Energy Forecasting: Past, Present, and Future

TAO HONG

PREVIEW *When you flick that switch, you expect the lights to go on – but the business of keeping them on is not nearly as straightforward. Dr. Tao Hong offers a practical overview of energy forecasting; it's an important task, one that electric utilities have been doing daily for over a century, but now with new challenges.*

INTRODUCTION

Our electric power systems are the most complex man-made objects ever, producing and delivering electricity to 5.6 billion people. Similar to airlines, consumer package goods, and oil and gas industries, the electric power industry needs forecasts of supply, demand, price, and energy forecasts to plan and operate the grid.

Figure 1a illustrates a typical short-term load forecasting problem, while **Figure 1b** depicts a typical long-term load forecasting problem.

Figure 1a. Three weeks of hourly loads of a small utility.

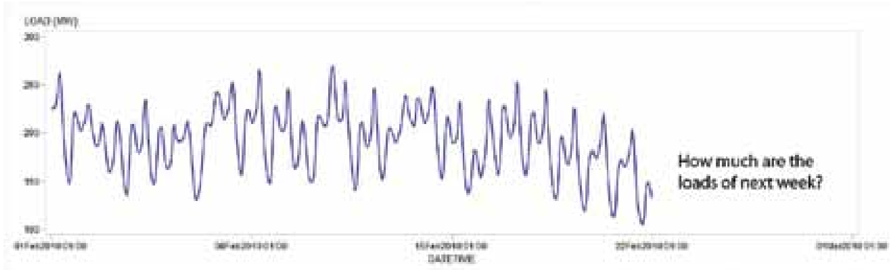

Figure 1b. Twelve years of daily peak load of a small utility.

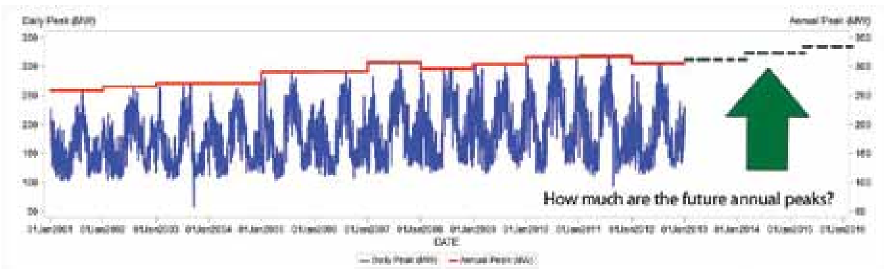

While other industries have some form of inventory to store and buffer their products and services, electricity cannot be massively stored using today's technologies. As a result, it has to be generated and delivered for immediate consumption; in short, utilities have to balance supply and demand *every moment*. Storage limitations and societal dependence on electricity lead to several interesting features of energy forecasting, including complex seasonal patterns, 24/7 data collection across the grid, and needs for precise accuracy.

This paper discusses the evolution of energy forecasting practices in chronological order, starting with short- and long-term load forecasting in the pre-PC era. It then summarizes the computer-based methods for spatial-load forecasting, short-term load forecasting, and electricity price forecasting, moving on to emerging smart-grid-era topics like demand response forecasting, renewable generation forecasting, and Global Energy Forecasting Competitions (GEFCom). The paper concludes with three lessons learned from 100-plus years of forecasting practice.

ORIGINS

Counting Lightbulbs

When Thomas Edison developed the Pearl Street Station in New York City in 1882, his motivation was simple: to promote the sale of lightbulbs. This first steam-powered generating station initially served about 3,000 lamps for 59 customers. When lighting was the sole end use of electricity, energy forecasting was straightforward; power companies counted how many lightbulbs they installed and planned to install. Then they roughly knew the level of load in the evening. This "ancient" method is still used in today's power systems planning for forecasting streetlight loads.

Key Points

- Energy forecasting in the utility industry has several distinct aspects: short-term load forecasting, long-term load forecasting, spatial-load forecasting, price forecasting, demand-response forecasting, and renewable-generation forecasting.
- Energy forecasting practices have gone through several important stages, from an engineering approach with charts and tables in the pre-PC era to today's more recent computer-based methods.
- Smart-grid investment and technologies have brought new challenges to the energy forecasting field, such as demand-response forecasting and renewable-generation forecasting. The century-old energy forecasting field has found new life in the smart-grid era.
- Advancement of energy forecasting relies on following rigorous out-of-sample tests, understanding business needs, and learning in an interdisciplinary manner across the fields of statistics, electrical engineering, meteorological science, and more.

The Engineering Approach

As more electrical appliances (irons, radios, washers) became popular, the complexity of the forecasting problem grew. A special event like a presidential radio address could now cause a spike in the load curve, as millions of people listened in at the same time.

Figure 2. Relationship between load and temperature

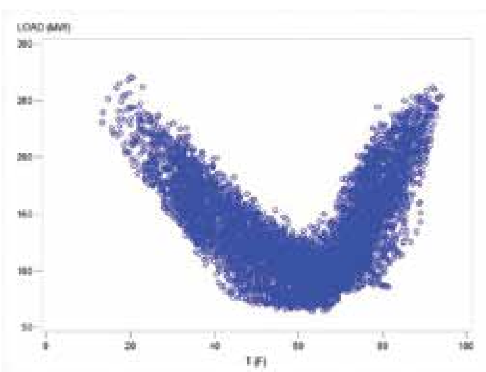

A scatter plot of hourly load and temperature. The scatter shows an asymmetric "U" shape. On the left, the lower the temperatures are, the higher the loads are, which is primarily due to heating needs. The opposite happens on the right, which is primarily due to space cooling needs.

In the 1940s, electricity demand became affected by weather, principally due to the penetration of air conditioners. **Figure 2** shows the relationship between load and temperature. In the winter, load and temperature are negatively correlated, primarily due to space-heating needs; in the summer, the correlation is positive, primarily due to cooling needs. Since then, weather variables such as temperature and humidity have been widely used to forecast load.

Because there were no statistical software packages at that time, an engineering approach based on charts and tables was developed that manually forecast future loads. Some of the elements considered – heating/cooling degree days, temperature-humidity index, and wind-chill factor – remain in today's load-forecasting models. The similar-day method, which derives a future load profile using the historical days with similar temperature profiles and day type (e.g., day of the week and holiday), is still used in many utility operations centers.

Spatial Load Forecasting

Computer applications ramped up in the 1980s. A significant amount of research was devoted to long-term spatial load forecasting regarding when, where, and how much load growth would occur (Willis, 2002). The forecasting horizon ranged from several years to several decades. These forecasts have been widely used since, in transmission and distribution planning.

Most of these forecasts were made using one of three methods: trending, simulation, and hybrid methods.

Trending methods use a mathematical function to fit the *past* load growth and then extrapolate to the future load. The most common function is a polynomial regression in which the load history is made a function of trend terms such as trend, trend-squared, etc.

The advantages of this method include ease of use, simplicity, and a short-range response to recent load-growth trends. However, the method often fails when used to estimate the long-range load, due to over-fitting or extrapolation of high-ordered polynomials, those with trend terms raised to more than two powers, such as trend-cubed.

Simulation methods attempt to model the load growth *process*, identifying the temporal, spatial, and magnitude patterns at play. They can also model the effect of a new urban development process based on land-use information from the government, customer rate classes from utilities, and load curve models of consumption patterns. Depending on data quality, this approach has had fair to very good short-range accuracy and good to excellent long-range usefulness for planning. Its drawback is its expensive development and training cost.

Table 1. Comparison among Short-Term Load Forecasting Techniques

	Pros	Cons
Multiple Regression	Interpretable; easy to implement, update, and automate; good accuracy	Needs explanatory variables, a designated functional form, and at least two years of history
ARIMA	Few parameters to estimate; does not require a long history; good accuracy in (very) short term	Low accuracy in longer term; high cost to implement, update and automate; often difficult to interpret
ANN	Minimum statistical or domain knowledge required; good accuracy during normal days	Heavy computation; over-parameterization; difficult to interpret; low accuracy during extreme weather conditions

Hybrid methods combine the favorable features of trending and simulation. An ideal hybrid should respond to the recent trend of load history in the short range and keep the long-range defensibility provided by the simulation methods, all without requiring much skill and interaction from the user.

Several years ago, I developed a modern hybrid method (2008) that divides an entire territory into thousands of 50-acre areas. These small areas are used to build a hierarchy with multiple levels. Load growth in each small area or region is projected by an S curve based on input information of historical loads, long-term load forecasts at the corporate level, and land use development plans. This method has been commercialized and deployed to many utilities in North America.

Short-Term Load Forecasting

Computers not only improved the practice of spatial-load forecasting, but also that of short-term load forecasting (STLF). Late in the last century, the power industry went through a major structural change, making accurate short-term load forecasts even more critical. Analysts first tried to apply statistical techniques, such as regression analysis and ARIMA (Box-Jenkins models). Then artificial intelligence (AI) became one of the hottest terms in the scientific community, resulting in hundreds of papers reporting AI-based approaches to STLF. In the 1990s, Electric Power Research Institute (EPRI) sponsored a project that developed several artificial-neural-networks-based short-term load forecasters (Khotanzad and colleagues, 1998). Some research generated from this EPRI project was later commercialized and became a popular STLF service provider in today's industry.

Artificial Intelligence. The models based on AI techniques such as artificial neural networks (ANN), fuzzy logic, and support-vector machines were *black-box models* that appealed to organizations unwilling to build an in-house team of forecasting analysts. However, other utilities weren't comfortable with black-box approaches and instead developed forecasts using more traditional methods such as the similar-day and regression methods. Some utilities with in-house forecasting teams also built black-box models or purchase forecasts for comparison purposes.

Regression Methods. Recently I reported on a comprehensive research study of regression-based short-term load forecasting. My methodology was soon adopted by many utilities, retailers, and trading firms worldwide and became part of the engine of a commercial energy-forecasting solution.

Table 1 summarizes the key approaches to short-term load forecasting.

Electricity-Price Forecasting

The dawning of electricity markets brought a new challenge: electricity-price forecasting. While load forecasts tell utilities how much power supply they need to balance the demand, price forecasts help them determine how much energy they should buy or sell.

Three categories of price forecasting methods have been employed: simulation, statistical, and artificial intelligence.

Figure 3. Electricity is the most volatile commodity in the world. Price spikes are hard to predict.

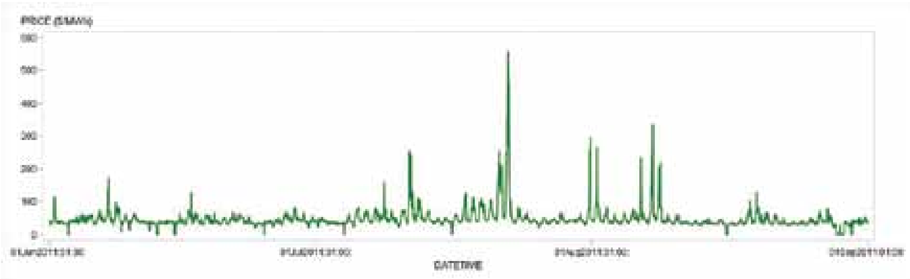

Simulation methods require a mathematical model of the electricity market, load forecasts, outage information, and bids from market participants. Conducting a simulation takes very specialized skills and knowledge of power systems as well as sophisticated software packages. Price forecasting accuracy is highly dependent upon the quality of the input information; the load forecast at each node of the market is a driver of the electricity-price forecast.

Statistical and AI-based methods do not require such comprehensive knowledge of market operations. They use historical prices, weather, outages, and loads to forecast future prices. Weron (2006) discussed applications of several statistical techniques for load and price forecasting. Their implementation cost is less than those of the simulation methods. However, they often find difficulties in forecasting price spikes as illustrated in **Figure 3**. These spikes are mostly due to congestions in the transmission network, which can sometimes be picked up through simulation methods.

THE "SMART GRID" ERA

Demand-Response Forecasting

In the past decade, the electric power industry has undergone a grid-modernization process, installing millions of smart meters, sensors, and communication devices. Smart-grid technologies offer great potential for a greener and more reliable grid, at reduced cost.

One way to achieve these goals is through demand response (DR), which is defined by the U.S. Federal Energy Regulatory Commission as "changes in electric usage by end-use customers from their normal consumption patterns in response to changes in the price of electricity over time, or to incentive payments designed to induce lower electricity use at times of high wholesale market prices or when system reliability is jeopardized."

To effectively design and implement the DR programs, utilities have to perform a series of analytical tasks such as forecasting electricity price, forecasting load at household level with or without DR programs, and after-the-fact estimation of the consumption pattern changes due to various demand response programs. A major challenge in this area is to estimate the normal (baseline) consumption pattern. Because time is irreversible, there's no way for a utility to conduct an experiment to get the actual value of normal consumption pattern for an event day (the day when DR programs are triggered).

There are over a hundred baseline estimation methods, most of them derived from a simple average of the load profiles of several days prior to the date when DR programs are triggered. These methods are shown to be insufficient at household level. We don't yet have a proven solution to this problem.

Renewable-Generation Forecasting

With increased power generation from wind turbines, rooftop PV panels, and solar farms, the century-old energy forecasting problem finds new life. Volatility of renewable generation is a major challenge to system operators and energy traders. **Figure 4** shows one week of solar generation at 5-minute intervals. While many large and medium utilities today operate their systems with the error of one-day-ahead load forecasts at 3% or lower, they can hardly achieve similar accuracy in forecasts of their wind or solar generation.

Major advances are being made in renewable-generation forecasting by meteorologists and forecasters. Meteorologists sample the state of the fluid at a given time and apply the equations of fluid dynamics and thermodynamics to estimate the state of the fluid at some time in the future. This numerical, weather prediction approach requires significant computing resources as well as working knowledge of meteorological science. Moreover, further analysis is needed to translate wind and solar forecasts into renewable-generation forecasts. The statistical alternative uses meteorological forecasts as one of its modeling inputs; other inputs include lagged wind power and calendar variables. A recent review of the state of the art is by Giebel and colleagues (2011).

GLOBAL ENERGY FORECASTING COMPETITION

To tackle the emerging challenges in energy forecasting, the IEEE Working Group on Energy Forecasting organized the Global Energy Forecasting Competition 2012 (GEFCom2012), bringing together many new ideas from data scientists. The competition consisted of two tracks: hierarchical load forecasting and wind-power forecasting. There were 2,000+ entries submitted by 200+ teams. Eight teams from eight countries are recognized as the winners.

Multiple linear regression and the gradient boosting method (GBM) were among the top four winning entries of each track. GBM is a relatively new data-mining approach that seeks to create a strong prediction model by combining weak prediction models. A more comprehensive introduction to GEFCom2012 is by Hong, Pinson and Fan (2014).

Difficulties in accurately predicting price, renewable generation, and long-term load have prompted interest in *probabilistic forecasts*. Probabilistic forecasts are the counterpart of point forecasts, supplementing point forecasts with probability information about their likely errors. As such, they offer a more comprehensive description of the future, making them more valuable for risk management.

There are still many research challenges with probabilistic energy forecasting, such as visualization, interpretation, and evaluation of the results, incorporating probabilistic inputs, and integration with business decision-making processes. With the success of GEFCom2012, the competition organizers will launch an upgraded version in 2014 with the

Figure 4. One week of solar power generation at 5-minute intervals. Volatility depends upon the mood of the clouds. The generation profile is a nice bell curve on a sunny day; on a cloudy or partially cloudy day, generation is quite unpredictable.

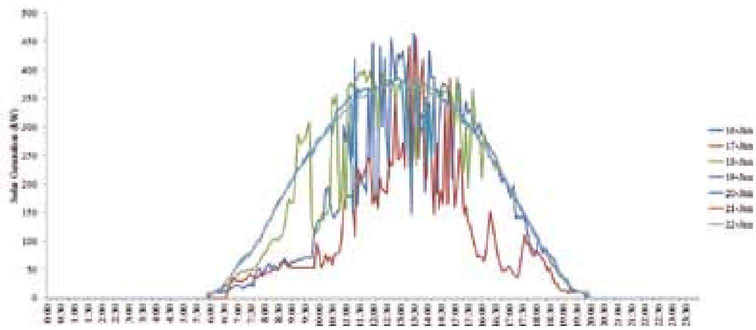

theme of probabilistic energy forecasting. GEFCom2014 will feature four tracks covering load, price, wind, and solar forecasting.

LESSONS LEARNED

Much has been learned as this fascinating field has advanced over the past century. Here are three major lessons:

Out-of-Sample Tests

Many energy forecasting papers reported amazingly low errors but failed miserably in practice. A primary reason is lack of rigorous out-of-sample tests. One obvious mistake is to use a model with 1,000+ parameters to fit a data set with a few hundred observations; other mistakes made in the literature are more difficult to identify.

For instance, when developing a regression model, the authors slice the data into two pieces, one piece for parameter estimation, the other as validation to calculate the Mean Absolute Percentage Error (MAPE). Seeing the MAPE value is too high, they change the model, recalculating the MAPE until its value is low enough to be impressive. When submitting the paper, they report the model with the lowest MAPE. Although this validated data was never used for model fit, its information is used for variable selection when building the model. In real life, we will never get a chance to use tomorrow's actual load to build the model. Peeking into the future produces nice results on paper, but not in practice. Tashman (2000) offered a comprehensive review of out-of-sample tests, of which sliding simulation is the most appropriate and effective for energy forecasting.

Understanding Business Needs

There are many reasons why utilities do things the way they do today. Some can be improved; some cannot. For example, the long-term-forecasting methodology used for rate cases has to be published and archived. Many parties – such as managers in the utility, regulatory commissions, and shareholders (if the utility is investor owned) – may review the document, even years after the case. Therefore, the long-term load forecasting methodology has to be interpretable and transparent for ease of communication and defense. Many utilities are simply not allowed to use black-box models for long-term forecasts for rate cases. Despite how fancy the models are, they are not useful for business.

Many utilities are simply not allowed to use black-box models for long-term forecasts for rate cases. Despite how fancy the models are, they are not useful for business.

Tao Hong is an Assistant Professor of Systems Engineering and Engineering Management at the University of North Carolina at Charlotte. He is Founding Chair of the IEEE Working Group on Energy Forecasting and General Chair of Global Energy Forecasting Competition (GEFCom.org). He is the author of the online book *Electric Load Forecasting: Fundamentals and Best Practices* (***otexts.org/elf***) and the blog "Energy Forecasting" (blog.drhongtao.com). Dr. Hong received his BEng in automation from Tsinghua University in Beijing and his PhD with co-majors in operations research and electrical engineering from North Carolina State University.

hongtao01@gmail.com

It Takes a Village

In today's dynamic environment, virtually all types of energy forecasts are connected. A short-term load forecasting model can be augmented to a long-term model by adding macroeconomic indicators (Hong, Wilson and Xie, 2013). Electricity prices are no longer driven by load only; the volatile renewable generation of wind and solar farms also affects price significantly. Price signals trigger demand response programs, which in turn affect loads. A best practice today for utilities is to build an in-house analytics center of excellence, where statisticians, data miners, meteorologists, business liaisons, IT analysts, and software developers can work together to tackle the emerging challenges of energy forecasting.

Energy forecasting is an interdisciplinary field. To further advance our knowledge, we have to involve various communities, such as statistical forecasting, artificial intelligence, meteorological science, and electrical engineering. To make the research meaningful to the industry, it is necessary to take inputs from utilities. It took a village to get to where we are today; it will take a village to be where we want to be in the future.

REFERENCES

Giebel, G., Brownsword, R., Kariniotakis, G., Denhard, M. & Draxl, C. (2011). *The State-Of-The-Art in Short-Term Prediction of Wind Power: A Literature Overview*, 2nd Ed., Project report for the Anemos.plus and SafeWind projects.

Hong, T., Pinson, P. & Fan, S. (2014). Global Energy Forecasting Competition 2012, *International Journal of Forecasting*.

Hong, T., Wilson, J. & Xie, J. (2013), Long Term Probabilistic Load Forecasting and Normalization with Hourly Information, IEEE Transactions on Smart Grid.

Khotanzad, A., Afkhami-Rohani, R. & Maratukulam, D. (1998). ANNSTLF-Artificial Neural Network Short-Term Load Forecaster Generation Three, IEEE Transactions on Power Systems.

Tashman, L. (2000). "Out-of-Sample Tests of Forecasting Accuracy: An Analysis and Review," *International Journal of Forecasting*, 16:4, 437-450.

Weron, R. (2006). *Modeling and Forecasting Electricity Loads and Prices: A Statistical Approach*, Wiley.

Willis, H. L. (2002). *Spatial Electric Load Forecasting*, 2nd Ed., CRC Press.