Blowing in the wind

Part of the answer to rising energy needs and costs may literally be blowing in the wind. Among sustainable sources of electricity, only wind energy has the capacity and technology needed to compete in the open marketplace. The largest onshore wind farm in Europe is being built in Scotland, the largest in the USA is planned for southern California, and the biggest offshore wind farm production in the world is slated for the Thames Estuary. But wind is intermittent. **Marc Genton** and **Amanda Hering** explain how advanced statistical techniques will enable wind energy to be more efficiently incorporated into the electrical grid.

Wind power basics

Harnessing the power of wind for the benefit of humans is not a new concept. Historically, windmills have been used to pump water from wells or to grind grain for centuries. But fastforwarding into the 21st century, "windmills" are being used to generate electricity. Wind turbines, as they are now commonly called, are enormous structures, generally up to 80 m tall, which is roughly the equivalent of a 26-storey building. With blades up to 40 m in length and costing up to US $2.5 million to manufacture and install a single unit, the science behind effective wind turbine design has evolved rapidly over the last two decades. Within the wind turbine housing is a gearbox to increase the rotational speed and a generator to convert the motion into electricity. A computer in the tower senses the wind direction, points the blades in the optimal direction and shuts the turbine off in dangerously high winds.

So, can these supercharged wind turbines actually produce enough energy to make a significant contribution to meeting demand? Most modern turbines installed onshore are rated to produce between 1.5 and 1.8 MW of electricity each, which is enough to power 1000 homes for an entire year (British Wind Energy Association data available at `http://www.bwea.org`). Depending on the size and number of turbines, clusters of them located in windy locations can produce electricity for many thousands of homes. These clusters, such as in Figure 1, are called wind farms. Construction of the largest onshore wind farm in Europe started in the autumn of 2006, south of Glasgow, Scotland. The construction will take 3 years to complete and will consist of 140 turbines producing 322 MW of electricity, enough for about 200 000 homes. The largest wind farm in the USA is planned for a region just north of Los Angeles in California and will produce over 1500 MW of power.

Figure 2 illustrates the amount of power that can be produced by a typical onshore turbine at various wind speeds. At the cut-in speed, the blades begin to

Figure 1. Typical wind farm in the state of Washington, USA

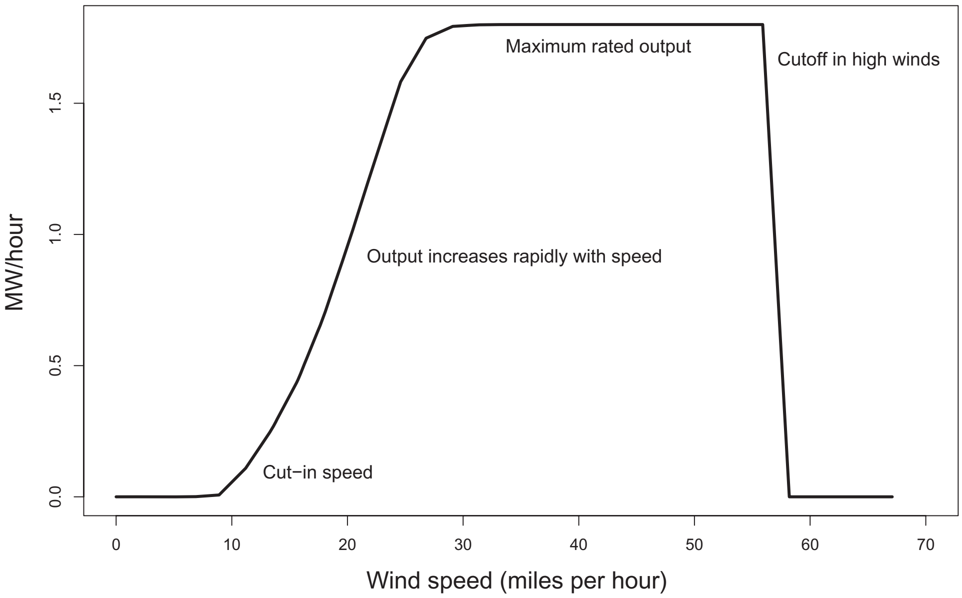

Figure 2. Potential power output of a wind turbine. The data are from a 1.8 MW Vesta V801800II turbine

rotate, but the power output increases rapidly even with very small increases in the wind speed. In this range, power is proportional to the cube of wind speed, so small differences in speed can make large differences in power output. The maximum power output of 1.8 MW for this particular turbine occurs at about 30 miles per hour and shuts down at just over 50 miles per hour. However, power depends not only on wind speed but also on variables such as the diameter of the blades, the density of the air and the direction from which the wind is blowing. Thus, wind power varies from one turbine make and model to another.

A large amount of growth and research is now being invested in offshore wind turbines, whose larger sizes (up to 3 MW with 5 and 7 MW machines in development) can take advantage of stronger ocean breezes. Just over 15 offshore wind farms are currently in operation, mainly off the coasts of Denmark, Sweden and the UK, but many more are in the planning stages. The Thames Estuary scheme announced by the government in December 2006 will use 341 turbines to generate a planned 1000 MW at a capital cost of £2 billion. Most offshore wind farms are located in water less than 30 m deep, but engineers feel that they can draw on their experiences with oil platforms and move these farms even farther from land and out of public view.

Compared with traditional power plants fired by coal, natural gas or nuclear reactions produce, averaged over the year, 50% of their maximum designed output, wind farms produce, on average, about 30% of their maximum rated output (`http://www.eia.doe.gov`). In the USA, the current cost for a kilowatt-hour of wind-generated electricity is between \$0.04 and \$0.06, very similar to traditional energy sources which cost between \$0.04 and \$0.055. Opponents of wind energy claim that greater start-up costs are involved. Transmission lines to move electricity from windy places, which tend to be remote, must be established; however, once a wind farm is operable, it should pay for itself in its first 6–8 months of operation (`http://www.bwea.org`). In addition, decommissioning a wind farm, whose turbines last 20–25 years, is simply a matter of disassembling the turbines, removing them and recycling the materials. This is a much simpler and more environmentally friendly process than decommissioning a nuclear power plant, for instance.

Wind farms have other tangible and intangible benefits. Once installed and operable, wind farms produce clean fuel, with no greenhouse gas pollutants or gas emissions. Quantifying the importance of this benefit is difficult, but it is recognised as significant. The Energy Information Administration projects that oil and gas prices will remain high for at least the next 20 years. Every hiccup in these prices can send economies into turmoil, so countries who invest in diversifying their energy portfolio will help to stabilise their economies. Not only will demand for oil and gas decrease, thereby causing a decrease in prices, but also, more importantly, volatility in energy prices will be reduced.

Worldwide, only 1% of electricity is generated from wind, but the growth rate has been rapid—24% overall in 2005, with a stunning 48% increase in Asian markets. The World Wind Energy Association expects that 70 000 MW of wind power will be installed by the end of 2006 (`http://www.wwindea.org`). Many countries already boast a large proportion of wind-generated electricity. The pioneering countries of Denmark and Germany, who generate over 20% and 8%, respectively, of their total electricity needs from wind, have set an example to others. Countries such as the USA and the UK (both currently generating 1% of their electricity needs from wind) are aggressively developing their abundant wind resources. Figure 3 shows how

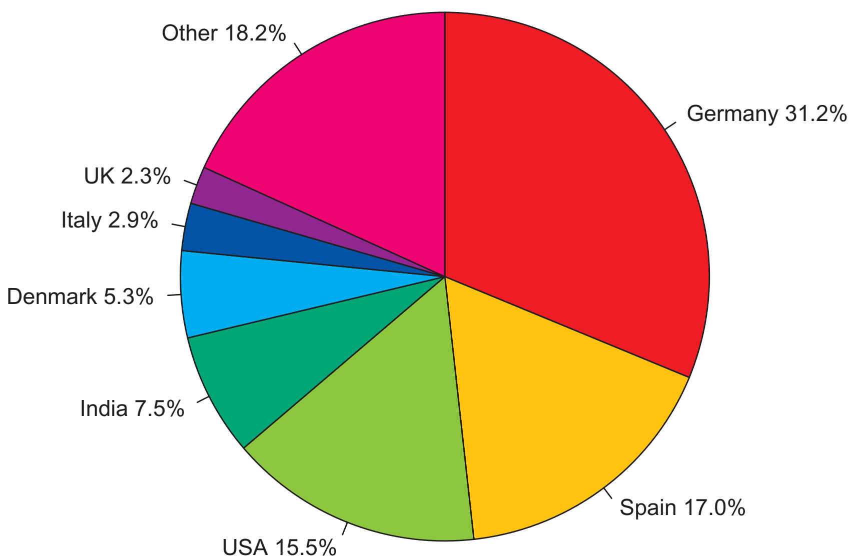

Figure 3. Percentage of worldwide wind capacity generated by each country in 2005. *Source*: World Wind Energy Association: `http://www.wwindea.org`

much electricity of the worldwide total is generated by each of the top seven wind-energy generating countries.

With all of the advantages of and interest in wind-produced electricity, barriers to widespread usage still exist. In particular, utility companies must manage a delicate balance between electricity supply and demand. In larger markets, excess electricity can be sold, and deficits can be bought. But, depending on regulations in each particular market, monetary penalties can be imposed when energy is wasted. In smaller markets, such as those on islands, with no one else available to buy or sell electricity, there is little room for error.

Electricity demand by consumers varies in a nearly deterministic fashion based on outdoor temperature, daylight hours and holidays. Thus, demand can be predicted, but, when wind power is added as a source of electricity, supply becomes unpredictable[1]. Wind power is intermittent—obviously the wind does not blow at a constant speed, but is variable. No cost-effective solution to storing wind energy has been found, so wind energy must be used immediately when it enters the grid. A utility company consequently needs to schedule how much energy it needs to "order" from its traditional plants so that supply will equal demand.

Gas turbine plants need at least 20 minutes notice to begin production, but large coal and oil plants require at least 8 hours to come online. Markets with slow start production units would benefit the most from accurate wind power forecasts.

Statistical solutions

We cannot be completely reliant on wind energy because of its uncontrollable and intermittent nature. However, intermittency is a close cousin of variability, which is any statistician's playground. Given certain information, electricity dispatchers do not have to fly blindly without any knowledge of how much electricity the wind will produce during a critical stage of decision-making. Statistical modelling to predict wind speeds or wind power can improve on our "best guess" estimate, which is the current wind speed, called the persistence model.

The number of hours ahead that a forecast is needed is called the forecast horizon and can vary depending on the reason for the prediction. The maximum horizon needed would be for 2–5 days ahead to schedule maintenance of the turbines during slow wind days. Otherwise, 48 hour forecasts are needed for trading in the electricity market. For scheduling and dispatch, a typical horizon is between 3 and 10 hours, but, in systems whose conventional sources generate electricity quickly, the horizon can be under 3 hours.

Both physical and statistical models for predicting wind power have been proposed and are currently in use, but both approaches follow similar strategies. The available data that the models will be built with must be scaled to the hub height of the turbine. For instance, the wind speeds for the past 24 hours may be available at a height of 10 m above ground level, but as the altitude increases, wind speed also increases in a logarithmic fashion. As a result, doubling the altitude can increase the wind speed by 10% and the power output by 34%.

The next step is deciding whether to predict the wind speed or to jump straight to predicting wind power output, i.e. the bottom line for utilities. If wind speed is predicted, then an additional step of translating that into power output for the particular types and numbers of turbines in use must be done. However, solely predicting power for a particular region may make it difficult to predict power output for a nearby wind farm with different turbines. In statistical models used to date, it has been found that modelling the wind speed itself is most efficient for horizons up to 8 hours and then modelling the power output thereafter is sufficient[1].

Finally, predictions can be scaled up for an entire region. This is especially important for areas like the UK and Europe where wind farms are geographically dense, and utility companies may manage several wind farms close to each other.

Most physical models used to predict wind speed or power incorporate output from numerical weather prediction (NWP) models. The basic premise of these models is the same–use an increasingly fine grid of information to get a more complete picture of terrain and air flow. NWP-based models can cover thousands of kilometers horizontally with grid resolutions from 5 to 25 km, but they are computationally extremely expensive to run. Models can require up to 4 hours of computer time and therefore cannot generate fast, reliable forecasts for short horizons[1]. These short horizons are the typical time needed to schedule transmissions and dispatch. Thus, physical models are more effective for 24-hour predictions. Ensemble models, averaging many different physical models together or combining them with statistical models, are also becoming popular.

Statistical models are the most competitive for short forecast lead times. Neural networks, fuzzy logic, local regression and time series methods have all been applied to the problem of wind speed prediction. Many of these models improve when additional information from the wind farm is included, such as wind direction, time of day, atmospheric pressure, and even physical model output[2]. The best statistical models, however, do not use a "black box" approach but also incorporate expert knowledge of the wind characteristics of a particular region[2]. It also makes sense that allowing parameters in these models to vary seasonally can result in improvements since a variable's influence may change throughout the year.

A growing area of emphasis has been to incorporate offsite observations into statistical models[2,3]. Changes in wind speeds may be detected at upwind locations before reaching the wind farm and can improve predictions. An argument against this methodology is that sites "upwind" of a wind farm can change as the wind direction changes[4], and no single offsite location may exist that has consistently high correlation with wind speeds at the prediction site. The ANEMOS project group (a consortium in Europe whose goal is to improve wind forecasting) found that with information on 23 offsite locations, predictions could be improved using between three and five of these sites whose meteorological conditions were most representative of the region (`http://anemos.cma.fr` and `http://anemos.cma.fr/modules.php?name=AvantGo&file=print&sid=33`). Even physical models have been shown to benefit from the use of additional spatial information.

The best statistical models incorporate expert knowledge of the wind characteristics of a particular region

With the plethora of models being proposed and tested, a consistent way to compare them is needed but not straightforward. Differences in complexity of the terrain, forecast and data resolution (10 minute, hourly, daily), and size and number of wind turbines at a farm can all affect model comparison. A common way to evaluate a model is to compute some function of the error—the actual observation minus what was forecast—such as the root-mean-square error (RMSE). RMSE will vary from one dataset to another; a skill score is used to remove the inherent variability in the observations. It is defined as the difference between the RMSE of a reference forecast and the RMSE of a model, divided by the RMSE of the reference forecast. The skill score can only be computed if a reference forecast (a model currently in place or the persistence forecast) is available. Even though RMSE is the most common measure to quantify error, it is not sufficiently sensitive to reflect improvements in prediction quality. In addition, comparisons made only against the persistence model may be overly optimistic since improving upon the persistence forecast can be accomplished with the simplest of statistical techniques. The ANEMOS project group has also suggested that errors be normalised with respect to the installed capacity of the wind farm.

Besides the most obvious problem of forecasting the wind speed or wind power for a particular horizon, more detailed information about the quality of the forecasts is also desired. Statistical forecasts have a built-in probabilistic error rate based on sampling distributions. These error bands around the predictions, or confidence bands, give dispatchers an idea of how certain the forecasts are. Very wide bands may indicate an unpredictable forecast, and smaller bands may indicate a more reliable estimate. Ensemble predictions can also give a sense of the forecast uncertainty (see the ANEMOS project web site). If the predictions from several different models are similar, then the collective prediction is more certain than if the forecasts vary dramatically. It is also of interest to identify conditions that lead to unpredictable power output or dramatic changes in power. When those conditions occur, utilities can protect themselves by carrying larger rolling reserves from traditional energy sources.

Future work

Predicting wind speed and power is a blossoming area of research. Besides the issues previously mentioned, predictions at offshore wind farms add another dimension to the process. The vertical wind profile (and thus the relationship between wind speed at an observed height and the turbine hub height) differs offshore owing to nonlinear interaction between the wind and waves, surface heating and the land–sea interface that modifies the air flow. Understanding the wake effect behind massive offshore turbines will influence turbine orientation and spacing. A wake is the decrease in wind speed since some energy is lost after moving through the turbine blades. They differ from one turbine to another and can decrease power output by up to 10% (see the ANEMOS project web site).

Predictions both onshore and offshore may benefit from the use of more advanced statistical techniques. Many statistical methods are built on the assumption that the variable of interest is normally (symmetric and bell shaped) distributed. This is decidedly untrue for wind speeds that are constrained to be positive and for which large values occur less frequently than small ones, as illustrated in Figure 4. Non-normality should be incorporated in statistical models. In addition, placement criteria for new wind farms and for turbines within a wind farm can be evaluated and aided with the use of spatial statistics.

Improved statistical forecasting has already had an influence in increasing wind energy production. As the industry continues to expand, the end wants of utilities will only grow in number and complexity; they will need longer forecasts, more accurate forecasts, measures of forecast quality and good tools for forecasting. As statisticians and scientists work together to provide these tools, we believe that the power blowing in the wind will be harnessed and become a mainstream solution to energy demands.

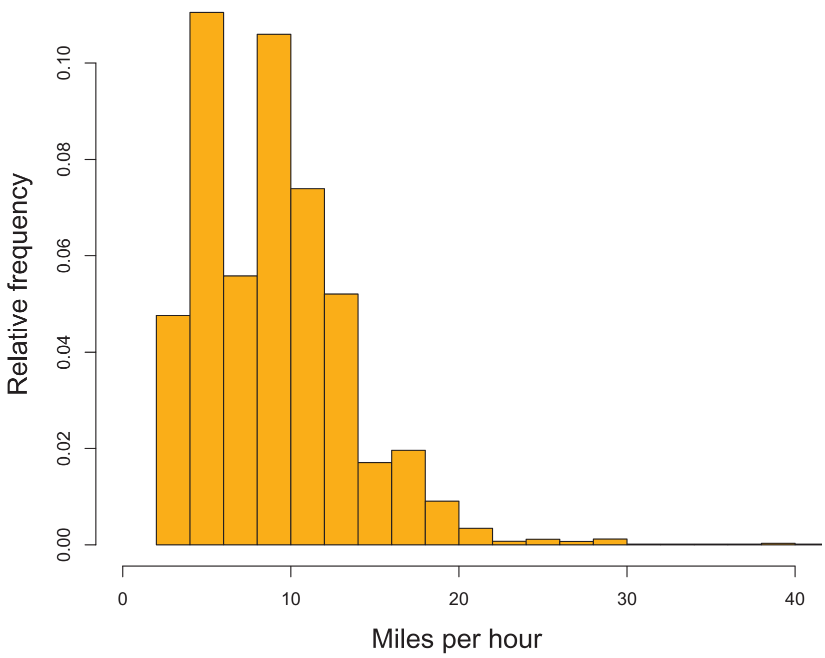

Figure 4. Wind speed recorded hourly during the year 2005 at Houston International Airport in Houston, TX, USA.

References

1. Giebel, G., Brownsword, R. and Karinotakis, G. (2003) *The State of the Art in Short Term Prediction of Wind Power: A Literature Overview*. ANEMOS Project.
2. Gneiting, T., Larson, K., Westrick, K., Genton, M. G. and Aldrich, E. (2006) Calibrated probabilistic forecasting at the Stateline wind energy center: the regime switching spacetime method. *Journal of the American Statistical Association*, **101**, 968–979.
3. Larson, K. and Westrick, K. (2006) Short term wind forecasting using offsite observations. *Wind Energy*, **9**, 55–62.
4. Kretzschmar, R., Eckert, P. and Cattani, D. (2004) Neural network classifiers for local wind prediction. *Journal of Applied Meteorology*, **43**, 727–738.

Marc G. Genton is Professor of Statistics in the Department of Econometrics at the University of Geneva and Associate Professor in the Department of Statistics at Texas A&M University. His research interests include spatiotemporal statistics, robustness, multivariate analysis, skewed multivariate distributions and data mining. This work was partially supported by National Science Foundation grant DMS0504896.

Amanda S. Hering is a doctoral student working under the supervision of Marc Genton in the Department of Statistics at Texas A&M University on the topic of forecasting wind power. Her research interests include circular statistics and spacetime modelling of environmental data, particularly with nonGaussian distributions.