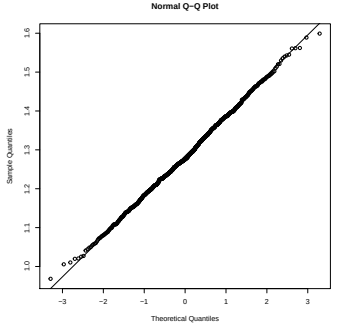

Should we be discouraged about the chemoterhapy + panitumumab option?

- The methodology was used to estimate an optimal ITR in an effort to extend the expected time spent in response within the first 18 months
- The tailoring variables were age, gender, disease stage, and exposure to prior treatment
- Results regarding the (estimated) expected duration of response:
  (i) Optimal ITR: **2.81** months
  (ii) Chemotherapy + panitumumab: **2.46** months
  p-value vs optimal ITR: 0.156
  (iii) Chemotherapy alone: **1.71** months
  p-value vs optimal ITR: **0.031**

Concluding remarks

- Estimation of optimal individualized interventions is crucial in heterogeneous chronic diseases (e.g., cancer)
- The proposed approach is **nonparametric** and relies on weak assumptions
- The validity of our estimation approach for multistate disease processes was justified both **theoretically** and via **simulation experiments**
- **Next steps**:
  - (i) Observational studies (violation of assumption A2)
  - (ii) Multiple decision points
  - (iii) Patient preferences and financial constraints
  - (iv) Missing values in tailoring variables
  - (v) Interval censoring issue

Acknowledgements

This project was supported by the National Institute Of Allergy And Infectious Diseases grant number R01AI140854. The content is solely the responsibility of the author and does not necessarily represent the official views of the National Institutes of Health.

This presentation is based on research using information obtained from `www.projectdatasphere.org`, which is maintained by Project Data Sphere. Neither Project Data Sphere nor the owner(s) of any information from the web site have contributed to, approved or are in any way responsible for the contents of this presentation.