Paper SAS729-2017

A Man with One Watch Always Knows What Time It Is

Tim Fairchild, SAS Institute Inc., Cary, NC

ABSTRACT

“A man with one watch always knows what time it is. A man with two watches is never sure.” Contrary to this adage, load forecasters at electric utilities would gladly wear an armful of watches. With only one model to choose from, it is certain that the forecasts will be wrong. But with multiple models, forecasters can have confidence about periods when the forecasts agree and can focus their attention on periods where the predictions diverge. Having a second opinion is preferred, and that’s one of the six classic rules for forecasters as per Dr. Tao Hong of the University of North Carolina at Charlotte. Dr. Hong is the premiere thought leader and practitioner in the field of energy forecasting. This presentation discusses Dr. Hong's six rules, how they relate to the increasingly complex problem of forecasting electricity consumption, and the role that predictive analytics plays.

INTRODUCTION

Electrification was named by the National Academy of Engineering as one of the greatest engineering achievements of the 20th century. The engineering marvel known as the electric grid has been described as the largest and most complicated machine ever built.

Economical large scale storage of electricity doesn’t exist yet, so throughout the world electric utilities face the challenge of precisely balancing electricity production with electricity consumption on a real-time basis. Accurately predicting consumption is essential for utilities as they make short-term decisions about how much electricity to generate or buy in order to maintain the equilibrium between supply and demand in the most economical and efficient manner.

Similarly, long-term forecasts of consumption are critical inputs as utilities make multi-million and multi-billion dollar decisions about constructing (and in some cases decommissioning) generation facilities.

Arguably there’s no other industry where forecasting plays a more important role. Miss a forecast in the retail industry and shelves are temporarily bare. Miss a forecast in the airline industry and passengers get bumped from an oversold flight. But come up short on the forecast of tomorrow’s electricity demand and the results could include brownouts or rolling blackouts with significant economic impacts. Over the last decade the already complex task of forecasting electric load has become even more difficult due to a number of factors including rooftop solar, electric vehicle penetration, and the proliferation of energy-efficient appliances and lighting. At the same time, the availability of hourly and in many cases sub-hourly interval data from smart meters has led to an explosion of data inputs.

Dr. Tao Hong is the premiere global thought leader and practitioner in the field of energy forecasting. He is currently the NCEMC Faculty Fellow of Energy Analytics, a Graduate Program Director and Assistant Professor of Systems Engineering and Engineering Management, an Adjunct Professor of Electrical and Computer Engineering, an Associate of the Energy Production and Infrastructure Center (EPIC), and the Director of BigDEAL (Big Data Energy Analytics Laboratory) at the University of North Carolina at

Charlotte. Dr. Hong is the Founding Chair of the IEEE Working Group on Energy Forecasting, the General Chair of the Global Energy Forecasting Competition, the lead author of the online book *Electric Load Forecasting: Fundamentals and Best Practices*, and the author of a blog focused on the topic of energy forecasting.

Prior to joining the University of North Carolina at Charlotte, Dr. Hong worked at SAS where he played a key role in the development of SAS® Energy Forecasting, an advanced forecasting solution designed for the critical and increasingly complex problem of forecasting electric load.

In a series of posts on his blog, Dr. Hong shared with his readers six basic principles of forecasting. In the sections that follow, each of Dr. Hong's principles will be explained in the context of electric forecasting and analytics with illustrations as appropriate from SAS Energy Forecasting.

SIX MUST-KNOW PRINCIPLES OF FORECASTING

Principle #1: Forecasting Is a Stochastic Problem

By its very nature, forecasting is a stochastic problem rather than a deterministic problem. But what's a stochastic problem? What's a deterministic problem?

Inherent randomness is the distinguishing mark of a stochastic problem. In a stochastic model, the same set of parameter values and initial conditions will lead to different outputs each time the model is run. There is randomness and uncertainty involved. For example, if you buy $1,000 worth of Apple® stock today, how much will it be worth in one year's time? Perhaps more…perhaps less…perhaps about the same. Clearly, there is an uncertainty that leads to a range of potential values.

Deterministic problems typically have one answer. In deterministic models, the output of the model is determined by the parameter values and the initial conditions. The present state completely determines the future state. For example, a guaranteed investment of $1,000 at 5% interest, compounded annually, will be worth $1,050 in one year, $1,102.50 in two years, and so on.

Forecasting is stochastic so that there's no such thing as certainty and there's always a range of uncertainty around each point.

Figure 1 below shows a long term load probabilistic load forecast. The red line at the top represents a 90% probability of loads being below that level while the blue line at the bottom represents a 10% probability of loads being below that level.

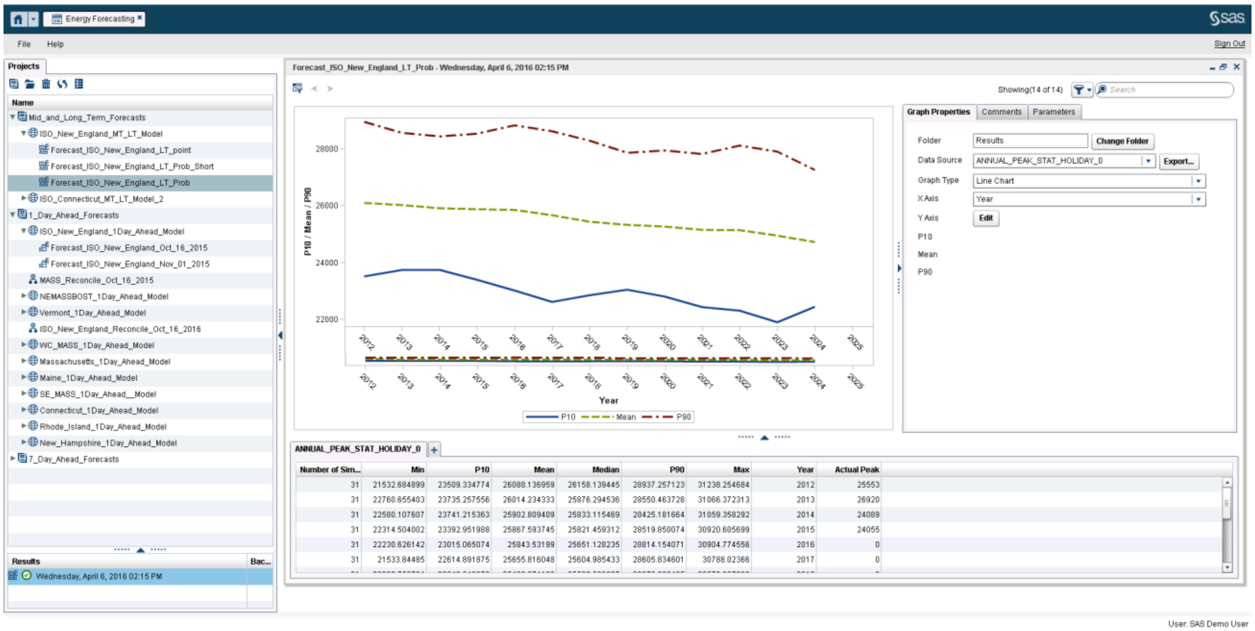

Figure 1 – Long term probabilistic forecast

Principle #2: All Forecasts Are Wrong

Because forecasting is a stochastic problem (see Principle #1), the result can never be determined with 100% certainty. You should always expect a forecast to be wrong and should never be surprised when it is. What should surprise you is when the forecasted value is exactly the same as the actual value.

The question "Why is your forecast different from the actual?" should never be asked, but sometimes it is and it usually comes from a person who doesn't understand the nature of forecasting. Unfortunately sometimes that person is a manager, a senior executive, a board member, or a regulator.

However, the extent to which a forecast is wrong is fair to question – especially when the error is due to poor data quality, inappropriate methodologies and algorithms, or inferior technology. While all forecasts are wrong, it's the goal of the forecaster to use the model with the lowest error rate.

Figure 2 below shows the results of applying 5 different models to the same data set. The red line at the bottom shows the dramatic reduction in error that resulted from using a two-stage model.

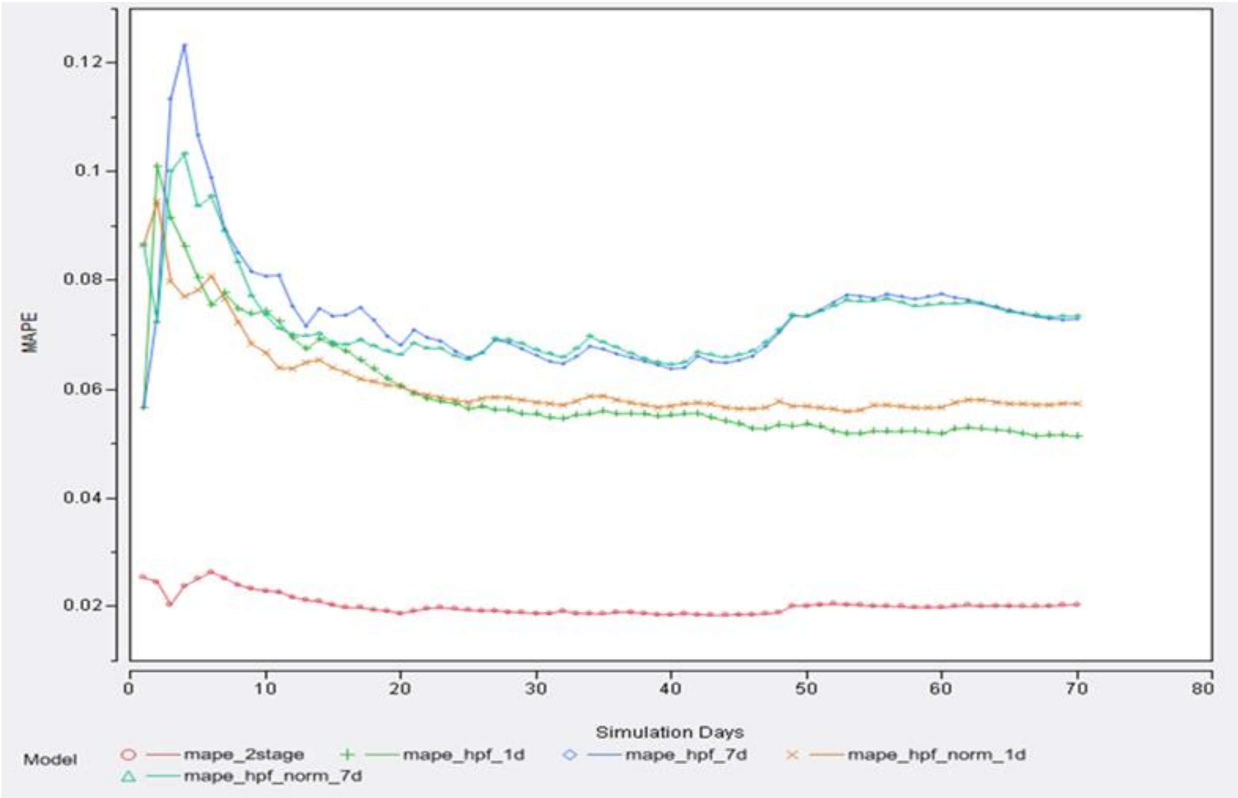

Figure 2 – Error rates for 5 different forecasting models

Principle #3: Some Forecasts Are Useful

Proctor & Gamble® needs SKU-level forecasts to optimize promotion offerings and manage inventory. American Airlines® needs passenger forecasts to plan and run their operations. Apple® needs forecasts to procure parts. And of course the utility industry needs load forecasts to plan and operate the electric grid.

Two factors determine the usefulness of a load forecast: accuracy and defensibility.

There are numerous ways to measure forecast accuracy. In fact, due to the unique nature of the electric utility industry, SAS Energy Forecasting includes 21 different built-in methods:

Annual Mean Absolute Error
Annual Mean Absolute Percentage Error
Annual Mean Error
Annual Peak Mean Absolute Error
Annual Peak Mean Absolute Percentage Error
Annual Peak Mean Error
Daily Mean Absolute Percentage Error
Daily Mean Error
Daily Peak Mean Absolute Error
Daily Peak Mean Absolute Percentage Error

Daily Peak Mean Error
Hourly Mean Absolute Error
Hourly Mean Absolute Percentage Error
Hourly Mean Error
Monthly Mean Absolute Error
Monthly Mean Absolute Percentage Error
Monthly Mean Error
Monthly Peak Mean Absolute Error
Monthly Peak Mean Absolute Percentage Error
Monthly Peak Mean Error
Weighted Mean Absolute Percentage Error

In some situations, the load forecaster will want to measure accuracy on an hourly or daily basis. In other situations such as where forecasting peaks is the primary objective, the appropriate accuracy measure will look at forecast performance for daily, monthly, or annual peaks.

Figure 3 below shows three error measures (Hourly MAPE, Daily MAPE, and Monthly MAPE) for nine different models and a 10th model that is a combined model. Scrolling across the table at the bottom in the application shows all 21 error measures. The different measures of errors allow the user to have the SAS Energy Forecasting solution automatically select the forecast based on the factors that matter the most according to the purpose for which the forecast will be used. For example, the user might want the model that gives the best hourly fit or the best fit for the day. Alternatively, the user might want the model that does the best job of forecasting the monthly or annual peak.

Figure 3 – Measuring forecast error in SAS Energy Forecasting

The defensibility of a forecast can be considered along a number of dimensions including interpretability and reproducibility.

Interpretability of a forecast is a function of the type of model used to develop the forecast and to the audience to whom the forecast needs to be explained. In some situations it might be appropriate to trade off accuracy for interpretability. For example, in a situation where a forecast needs to be explained to a regulator who lacks an in-depth technical understanding of forecasting, it might be better to use a statistical approach such as regression or time series rather than a black-box approach such as an artificial neural network.

The reproducibility of a forecast depends on the process and workflows involved in producing the forecast. Will the same models applied to the same data produce identical results a week, a month, or a year from now? The design and workflows of SAS Energy Forecasting (see Figure 4 below) ensure that the answer to that question is yes.

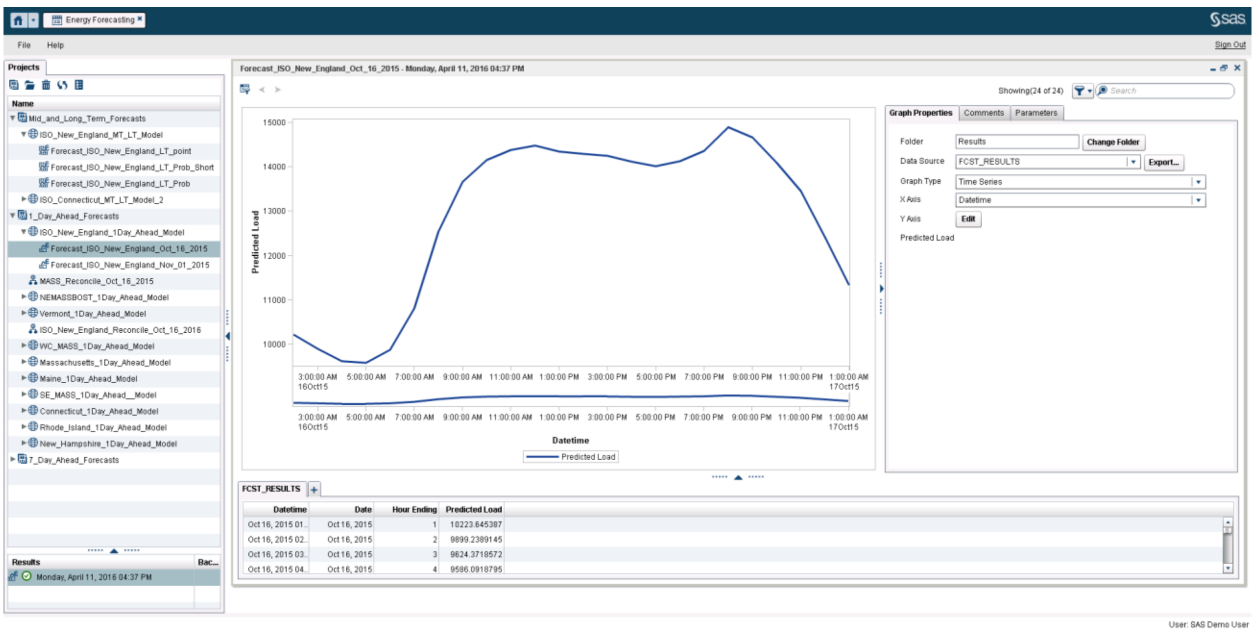

Figure 4 – The forecast workspace in SAS Energy Forecasting

The two factors that make a forecast useful – accuracy and defensibility – should be prioritized depending on the exact business need. For example, even though an artificial neural network model might provide a more accurate forecast result, the difficulty of explaining the model to a non-technical regulator might make it prudent to trade off accuracy for defensibility. This makes it essential for the forecaster to have a clear understanding of the business needs and the forecast review process before developing the forecast.

Principle #4: All Forecasts Can Be Improved

When you improve a forecast, you make it more useful. And what determines the usefulness of a forecast? Accuracy and defensibility.

Two specific areas where forecasts can be improved are in the spread of errors and the interpretability of errors.

Reducing the spread of errors reduces the uncertainty of the forecast and sometimes the forecaster might choose to sacrifice some degree of accuracy for a smaller range of errors. For example, one forecast might have a Mean Absolute Percentage Error (MAPE) of 1.5% while another has a MAPE of 1.7%. But the forecast with the slightly higher MAPE might have a significantly tighter distribution of error (as measured by the standard deviation of the Absolute Percentage Error) and in some circumstances this could be preferred.

Interpretability of errors comes into play particularly for long term load forecasting due to the inherent uncertainty in weather and economic scenarios over time periods that can span decades. It’s a best practice for the forecaster to help the consumer of the forecast (whether it is a business user or a regulator) understand how much of the error is accounted for by model error versus error in the weather forecasts versus error in the economic scenarios. Deconstructing the error in this way increases the interpretability of the forecast and therefore its usefulness.

A third way to improve a forecast involves resources. Technology can be used to make the forecasting process faster and more efficient allowing the forecaster to produce more forecasts and to spend more time on interpreting the results of the forecasts and less time producing the forecasts themselves.

Principle #5: Forecast Accuracy Is Never Guaranteed

Because forecasting is a stochastic process (see Principle #1) you should never expect that past performance of the model – no matter how good it is – will necessarily continue into the future.

If you’ve invested in the stock market you have most certainly heard the phase “Past performance is not indicative of future results.” Investment companies use this phrase to comply with Rule 156 of the Securities Act of 1933 that prohibits “Portrayals of past performance, made in a manner that would imply that gains or income realized in the past would be repeated in the future.”

But even though forecast accuracy is never guaranteed, remember that all forecasts can be improved. (See Principle #4.)

Principle #6: Having a Second Opinion Is Preferred

“A man with one watch always knows what time it is. A man with two watches is never sure.”

Sometimes referred to as Segal's Law, this adage refers to what happens when you have too much conflicting information when you’re trying to make a decision, and it inspired the title of this paper.

The phrase dates back to 1961 when a column in a New York newspaper attributed the saying to Lee Segal of KIXL in Dallas, Texas.

Or maybe not. An article in the San Diego Union in September 1930 stated that “Retail jewelers assert that every man should carry two watches. But a man with one watch knows what time it is, and a man with two watches could never be sure.”

So perhaps Lee Segal just used a long-existing saying instead of coining the phrase himself. It’s interesting how the derivation of the phrase illustrates its meaning.

But load forecasters don’t believe in Segal’s Law. The more information they have, the better. They would wear an arm-ful of watches.

If the forecaster only uses one model, it will occasionally produce bad forecasts. But with multiple models, the forecaster will have high confidence when the models agree with each other and they can focus their attention on those times when the models exhibit significant disagreement.

In the example shown in Figure 5 below, the load forecaster is looking back at the forecast produced by three different models and comparing them to the actual load. Note how for most of the hourly time periods, the three models produce forecasted loads that are relatively similar to each other. But there are several periods (such as span of dates between October 30 and November 3) where the forecasted results diverge significantly. These are the times periods that the forecaster will focus on in their analysis.

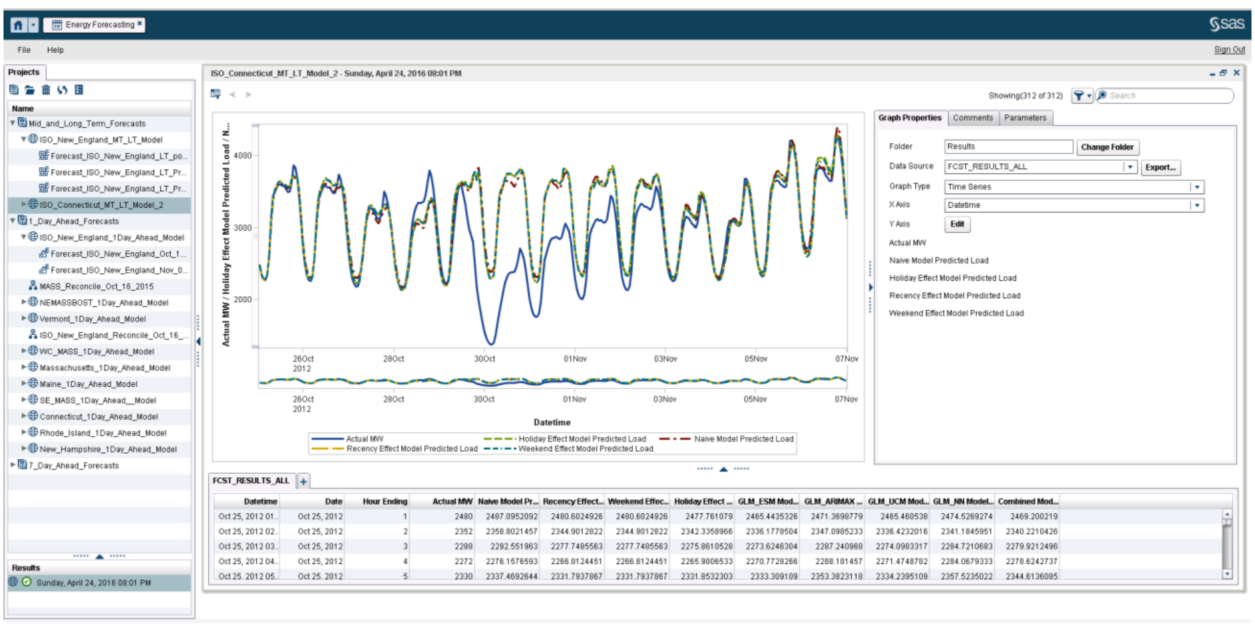

Figure 5: Forecast results workspace in SAS Energy Forecasting

SUMMARY

In this paper we’ve examined what Dr. Tao Hong describes as the six must-know basics of forecasting:

- Forecasting is a stochastic problem
- All forecasts are wrong
- Some forecasts are useful
- All forecasts can be improved
- Forecast accuracy is never guaranteed
- Having a second opinion is preferred

These principles are well known and understood by experienced load forecasters at electric utilities and were used as design points for SAS Energy Forecasting.

Many of the business users within a utility who depend on load forecasts as essential inputs for their planning and operational processes – as well as senior executives, board members, and regulators – are not forecasters themselves and don't understand the models, algorithms, and statistical procedures that produce the forecast. The same goes for executives and board members.

Understanding these six principles establishes the groundwork for productive discussions between load forecasters and the often non-technical constituencies who make decisions based on the forecast.

A man with one watch always knows what time it is. A man with two watches is never sure. A load forecaster with an arm-ful of watches is confident when he tells you what time it is.

ACKNOWLEDGMENTS

The author wishes to express his appreciation to Dr. Tao Hong.

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author:

Tim Fairchild

100 SAS Campus Drive

Cary, NC 27513

SAS Institute Inc.

Tim.Fairchild@sas.com

http://www.sas.com/