Machine Learning

Using Machine Learning to Improve Demand Response Forecasts

CONTENTS

INTRODUCTION

The utility industry, like many industries, is currently undergoing a data revolution. Metering infrastructure is dramatically improving, data storage costs are plummeting and computing power is still following Moore's famous law of exponential improvement. Yet many utilities are unable or unsure about how to take advantage of this modern-day gold. At Itron, we are tackling ways to turn this flood of data into insights that promise to help our customers improve forecasting of demand response (DR) control events. This will increase efficiency, help with targeting and ultimately improve utilities' bottom lines.

This white paper gives an overview of how Itron is using current techniques in data science and machine learning to bring improved demand response forecasts to our customers, primarily using our IntelliSOURCE®- Analytics forecasting tool.

WHAT IS MACHINE LEARNING?

Machine learning is a rapidly growing field that has its roots in Artificial Intelligence. It brings together the fields of mathematics and computer science, and may be thought of as creating computer-generated predictions based on past observations and future expected parameters. Over time, as the model encounters more and more examples of a behavior, its predictions should become better and better. Some examples of machine learning that you may be familiar with are:

» Your favorite streaming movie service giving movie recommendations based on past movies watched, browsed and rated.

» Online retailers giving product recommendations based on past purchases and purchases of people similar to you.

» Self-driving cars! The more they drive, and the more diverse situations they encounter, reportedly the safer they become.

But this isn't (yet) invasion of the robots! The machine learning models take a significant amount of human interaction to develop—from curating the data and choosing the best input "features," to selecting a model, then evaluating and tuning the model—repeating those steps several times and finally deploying a production-ready model to a larger system.

HOW DOES ITRON USE MACHINE LEARNING?

Machine Learning is also proving to be extremely valuable for demand response. Historically, utilities have forecasted "curtailable load" during a control event by using look-up tables or rules-of-thumb, where the forecast relies upon weather variables, day-of-week variables, time of year or holiday variables and perhaps some estimates using recent similar events. While this method of forecasting has its strengths, such as close utility employee engagement, it can also be time-consuming, difficult to do in an emergency situation and slow to respond to systematic changes like changes in energy efficiency or installation of more rooftop solar.

Itron is changing that paradigm by integrating machine learning into our IntelliSOURCE-Analytics forecasting tool. Rather than relying upon specific rules to forecast curtailable load during a DR event, we are using data to "train" the models to find the rules themselves. Over time, as the models "learn" from more experience (more DR events), the forecasts become more accurate. As with all forecasts, there is some amount of uncertainty: uncertainty with the weather forecasts, natural variation in the utility customers' behavior, etc. So we can't expect our forecasts to be perfect, even if that is the goal. But, by using machine learning to forecast curtailable load during a DR event, we are making forecasts more accurate, more reliable, more scalable and ultimately faster to respond to system-wide changes.

We have experimented with a variety of methods, including linear regression models with dozens of factor features, random forests, gradient boosted machines, neural networks, ridge regression and others. Experimenting with many different algorithm and feature combinations, modeling and evaluating the models takes months of analytic work on our part; yet, we are always keeping an eye out for methods that may give us an additional edge.

AN EXAMPLE: IMPROVING DEMAND RESPONSE FORECASTS

Here is an example of how we create, evaluate, tune and deploy a machine learning model for demand response forecasting.

For illustration purposes, we'll take a subset of approximately 1,900 premises enrolled in a DR program for one of our customers—a medium-sized municipal utility. Each of the premises has a smart meter, taking energy usage readings every 15 minutes, and we have approximately three years' worth of data for each meter. This gives us over 180 million data points—a modest, but not overwhelming data set size. The first step is to inspect and clean the data, removing faulty readings and interpolating during meter "offline" times. Then we aggregate the meter data, and merge with weather data and other calendar variables, to create a time series of actual energy readings over time. For this particular utility, we also import data from IntelliSOURCE® Enterprise™ to give us information about when previous DR events occurred, and which premises participated (not all 1,900 premises will participate in every event).

Then, we create two data sets—one "training" set, and one "test" set. The test set is the time around a known DR event, and the training set is all the time before the DR event. After that, we feed the training set to an algorithm, which creates a model. This model uses all of the data to "learn" the tendencies of the DR events— perhaps there is more curtailable load at 5pm than at 2pm, or more on a Thursday afternoon than a Saturday afternoon, or less on a 90-degree day in May than a 90-degree day in July. The more examples of DR events that the model can train on and the more diversity in the conditions around the events, the more ability it has to create accurate forecasts.

After creating the model, we use the test set to generate a "forecast." Forecast is in quotes there because it's not a true forecast, since the event already happened. However, the model does not use the test set's actual energy levels in its prediction, but uses the actual energy level to evaluate the accuracy of the prediction.

The annotated figure below shows one such early-run model, with the predictions plotted against actual average load (averaged over 15-minute intervals). From this plot, we can see our predicted load (crimson bars) is matching the actual load (black line) fairly well prior to the DR event, which started at 15:00 local time. Our curtailable load prediction (dark blue bars), however, only matches about 2/3 of the actual curtailable load during the event.

Model 1: Predicted Curtailable Load During 2.5-hr DR Event, Approx. 1,900 Premises
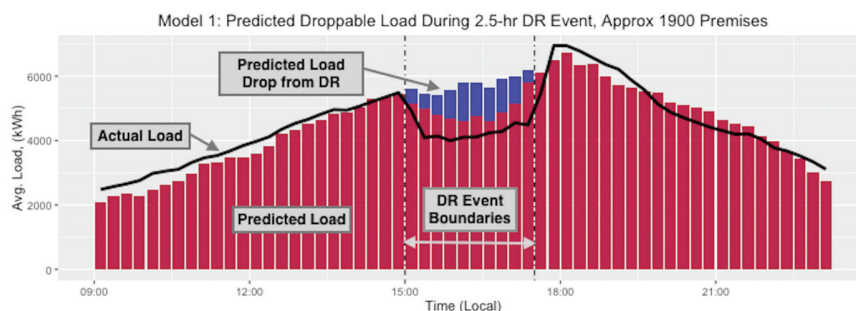
This forecast is fairly good, but we believe it could be better. After adjusting some of the model parameters and inputs, we get a second model and its predictions, shown below. We can see that the model does a much better job predicting the curtailable load.

Model 2: Predicted Curtailable Load During 2.5-hr DR Event, Approx. 1,900 Premises
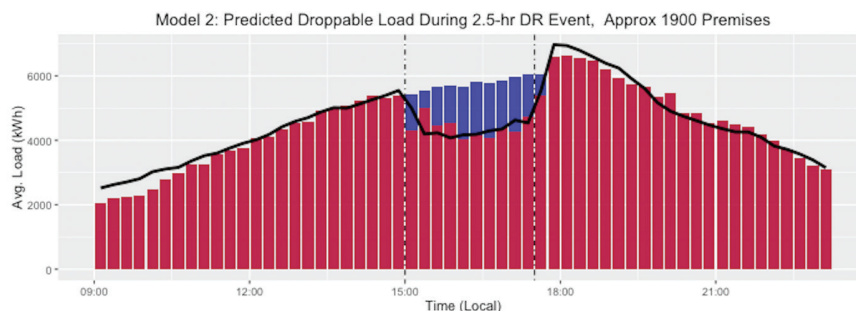
But we don't stop here. Just because a model works best in a single instance doesn't mean it is always the best model. We continue to run models—creating dozens or hundreds of them—and compare the results across several different test sets. Eventually, we must settle on a model that generalizes well to many situations and provides useful information to the users.

INTEGRATING MODELS INTO INTELLISOURCE-ANALYTICS

After the intensive process of modeling, Itron data scientists work on developing production-ready model code to integrate with the rest of our application and work with the application developers to help them understand what data is needed, and in what format. Necessarily, the analytics and modeling is sometimes several months in front of the progress of the application developers, because in some ways our results guide the direction of product development (it's a two-way street that requires a lot of communication, however).

The IntelliSOURCE-Analytics forecasting tool includes two major forms of forecasting models. First, we have system load forecasts, which are also machine learning models that we have developed. Secondly, our demand response forecasts are built as a second layer on top of the system load forecasts—and curtailable load is subtracted from the total forecasted system load. We keep the models separate because our goal with the application is

to optimize DR events, meaning we want to be able to split the participant populations in whatever way we choose. This requires the DR model to sit on top of the system load forecast model, instead of being completely integrated.

Two screenshots of the IntelliSOURCE-Analytics forecasting tool are shown below. This particular configuration shows actual load, actual and forecast average hourly power (labeled "energy" in the app), current and forecast temperature and forecast curtailable load. The curtailable load forecasts for each are the large-font numbers on top of the forecast bars in the second screen shot. The app gives the user the ability to zoom to any part of a one-week forecast, and planned future releases will allow for more custom configurations, such as user-set load thresholds, selectable DR sources (such as air conditioners or water heaters), DR forecasting by geographic subpopulation and significantly tighter integration with the rest of the IntelliSOURCE Enterprise architecture.
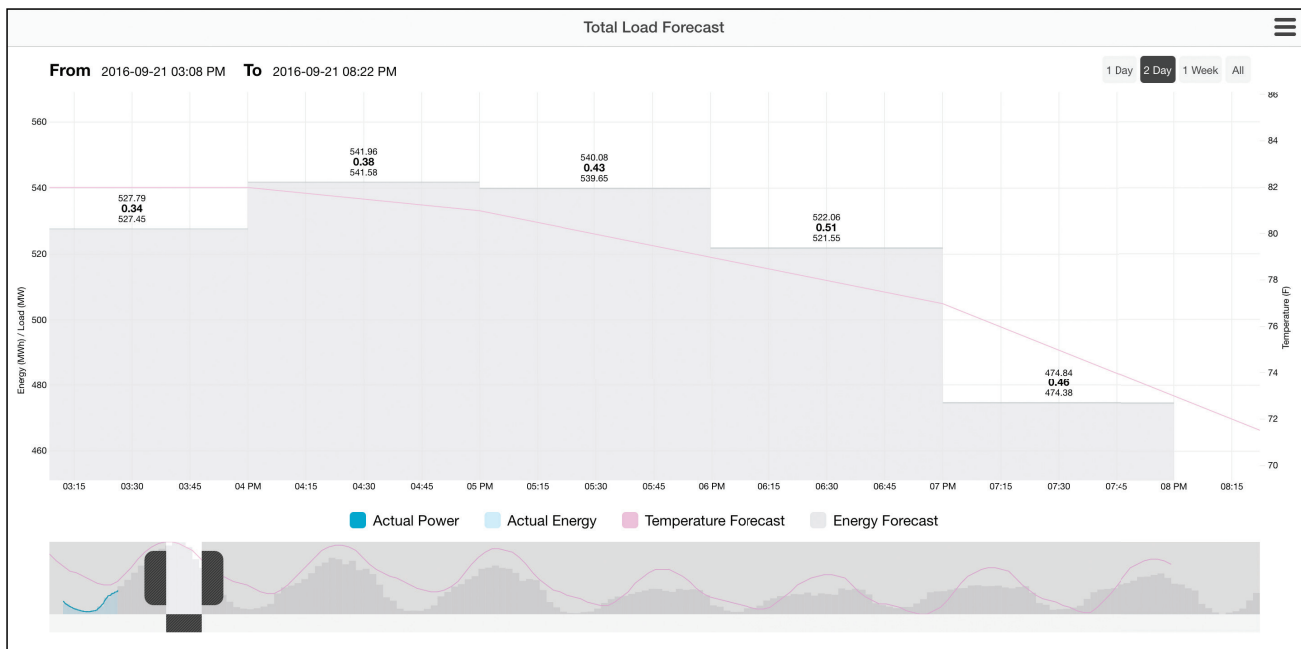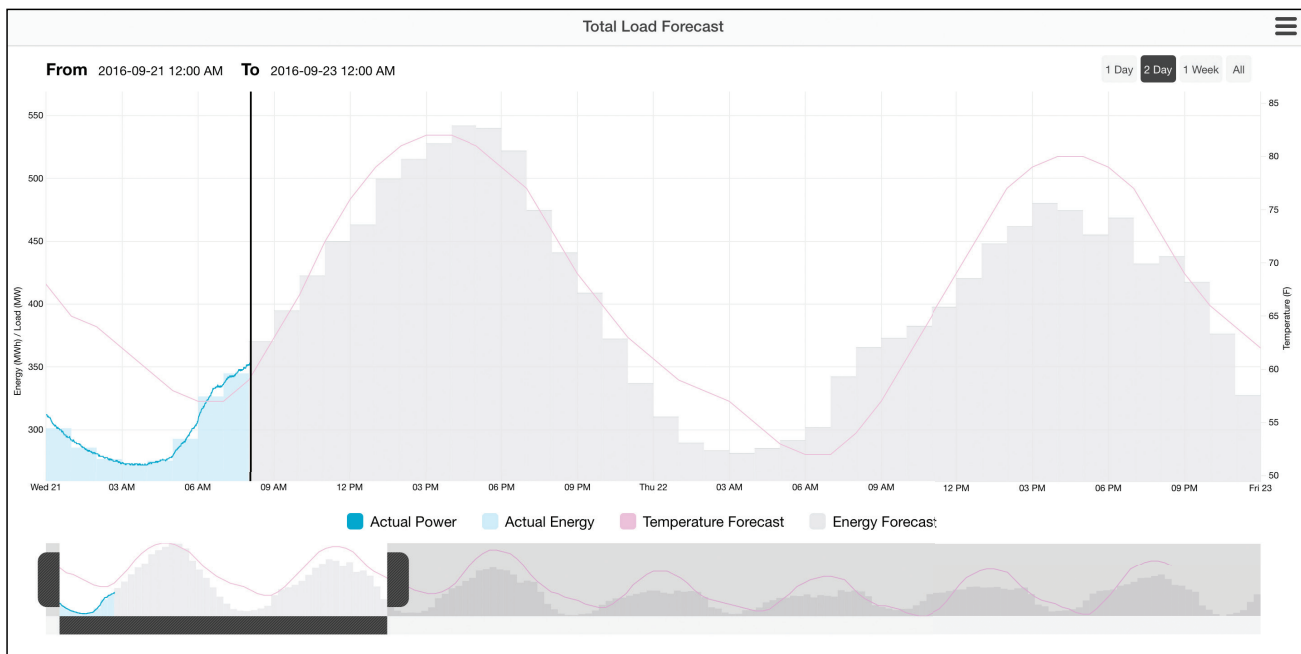
CONCLUSION

This white paper only scratches the surface of how we are using machine learning at Itron. In addition to the example of forecasting curtailable load from DR for residential thermostats (connected to air conditioners), we are also integrating forecasts from water heaters and pool pumps. We are continually refining our methods and experimenting with ways to integrate telemetry data from our thermostats into our models. With the explosion of data in the utilities sector, we forecast that we will have no shortage of work to do, nor shortage of benefits to bring to our customers, in the near future.

Join us in creating a more **resourceful world**.
To learn more visit **itron.com**

CORPORATE HQ

2111 North Molter Road
Liberty Lake, WA 99019 USA

Phone: 1.800.635.5461
Fax: 1.509.891.3355