Interpretable Treatment Regimes

Yichi Zhang
Joint work with Dr. Eric Laber,
Dr. Anastasios Tsiatis and Dr. Marie Davidian

Department of Statistics, North Carolina State University

IMPACT 2014

Personalized Medicine

- One size fits all?
- What treatment should he/she receive?

Personalized Medicine

- One size fits all?
- What treatment should he/she receive?
- Data driven and scientifically valid treatment regimes
    - Treatment regime: a function that maps patient covariates to treatment options, and thus provides treatment recommendations from individual patient characteristics
    - Patient covariates: demographic, genetic, clinical measurements, medical history ...
- Aims to optimize some clinical outcome
- Requires joint effort of clinical scientists and statisticians

Example of Treatment Regimes

A breast cancer clinical trial (Fisher et al., 1983)

- Treatments after surgery:
  - chemotherapy alone
  - chemotherapy with tamoxifen
- Patient covariates:
  - age (year)
  - estrogen receptor level (ER, fmol)
  - progesterone receptor level (PR, fmol)
  - tumor size (cm)
  - number of histologically positive nodes
- Outcome: three-year disease-free survival

Example of Treatment Regimes

Give chemotherapy alone if
age ≤ 50 and PR ≤ 10
(Gail and Simon, 1985)

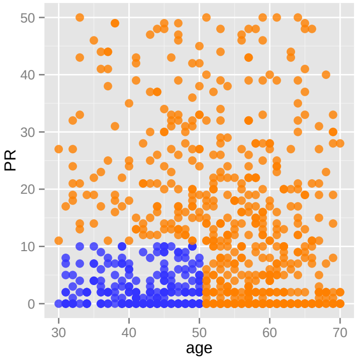

Give chemotherapy alone if
age $+ 7.98 \log(1 + \text{PR}) \leq 60$
(Zhang et al., 2012)

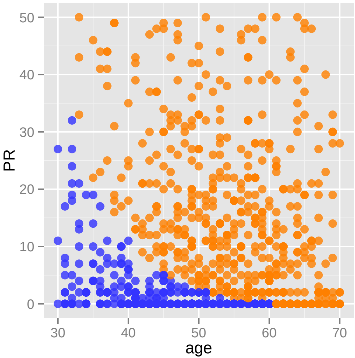

Estimation of Treatment Regimes

- Using the conditional mean of outcome given treatment and patient covariates, e.g. Q- and A-learning (Murphy, 2003; Robins, 2004; Murphy, 2005; Moodie et al., 2007; Henderson et al., 2009; Zhao et al., 2009; Qian and Murphy, 2011; Laber et al., 2014)
- Using the marginal mean of outcome under a treatment regime, e.g. policy search or classification perspective (Robins et al., 2008; Orellana et al., 2010; Zhang et al., 2012, 2013; Zhao et al., 2012, 2014)

Estimation of Treatment Regimes

Estimated regimes often in the form of linear combinations of many variables or basis functions, difficult to interpret

Estimation of Treatment Regimes

Estimated regimes often in the form of linear combinations of many variables or basis functions, difficult to interpret

Besides interpretability, ...

- Identify the most relevant covariates
- Account for the cost of applying the regime

Can we construct high-quality treatment regimes that are much more interpretable?

Framework

- Single decision point
- n i.i.d. samples
- (Y_i, A_i, X_i) for each patient
  - Y_i: scalar outcome of interest, the larger the better
  - A_i: one of m treatment options
  - X_i: vector of patient covariates at baseline

Framework

- Single decision point
- n i.i.d. samples
- (Y_i, A_i, X_i) for each patient
  - Y_i: scalar outcome of interest, the larger the better
  - A_i: one of m treatment options
  - X_i: vector of patient covariates at baseline
- A treatment regime π recommends treatment $\pi(x)$ to a patient with covariates x
  - Easy to interpret
  - Cheap to apply

Constructing Interpretable Treatment Regimes

- A list of if-then statements (decision lists, Rivest, 1987; Marchand and Sokolova, 2005)

If c_1 **then** a_1;
else if c_2 **then** a_2;
...
else if c_L **then** a_L;
else a_0.

- a_ℓs are treatments; c_ℓs are logical conditions of form:

$x_{k_1} \leq \tau_1$,	$x_{k_1} > \tau_1$,
$x_{k_1} \leq \tau_1$ and $x_{k_2} \leq \tau_2$,	$x_{k_1} \leq \tau_1$ or $x_{k_2} \leq \tau_2$,
$x_{k_1} \leq \tau_1$ and $x_{k_2} > \tau_2$,	$x_{k_1} \leq \tau_1$ or $x_{k_2} > \tau_2$,
$x_{k_1} > \tau_1$ and $x_{k_2} \leq \tau_2$,	$x_{k_1} > \tau_1$ or $x_{k_2} \leq \tau_2$,
$x_{k_1} > \tau_1$ and $x_{k_2} > \tau_2$,	$x_{k_1} > \tau_1$ or $x_{k_2} > \tau_2$,

Constructing Interpretable Treatment Regimes

- Readily interpretable

Constructing Interpretable Treatment Regimes

- Readily interpretable

If $x_2 > 7$ **then** Trt A;
else if $x_1 \leq 4$ and $x_2 > 5$ **then** Trt A;
else Trt B.

Constructing Interpretable Treatment Regimes

- Readily interpretable

If $x_2 > 7$ **then** Trt A;
else if $x_1 \leq 4$ and $x_2 > 5$ **then** Trt A;
else Trt B.

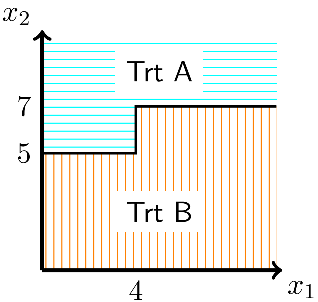

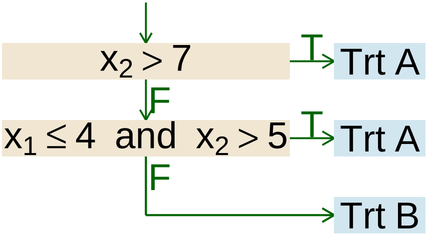

Constructing Interpretable Treatment Regimes

- Expressive

Constructing Interpretable Treatment Regimes

- Expressive

Give Trt A if and only if at least one of the conditions $c_1, \ldots, c_L$ is met:

If c_1 **then** Trt A;
else if c_2 **then** Trt A;
...
else if c_L **then** Trt A;
else Trt B.

Constructing Interpretable Treatment Regimes

- Expressive

Give Trt A if and only if at least one of the conditions $c_1, \ldots, c_L$ is met:

If c_1 **then** Trt A;
else if c_2 **then** Trt A;
...
else if c_L **then** Trt A;
else Trt B.

Give Trt A if and only if all of the conditions $c_1, \ldots, c_L$ are met:

If not c_1 **then** Trt B;
else if not c_2 **then** Trt B;
...
else if not c_L **then** Trt B;
else Trt A.

Constructing Interpretable Treatment Regimes

- Cost effective
  - Suppose a treatment regime is representable by

    If $x_2 > 7$ **then** Trt A;
    else if $x_1 \leq 4$ and $x_2 > 5$ **then** Trt A;
    else Trt B.

  - Only two variables are involved
  - Short-circuited recommendation is possible: Measurement of x_1 is unnecessary for patients with $x_2 > 7$

Non-uniqueness

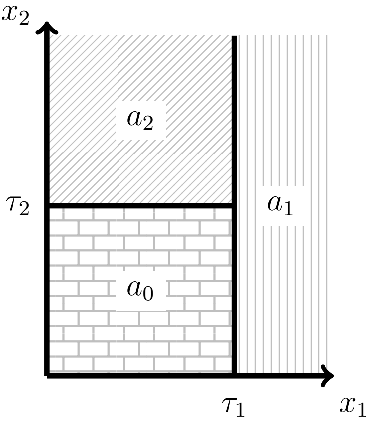

Treatment regime π_1:

If $x_1 > \tau_1$ **then** a_1;
else if $x_2 \geq \tau_2$ **then** a_2;
else a_0.

Treatment regime π_2:

If $x_1 \leq \tau_1$ **and** $x_2 \geq \tau_2$ **then** a_2;
else if $x_1 > \tau_1$ **then** a_1;
else a_0.

Non-uniqueness

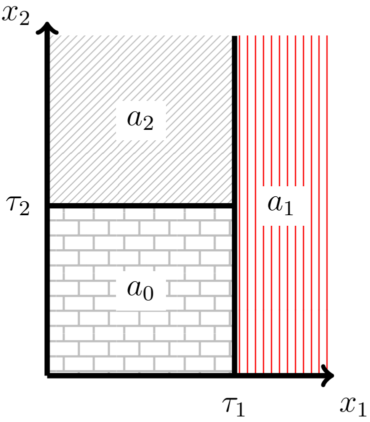

Treatment regime π_1:

If $x_1 > \tau_1$ **then** a_1;

else if $x_2 \geq \tau_2$ **then** a_2;

else a_0.

Treatment regime π_2:

If $x_1 \leq \tau_1$ **and** $x_2 \geq \tau_2$ **then** a_2;

else if $x_1 > \tau_1$ **then** a_1;

else a_0.

Non-uniqueness

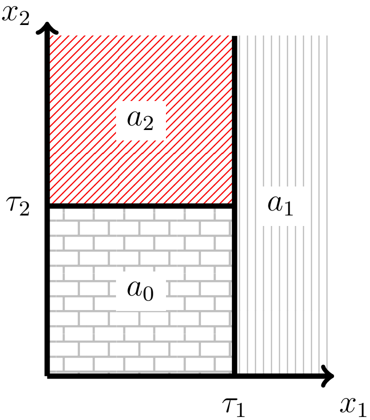

Treatment regime π_1:

If $x_1 > \tau_1$ **then** a_1;

else if $x_2 \geq \tau_2$ **then** a_2;

else a_0.

Treatment regime π_2:

If $x_1 \leq \tau_1$ **and** $x_2 \geq \tau_2$ **then** a_2;

else if $x_1 > \tau_1$ **then** a_1;

else a_0.

Non-uniqueness

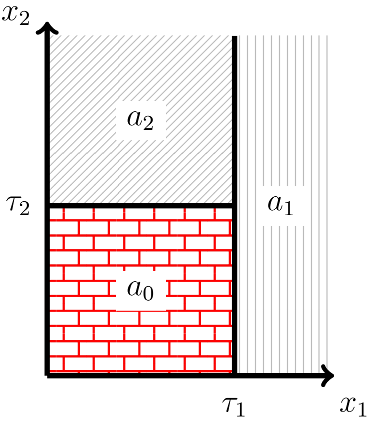

Treatment regime π_1:

If $x_1 > \tau_1$ **then** a_1;
else if $x_2 \geq \tau_2$ **then** a_2;
else a_0.

Treatment regime π_2:

If $x_1 \leq \tau_1$ **and** $x_2 \geq \tau_2$ **then** a_2;
else if $x_1 > \tau_1$ **then** a_1;
else a_0.

Non-uniqueness

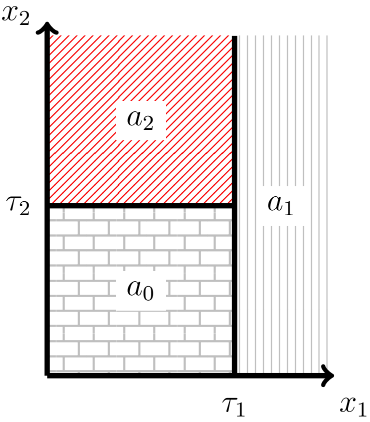

Treatment regime π_1:

If $x_1 > \tau_1$ **then** a_1;
else if $x_2 \geq \tau_2$ **then** a_2;
else a_0.

Treatment regime π_2:

If $x_1 \leq \tau_1$ **and** $x_2 \geq \tau_2$ **then** a_2;
else if $x_1 > \tau_1$ **then** a_1;
else a_0.

Non-uniqueness

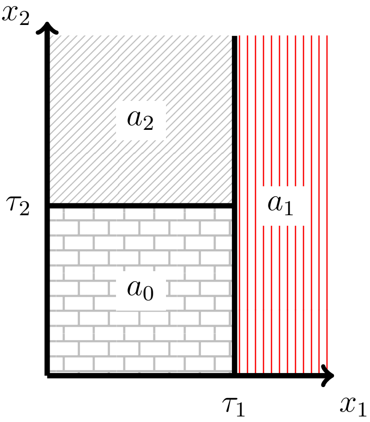

Treatment regime π_1:

If $x_1 > \tau_1$ **then** a_1;
else if $x_2 \geq \tau_2$ **then** a_2;
else a_0.

Treatment regime π_2:

If $x_1 \leq \tau_1$ **and** $x_2 \geq \tau_2$ **then** a_2;
else if $x_1 > \tau_1$ **then** a_1;
else a_0.

Non-uniqueness

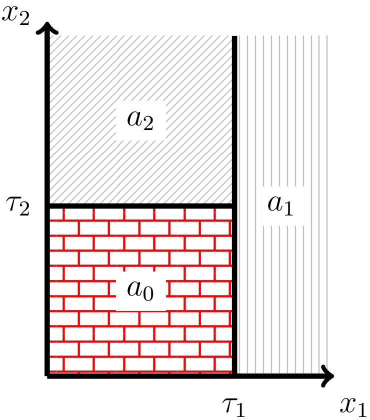

Treatment regime π_1:

If $x_1 > \tau_1$ **then** a_1;
else if $x_2 \geq \tau_2$ **then** a_2;
else a_0.

Treatment regime π_2:

If $x_1 \leq \tau_1$ **and** $x_2 \geq \tau_2$ **then** a_2;
else if $x_1 > \tau_1$ **then** a_1;
else a_0.

Non-uniqueness

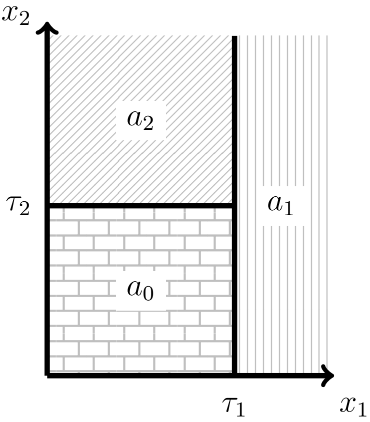

Treatment regime π_1:

If $x_1 > \tau_1$ **then** a_1;
else if $x_2 \geq \tau_2$ **then** a_2;
else a_0.

Treatment regime π_2:

If $x_1 \leq \tau_1$ **and** $x_2 \geq \tau_2$ **then** a_2;
else if $x_1 > \tau_1$ **then** a_1;
else a_0.

π_1 is better is terms of cost

Optimizing Interpretable Treatment Regimes

If c_1 **then** a_1;
else if c_2 **then** a_2;
...
else if c_L **then** a_L;
else a_0.

Many choices for c_ℓs and a_ℓs! How to choose?

1. Improve outcome
   Which regime leads to the largest expected value of outcome, if all patients followed that regime?
2. Reduce cost
   Which regime incurs the smallest cost to implement?

Optimizing Interpretable Treatment Regimes

- Π denotes the class of interpretable treatment regimes
- Step 1 (Maximize outcome): Find

$$\widetilde{\pi} \in \arg\max_{\pi \in \Pi} \widehat{R}(\pi)$$

where $\widehat{R}(\pi)$ estimates the expected value of outcome if the entire population of patients followed the regime π

Optimizing Interpretable Treatment Regimes

- Π denotes the class of interpretable treatment regimes
- Step 1 (Maximize outcome): Find

$$\widetilde{\pi} \in \arg\max_{\pi \in \Pi} \widehat{R}(\pi)$$

where $\widehat{R}(\pi)$ estimates the expected value of outcome if the entire population of patients followed the regime π

  - Use the doubly robust estimator (Tsiatis, 2006) as $\widehat{R}(\pi)$
  - Build if-then statements one by one to form π

Optimizing Interpretable Treatment Regimes

- Π denotes the class of interpretable treatment regimes
- Step 1 (Maximize outcome): Find

$$\widetilde{\pi} \in \arg\max_{\pi \in \Pi} \widehat{R}(\pi)$$

  where $\widehat{R}(\pi)$ estimates the expected value of outcome if the entire population of patients followed the regime π

- Step 2 (Minimize cost): Find

$$\widehat{\pi} \in \arg\min_{\pi \in \Pi} \widehat{S}(\pi) \text{ subject to } \widehat{R}(\pi) = \widehat{R}(\widetilde{\pi})$$

  where $\widehat{S}(\pi)$ estimates the expected cost of measurements required to apply the regime π

Simulation Study

- Randomized clinical trial with two or three treatment options
- Patient covariates are 50-dimensional, multivariate normal, and weakly correlated
- Only the first two covariates are important
- Outcome is continuous or binary

Goal:

- To show that the proposed method, compared with Q-learning, constructs much more interpretable treatment regimes with comparable quality and much lower cost

Simulation Study: Optimal Regime is Decision List

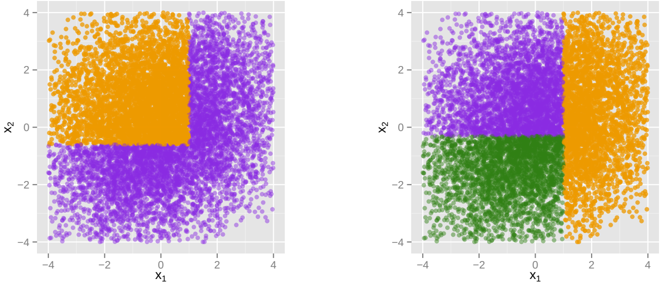

		n	$R(\widehat{\pi})$	$R(\widehat{\pi}_{\text{glm}})$	$R(\widehat{\pi}_{\text{svm}})$	$S(\widehat{\pi})$	$S(\widehat{\pi}_{\text{glm}})$
Cont	L	500	2.76	2.51	2.36	1.81	21.3
	R	750	2.87	2.63	2.33	2.14	28.5
Bin	L	1000	0.76	0.73	0.69	2.64	21.9
	R	1500	0.74	0.72	0.61	3.16	30.4

Simulation Study: Optimal Regime is *not* Decision List

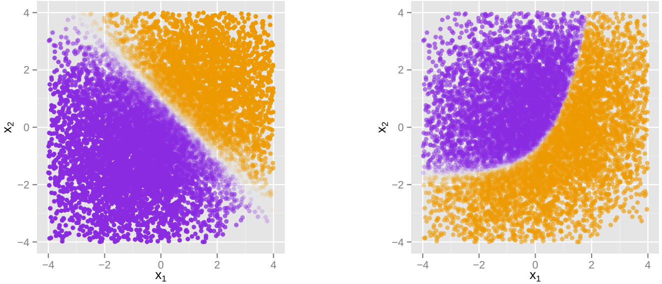

		n	$R(\widehat{\pi})$	$R(\widehat{\pi}_{\text{glm}})$	$R(\widehat{\pi}_{\text{svm}})$	$S(\widehat{\pi})$	$S(\widehat{\pi}_{\text{glm}})$
Cont	L	500	2.70	2.79	2.73	1.64	21.4
	R	500	2.59	2.52	2.35	1.71	23.1
Bin	L	1000	0.71	0.72	0.60	1.87	26.2
	R	1000	0.73	0.72	0.67	2.53	25.0

Breast Cancer Clinical Trial

- Treatments after surgery:
  - chemotherapy alone
  - chemotherapy with tamoxifen
- Patient covariates:
  - age (year)
  - estrogen receptor level (ER, fmol)
  - progesterone receptor level (PR, fmol)
  - tumor size (cm)
  - number of histologically positive nodes
- Outcome: three-year disease-free survival
- 1164 patients

Breast Cancer Clinical Trial

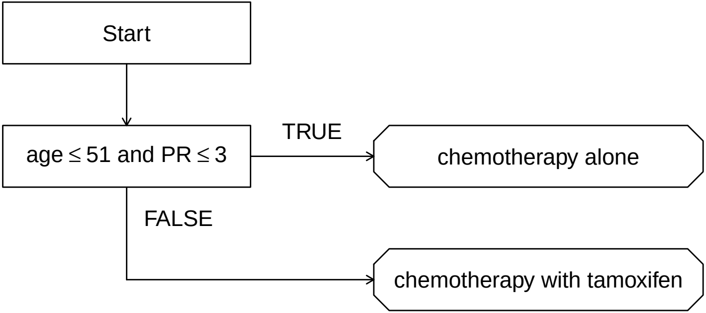

Figure 1 : Estimated regime. Replacing the condition with age ≤ 50 and PR ≤ 10 leads to the regime suggested by Gail and Simon (1985). These two regimes agree for 92% of the patients in the data.

Breast Cancer Clinical Trial

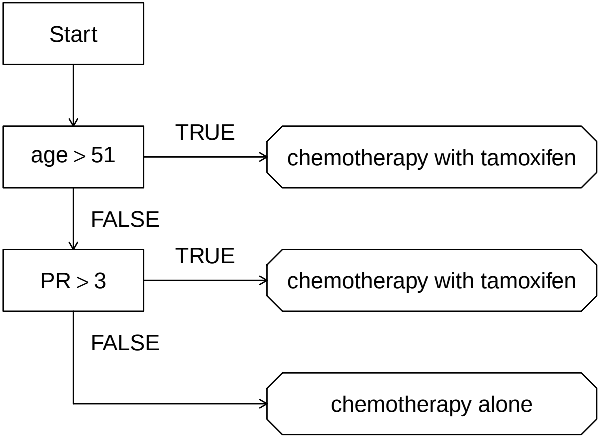

Figure 2 : Estimated regime with minimal cost.

Summary

- Construct interpretable treatment regimes using a list of if-then statements
  - Use simple inequalities as logical conditions
  - Reduce cost and identify important covariates
  - Protect against model misspecification
- May extend to multiple decisions, e.g. chronic diseases
- May apply to other areas, e.g. risk prediction, to handle costly measurements
- R Package: http://www4.ncsu.edu/~yzhang52/

Thank you!

yzhang52@ncsu.edu

http://www4.ncsu.edu/~yzhang52/

References I

Fisher, B., Redmond, C., Brown, A., Wickerham, D. L., Wolmark, N., Allegra, J., Escher, G., Lippman, M., Savlov, E., and Wittliff, J. (1983). Influence of tumor estrogen and progesterone receptor levels on the response to tamoxifen and chemotherapy in primary breast cancer. *Journal of Clinical Oncology*, 1(4):227–241.

Gail, M. and Simon, R. (1985). Testing for qualitative interactions between treatment effects and patient subsets. *Biometrics*, 41(2):361–372.

Henderson, R., Ansell, P., and Alshibani, D. (2009). Regret-regression for optimal dynamic treatment regimes. *Biometrics*, 66(4):1192–1201.

References II

Laber, E. B., Linn, L. A., and Stefanski, L. A. (2014). Interactive model-building for Q-learning. *Biometrika*. in press.

Marchand, M. and Sokolova, M. (2005). Learning with decision lists of data-dependent features. *Journal of Machine Learning Research*, 6:427–451.

Moodie, E. E. M., Richardson, T. S., and Stephens, D. A. (2007). Demystifying optimal dynamic treatment regimes. *Biometrics*, 63(2):447–455.

Murphy, S. A. (2003). Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65(2):331–355.

Murphy, S. A. (2005). A generalization error for Q-learning. *Journal of Machine Learning Research*, 6:1073–1097.

References III

Orellana, L., Rotnitzky, A., and Robins, J. M. (2010). Dynamic regime marginal structural mean models for estimation of optimal dynamic treatment regimes, Part I: Main content. *The International Journal of Biostatistics*, 6(2).

Qian, M. and Murphy, S. A. (2011). Performance guarantees for individualized treatment rules. *Annals of Statistics*, 39(2):1180–1210.

Rivest, R. L. (1987). Learning decision lists. *Machine Learning*, 2(3):229–246.

Robins, J. M. (2004). Optimal structural nested models for optimal sequential decisions. In Lin, D. Y. and Heagerty, P. J., editors, *Proceedings of the Second Seattle Symposium in Biostatistics*, volume 179 of *Lecture Notes in Statistics*, pages 189–326. Springer.

References IV

Robins, J. M., Orellana, L., and Rotnitzky, A. (2008). Estimation and extrapolation of optimal treatment and testing strategies. *Statistics in Medicine*, 27(23):4678–4721.

Tsiatis, A. A. (2006). *Semiparametric theory and missing data*. Springer.

Zhang, B., Tsiatis, A. A., Laber, E. B., and Davidian, M. (2012). A robust method for estimating optimal treatment regimes. *Biometrics*, 68(4):1010–1018.

Zhang, B., Tsiatis, A. A., Laber, E. B., and Davidian, M. (2013). Robust estimation of optimal dynamic treatment regimes for sequential treatment decisions. *Biometrika*, 100(3):681–694.

References V

Zhao, Y., Kosorok, M. R., and Zeng, D. (2009). Reinforcement learning design for cancer clinical trials. *Statistics in Medicine*, 28(26):3294–3315.

Zhao, Y., Zeng, D., Laber, E. B., and Kosorok, M. R. (2014). New statistical learning methods for estimating optimal dynamic treatment regimes. *Journal of the American Statistical Association*. in press.

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.