Residential Demand Response Targeting Using Machine Learning with Observational Data

Datong-Paul Zhou, Maximilian Balandat, and Claire Tomlin

University of California, Berkeley

[datong.zhou, balandat, tomlin]@eecs.berkeley.edu

December 14, 2016

Residential Demand Response

Background

- Reliable grid operation is dependent on adequate supply and flexible resources
- 2009: CAISO proposes "**Proxy Demand Resource Product**" to "facilitate the participation of existing retail demand programs in the ISO market"
- Proxy Demand Resources (PDRs) offer bids to reflect flexibility to adjust load in response to market schedules
- July 2015: CPUC Resolution E-4728 approves an auction mechanism for demand response capacity, called the "**Demand Response Auction Mechanism** (DRAM)"
- (Residential) Demand Response Providers with the ability to aggregate customers capable of reducing load participate in the ISO day-ahead, real-time ancillary services market
- Utilities can obtain resource adequacy through residential resources
- Benchmarking of reduction potential by using consumption "baselines"

Residential Demand Response

Background

- Reliable grid operation is dependent on adequate supply and flexible resources
- 2009: CAISO proposes "**Proxy Demand Resource Product**" to "facilitate the participation of existing retail demand programs in the ISO market"
- Proxy Demand Resources (PDRs) offer bids to reflect flexibility to adjust load in response to market schedules
- July 2015: CPUC Resolution E-4728 approves an auction mechanism for demand response capacity, called the "**Demand Response Auction Mechanism** (DRAM)"
- (Residential) Demand Response Providers with the ability to aggregate customers capable of reducing load participate in the ISO day-ahead, real-time ancillary services market
- Utilities can obtain resource adequacy through residential resources
- Benchmarking of reduction potential by using consumption "baselines"

Residential Demand Response

Background

- Reliable grid operation is dependent on adequate supply and flexible resources
- 2009: CAISO proposes "**Proxy Demand Resource Product**" to "facilitate the participation of existing retail demand programs in the ISO market"
- Proxy Demand Resources (PDRs) offer bids to reflect flexibility to adjust load in response to market schedules
- July 2015: CPUC Resolution E-4728 approves an auction mechanism for demand response capacity, called the "**Demand Response Auction Mechanism** (DRAM)"
- (Residential) Demand Response Providers with the ability to aggregate customers capable of reducing load participate in the ISO day-ahead, real-time ancillary services market
- Utilities can obtain resource adequacy through residential resources
- Benchmarking of reduction potential by using consumption "baselines"

Residential Demand Response

Background

- Reliable grid operation is dependent on adequate supply and flexible resources
- 2009: CAISO proposes "**Proxy Demand Resource Product**" to "facilitate the participation of existing retail demand programs in the ISO market"
- Proxy Demand Resources (PDRs) offer bids to reflect flexibility to adjust load in response to market schedules
- July 2015: CPUC Resolution E-4728 approves an auction mechanism for demand response capacity, called the "**Demand Response Auction Mechanism** (DRAM)"
- (Residential) Demand Response Providers with the ability to aggregate customers capable of reducing load participate in the ISO day-ahead, real-time ancillary services market
- Utilities can obtain resource adequacy through residential resources
- Benchmarking of reduction potential by using consumption "baselines"

Estimating the Counterfactual

- Causal Inference: Set of users $\mathcal{I} = \{1, \ldots, N\}$, each user has observed outcomes $\mathbf{y}_i = \{y_{i1}, \ldots, y_{i\tau}\}$, corresponding covariates $X_i = \{\mathbf{x}_{i1}, \ldots, \mathbf{x}_{i\tau}\}$, and binary treatment indicator D_{it}
- Potential outcomes framework (Rubin, 1974) and fundamental problem of causal inference:

$$y_{it} = y_{it}^0 + D_{it}(y_{it}^1 - y_{it}^0) \quad \forall\, i \in \mathcal{I}, t \in \{1, \ldots, \tau\}$$

  - That is, either the treatment outcome y_{it}^1 or the control outcome y_{it}^0 can be observed, but never both.
- Individual treatment effect β_i (ITE) and average treatment effect (ATE) μ:

$$\beta_i = \mathbb{E}[y_{it}^1 - y_{it}^0] = \frac{1}{|\mathcal{T}_i|} \sum_{j \in \mathcal{T}_i} (y_{ij}^1 - y_{ij}^0) \qquad \mu = \frac{1}{N} \sum_{i \in \mathcal{I}} \beta_i$$

- How can we estimate ITEs, given the fundamental problem of causal inference?
- $\Rightarrow$ Estimate counterfactuals!

Estimating the Counterfactual

- Causal Inference: Set of users $\mathcal{I} = \{1, \ldots, N\}$, each user has observed outcomes $\mathbf{y}_i = \{y_{i1}, \ldots, y_{i\tau}\}$, corresponding covariates $X_i = \{\mathbf{x}_{i1}, \ldots, \mathbf{x}_{i\tau}\}$, and binary treatment indicator D_{it}
- Potential outcomes framework (Rubin, 1974) and fundamental problem of causal inference:

$$y_{it} = y_{it}^0 + D_{it}(y_{it}^1 - y_{it}^0) \quad \forall\, i \in \mathcal{I}, t \in \{1, \ldots, \tau\}$$

  - That is, either the treatment outcome y_{it}^1 or the control outcome y_{it}^0 can be observed, but never both.
- Individual treatment effect β_i (ITE) and average treatment effect (ATE) μ:

$$\beta_i = \mathbb{E}[y_{it}^1 - y_{it}^0] = \frac{1}{|\mathcal{T}_i|} \sum_{j \in \mathcal{T}_i} (y_{ij}^1 - y_{ij}^0) \qquad \mu = \frac{1}{N} \sum_{i \in \mathcal{I}} \beta_i$$

- How can we estimate ITEs, given the fundamental problem of causal inference?
- $\Rightarrow$ Estimate counterfactuals!

Estimating the Counterfactual

- Causal Inference: Set of users $\mathcal{I} = \{1, \ldots, N\}$, each user has observed outcomes $\mathbf{y}_i = \{y_{i1}, \ldots, y_{i\tau}\}$, corresponding covariates $X_i = \{\mathbf{x}_{i1}, \ldots, \mathbf{x}_{i\tau}\}$, and binary treatment indicator D_{it}
- Potential outcomes framework (Rubin, 1974) and fundamental problem of causal inference:

$$y_{it} = y_{it}^0 + D_{it}(y_{it}^1 - y_{it}^0) \quad \forall\, i \in \mathcal{I}, t \in \{1, \ldots, \tau\}$$

  - That is, either the treatment outcome y_{it}^1 or the control outcome y_{it}^0 can be observed, but never both.

- Individual treatment effect β_i (ITE) and average treatment effect (ATE) μ:

$$\beta_i = \mathbb{E}[y_{it}^1 - y_{it}^0] = \frac{1}{|\mathcal{T}_i|} \sum_{j \in \mathcal{T}_i} (y_{ij}^1 - y_{ij}^0) \qquad \mu = \frac{1}{N} \sum_{i \in \mathcal{I}} \beta_i$$

- How can we estimate ITEs, given the fundamental problem of causal inference?
- $\Rightarrow$ Estimate counterfactuals!

Estimating the Counterfactual

- Causal Inference: Set of users $\mathcal{I} = \{1, \ldots, N\}$, each user has observed outcomes $\mathbf{y}_i = \{y_{i1}, \ldots, y_{i\tau}\}$, corresponding covariates $X_i = \{\mathbf{x}_{i1}, \ldots, \mathbf{x}_{i\tau}\}$, and binary treatment indicator D_{it}
- Potential outcomes framework (Rubin, 1974) and fundamental problem of causal inference:

$$y_{it} = y_{it}^0 + D_{it}(y_{it}^1 - y_{it}^0) \quad \forall\, i \in \mathcal{I}, t \in \{1, \ldots, \tau\}$$

  - That is, either the treatment outcome y_{it}^1 or the control outcome y_{it}^0 can be observed, but never both.

- Individual treatment effect β_i (ITE) and average treatment effect (ATE) μ:

$$\beta_i = \mathbb{E}[y_{it}^1 - y_{it}^0] = \frac{1}{|\mathcal{T}_i|} \sum_{j \in \mathcal{T}_i} (y_{ij}^1 - y_{ij}^0) \qquad \mu = \frac{1}{N} \sum_{i \in \mathcal{I}} \beta_i$$

- How can we estimate ITEs, given the fundamental problem of causal inference?
- $\Rightarrow$ Estimate counterfactuals!

Measuring Reduction in Consumption

- Estimate the demand reduction by using suitable baselines (counterfactuals)

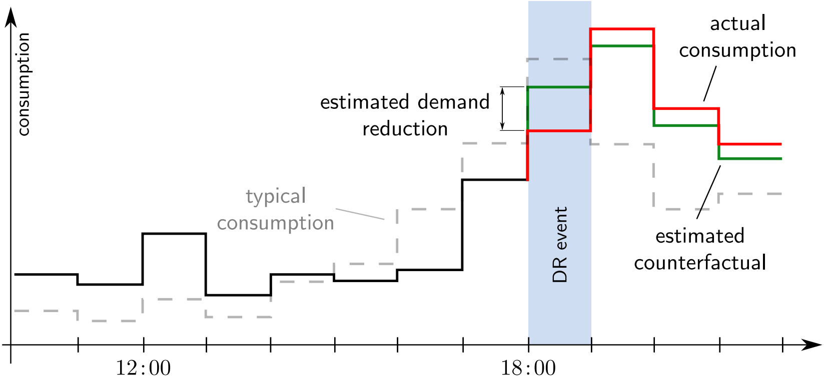

Figure: Estimated counterfactual vs. actual consumption[1]

[1] Courtesy of Maximilian Balandat

Today's Talk

Agenda

- Motivation for Residential Demand Response (done)
- Related work
- Methods for estimating the counterfactual consumption
  - (Black box) Machine Learning methods
- Estimation of average treatment effects
- User-specific analysis
  - Typical load shapes
  - Variation of ITE across users
- Case study on DR program in California
- Conclusion and further work

Today's Talk

Agenda

- Motivation for Residential Demand Response (done)
- Related work
- Methods for estimating the counterfactual consumption
  - (Black box) Machine Learning methods
- Estimation of average treatment effects
- User-specific analysis
  - Typical load shapes
  - Variation of ITE across users
- Case study on DR program in California
- Conclusion and further work

Today's Talk

Agenda

- Motivation for Residential Demand Response (done)
- Related work
- Methods for estimating the counterfactual consumption
  - (Black box) Machine Learning methods
- Estimation of average treatment effects
- User-specific analysis
  - Typical load shapes
  - Variation of ITE across users
- Case study on DR program in California
- Conclusion and further work

Today's Talk

Agenda

- Motivation for Residential Demand Response (done)
- Related work
- Methods for estimating the counterfactual consumption
  - (Black box) Machine Learning methods
- Estimation of average treatment effects
- User-specific analysis
  - Typical load shapes
  - Variation of ITE across users
- Case study on DR program in California
- Conclusion and further work

Today's Talk

Agenda

- Motivation for Residential Demand Response (done)
- Related work
- Methods for estimating the counterfactual consumption
  - (Black box) Machine Learning methods
- Estimation of average treatment effects
- User-specific analysis
  - Typical load shapes
  - Variation of ITE across users
- Case study on DR program in California
- Conclusion and further work

Today's Talk

Agenda

- Motivation for Residential Demand Response (done)
- Related work
- Methods for estimating the counterfactual consumption
  - (Black box) Machine Learning methods
- Estimation of average treatment effects
- User-specific analysis
  - Typical load shapes
  - Variation of ITE across users
- Case study on DR program in California
- Conclusion and further work

Related Work

Economics

- Experimental Settings (Randomized Control Trials): Measure the counterfactual
- Non-experimental settings: Obtain non-experimental estimates of the counterfactual[2][3]

Engineering

- Short-term load forecasting (STLF):
  - Effect of aggregation size
  - Estimation accuracy
- Load shape analysis
- Load scheduling
- ...

[2] Abadie et al.: Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program, Journal of the American Statistical Association, 2012

[3] Bollinger et al.: Welfare Effects of Home Automation Technology with Dynamic Pricing, 2015

Related Work

Economics

- Experimental Settings (Randomized Control Trials): Measure the counterfactual
- Non-experimental settings: Obtain non-experimental estimates of the counterfactual[2][3]

Engineering

- Short-term load forecasting (STLF):
  - Effect of aggregation size
  - Estimation accuracy
- Load shape analysis
- Load scheduling
- ...

[2] Abadie et al.: Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program, Journal of the American Statistical Association, 2012

[3] Bollinger et al.: Welfare Effects of Home Automation Technology with Dynamic Pricing, 2015

Related Work

Economics

- Experimental Settings (Randomized Control Trials): Measure the counterfactual
- Non-experimental settings: Obtain non-experimental estimates of the counterfactual[2][3]

Engineering

- Short-term load forecasting (STLF):
  - Effect of aggregation size
  - Estimation accuracy
- Load shape analysis
- Load scheduling
- ...

[2] Abadie et al.: Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program, Journal of the American Statistical Association, 2012

[3] Bollinger et al.: Welfare Effects of Home Automation Technology with Dynamic Pricing, 2015

Related Work

Economics

- Experimental Settings (Randomized Control Trials): Measure the counterfactual
- Non-experimental settings: Obtain non-experimental estimates of the counterfactual[2][3]

Engineering

- Short-term load forecasting (STLF):
  - Effect of aggregation size
  - Estimation accuracy
- Load shape analysis
- Load scheduling
- ...

[2] Abadie et al.: Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program, Journal of the American Statistical Association, 2012

[3] Bollinger et al.: Welfare Effects of Home Automation Technology with Dynamic Pricing, 2015

Machine Learning for Short-Term Load Forecasting

- Basic outcome model:

$$y_{it} = f_i(\mathbf{x}_{it}) + D_{it}\beta_i + \epsilon_{it}$$

- Treatment and control data:

$$\mathcal{D}_{i,t} = \{(\mathbf{x}_{it}, y_{it}) \mid t \in \mathcal{T}_i\} \qquad \mathcal{T}_i = \{t \in \mathbb{T} \mid D_{it} = 1\}$$
$$\mathcal{D}_{i,c} = \{(\mathbf{x}_{it}, y_{it}) \mid t \in \mathcal{C}_i\} \qquad \mathcal{C}_i = \{t \in \mathbb{T} \mid D_{it} = 0\}$$

- Two-step strategy to estimate counterfactuals:
  - Estimate conditional mean function $\hat{f}_i(\cdot) : \mathbb{R}^n \mapsto \mathbb{R}$ on control data $\mathcal{D}_{i,c}$ with *any* regression method
  - Obtain counterfactuals: $\hat{y}^0_{ij} = \hat{f}_i(\mathbf{x}^1_{ij}) \quad \forall\, j \in \mathcal{T}_i \Longrightarrow \hat{\beta}_i = \frac{1}{|\mathcal{T}_i|}\sum_{j\in\mathcal{T}_i}(\hat{f}_i(\mathbf{x}^1_{it}) - y^1_{ij})$
- Regression methods used:
  - CAISO 10-in-10 Baseline (BL)
  - Ordinary Least Squares Regression (OLS)
  - LASSO (L1) and Ridge Regression (L2)
  - *k*-Nearest Neighbors Regression (KNN)
  - Decision Tree Regression (DT)
  - Random Forest Regression (RFR)
  - ...

Machine Learning for Short-Term Load Forecasting

- Basic outcome model:

$$y_{it} = f_i(\mathbf{x}_{it}) + D_{it}\beta_i + \epsilon_{it}$$

- Treatment and control data:

$$\mathcal{D}_{i,t} = \{(\mathbf{x}_{it}, y_{it}) \mid t \in \mathcal{T}_i\} \qquad \mathcal{T}_i = \{t \in \mathbb{T} \mid D_{it} = 1\}$$
$$\mathcal{D}_{i,c} = \{(\mathbf{x}_{it}, y_{it}) \mid t \in \mathcal{C}_i\} \qquad \mathcal{C}_i = \{t \in \mathbb{T} \mid D_{it} = 0\}$$

- Two-step strategy to estimate counterfactuals:
  - Estimate conditional mean function $\hat{f}_i(\cdot) : \mathbb{R}^n \mapsto \mathbb{R}$ on control data $\mathcal{D}_{i,c}$ with *any* regression method
  - Obtain counterfactuals: $\hat{y}^0_{ij} = \hat{f}_i(\mathbf{x}^1_{ij}) \quad \forall\, j \in \mathcal{T}_i \Longrightarrow \hat{\beta}_i = \frac{1}{|\mathcal{T}_i|}\sum_{j\in\mathcal{T}_i}(\hat{f}_i(\mathbf{x}^1_{it}) - y^1_{ij})$
- Regression methods used:
  - CAISO 10-in-10 Baseline (BL)
  - Ordinary Least Squares Regression (OLS)
  - LASSO (L1) and Ridge Regression (L2)
  - *k*-Nearest Neighbors Regression (KNN)
  - Decision Tree Regression (DT)
  - Random Forest Regression (RFR)
  - ...

Machine Learning for Short-Term Load Forecasting

- Basic outcome model:

$$y_{it} = f_i(\mathbf{x}_{it}) + D_{it}\beta_i + \epsilon_{it}$$

- Treatment and control data:

$$\mathcal{D}_{i,t} = \{(\mathbf{x}_{it}, y_{it}) \mid t \in \mathcal{T}_i\} \qquad \mathcal{T}_i = \{t \in \mathbb{T} \mid D_{it} = 1\}$$
$$\mathcal{D}_{i,c} = \{(\mathbf{x}_{it}, y_{it}) \mid t \in \mathcal{C}_i\} \qquad \mathcal{C}_i = \{t \in \mathbb{T} \mid D_{it} = 0\}$$

- Two-step strategy to estimate counterfactuals:
  - Estimate conditional mean function $\hat{f}_i(\cdot) : \mathbb{R}^n \mapsto \mathbb{R}$ on control data $\mathcal{D}_{i,c}$ with *any* regression method
  - Obtain counterfactuals: $\hat{y}_{ij}^0 = \hat{f}_i(\mathbf{x}_{ij}^1) \quad \forall\, j \in \mathcal{T}_i \Longrightarrow \hat{\beta}_i = \frac{1}{|\mathcal{T}_i|}\sum_{j\in\mathcal{T}_i}(\hat{f}_i(\mathbf{x}_{it}^1) - y_{ij}^1)$
- Regression methods used:
  - CAISO 10-in-10 Baseline (BL)
  - Ordinary Least Squares Regression (OLS)
  - LASSO (L1) and Ridge Regression (L2)
  - k-Nearest Neighbors Regression (KNN)
  - Decision Tree Regression (DT)
  - Random Forest Regression (RFR)
  - ...

Machine Learning for Short-Term Load Forecasting

- Basic outcome model:

$$y_{it} = f_i(\mathbf{x}_{it}) + D_{it}\beta_i + \epsilon_{it}$$

- Treatment and control data:

$$\mathcal{D}_{i,t} = \{(\mathbf{x}_{it}, y_{it}) \mid t \in \mathcal{T}_i\} \qquad \mathcal{T}_i = \{t \in \mathbb{T} \mid D_{it} = 1\}$$
$$\mathcal{D}_{i,c} = \{(\mathbf{x}_{it}, y_{it}) \mid t \in \mathcal{C}_i\} \qquad \mathcal{C}_i = \{t \in \mathbb{T} \mid D_{it} = 0\}$$

- Two-step strategy to estimate counterfactuals:
  - Estimate conditional mean function $\hat{f}_i(\cdot) : \mathbb{R}^n \mapsto \mathbb{R}$ on control data $\mathcal{D}_{i,c}$ with *any* regression method
  - Obtain counterfactuals: $\hat{y}_{ij}^0 = \hat{f}_i(\mathbf{x}_{ij}^1) \quad \forall\, j \in \mathcal{T}_i \Longrightarrow \hat{\beta}_i = \frac{1}{|\mathcal{T}_i|}\sum_{j\in\mathcal{T}_i}(\hat{f}_i(\mathbf{x}_{it}^1) - y_{ij}^1)$
- Regression methods used:
  - CAISO 10-in-10 Baseline (BL)
  - Ordinary Least Squares Regression (OLS)
  - LASSO (L1) and Ridge Regression (L2)
  - k-Nearest Neighbors Regression (KNN)
  - Decision Tree Regression (DT)
  - Random Forest Regression (RFR)
  - ...

Further Analyses

Statistical Hypothesis Testing

- Wilcoxon Signed Rank Test (paired difference test) to compare $\{y_{ij}^1 \mid j \in \mathcal{T}_i\}$ and $\{\hat{y}_{ij}^0 \mid j \in \mathcal{T}_i\}$
- Null hypothesis H_0: The samples $\{y_{ij}^1 \mid j \in \mathcal{T}_i\}$ and $\{\hat{y}_{ij}^0 \mid j \in \mathcal{T}_i\}$ are drawn from the *same* distribution

$$U \sim \mathcal{N}(\mu, \sigma^2) \qquad \mu = 0, \sigma^2 = \frac{|\mathcal{T}_i|(|\mathcal{T}_i| + 1)(2|\mathcal{T}_i| + 1)}{6} \qquad \text{if } |\mathcal{T}_i| \text{ large enough}$$

- For each user $i \in \mathcal{I}$, determine rank and test statistic to determine a p-value

Variability of Load Shapes

- Use k-means clustering on daily load shapes to find characteristic profiles $\{C_1, \ldots, C_k\}$
- Use *entropy* measure H_i to characterize variability of consumption behavior:

$$H_i = -\sum_{s=1}^{k} p_s(C_s) \log(p_s(C_s))$$

Further Analyses

Statistical Hypothesis Testing

- Wilcoxon Signed Rank Test (paired difference test) to compare $\{y_{ij}^1 \mid j \in \mathcal{T}_i\}$ and $\{\hat{y}_{ij}^0 \mid j \in \mathcal{T}_i\}$
- Null hypothesis H_0: The samples $\{y_{ij}^1 \mid j \in \mathcal{T}_i\}$ and $\{\hat{y}_{ij}^0 \mid j \in \mathcal{T}_i\}$ are drawn from the *same* distribution

$$U \sim \mathcal{N}(\mu, \sigma^2) \qquad \mu = 0, \sigma^2 = \frac{|\mathcal{T}_i|(|\mathcal{T}_i| + 1)(2|\mathcal{T}_i| + 1)}{6} \qquad \text{if } |\mathcal{T}_i| \text{ large enough}$$

- For each user $i \in \mathcal{I}$, determine rank and test statistic to determine a p-value

Variability of Load Shapes

- Use k-means clustering on daily load shapes to find characteristic profiles $\{C_1, \ldots, C_k\}$
- Use *entropy* measure H_i to characterize variability of consumption behavior:

$$H_i = -\sum_{s=1}^{k} p_s(C_s) \log(p_s(C_s))$$

Further Analyses

Statistical Hypothesis Testing

- Wilcoxon Signed Rank Test (paired difference test) to compare $\{y_{ij}^1 \mid j \in \mathcal{T}_i\}$ and $\{\hat{y}_{ij}^0 \mid j \in \mathcal{T}_i\}$
- Null hypothesis H_0: The samples $\{y_{ij}^1 \mid j \in \mathcal{T}_i\}$ and $\{\hat{y}_{ij}^0 \mid j \in \mathcal{T}_i\}$ are drawn from the *same* distribution

$$U \sim \mathcal{N}(\mu, \sigma^2) \qquad \mu = 0, \sigma^2 = \frac{|\mathcal{T}_i|(|\mathcal{T}_i| + 1)(2|\mathcal{T}_i| + 1)}{6} \qquad \text{if } |\mathcal{T}_i| \text{ large enough}$$

- For each user $i \in \mathcal{I}$, determine rank and test statistic to determine a p-value

Variability of Load Shapes

- Use k-means clustering on daily load shapes to find characteristic profiles $\{C_1, \ldots, C_k\}$
- Use *entropy* measure H_i to characterize variability of consumption behavior:

$$H_i = -\sum_{s=1}^{k} p_s(C_s) \log(p_s(C_s))$$

Further Analyses

Statistical Hypothesis Testing

- Wilcoxon Signed Rank Test (paired difference test) to compare $\{y_{ij}^1 \mid j \in \mathcal{T}_i\}$ and $\{\hat{y}_{ij}^0 \mid j \in \mathcal{T}_i\}$
- Null hypothesis H_0: The samples $\{y_{ij}^1 \mid j \in \mathcal{T}_i\}$ and $\{\hat{y}_{ij}^0 \mid j \in \mathcal{T}_i\}$ are drawn from the *same* distribution

$$U \sim \mathcal{N}(\mu, \sigma^2) \qquad \mu = 0, \sigma^2 = \frac{|\mathcal{T}_i|(|\mathcal{T}_i| + 1)(2|\mathcal{T}_i| + 1)}{6} \qquad \text{if } |\mathcal{T}_i| \text{ large enough}$$

- For each user $i \in \mathcal{I}$, determine rank and test statistic to determine a p-value

Variability of Load Shapes

- Use k-means clustering on daily load shapes to find characteristic profiles $\{C_1, \ldots, C_k\}$
- Use *entropy* measure H_i to characterize variability of consumption behavior:

$$H_i = -\sum_{s=1}^{k} p_s(C_s) \log(p_s(C_s))$$

Further Analyses

Statistical Hypothesis Testing

- Wilcoxon Signed Rank Test (paired difference test) to compare $\{y_{ij}^1 \mid j \in \mathcal{T}_i\}$ and $\{\hat{y}_{ij}^0 \mid j \in \mathcal{T}_i\}$
- Null hypothesis H_0: The samples $\{y_{ij}^1 \mid j \in \mathcal{T}_i\}$ and $\{\hat{y}_{ij}^0 \mid j \in \mathcal{T}_i\}$ are drawn from the *same* distribution

$$U \sim \mathcal{N}(\mu, \sigma^2) \qquad \mu = 0, \sigma^2 = \frac{|\mathcal{T}_i|(|\mathcal{T}_i| + 1)(2|\mathcal{T}_i| + 1)}{6} \qquad \text{if } |\mathcal{T}_i| \text{ large enough}$$

- For each user $i \in \mathcal{I}$, determine rank and test statistic to determine a p-value

Variability of Load Shapes

- Use k-means clustering on daily load shapes to find characteristic profiles $\{C_1, \dots, C_k\}$
- Use *entropy* measure H_i to characterize variability of consumption behavior:

$$H_i = -\sum_{s=1}^{k} p_s(C_s) \log(p_s(C_s))$$

Case Study on DR Program

- Hourly smart meter time series data on residential customers in California
- Hourly ambient air temperature scraped from public data sources
- Data preprocessing
  - Drop obvious outliers
  - Drop negative consumption data (and all users on net energy metering)
- Split cleaned hourly consumption data into three sets:
  - $\mathcal{D}_{i,t}$, i.e. the treatment data set
  - $\mathcal{D}_{i,p} \subset \mathcal{D}_{i,c}$, the placebo treatment set
  - $\mathcal{D}_{i,tr} = \mathcal{D}_{i,c} \setminus \mathcal{D}_{i,p}$, the training data set
- Choose covariates:
  - Hour of the day as a categorical variable
  - Ambient air temperature
  - Previous n_{ar} hourly consumption values

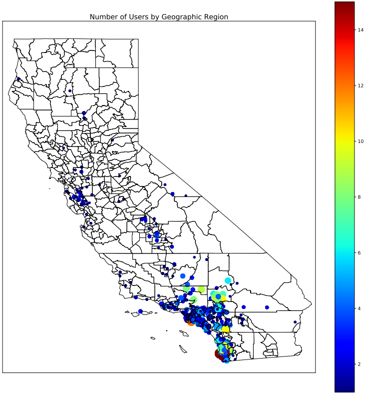

Figure: Distribution of Users

Case Study on DR Program

- Hourly smart meter time series data on residential customers in California
- Hourly ambient air temperature scraped from public data sources
- Data preprocessing
  - Drop obvious outliers
  - Drop negative consumption data (and all users on net energy metering)

- Split cleaned hourly consumption data into three sets:
  - $\mathcal{D}_{i,t}$, i.e. the treatment data set
  - $\mathcal{D}_{i,p} \subset \mathcal{D}_{i,c}$, the placebo treatment set
  - $\mathcal{D}_{i,tr} = \mathcal{D}_{i,c} \setminus \mathcal{D}_{i,p}$, the training data set
- Choose covariates:
  - Hour of the day as a categorical variable
  - Ambient air temperature
  - Previous n_{ar} hourly consumption values

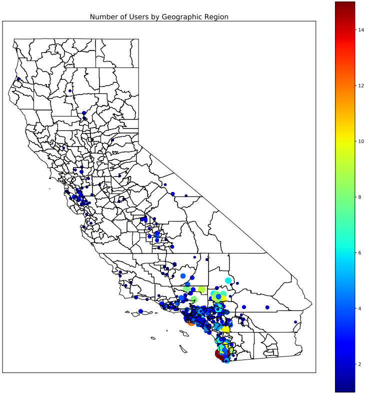

Figure: Distribution of Users

Case Study on DR Program

- Hourly smart meter time series data on residential customers in California
- Hourly ambient air temperature scraped from public data sources
- Data preprocessing
  - Drop obvious outliers
  - Drop negative consumption data (and all users on net energy metering)

- Split cleaned hourly consumption data into three sets:
  - $\mathcal{D}_{i,t}$, i.e. the treatment data set
  - $\mathcal{D}_{i,p} \subset \mathcal{D}_{i,c}$, the placebo treatment set
  - $\mathcal{D}_{i,tr} = \mathcal{D}_{i,c} \setminus \mathcal{D}_{i,p}$, the training data set
- Choose covariates:
  - Hour of the day as a categorical variable
  - Ambient air temperature
  - Previous n_{ar} hourly consumption values

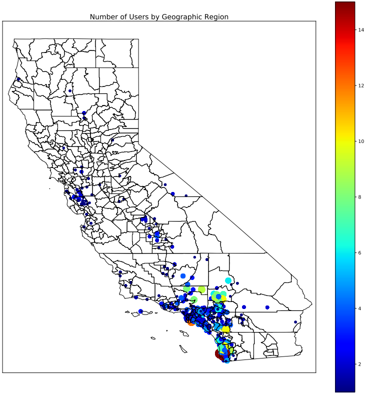

Figure: Distribution of Users

Case Study on DR Program (cont'd)

Prediction Accuracy

- Define Mean Absolute Percentage Error (MAPE) as metric for prediction accuracy:

$$\text{MAPE} = \frac{1}{|\mathcal{D}_{i,v}|} \sum_{j \in \mathcal{D}_{i,v}} \left| \frac{\hat{f}_i(\mathbf{x}_{ij}) - y_{ij}}{y_{ij}} \right| \cdot 100\%$$

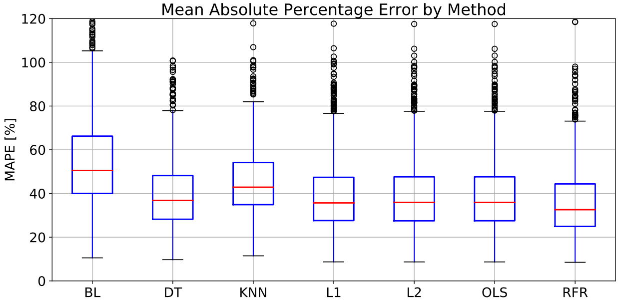

Figure: MAPEs by Prediction Method

- CAISO baseline performs worst

Case Study on DR Program (cont'd)

ITEs by Geographic Region

- Larger reductions in warmer regions

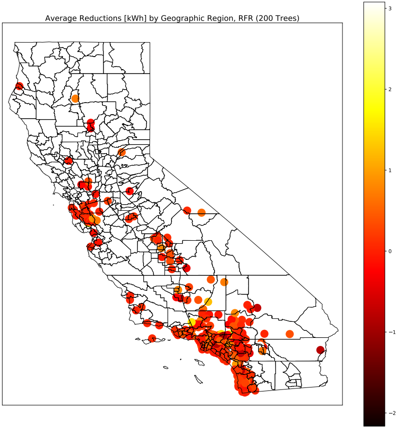

Figure: Geographic Distribution of ITEs

Case Study on DR Program (cont'd)

Average Treatment Effect Estimation

- Random Forest performs best
- Highest reductions in the evening
- Placebo events show ≈ zero reduction

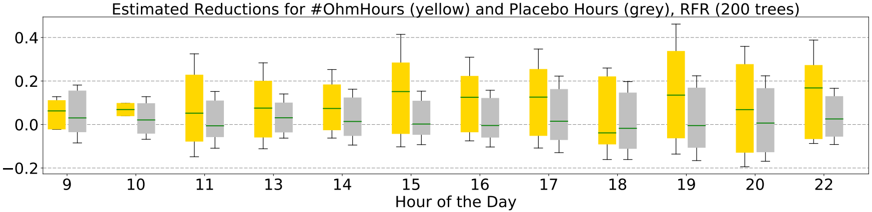

Figure: Estimated Reductions by Hour of the Day

Types of Load Shapes

- Morning + evening peak
- Daytime peak
- Night peak
- Evening peak

Interesting Observation

- Users with higher entropy reduce more

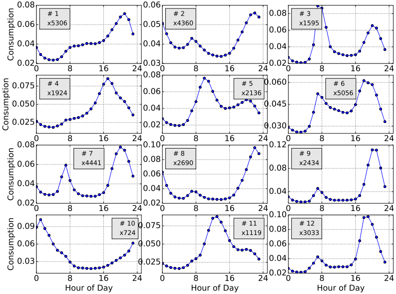

Figure: Characteristic Load Shapes

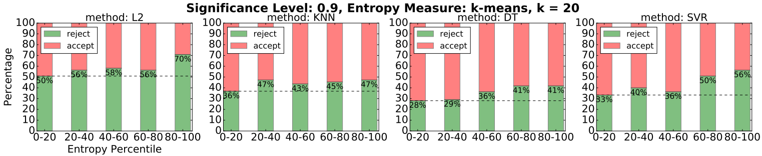

Figure: Percentage of Rejected Nulls vs. Forecasting Method

Conclusion

- Description of Residential Demand Response using ordinary machine learning methods
- Estimation of counterfactual consumption during DR events
- Presented Black-box Machine Learning Methods
  - OLS, L1, L2, KNN, DT, RFR
  - Random Forest has lowest MAPE
- Identified “dictionary” of load shapes to compute variability of consumption
- Discovered a higher percentage of reduction among more variable users

Further Work

Completed/In Progress

- Improve estimation of counterfactual by using **latent variables**[4]
  - Hidden Markov Model
  - Mixture Models
- Nonparametric Estimators[5]
- Analysis of bias and variance in estimation process[5]
- Mechanism Design for DR elicitation

Upcoming

- Time series modeling to estimate causal impact of DR interventions
- Run **Randomized Control Trial (RCT)** to
  - Validate non-experimental estimates of DR reduction
  - Target most susceptible users for DR incentives

[4] D. Zhou, M. Balandat, C. Tomlin: A Bayesian Perspective on Residential Demand Response Using Smart Meter Data. 54th Annual Allerton Conference on Communication, Control, and Computing, September 2016

[5] M. Balandat: PhD Thesis. University of California, Berkeley, 2016

Thank You!
Questions?