Smoothers, harmonic autoregression, and gradient boosting for seasonal time series:

Predicting hourly electricity demand in Ontario

UNIVERSITY OF ALBERTA

Xilai Fu, Danyi Liu, Ree Yang, Yijia Yu

Department of Mathematical and Statistical Sciences, University of Alberta, Edmonton, AB, Canada

Background and Objectives

Models able to accurately predict hourly electricity demand can help energy production meet trending energy demands, both in terms of long term trends over years, and finely across seasons and days of the week. Where new, emerging sources of energy require greater consideration for energy storage, properly anticipating energy demand can ensure that residents and industries have a stable electricity supply.

In this case study, we build predictive models for sector-specific hourly electricity demand that are able to predict annual demand. We do so using 14 years of real-world data from Ontario, including sector-aggregated hourly demand, sector-specific annual totals, and population-weighted hourly weather data. Our methods focus specifically on seasonal trend estimation using a variety of techniques including smoothers, harmonic regression, autoregression, and regression trees to address real-world difficulties in the data.

Data

Target: Sector-Specific Annual Electricity Demand

- Five sectors: residential, industrial, commercial, agriculture, and transportation
- Data includes annual totals from 2003 to 2016

Hourly Total Electricity Demand

- Total energy demand (aggregated over sectors), measured every hour during 2003 to 2016

Hourly Weather

- Population-weighted air and weather data (precipitation, snowfall, air density, solar irradiation, cloud cover fraction), measured every hour

Seasonality Exploration

- Strong yearly, weekly, and daily trends evident from autocorrelation plots of hourly demand.

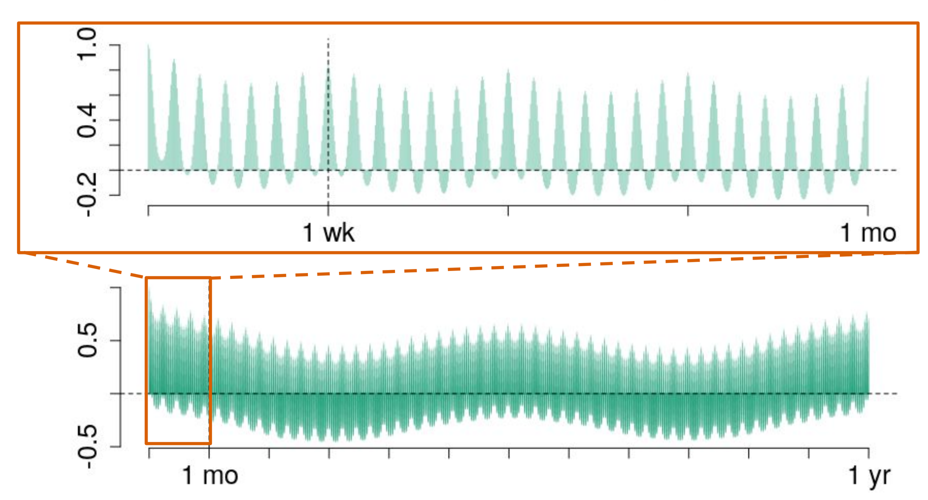

- Seasonal trend modelling is clearly central to modelling long-term hourly electricity demand

Pre-processing

Power Outage Imputation

- Observed electricity usage during the August 2003 "Northeast Blackout" adds noise to the data.
- We imputed data from 4:00pm August 14 to 8:00am August 18, 2003 by averaging demand at the same hour in the surrounding 3 days.

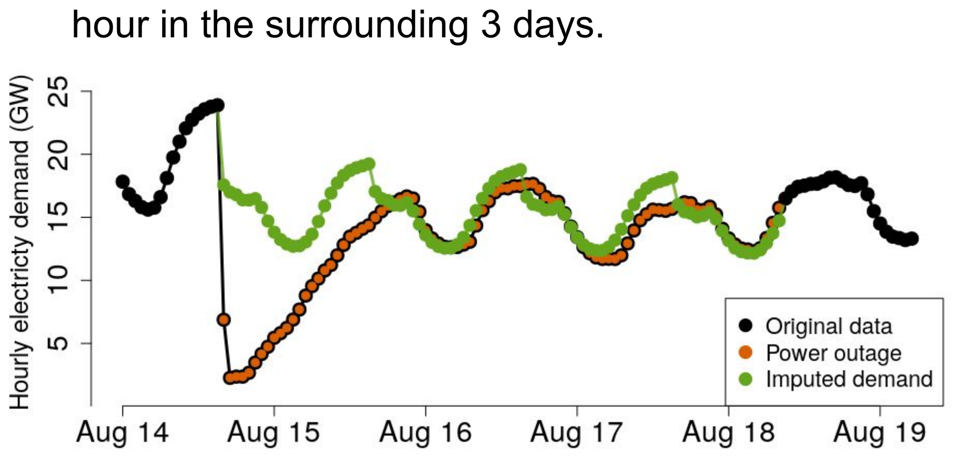

- Intuitively, imputation lessens the effect on seasonal trend estimates with little effect on model predictions

Annual / Total Hourly Demand Discrepancy

- A discrepancy exists between summed hourly demand and annual demand targets that changes over time.
- We adjust hourly demand by the discrepancy, assumed uniform over the year.

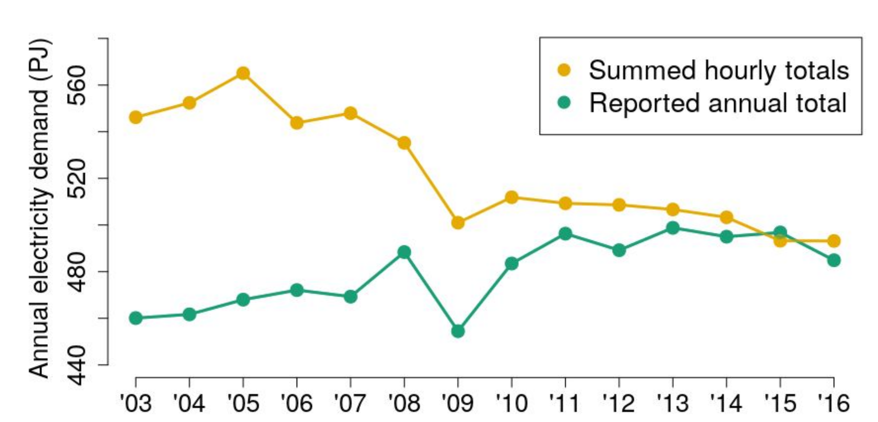

Sector Usage Proportions

- To predict sector specific usage from hourly total for the test year, we apply a moving average (with window 1) to annual usage targets.

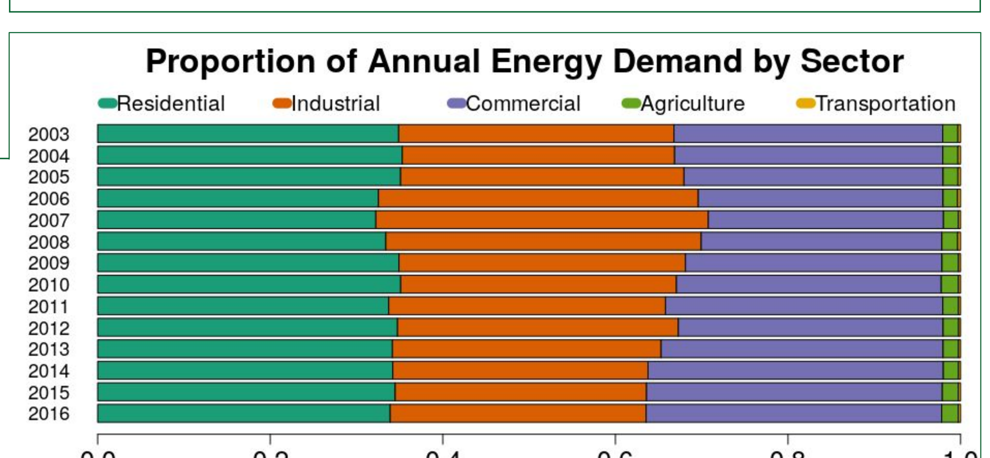

Methods

MSTL Model

- The **m**ultiple **s**easonality **t**ime **s**eries decomposition via **L**OESS method flexibly estimates time series components using smoothing estimators.
- LOESS and Friedman's super-smoother estimate long-term and seasonal trends (with pre-specified periods).
- Forward and backward AR(2) models are fit to residuals to improve estimates.

Dynamic Linear (Quantile) Models (DYNLM/DYNQR)

- Linear regression framework allowing lagged predictors and harmonic terms.
- A harmonic predictor with order p and order K includes terms $\sin(2k\pi t/p)$, $\cos(2k\pi t/p)$, $k = 1, \ldots, K$
- We fit both a least squares (DYNLM) and median (quantile) regression (DYNQR) model.

Harmonic Autoregressive Model (AR Model)

- Estimates ARIMA time series parameters after accounting for model covariates in a linear regression model.

XGBOOST Model

- A popular tree-based supervised learning method known for its high predictive power
- Uses boosting to build an ensemble of (regression) trees that balances model complexity by penalizing tree depth, the importance of each split

Results

Evaluation

- Models are evaluated using leave-one-out absolute error (MAE) in predicting sector-specific energy demand, averaged over 14 years. As a baseline, we also consider a Naive model fit to annual demand.
- All models (except XGBOOST) average predictions over a forward and backward predictive model when training data comes before and after the test year.

MSTL Model

- Predictors for residuals: time, temperature^2. Residuals showed no relationship with other weather variables.
- Optimal window size (77760) was tuned via cross-validation.
- Sudden decrease in demand in 2009 was not detected by smoothers.
- Flexibility of smoother estimators is a benefit, but its variability may negatively impact test set predictions for some years. The AR residual model reduced the impact of smoother variability.

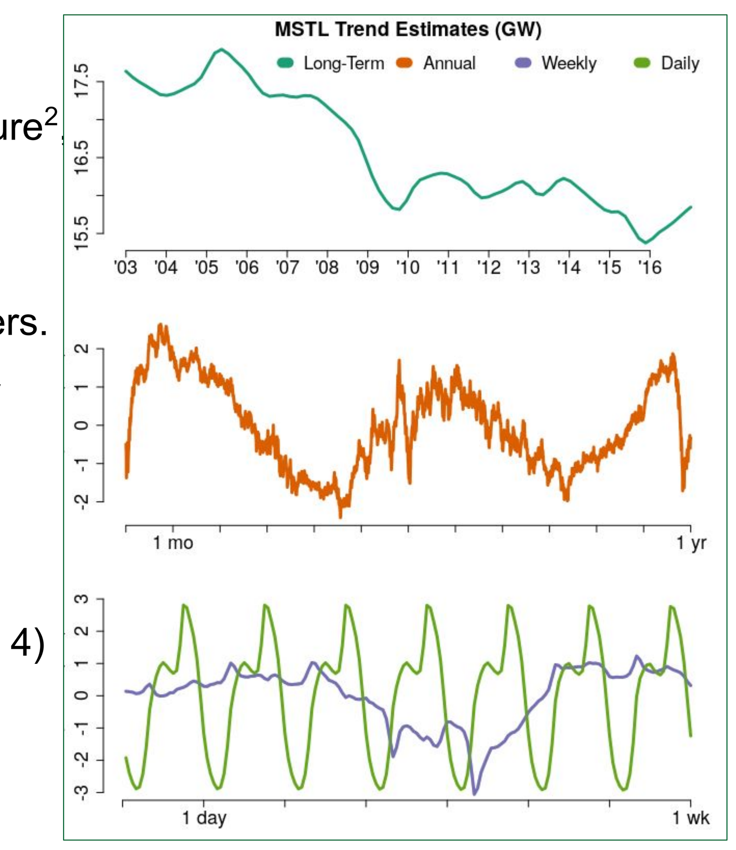

AR Model

- Harmonic terms: yearly, weekly, and daily periods (orders 2, 6, and 4)
- Residual diagnostics showed that an ARMA(3,0) model with harmonic terms, time, temperature^2 air density^2 was suitable

DYNLM/DYNQR Models

- Weather predictors: linear deviation in temperature from 18° and air density below 1.145 kg/m^3. Harmonic terms: yearly, weekly, and daily periods with orders 2, 1, and 6. Other predictors: weekend indicator and 1-year-lagged electricity demand.
- Other weather predictors could not be included, likely due to high correlation and overlap in variability explained by the harmonic terms and weather predictors.

XGBOOST Model

- Predictors: time (year, month, hour, and holiday indicator), temperature, and lagged electricity demand
- The first two lags had greatest predictive power while inclusion of the third lag decreases model performance. Other weather predictors are not informative in prediction.

Test MAE by Sector	MSTL	AR	DYNQR	DYNLM	XGBOOST	Naive
Residential	**3.281**	3.903	3.660	3.678	3.375	4.721
Industrial	4.590	4.647	5.243	5.268	**4.106**	6.244
Commercial	3.570	3.974	3.473	**3.431**	4.320	3.739
Agriculture	0.238	0.218	**0.196**	**0.196**	0.247	0.228
Transportation	0.129	0.128	0.130	0.130	0.137	**0.084**
Weighted MAE	**3.735**	4.091	4.047	4.048	3.842	4.820

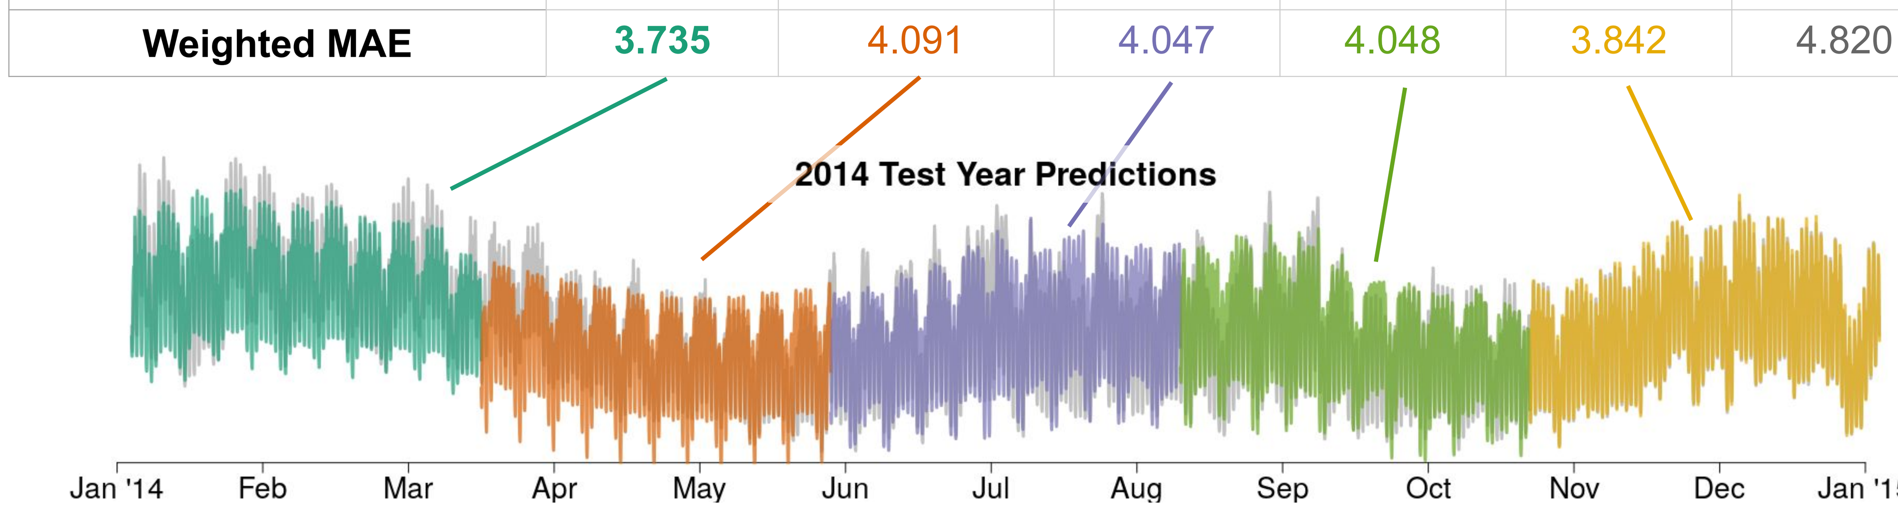

Conclusions

Electricity Demand Prediction

- **Overall, the MSTL model seems to perform best** relative to other models. We suspect this is due to the flexibility of the smoothing estimators and our residual AR model that accounts for high smoother variability. DYNQR performed reliably across all sectors, though did not greatly outperform other models in any.
- Accounting for harmonic and lagged predictors, **weather covariates were generally found to not be useful**. However, the dynamic linear and quantile models found extreme temperature and air density values useful in predicting spikes in hourly energy usage.
- Comparable performance of statistical models with XGBOOST suggests **continued utility of traditional methods**.

Challenges and Limitations

- We relied on a moving average of annual totals to predict sector-specific demand from total demand. Other methods can be explored to improve the quality of sector-specific proportion estimates.
- The discrepancy between annual targets and summed hourly demand added further noise and degraded model performance. The source of this discrepancy (presumably transmission loss) is unconfirmed.

References and Acknowledgements

The authors acknowledge the support of the Training Consulting Centre (TCC) at the University of Alberta in organizing participation in the case study and the supervision of TCC consultant Matthew Pietrosanu throughout.

References relevant to this work are available through the accompanying QR code.