Learning Optimal Treatment Rules

Mentee: Max Bi
Mentor: Nina Galanter

Motivation

- Some subset of patients may respond better to one treatment as opposed to another
- How do we find which patients respond better to which treatments?
- Optimal Treatment Rule (OTR)
- Learned based on data, can predict future treatments

Example Scenario

- Based on a measurable covariate, determine whether Benadryl or Melatonin is a better sleep aid
- Y = …
  - Some real number(representing time slept)
- A = …
  - 0: Benadryl
  - 1: Melatonin
- X_1, X_2, … X_d: many number of covariates that may affect drug effectiveness
  - Are some covariates more relevant than others?
  - Are there covariates that are irrelevant?

Assumptions for Causal Inference

- Scenario 1: Experiment
    - Randomized treatments
    - No confounding, covariates don't affect probability of treatment

Assumptions for Causal Inference

- Scenario 2: Observational study
  - Treatments are distinct
    - E.g. medicine or placebo
  - No unaccounted for confounding
  - Nonzero probability of each treatment value

Q-Learning

- Propose model for expectation of outcome given covariates, treatments
- Obtain coefficients through ordinary least squares or LASSO

$$E(Y|X, A) = b_0 + b_1 X + A(b_2 + b_3 * X)$$

- Solve for coefficients to obtain rule
- Interpretable for simpler examples

Treatment Rule Example

- $A = 1, 0$
- $B = [b_0, b_1, b_2, b_3]$
- X is blood melatonin level

$$E(Y|X, A) = b_0 + b_1X + A(b_2 + b_3 * X)$$

Treatment Rule Example

$$
\begin{aligned}
b_2 + b_3 X &> 0 \\
b_3 X &> -b_2 \\
X &< \frac{b_2}{b_3}
\end{aligned}
$$

- Interpretation: treat if blood-melatonin level is below b_2 / b_3

Least Squares

$$argmin_B \sum_{i=1}^{n}(B^T x_i - y_i)^2$$

- Find coefficients that minimize the following

Least Squares Visualization

- “Bowl” shaped, convex function
- Objective: find minimum value of function

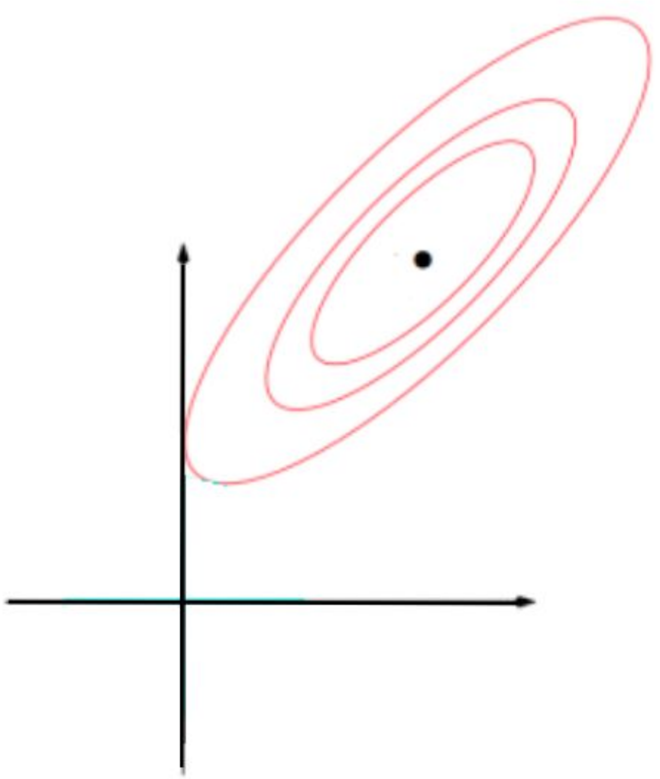

LASSO

$$argmin_B \sum_{i=1}^{n} (B^T x_i - y_i)^2 + \lambda ||B||_1$$

- Add regularization to least squares function
- Now, find coefficients that minimize the function
- As regularization increases, the L1 norm matters more, thus increasing sparsity

LASSO Visualization

- Find another solution on the "bowl"
  - Still have the "bowl"
  - Add a search space
  - Find intersection between search space and least squares function

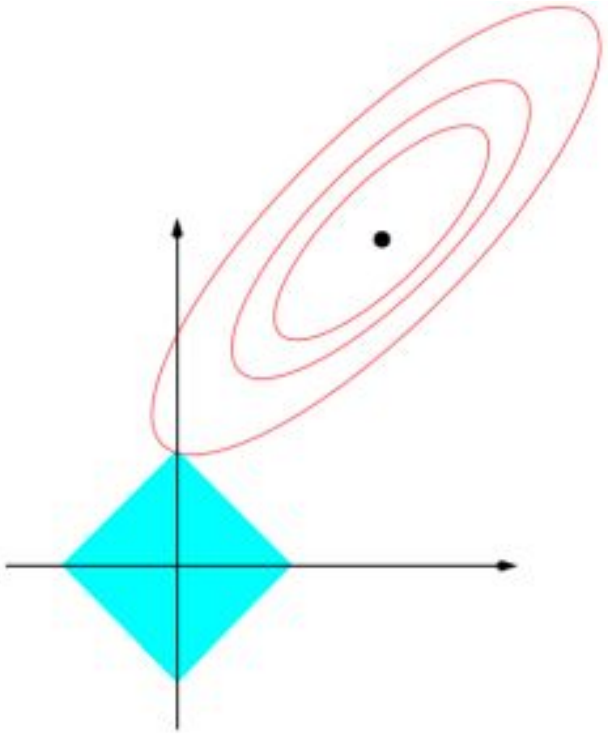

Let's Simulate! But … Why?

- Accurately assess model accuracy
  - True function for real data UNKNOWN
  - Easily test multiple methods

Simulation

Setup

- Generate X, 1000 random observations
- Generate sparse B
- Nonzero coefficients are arbitrarily chosen
- Plug into equation and add random noise to generate Y (training data)

Simulation

Estimate

- Train plain Q-learning with OLS (lm) and with LASSO (glmnet)

Simulation

Evaluate

- Generate more test data with true coefficient values
- Predict from test observations
- Compare predictions to true values

Simulation

Results

- Average accuracies:
  - OLS: 97.17%
  - LASSO: 97.45%

Red dots correspond to true coefficient values

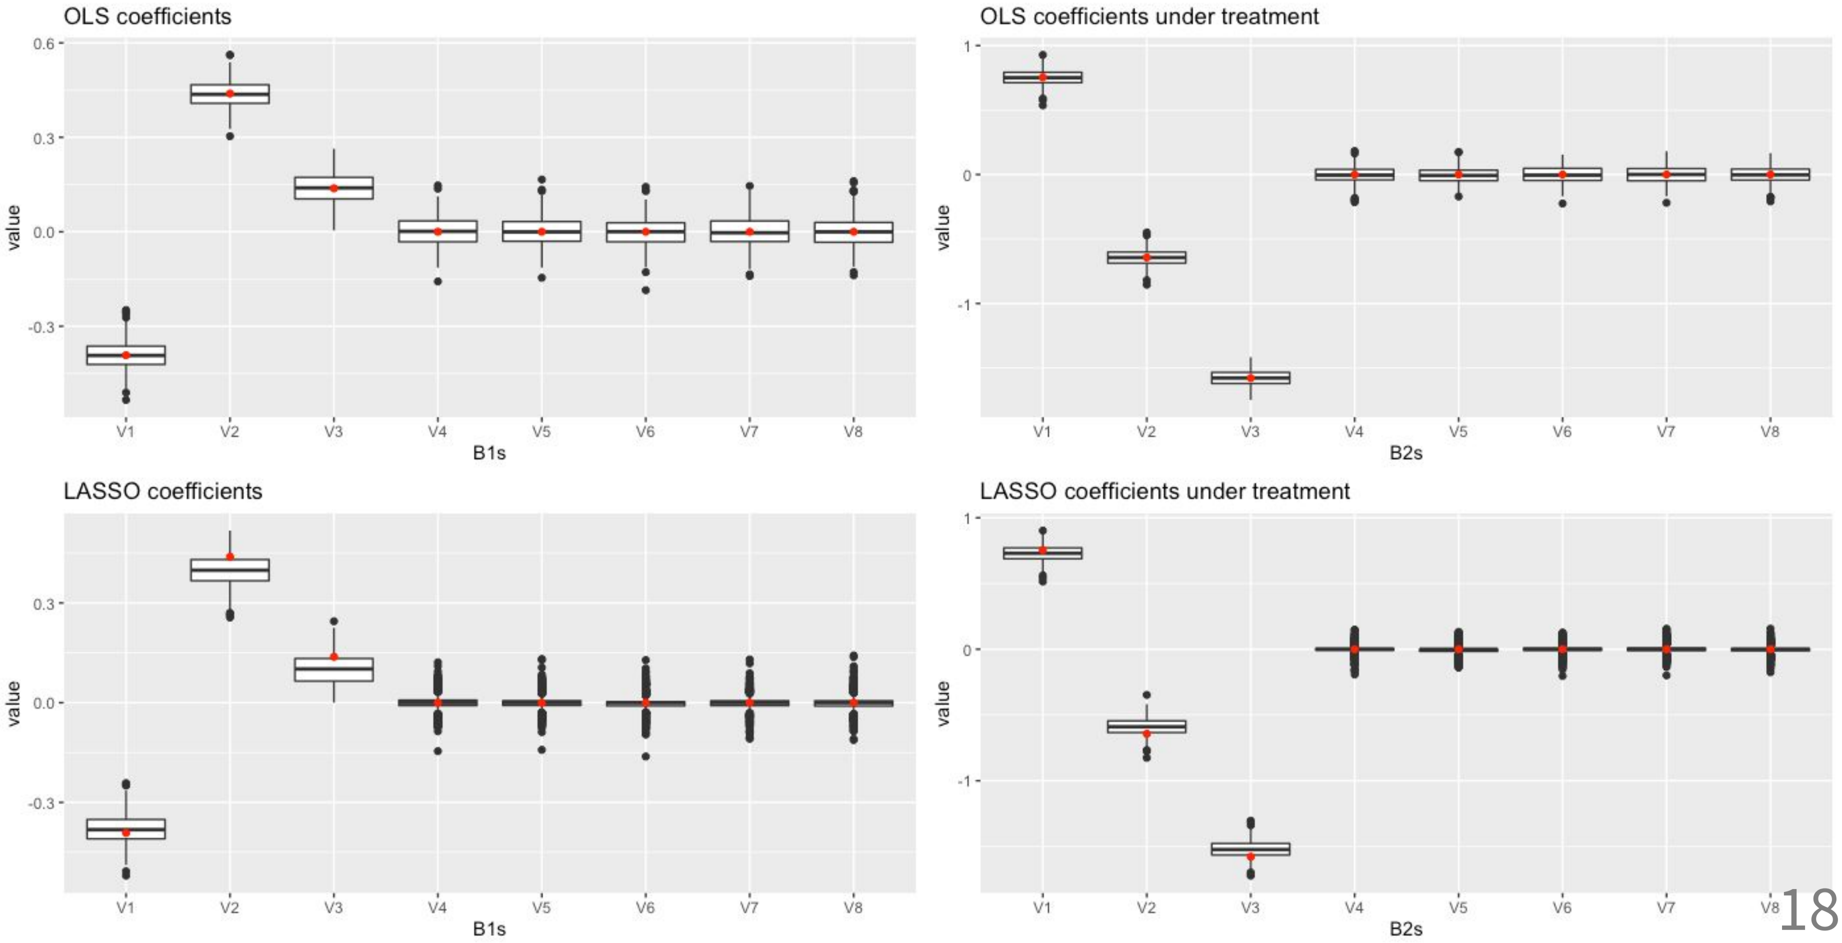

Conclusion

Extensions:

- Multi-stage treatments
  - What happens if we want different, consecutive treatments?
- Adaptations to Observational Data
  - How do we account for confounding mathematically?

References

1. Hernán MA, Robins JM (2020). Causal Inference: What If. Boca Raton: Chapman & Hall/CRC.
2. Kosorok, M. R., & Moodie, E. E. M. (2016). Adaptive Treatment Strategies in Practice: Planning Trials and Analyzing Data for Personalized Medicine (ASA-SIAM Series on Statistics and Applied Probability, Series Number 21). Society for Industrial and Applied Mathematics.
3. Murphy, K. P. (2022). Probabilistic Machine Learning: An Introduction (Adaptive Computation and Machine Learning series). The MIT Press.