What’s going on in there? Interpretable Machine Learning for Deeper Energy Savings

EFFICIENCY VERMONT R&D PROJECT: DEEPER ENERGY SAVINGS THROUGH ADVANCED REGRESSION MODELING

Abiodun Iwayemi, Jeff Hullstrung, Dan Fredman, and Nick Neverisky
March 30, 2021

Introduction

Why is artificial intelligence being woven into the fabric of everyday life and work, and what are the benefits and challenges that come with this ever-increasing usage? Due to the power and flexibility of artificial intelligence, it is used for mundane tasks such as asking a voice assistant to turn off your living room lights, automatically tagging pictures of friends and family on Facebook, helping autonomous vehicles navigate city roads and pedestrians, translating text from one language to another, or powering automated customer service “bots” on websites. This is all possible via machine learning, a process by which computer algorithms learn patterns and relationships from large amounts of data, and then use the learned models to make predictions or decisions on new or unseen data.

These more powerful models are capable of amazing feats but the growth in their capabilities is accompanied by increasing model complexity and opacity of model decision-making. This is a challenge in regulated sectors like the energy efficiency industry, where it’s important to get the correct answer, and to know what factors influenced the decisions or outputs of a model. This is especially important when the models can learn spurious relationships, contain hidden biases1 or fail in unexpected ways. The adoption of machine learning in such sectors is a function of how trustworthy these models are, and that trust hinges on being able to explain how the models make their decisions.

Rather than stay with the status quo and default to using less powerful but more explainable machine learning models, Efficiency Vermont set out to learn if and how explainable artificial intelligence (XAI)2 approaches could drive adoption of cutting-edge machine learning models for energy modeling at customer facilities. Achieving these outcomes requires demonstrating that these models provide consistently superior performance than current approaches, and that Efficiency Vermont can communicate the model outputs to energy modelers, evaluators, and customers, in simple but more insightful ways than the current visualizations.

[1] Julia Angwin Mattu Jeff Larson,Lauren Kirchner,Surya, “Machine Bias,” ProPublica, May 23, 2016, https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing.
2 David Gunning and David Aha, “DARPA’s Explainable Artificial Intelligence (XAI) Program,” *AI Magazine* 40, no. 2 (June 24, 2019): 44–58, https://ojs.aaai.org//index.php/aimagazine/article/view/2850.

Why Machine Learning Matters in the Energy Efficiency Sector

Regulatory frameworks are the structure for energy efficiency program accountability and the proving ground for the cost-effective use of ratepayer funds. The energy efficiency sector offers an excellent platform for demonstrating how well artificial intelligence can improve program performance. State-of-the-art models such as deep neural networks and ensemble models like Gradient Boosted Regression Trees[3] (GBRTs) are among the most powerful and complex machine learning models. These models capture greater energy savings for energy efficiency utilities and energy service companies via more accurate and precise regression energy modeling. Unfortunately, these classes of machine learning models are black box models.

In black box models, machine learning algorithms map input values (for example, temperature or day of week) to output predictions (such as power consumption) in non-linear, complex ways that are difficult for humans to understand. Unlike interpretable or explainable models, like ordinary least squares (OLS) regression[4], the input to output relationship of a black box model cannot be expressed in a simple equation or decision tree. This leads to a tradeoff between model interpretability and accuracy. The most interpretable machine learning models—like OLS, the linchpin of energy regression modeling —are less flexible and have lower accuracy while more accurate methods, like the deep neural networks and ensemble methods[5], are less interpretable.

What Is the Machine Learning Problem that Efficiency Vermont Has Solved?

Energy efficiency is a field in which model explainability is very valuable. Efficiency Vermont engineers provide technical assistance to building energy managers by creating energy models of customer facilities. These models use historical consumption data to predict future energy use from equipment and processes, and the drivers of that energy consumption. Predicting, however, is not the same as explaining. The current approach uses OLS-based time-of-week and temperature regression (TOWT) [6] with tried and tested algorithms developed by the Lawrence Berkeley National Laboratory. It provides good results, but does so at the expense of explainability - the regression equation has 176 parameters. Blind trust in any model is risky, and simply telling a customer (or regulator) to trust the model is not an option. How does one go about predicting *and* explaining the drivers of facility energy usage in a compelling way so that building managers can confidently make business decisions about energy efficiency investments?

Efficiency Vermont researchers piloted a hybrid model that supports the use of cutting-edge algorithms and the design of insightful, informative visualizations. The researchers produced predictions that are highly accurate, intuitive, and easy to understand.

Project Objectives
The goal of the pilot project's research was to answer two questions:

---

[3] Peter Prettenhofer and Gilles Louppe, "Gradient Boosted Regression Trees in Scikit-Learn," 2014, https://orbi.uliege.be/bitstream/2268/163521/1/slides.pdf.
[4] Bob Nau, "Notes on Linear Regression Analysis," accessed March 29, 2021, http://people.duke.edu/%7Ernau/Notes_on_linear_regression_analysis--Robert_Nau.pdf.
[5] "Ensemble Learning," in *Wikipedia*, March 21, 2021, https://en.wikipedia.org/w/index.php?title=Ensemble_learning.
[6] Mathieu, Johanna L., Philip N. Price, Sila Kiliccote, and Mary Ann Piette, 2011. "Quantifying Changes in Building Electricity Use, With Application to Demand Response." Berkeley, CA: Lawrence Berkeley National Laboratory. https://www.osti.gov/servlets/purl/1048308.

- To what extent do advanced, black box machine learning models such as GBRTs outperform OLS-based time-of-week and temperature regression models?
- What are the best methods and charts for effectively and simply communicating the outputs of these machine learning models of a facility's power use to broad audiences with varying levels of visualization expertise so that those audiences can make energy efficiency decisions?

Methods

The team selected eight grocery stores in Vermont for validating the proposed approach. The grocery store average size is 50,000 square feet, and the total annual energy consumption for all stores was approximately 14,000 MWh. Efficiency Vermont analyzed the stores previously for large multi-measure projects using it's standard weather normalization process, making the stores good candidates for this new approach. Efficiency Vermont evaluated both TOWT, and GBRT's (using Microsoft Research's LightGBM[7] software library) on this data and compared model performance in terms of prediction accuracy (R^2) and precision (CV-RMSE).

Results

The GBRT models provided more accurate and precise energy modeling results (approximately 4 percent average improvement in adjusted R^2, and a 15 percent lower CV-RMSE)[8] than TOWT regression modeling.

Using the SHAP[9] explainable artificial intelligence framework, in conjunction with focus group interviews of 21 VEIC[10] employees, the team selected the most effective visualizations and key chart elements for communicating the model results. The focus group represented VEIC's marketing, engineering, customer service, and account management departments. The interviews sought information about the data visualizations that most clearly explained the outputs of the regression models. This exploration yielded general agreement on the most effective visualizations and chart elements for successfully communicating model performance to a likely audience of facilities managers, building owners, and operations staffs. It also provided insights into the drivers of energy use in a facility. Findings from the focus groups:

- Visually easy-to-understand charts offer good representations of data about power use. The charts are especially useful for facility managers who want a tool to understand how and where the facility is using power, and to communicate that information to their organizations; and for Efficiency Vermont technical and account management staff who are seeking information about a customer site's power use.
    - Clarity of chart labels and design elements, such as helper text and color-blind-friendly color schemes, are important for ensuring that charts can be easily understood and are accessible to as broad an audience as possible.
- Presenting power use in terms of the customer's dollar costs (rather than energy use costs in kWh) makes it easier for the customer to relate the data to energy behaviors and to take action on energy efficiency investments.

---

[7] *Microsoft LightGBM* (2016; repr., Microsoft, 2021), https://github.com/microsoft/LightGBM.
[8] Co-efficient of variation of root-mean squared error, a measure of the differences between a model's predicted values and the values actually achieved or observed.
[9] Scott M. Lundberg and Su-In Lee, "A Unified Approach to Interpreting Model Predictions," *Advances in Neural Information Processing Systems* 30 (2017): 4765–74. https://arxiv.org/pdf/1705.07874.pdf.
[10] VEIC operates Efficiency Vermont.

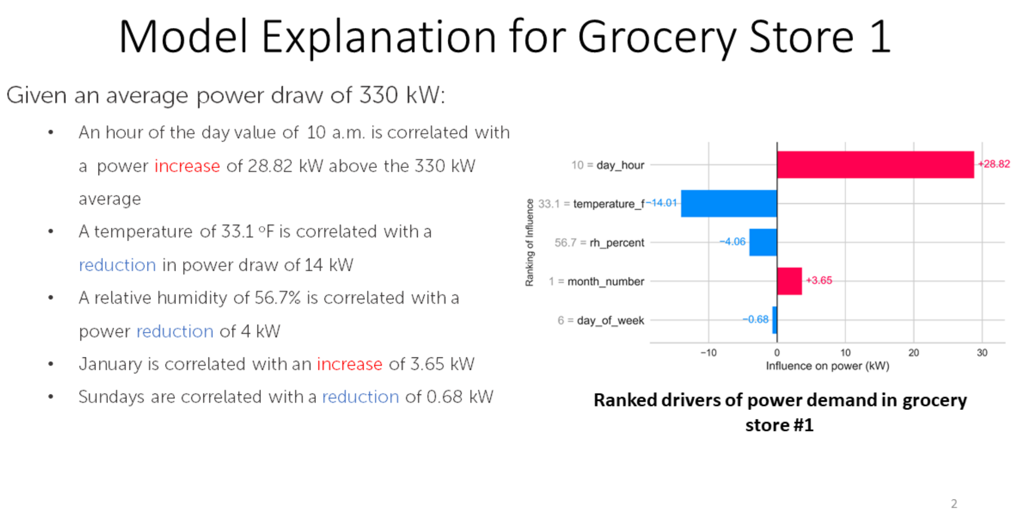

Figure 1. How the SHAP library explains the model's inputs and their effects on prediction of power draw at grocery store.

The project's hybrid approach—combining advanced algorithms and customer-centric data visualization—produced more accurate predictions of facility energy consumption and detailed useful insights about the factors influencing their energy consumption. Efficiency Vermont will leverage this new template to create high-value energy reports that help commercial and industrial customers more quickly make decisions about energy efficiency investments. The research project set the stage for designing automated processes and reporting templates for Efficiency Vermont staff in their reports to customers and evaluators.