Machine Learning Methods for Causal Inference on High-frequency Observational Data

The Case of Residential Demand Response

Maximilian Balandat
work with James Gillan, Datong Zhou and Claire Tomlin

FORCES All Hands Meeting, MIT
June 9, 2016

Demand Response (DR)

Projected net load in CAISO system:

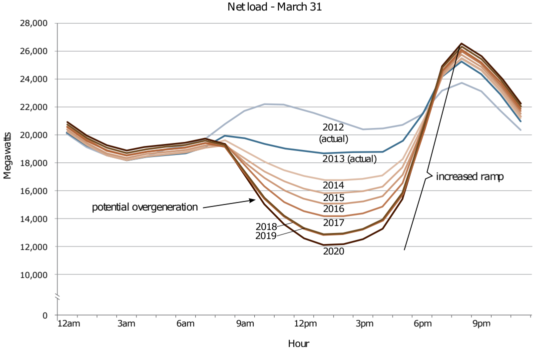

- Growing need for operational flexibility due to more renewables
- DR as "behavior-modifying mechanisms to change the net load shape"

Objective: Causal inference on participant behavior in Residential DR

Households receive messages (via text/app, email, etc.) incentivizing them to reduce electricity consumption.

Econometrics and Machine Learning

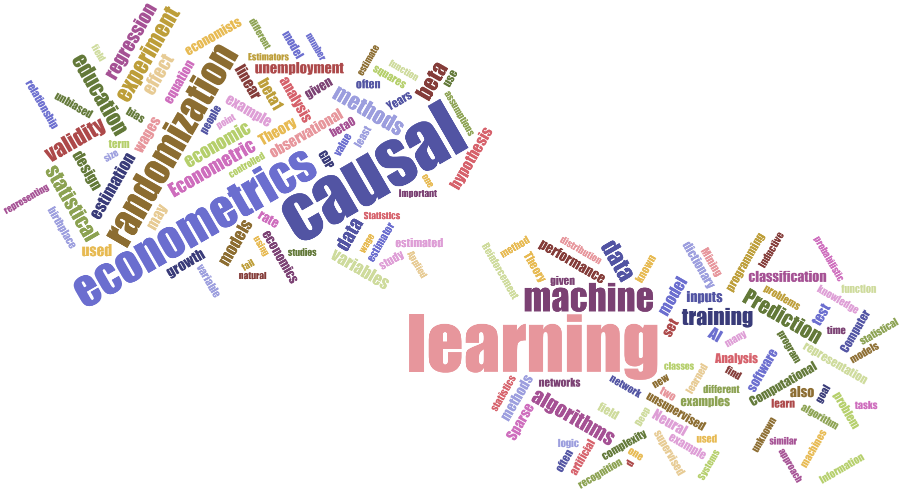

Little (but very interesting) work bringing the two fields together, e.g.:

- S. Athey and G. W. Imbens. Recursive Partitioning for Heterogeneous Causal Effects. ArXiV, Dec 2015.
- L. Tian et al. A simple method for estimating interactions between a treatment and a large number of covariates. Journal of the American Statistical Association, 2014.

Estimating the Counterfactual

Economics gold standard: Randomized Experiment

- Randomly split population in treatment and control groups
- Compare average outcome in the two groups

Estimating the Counterfactual

Economics gold standard: Randomized Experiment

- Randomly split population in treatment and control groups
- Compare average outcome in the two groups

Our goal: Use within-subject variation in repeated treatment assignment

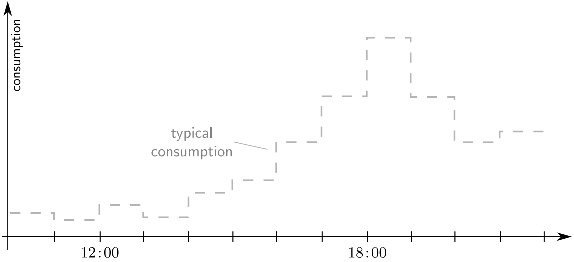

Estimating the Counterfactual

Economics gold standard: Randomized Experiment

- Randomly split population in treatment and control groups
- Compare average outcome in the two groups

Our goal: Use within-subject variation in repeated treatment assignment

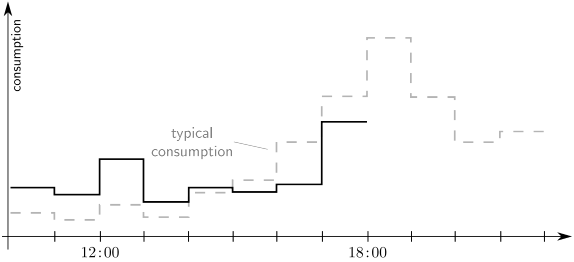

Estimating the Counterfactual

Economics gold standard: Randomized Experiment

- Randomly split population in treatment and control groups
- Compare average outcome in the two groups

Our goal: Use within-subject variation in repeated treatment assignment

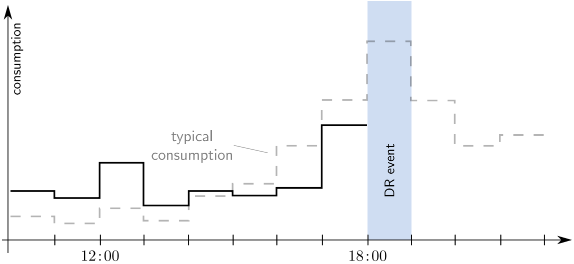

Estimating the Counterfactual

Economics gold standard: Randomized Experiment

- Randomly split population in treatment and control groups
- Compare average outcome in the two groups

Our goal: Use within-subject variation in repeated treatment assignment

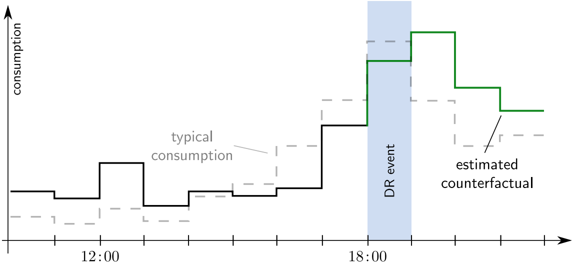

Estimating the Counterfactual

Economics gold standard: Randomized Experiment

- Randomly split population in treatment and control groups
- Compare average outcome in the two groups

Our goal: Use within-subject variation in repeated treatment assignment

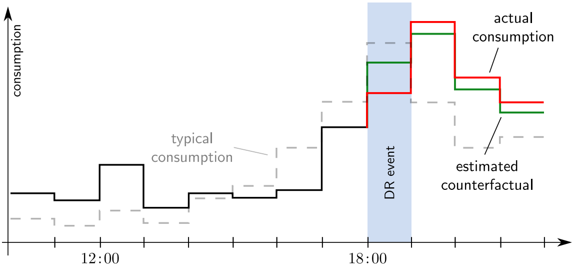

Estimating the Counterfactual

Economics gold standard: Randomized Experiment

- Randomly split population in treatment and control groups
- Compare average outcome in the two groups

Our goal: Use within-subject variation in repeated treatment assignment

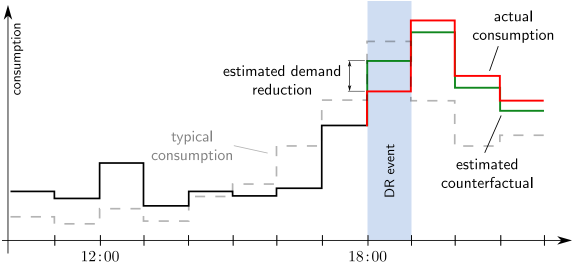

Estimating the Counterfactual

Economics gold standard: Randomized Experiment

- Randomly split population in treatment and control groups
- Compare average outcome in the two groups

Our goal: Use within-subject variation in repeated treatment assignment

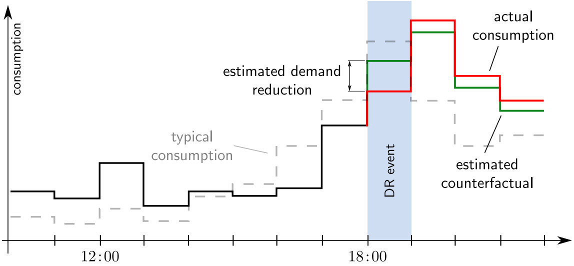

Assumption: Potential outcomes* and assignment of DR events are conditionally independent given covariates of the model.

*Consumption during potential event hour in case of event / no event.

Goals and Challenges

Individual Treatment Effect (ITE): Response of a single household to receiving a DR event*

Average Treatment Effect (ATE): Average response of population*

*possibly conditional on reward level, time of day, etc.

Goals

- Estimate ITE based solely on within-user variation
- Obtain estimate of ATE by marginalizing ITE estimates
  - benchmark against the experimental estimates
- Use ITE estimates for targeting to improve allocative efficiency

Challenges

- Very high residual variance in household-level consumption
- Limited number of DR events per household
- Validity relies on unconfoundedness assumption

Goals and Challenges

Individual Treatment Effect (ITE): Response of a single household to receiving a DR event*

Average Treatment Effect (ATE): Average response of population*

*possibly conditional on reward level, time of day, etc.

Goals

- Estimate ITE based solely on within-user variation
- Obtain estimate of ATE by marginalizing ITE estimates
  - benchmark against the experimental estimates
- Use ITE estimates for targeting to improve allocative efficiency

Challenges

- Very high residual variance in household-level consumption
- Limited number of DR events per household
- Validity relies on unconfoundedness assumption

A Large-Scale Field Experiment

Randomized Controlled Trial among customers of ohmconnect

- Estimate demand curve, effect of adopting home automation technology and targeting algorithm
- $\approx$ 12,500 participants over 15 month duration
- Funding from California Energy Commission's EPIC grant
- Rollout: July 2016

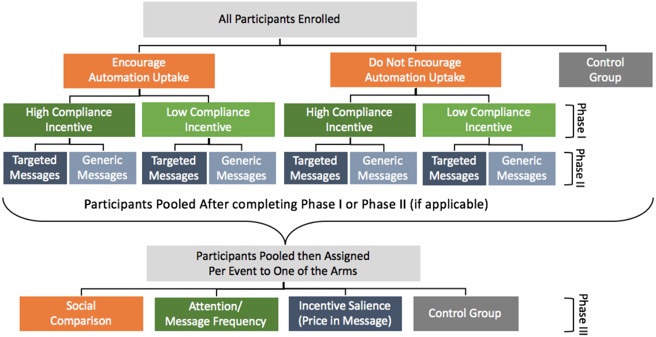

Synthetic Experiments: ITE Estimation Errors

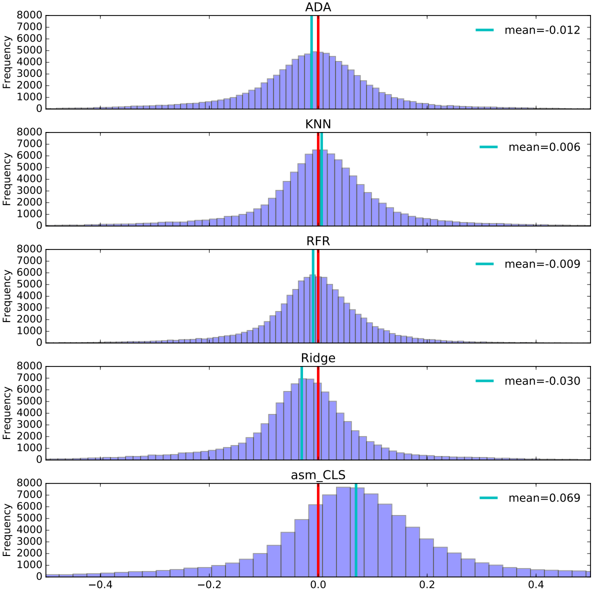

Synthetic Experiments: ATE estimates

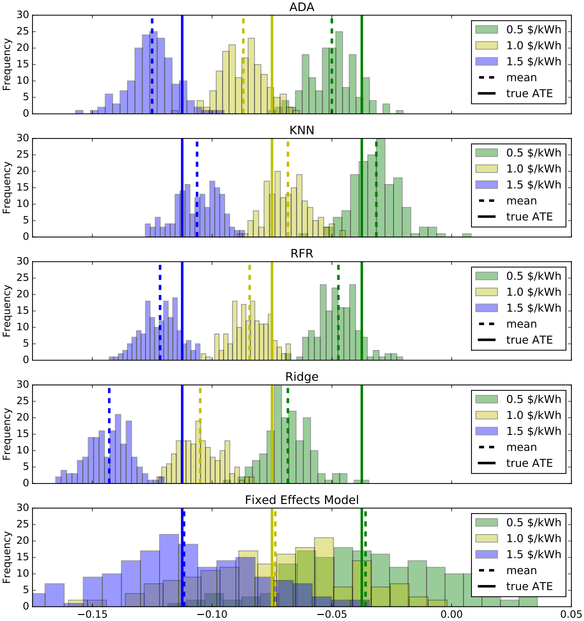

Synthetic Experiments: Targeting

Goal: Find participants with largest (most negative) ITEs

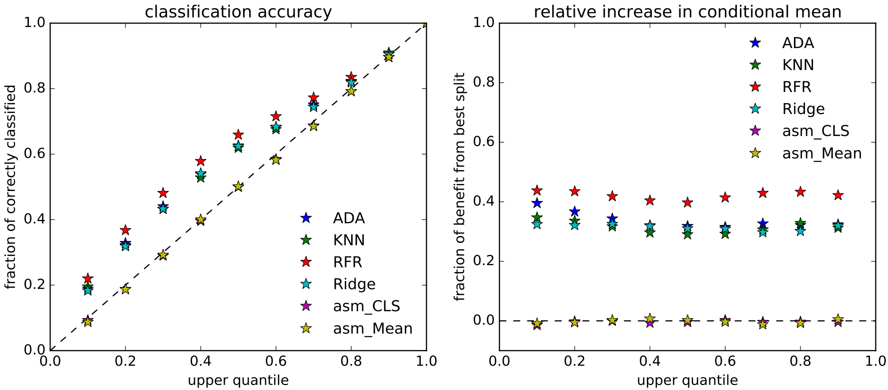

Summary

Methodology

Cross-pollination between Econometrics and Machine Learning

- Proposed novel algorithms for causal inference
  - theoretical properties (under some assumptions)
  - performance demonstrated in synthetic experiments using real data
- Developed targeting algorithms for Demand Response
  - improve allocative efficiency by exploiting heterogeneity across customers

Empirical Evaluation

Upcoming large-scale Randomized Controlled Trial in California

- Potential implications for technology and policy
- Benchmark our estimators against the experimental gold standard
- Evaluate effectiveness of our targeting algorithms