Wind Energy: Forecasting Challenges for Its Operational Management

Pierre Pinson

Abstract. Renewable energy sources, especially wind energy, are to play a larger role in providing electricity to industrial and domestic consumers. This is already the case today for a number of European countries, closely followed by the US and high growth countries, for example, Brazil, India and China. There exist a number of technological, environmental and political challenges linked to supplementing existing electricity generation capacities with wind energy. Here, mathematicians and statisticians could make a substantial contribution at the interface of meteorology and decision-making, in connection with the generation of forecasts tailored to the various operational decision problems involved. Indeed, while wind energy may be seen as an environmentally friendly source of energy, full benefits from its usage can only be obtained if one is able to accommodate its variability and limited predictability. Based on a short presentation of its physical basics, the importance of considering wind power generation as a stochastic process is motivated. After describing representative operational decision-making problems for both market participants and system operators, it is underlined that forecasts should be issued in a probabilistic framework. Even though, eventually, the forecaster may only communicate single-valued predictions. The existing approaches to wind power forecasting are subsequently described, with focus on single-valued predictions, predictive marginal densities and space–time trajectories. Upcoming challenges related to generating improved and new types of forecasts, as well as their verification and value to forecast users, are finally discussed.

Key words and phrases: Decision-making, electricity markets, forecast verification, Gaussian copula, linear and nonlinear regression, quantile regression, power systems operations, parametric and nonparametric predictive densities, renewable energy, space–time trajectories, stochastic optimization.

1. INTRODUCTION

Increased concerns related to climate evolution and energetic independence have supported the necessary technological and regulatory developments to broaden the energy mix all around the world, with a particular emphasis placed on renewable energy sources (Letcher, 2008). Among the various candidates, wind energy showed the most rapid and consistent deployment of power generating capacities. By June 2012, the cumulative installed wind power capacity worldwide had reached 254 GW, and it is still increasing at a rapid pace [WWEA (World Wind Energy Association), 2012]. Besides all the mathematical and statistical challenges in the development of turbines (aerodynamics, materials and structures, etc.), and in the deployment of wind energy capacities (e.g., wind resource estimation, logistics optimization), those relating to power systems operations and electricity markets have attracted substantial and growing interest over the last 3 decades. This is since, in contrast with conven-

Pierre Pinson is Professor in the Modeling of Electricity Markets, Centre for Electric Power and Energy, Department of Electrical Engineering, Technical University of Denmark, Elektrovej 325(058), 2800 Kgs. Lyngby, Denmark (e-mail: ppin@dtu.dk).

tional generation means, wind power generation cannot be scheduled at will, except maybe by curtailing the output of the wind turbines. Wind power is produced as the wind blows: the dynamics of wind power generation are the result of a nonlinear conversion and filtering of wind dynamics through the turbines' rotor and electric generator. It makes that the traditional operational methods used for conventional generators cannot directly apply to wind energy. For that reason, of the various renewable energy sources, wind, solar, wave and tidal energy are often referred to as stochastic power generation, owing to their inherent variability and uncertainty.

Wind energy is by far the renewable energy source that has attracted the most attention of researchers and practitioners. It is clear, however, that a number of operational and economic issues will be the same for the other forms of renewable energy sources. In practice, such challenges require the modeling and forecasting of the wind power generation process at various temporal and spatial scales, to be subsequently used as input to decision-making. Our objective here is to give an overview of these forecasts and of challenges stemming from their generation and verification. It is to be noted that forecasting is only one aspect of better accommodating renewable energy generation, such as that from the wind into existing power systems and electricity markets. For instance, from a more general perspective of investment, regulation and policy, even the way wind energy should be compared to conventional technologies challenges traditional practice (Joskow, 2011). Similarly, when assessing resource adequacy (i.e., making sure that the overall generating capacity is sufficient to meet demand at all times) and competition in electricity markets, it is argued that the impact of renewable energy sources on market dynamics ought to be accounted for (Wolak, 2013).

The most classical statistical problem involving wind energy is that of resource assessment, that is, focusing on unconditional distributions of wind speed and the corresponding potential power generation. In practice, this is based on estimating marginal wind distributions given a (potentially limited) sample of wind measurements on site and/or in the vicinity of the sites of interest. Even though these marginal distributions are highly valuable for the optimal siting and design of wind farms, they have nearly no value for the operational management of wind power generation: they give an unconditional picture only, hence, they do not give information on the volatile and conditional characteristics of wind and power dynamics at the relevant spatial and temporal scales. A succession of two

papers published in the *Journal of Applied Meteorology* in 1984 is a symbol of the transition from models for limiting distributions only to dynamic models. There, the seminal work of Conradsen, Nielsen and Prahm (1984) on fitting Weibull distributions to samples of wind speed measurements of various lengths is literally followed by that of Brown, Katz and Murphy (1984), which certainly was the first paper looking at dynamic (linear time-series) models for the prediction of wind speed and corresponding power generation. Not so long after, Haslett and Raftery (1989) bridged the gap between the two by focusing on the dynamic spatio-temporal structure of wind speed over Ireland and its implications for the wind energy resource. Since then, ample research was performed on stochastic dynamic models for the prediction of wind power generation at lead times between a few minutes and up to several days ahead, accounting or not for spatial effects. For an exhaustive review of the state of the art in that field, the reader is referred to Giebel et al. (2011), while a solid introduction to the physical concepts involved can be found in Lange and Focken (2006). Our state of knowledge today is that optimal decision-making involving wind power generation calls for predictions generated in a probabilistic framework. These should inform of uncertainties through predictive marginal densities, but also potentially of spatio-temporal dependencies through trajectories, which are known as scenarios in the operations research literature. As a very recent example of how forecasts in their most simple deterministic form, or as space–time trajectories, may be used as input to operational problems, the reader is referred to Papavasiliou and Oren (2013), focusing on a unit commitment problem (i.e., the least-cost dispatch of available generation units) under transmission network constraints.

Wind power generation is first introduced in Section 2 as a stochastic process observed at discrete points in space and in time. Subsequently, in order to underline the importance of probabilistic forecasts (in contrast to deterministic, single-valued forecasts), Section 3 describes representative decision problems involving wind energy in power systems operations and its participation in liberalized electricity markets. Section 4 then covers the various types of forecasts used today and to be employed in the future for optimal decision-making. The paper ends in Section 5 with a discussion that covers (i) the current and foreseen challenges for forecast improvement, (ii) the proposal of thorough and appropriate verification frameworks, and (iii) the importance of bridging the gap between forecast quality and value.

2. WIND POWER GENERATION AS A STOCHASTIC PROCESS

Some of the early works on dynamic modeling and forecasting of wind power generation were cast in a physical deterministic framework, as, for instance, Landberg and Watson (1994) on local wind conditions, and similarly for the follow-up study (Landberg, 1999) on power generation. Today however, there is a broad consensus that wind power generation should be modeled as a stochastic process, whatever the spatial and temporal scales involved. A part of uncertainty comes from our lack of knowledge of all the physical processes involved, combined to our limited ability to account for all of them in mathematical and statistical models. There may also be some inherent uncertainty in the data generating process. The choice for appropriate distributions may not be straightforward.

The physical basics of wind power generation are presented in Section 2.1. Definitions and notation are introduced subsequently in Section 2.2. Finally, the Western Denmark data set is described in Section 2.3. It will be used for illustrating the different forms of forecasts that will be described throughout the paper.

2.1 Physical Basics of Wind Power Generation

The generation of electric power from the wind relies on atmospheric processes. The power output of a single wind turbine is a direct function of the strength of the wind over the rotor swept area. Coarsely simplifying the meteorological aspects involved, winds originate from the movement of air masses from high to low pressure areas: the larger the difference in pressure, the stronger the resulting winds. On top of that come the boundary layer effects, complexifying wind behavior due to natural obstacles, friction effects, the nature of the surface itself, temperature gradients, etc. The boundary layer is formally defined as the lower part of the atmosphere where wind speed is affected by the surface. The resulting level of complexity makes that the characteristic features of wind variability may be better described in the frequency domain (Mur Amada and Bayod Rújula, 2010). Our state of the knowledge on wind dynamics in the boundary layer, and, more generally, mesoscale meteorology, is today still limited: resulting models of wind characteristics have systematic and random errors.

Wind speed exhibits fluctuations over a wide range of frequencies. Those in the order of days are induced by the movement of synoptic weather patterns, that is, by general changes in weather situations. These are modeled within global weather models such as those run at the European Centre for Medium-range Weather Forecasts (ECMWF, in the UK) and at the National Centers for Environmental Prediction (NCEP, in the US), among others. Those models encompass well-known dynamics of state variables for the global weather, while wind components are a by-product derived from the evolution of these state variables. In terms of forecasting, several directions are thought of today for improving the estimation of the initial state of the Atmosphere and also to better account for potential uncertainties in the model and its parametrization (Palmer, 2012).

Fluctuations referred to as diurnal and semi-diurnal cycles (with periods of 24 and 12 hours) are mainly induced by thermal exchanges between the surface (land or sea) and the Atmosphere. Their magnitude varies as a function of local climate and seasons. At these time scales, the phenomena involved are fairly well known, though certain aspects like their impact on wind profiles (i.e., the way wind evolves with height) still are a subject of active research, for example, Peña Diaz, Gryning and Mann (2010). At frequencies in the order of minutes to hours, local effects potentially including the presence of cumulus clouds, convective cells, precipitation, waves (for offshore sites), etc. are the drivers of wind speed variations. Here, the physical and mathematical aspects may become more challenging owing to the combination of a substantial number of interacting physical processes. Higher frequencies (seconds to a few minutes, not considered in the present paper) see a dominance of turbulence effects, which are a particular concern for the structural design of turbines, fatigues studies and, potentially, control. Finally, at the other end of the spectrum, very low frequencies also seen as long-term wind trends, have attracted increased attention recently since human activity and climate evolution may potentially impact surface winds at these time scales; see Vautard et al. (2010), for instance. In the following, emphasis is placed on time scales in the order of minutes to days, where existing meteorological challenges include the better understanding of the physical processes and their interaction, as well as their modeling.

Wind speed is the meteorological variable of most relevance to power generation. The process of the conversion of wind to electric power for a single wind turbine is described by its power curve. It is also influenced by air density (being a function of temperature, pressure and humidity) to a minor extent. Power curves for different turbines roughly have the same shape for
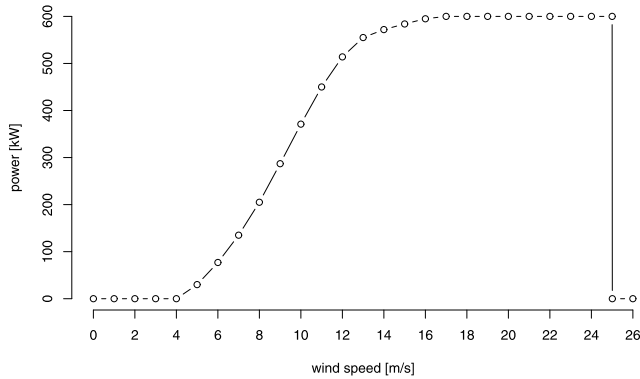
FIG. 1. *Power curve of the Vestas V44 (600 kW) wind turbines installed at the Klim wind farm, for an air density of* 1.225 *kg/m*^{3}.

all manufacturers and turbine types. In order to discuss and illustrate what manufacturer's (i.e., theoretical) and observed power curves may look like, let us take the example of the Klim wind farm in Western Denmark. It is composed of 35 Vestas V44 wind turbines having a capacity of 600 kW each, yielding a nominal capacity of 21 MW. The nominal capacity of a wind turbine or of a wind farm is the power output it generates within the range of wind conditions over which it was designed to operate, ideally. Figure 1 depicts the power curve for a V44 turbine. The power production is null below the cut-in wind speed (4 m/s), then sharply augments between the cut-in and rated wind speeds (16 m/s). At rated speed, it reaches its nominal power P_n. The power production is nearly constant between rated and cut-off wind speeds (here 25 m/s). At cut-off speed, the turbine stops for security reasons. This power curve example is for a fairly old wind turbine model, since this wind farm started operating in 1996. Various technological improvements have been permitted to lower cut-in and rated wind speeds, which are today between 2 and 4 m/s for the former one and between 12 and 15 m/s for the latter one. Moreover, cut-off wind speeds may reach up to 34 m/s. In a general manner there may also be a difference between the maximum (peak) and nominal power values (up to 10–20%). Most importantly, the nominal capacity of today's wind turbines is up to 7–8 MW.

A power curve such as in Figure 1 is a theoretical one, since it gives the power output of a single turbine exposed to ideal wind conditions as if in a wind tunnel (i.e., not altered by obstacles, without turbulence and for the turbine always perfectly facing the wind), for a given air density. In practice, however, wind turbines are almost always gathered in wind farms with potentially a mix of different turbine types. The combination

of these individual power curves will not be the same as that of any of the individual turbine types. Besides, depending upon the prevailing wind direction, some of the turbines within a wind farm may mask the others— the so-called shadowing effect, therefore reducing the wind seen by these turbines. This combines with additional surrounding topographic and orographic effects (i.e., hills, forest, etc.), making that the various turbines within a wind farm are constantly seeing different wind conditions, which also are different from the free-stream wind at a reasonable distance away from the wind farm. Consequently, the resulting wind farm power curve has features far more complex than the theoretical power curves provided by the manufacturers for individual wind turbines.

Figure 2 depicts the empirical power curve of the Klim wind farm based on hourly wind speed (at 10 m above ground level) and power measurements collected over the first 6 months of 2002. For both types of measurements, a record for a given point in time corresponds to the average value over the preceding hour. Measurement errors in power and wind speed observations certainly contribute to the scatter of data observed. However, the main reason for that scatter is the impact of other meteorological variables, as, for example, wind direction and air density, on the power generation from the wind farm. For measured wind speeds of 5 m/s the observed power output of the wind farm varies between 0 and 7 MW, while for wind speeds
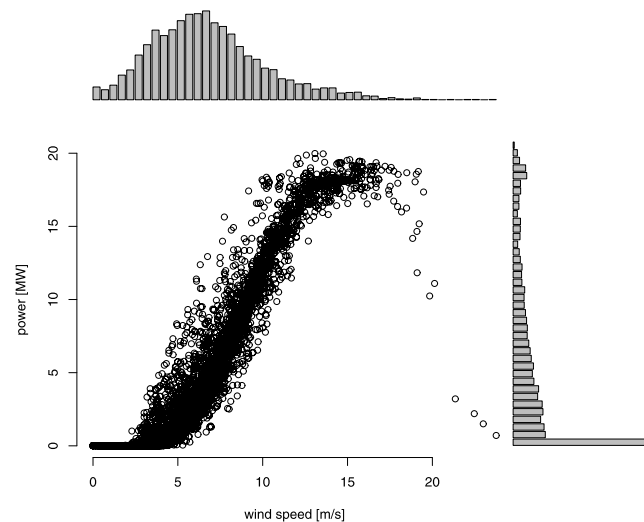

FIG. 2. *Example empirical power curve for the Klim wind farm over a* 6-*month period in* 2002, *based on hourly measurements of wind speed and corresponding power output. Marginal distributions of wind speed and power are also represented above, and, respectively, right of, the power curve itself.*

of 10 m/s, that same power output may be between 6 and 15 MW. Other reasons for these variations include natural changes in the environment of the wind farm, aging of turbine components, etc. At the turbine, wind farm or portfolio (i.e., a group of geographically distributed wind farms, though jointly operated) level, all empirical power curves exhibit characteristics differing from those of theoretical ones, also with a substantial scatter of observations. Other interesting empirical power curves for wind farms in Crete, as well as challenges related to their modeling, were recently discussed by Jeon and Taylor (2012).

2.2 Preliminaries and Definitions

Owing to the combination of complex physical processes, and since we may not have a perfect understanding of all these processes anyway, it is acknowledged that one should account for a random uncertainty component in the modeling of energy generation from wind turbines. Wind power is therefore considered as a discrete stochastic process, that is, as a set of random variables $Y_{s,t}$ observed at discrete points in time t and in space s. Depending upon the practical setup, it may reduce to a temporal process with a set of random variables Y_t for successive times, for instance, if concentrating on a single wind farm or on a fixed (geographically spread) portfolio, or to a spatial process with a set of random variables Y_s for a given time but for various locations, for instance, if looking at maps of wind energy resource over a region. The corresponding realizations of the process are denoted by $y_{s,t}$ in the more general spatio-temporal case, or, more simply, by y_t and y_s in the temporal and spatial cases, respectively. The notation f and F are used for probability density and cumulative distribution functions (abbreviated p.d.f. and c.d.f.) of the random variables involved, with appropriate indices.

Wind power generation as a stochastic process exhibits features that can be seen as fairly unique, even though relevant parallels with stochastic processes for other renewable energy sources, in meteorology and hydrology or in economics and finance, exist. Some of these characteristic features come from the very nature of wind, while some others are directly linked to the process of converting the energy in the wind to electric power. First of all, wind components and resulting wind speed have a combination of dynamic and seasonal features, which may vary depending on local wind climates and regions of the world. Besides, when focusing on spatial and temporal scales of relevance to power systems operations and electricity markets, the various meteorological phenomena involved

induce switches in the dynamic behavior of wind fluctuations and in their predictability, yielding a nonstationary process [see the discussion by Vincent et al. (2010), for instance]. Inspired by models developed in the econometrics literature, the existence of successive periods with different levels of predictability of wind speeds was first captured with a Generalised Auto Regressive Conditional Heteroscedastic (GARCH) model by Tol (1997), though focusing on coarser daily wind records.

In parallel, the conversion of the energy in the wind to electric power acts as a nonlinear transfer function (as represented in Figure 2) making wind power generation a nonlinear and bounded stochastic process. There may even be smooth temporal changes in this nonlinear transfer function owing to, for example, aging of equipment, changes in external environment, etc. The transfer function shapes the predictability of wind power generation. Consequently, conditional densities of wind power generation should be seen as non-Gaussian, with their moments of order greater than 1 directly influenced by their mean (Lange, 2005; Bludszuweit, Domínguez-Navarro and Llombart, 2008). Truncated Gaussian, Censored Gaussian and Generalized Logit–Normal distributions were proposed as relevant candidates for the modeling of conditional densities of wind power generation (Pinson, 2012). In terms of stochastic differential equations, this would translate to having a state-dependent diffusion component. The flat parts of the transfer function also yield concentration of probability mass at the boundaries, potentially requiring to consider wind power generation as a discrete-continuous mixture, similar to precipitation, for instance.

After proposing a suitable model structure, and estimating its parameters, such a model may be employed to simulate time-series of wind power generation for one or several locations, for instance, as input to power systems and market-related analysis. In most cases, however, forecasting is the final application. Predictions fed into operational decision problems always are for future points in time and rarely for new locations at which no observations are available. Consequently, even though spatial aspects are of crucial interest, the problem at hand is mainly seen as a temporal forecasting problem. The set of m locations is denoted by

$$(2.1) \qquad \mathbf{s} = \{s_1, s_2, \ldots, s_m\}.$$

In parallel, the set of n lead times is

$$(2.2) \qquad t + \mathbf{k} = \{t+1, t+2, \ldots, t+n\},$$

where n is the forecast length. Lead times are spaced regularly and with a temporal resolution equal to the sampling time of the process observations. Since the power generation process is bounded, it can be marginally normalized, so that

$$(2.3) \qquad \mathbf{y}_{\mathbf{s},t+\mathbf{k}} \in [0, 1]^{mn}.$$

At time t the aim is to predict some of the characteristics of

$$(2.4) \quad \begin{aligned} &\mathbf{Y}_{\mathbf{s},t+\mathbf{k}} \\ &= \{Y_{s,t+k}; s = s_1, \dots, s_m, k = 1, \dots, n\}, \end{aligned}$$

a multivariate random variable of dimension $m \times n$ in the complete spatio-temporal case, or of

$$(2.5) \qquad \mathbf{Y}_{t+\mathbf{k}} = \{Y_{t+k}; k = 1, \dots, n\},$$

a multivariate random variable of dimension n, in the simpler setup where spatial considerations are disregarded.

In the most general case, the forecaster issues at time t for the set of lead times $t + \mathbf{k}$ a probabilistic forecast $\hat{F}_{\mathbf{s},t+\mathbf{k}|t}$ (here a predictive c.d.f.) describing as faithfully as possible the c.d.f. $F_{\mathbf{s},t+\mathbf{k}}$ of the random variable $\mathbf{Y}_{\mathbf{s},t+\mathbf{k}}$, given the information available up to time t. It hence translates to a full description of marginal densities for every location and lead time, as well as spatio-temporal dependencies among the set of m locations and n lead times. This clearly comprises a difficult problem, both in terms of generating such forecasts and also for their verification. Consequently, since degenerate versions of that problem may be more tractable, a number of them have been dealt with in the literature, for instance, for the forecasting of marginal densities for each location and lead time individually, or even by forecasting some summary statistics (more precisely, mean and quantiles) of these marginal densities only.

The combination of all uncertainties, related to physical aspects to be accounted for in the models, but also in connection with the data-generating process, obviously is going to impact the quality of the resulting forecasts. In Section 4 some of the most common approaches to forecasting will be reviewed. They all tend to disregard the specific contributions of physical and data-generating processes to forecast quality. Alternative proposals in a robust forecasting framework could therefore be beneficial.

2.3 The Western Denmark Data Set

A data set for the Western Denmark area is used as a basis for illustration. It consists of wind power measurements as collected by Energinet.dk, the transmission system operator in Denmark. This region has one of the highest wind power penetrations (i.e., the share of wind power in meeting the electric energy demand) in the world, consistently between 25 and 30% over the last few years.

Wind power measurements are originally available at more than 400 geographically distributed grid-connection points. Observations have an hourly resolution over a period between 1 January 2006 and 24 October 2007. They represent average hourly power values. For operational purposes, these are gathered in 15 so-called control zones depicted in Figure 3 along with their identification numbers. The total nominal capacity slightly evolved during this period though generally being around 2.5 GW. In order to additionally simplify this case-study, the original 15 control zones are aggregated into 5 zones only (see Figure 3), each having a different share of the overall wind power capacity for that region. All power measurements are normalized by the respective nominal capacities of the 5 aggregated zones. This aggregation is for the sake of example only and could be seen as wind power generation portfolios operated by a set of power producers in that region. Working at such a coarse spatial resolution certainly is sufficient for some decision problems, also simplifying modeling and estimation challenges. However, it may be that for some applications the statistician and forecaster has to work with the original 400-location data set, so that he has to finely analyze and model the observed spatio-temporal dynamics; see Girard and Allard (2013), for instance. This would be the case if all the owners/operators of these individual wind farms ask for predictions in order to design market offering strategies or for the network operator to perform very detailed system simulations based on the impact of spatially distributed wind power generation.

Some of the features of this data at such temporal and spatial scales can be observed from the example episode with 24 hours of hourly wind power measurements in Figure 4, for the 5 aggregated zones of Western Denmark. The spatio-temporal interdependence structure of the wind power generation process, as induced by the inertia in weather phenomena and resulting winds, especially results in smooth temporal variations at each zone, individually, as well as in similarities in the patterns observed at the various
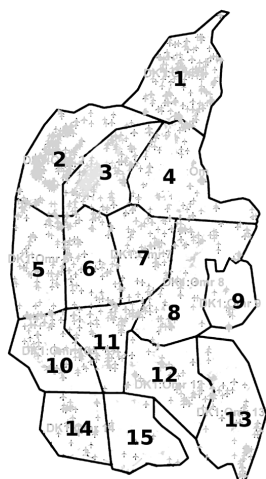
Agg. zone	Orig. zones	% of capacity
1	1, 2, 3	31
2	5, 6, 7	18
3	4, 8, 9	17
4	10, 11, 14, 15	23
5	12, 13	10

FIG. 3.   *The Western Denmark data set*: *original locations for which measurements are available*, 15 *control zones defined by Energinet.dk, as well as the* 5 *aggregated zones. The total nominal capacity for Western Denmark was* 2.5 *GW over the period covered by this data set.*

zones. These spatio-temporal dependencies are necessarily strengthened by the aggregation procedure employed. For instance, the drop in power generation observed in zone 4 on 19 March 2007 at 8:00 UTC (i.e., the 20th time step) is also visible for zone 5, at the same time and with a similar magnitude, while it may potentially be related to a drop of lesser magnitude observed in zones 2 and 3 around the same time. Note that UTC (for Coordinated Universal Time) is the most common standard for referring to time in the meteorological and wind energy communities.

3. SOME REPRESENTATIVE OPERATIONAL DECISION-MAKING PROBLEMS INVOLVING WIND ENERGY

Some of the representative operational decision problems are described here, while a more extensive overview of such problems may be found in Ackermann (2012). The side of power producers is taken first, by considering the issue of designing optimal offering strategies in electricity markets. Subsequently taking the side of the system operator instead (like Energinet.dk, the transmission system operator for Western Denmark), an issue of rising importance is that of quantifying the necessary backup generation to accommodate variability and limited predictability of wind power generation. These two decision-making problems are somehow interrelated, since the quantification of necessary backup capacities is performed in a dynamic way, conditional on the clearing of the electricity market. For both types of problems, forecasts for other quantities than wind power generation may be

necessary, like load and prices. There exists substantial literature on the statistical modeling and forecasting of these market variables. The interested reader is referred to Weron (2006) for an overview.

3.1 Participation of Wind Energy in Electricity Markets

In a number of countries with significant wind power generation, electricity markets are organized as electricity pools, gathering production and consumption offers in order to dynamically find the quantities and prices for electricity generation and consumption that permit to maximize social welfare. These electricity pools typically have two major stages which are the day-ahead and the balancing markets. The electricity pool for Scandinavia, used as an example here, is commonly known as the Nord Pool. For an overview of some the European electricity markets and of the way they deal with wind power generation, see Morthorst (2003). A parallel overview for the case of US electricity markets can be found in Botterud et al. (2010).

Electricity exchanges first take place in the day-ahead market for the next delivery period, that is, the next day. Production offers and consumption bids are to be placed for every time unit before gate closure, occurring 12 hours before delivery in the Nord Pool, where market time units are hourly. At the time t of gate closure, wind power producers shall propose energy offers based on forecasts with lead times $t + k$, $k \in \{13, 14, \ldots, 37\}$. The market clearing is there to match production offers and consumption bids through a single auction process, yielding the system price π_{t+k}^{c}
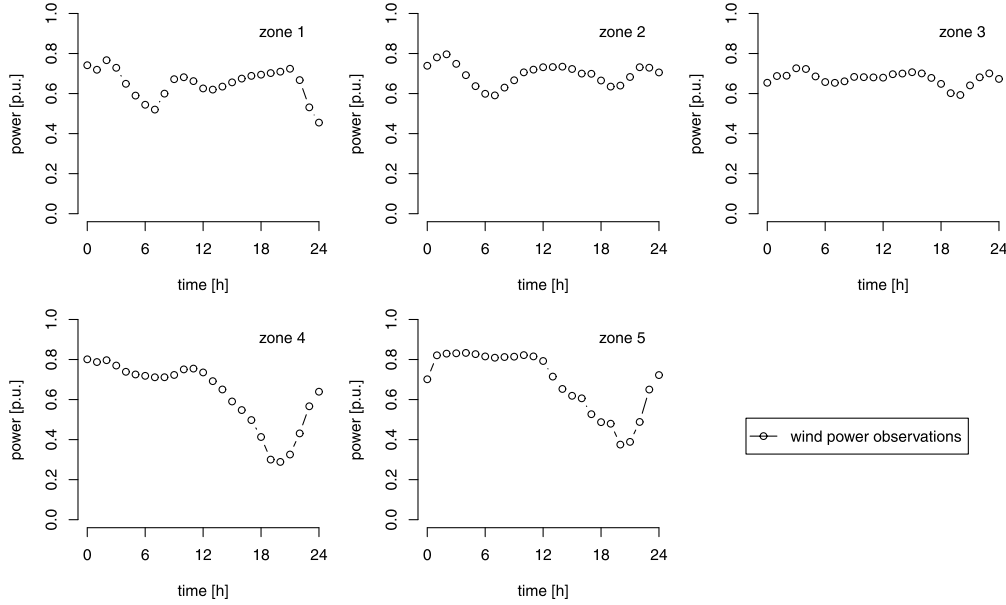
FIG. 4. *Example episode with normalized wind power measurements for the 5 zones of the Western Denmark data set over 24 hours, starting from the 18 March 2007 at 12 UTC.*

and the program of the market participants, that is, a set of energy blocks y_{t+k}^c to be delivered by wind power producers,[1] for every market time unit. The superscript c indicates that this combination of energy quantity and price defines a contract. Power producers are financially responsible for any deviation from this contract. Indeed, in a second stage, the balancing market managed by the system operator ensures the real-time balance between generation and load, while translating to financial penalties for those who deviate from their contracted schedule. The prices for buying and selling on the balancing market are denoted by π_{t+k}^b and π_{t+k}^s, respectively. They are generally less advantageous than those in the day-ahead market, fairly volatile and substantially different from one another in a two-price settlement system like that of the Nord Pool. The combination of the inherent uncertainty of wind power predictions and of the asymmetry of balancing prices encourages market participants to be more strategic when designing offering strategies (Skytte, 1999).

Simplifying the decision problem for clarity, potential dependencies among time units and in space throughout the network are disregarded. A wind power producer is seen as participating with a global portfolio of wind power generation in the electricity market.

---

[1]Note that the notation y_{t+k}^c is used abusively for simplification. This is since the energy block for hour $t + k$ is necessarily equal to the average power production value y_{t+k}^c over that one-hour period.

The overall market revenue R_{t+k} is a random variable, which, given the decision variable y_{t+k}^c and the random variable Y_{t+k}, can be defined as

$$(3.1) \quad R_{t+k} = s_{t+k}(y_{t+k}^c) + B_{t+k}(Y_{t+k}, y_{t+k}^c),$$

where the first part corresponds to the revenue from the day-ahead market, $s_{t+k}(y_{t+k}^c) = \pi_{t+k}^c y_{t+k}^c$, while the second is that from the balancing market, to be detailed below. Following Pinson, Chevallier and Kariniotakis (2007) (among others), this revenue can be reformulated as a combination of revenues and costs in a way that the decision variable appears in the balancing market term only

$$(3.2) \quad R_{t+k} = \tilde{S}_{t+k}(Y_{t+k}) - \tilde{B}_{t+k}(Y_{t+k}, y_{t+k}^c),$$

that is, as the sum of a stochastic, though fatal since out of the control of the decision-maker, component $\tilde{S}_{t+k}$ from selling of the energy actually produced through the day-ahead market, minus another stochastic component $\tilde{B}_{t+k}$, whose characteristics may be altered through the choice of a contract y_{t+k}^c. The imbalance is also a random variable, given by $Y_{t+k} - y_{t+k}^c$, yielding the following definition for $\tilde{B}_{t+k}$:

$$(3.3) \quad \begin{aligned} &\tilde{B}_{t+k}(Y_{t+k}, y_{t+k}^c) \\ &= \begin{cases} \pi_{t+k}^{\downarrow}(Y_{t+k} - y_{t+k}^c), & Y_{t+k} - y_{t+k}^c \geq 0, \\ -\pi_{t+k}^{\uparrow}(Y_{t+k} - y_{t+k}^c), & Y_{t+k} - y_{t+k}^c < 0, \end{cases} \end{aligned}$$

where $\pi_{t+k}^{\downarrow}$ and $\pi_{t+k}^{\uparrow}$ are referred to as the regulation unit costs for downward and upward balancing, respectively. They are readily given by

$$(3.4) \qquad \pi_{t+k}^{\downarrow} = \pi_{t+k}^c - \pi_{t+k}^s,$$

$$(3.5) \qquad \pi_{t+k}^{\uparrow} = \pi_{t+k}^b - \pi_{t+k}^c.$$

For most electricity markets regulation unit costs are always positive, making $\tilde{B}_{t+k} \geq 0$, while the overall market revenue has an upper bound obtained when placing an offer corresponding to a perfect point prediction, $y_{t+k}^c = \hat{y}_{t+k|t} = y_{t+k}$. As this is not realistic, and accounting for the uncertainty in wind power forecasts, optimal offering strategies are to be derived in a stochastic optimization framework. Assuming that the wind power producer is rational, his objective is to maximize the expected value of his revenue for every single market time unit, since this permits to maximize revenues in the long run. Additionally considering the market participant as a price-taker (i.e., not influencing the market outcome by his own decision), and having access to forecasts of the regulation unit costs ($\hat{\pi}_{t+k|t}^{\downarrow}$ and $\hat{\pi}_{t+k|t}^{\uparrow}$), the optimal production offer y_{t+k}^* at the day-ahead market is given by

$$(3.6) \qquad y_{t+k}^* = \arg\min_{y_{t+k}^c} \mathbb{E}\big[B_{t+k}\big(Y_{t+k}, y_{t+k}^c\big)\big].$$

This stochastic optimization problem has a closed-form solution, as first described by Bremnes (2004), that is, for any market time unit $t + k$, the optimal wind power production offer is given by

$$(3.7) \qquad y_{t+k}^* = \hat{F}_{t+k|t}^{-1}\left(\frac{\hat{\pi}_{t+k|t}^{\downarrow}}{\hat{\pi}_{t+k|t}^{\downarrow} + \hat{\pi}_{t+k|t}^{\uparrow}}\right),$$

where $\hat{F}_{t+k|t}$ is the predictive c.d.f. issued at time t (the decision instant) for time $t + k$. In other words, the optimal offer corresponds to a specific quantile of predictive densities, whose nominal level α is a direct function of the predicted regulation unit costs for this market time unit. That problem is a variant of the well-known linear terminal loss problem, also called the newsvendor problem (Raiffa and Schaifer, 1961). It was recently revisited by Gneiting (2010), who showed that for a more general class of cost functions [i.e., generalizing that in (3.3)] optimal point forecasts are quantiles of predictive densities with nominal levels readily determined from the utility function itself, analytically or numerically. Note that appropriate forecasts of regulation unit costs are also needed here. It was shown

by Zugno, Jónsson and Pinson (2013) and the references therein that these may be obtained from variants of exponential smoothing (in its basic form or as a conditional parametric generalization) and then directly embedded in offering strategies such as those given by (3.7).

In their simplest form, market participation problems involving wind energy rely on a family of piecewise linear and convex loss functions, for which optimal offering strategies are obtained in a straightforward manner, as in the above. These only require quantile forecasts for a given nominal level or maybe predictive densities of wind power generation for each lead time individually. However, when complexifying the decision problem by adding dependencies in space (e.g., spatial correlation in wind power generation, network considerations) and in time (e.g., accounting for the temporal structure of forecast errors), it then requires a full description of $\mathbf{Y}_{s,t+k}$ (ideally in the form of trajectories), instead of marginal densities for the whole portfolio and for each lead time individually. The same goes for alternative strategies of the decision-makers, for instance, if one aims to account for risk aversion. The resulting mathematical problems do not rely on studying specific families of cost functions, but instead translate to formulating large scenario-based optimization problems, in a classical operations research framework. Some of the resulting stochastic optimization problems may be found in Conejo, Carrion and Morales (2010). The price-taker assumption is also to be relaxed to a more general stochastic optimization framework, where wind-market dependencies are to be described and accounted for (Zugno et al., 2013).

3.2 Quantification of Necessary Power Systems' Reserves

On the other side, the electric network operator has the responsibility to ensure a constant match of electricity generation and consumption, outside of the market framework described before. It involves the quantification of so-called reserve capacities, prior to actual operations, to be readily available if needed. This may be either for supplementing generation lacking in the system, for example, in case of asset outages, general loss of production and unforeseen increase in electricity demand, or, alternatively, for lowering the overall level of generation in the system when demand is less than production. For an overview, see Doherty and O'Malley (2005).

For simplicity and clarity, the timeline here is the same as for the market participation problem described

earlier. Potential dependencies among time units and in space throughout the network (as induced by potential network congestion) are disregarded. The system operator has to make a decision at time t (market gate closure) for all time units $t + k$ of the following day. Reserves are to be quantified as two numbers $q_{t+k}^{\downarrow}$ and $q_{t+k}^{\uparrow}$ for the whole power system, for downward (when consumption is less than production) and upward (conversely) balancing, respectively. The choice of optimal reserve levels is linked to a random variable O_{t+k} describing all potential deviations from the chosen dispatch (consisting in the reference values for generation and consumption at every time $t + k$). This random variable is commonly referred to as the system generation margin.

O_{t+k} can be defined as a sum of random variables representing all uncertainties involved. These include (i) potential forecast errors ε_L for the electric load, (ii) the probability of generation loss through asset outages (assets being conventional generators, transmission lines and other equipment), and (iii) potential forecast errors ε_Y for the various forms of stochastic power generation. For simplicity, we only consider wind power here, corresponding to the operational situation where, as in most countries, wind power is the prominent form of stochastic power generation. In a more general setup the combination of uncertainties with, for example, solar and wave energy, should also be accounted for. These various uncertainties are fully characterized by probabilistic forecasts available at time t: $\hat{f}_{t+k|t}^{\varepsilon_L}$ for the load, $\hat{f}_{t+k|t}^G$ for generation losses, and $\hat{f}_{t+k|t}^{\varepsilon_Y}$ for wind power generation. This means that, besides the wind generation forecasts discussed in this paper, additional predictions of potential generation losses (e.g., the probability of failure of various equipment) are to be issued, for instance, based on reliability models in the spirit of Billinton and Allan (1984). Forecasts for the electric load can in addition be obtained from one of the numerous methods recently surveyed by Hahn, Meyer-Nieberg and Pickl (2009), though very few of them look at full predictive densities.

Assuming independence of the various random variables, the overall uncertainty, represented by a predictive marginal density $\hat{f}_{t+k|t}^O$, is obtained through convolution,

$$(3.8) \qquad \hat{f}_{t+k|t}^O = \hat{f}_{t+k|t}^{\varepsilon_L} * \hat{f}_{t+k|t}^G * \hat{f}_{t+k|t}^{\varepsilon_Y}.$$

This predictive density is split into its positive and negative parts, yielding $\hat{f}_{t+k|t}^{O^+}$ and $\hat{f}_{t+k|t}^{O^-}$, since decisions about downward and upward reserve capacities are to be made separately.

After such a description of system-wide uncertainties, the system operator can plug this density into a cost-loss analysis (Matos and Bessa, 2010), similar in essence to the market participation problem presented in the above. Based on cost functions $g^{\downarrow}$ and $g^{\uparrow}$ for the downward and upward cases, the optimal amounts of reserve capacities (in an expected utility maximization sense) are the solution of stochastic optimization problems of the form

$$(3.9) \qquad q_{t+k}^{\uparrow*} = \underset{q_{t+k}^{\uparrow}}{\arg\min} \, \mathbb{E}[g^{\uparrow}(O_{t+k}^-, q_{t+k}^{\uparrow})]$$

and

$$(3.10) \qquad q_{t+k}^{\downarrow*} = \underset{q_{t+k}^{\downarrow}}{\arg\min} \, \mathbb{E}[g^{\downarrow}(O_{t+k}^+, q_{t+k}^{\downarrow})],$$

which may be solved analytically or numerically, depending upon the complexity of the cost functions. Here the optimal reserve levels relate to specific quantiles of the predictive densities for the system margin O_{t+k}. However, it would be difficult to link the optimal reserve levels to specific quantiles of the input predictive densities of wind power generation.

In its more complex form the reserve quantification problem requires accounting for dependencies in space and in time, similar to the trading problems, with many more considerations relating to operational constraints, for example, unit characteristics (capability to increase or decrease power output over a predefined period of time—so-called ramping characteristics, nonconvexities in costs, etc.), and, potentially, risk aversion. The resulting stochastic optimization problems take the general form of those described and analyzed in Ortega-Vazquez and Kirschen (2009). They require space–time trajectories for all input variables.

4. MODELING AND FORECASTING WIND POWER IN A PROBABILISTIC FRAMEWORK

Decision-making problems relating to an optimal management of wind power generation in power systems and electricity markets require different types of forecasts as input. The lead forecast range considered in the above is between 13 and 37 hours ahead, with an hourly temporal resolution for the forecasts. In practice, various forecast ranges, spatial and temporal resolutions, are of relevance depending upon the decision problem. For instance, the shorter lead times, say, between 10 minutes and 2 hours ahead, are also crucial for a number of dispatch and control problems. Below are presented the leading forms of forecasts for wind power generation, as well as example approaches to generate them.

4.1 Point Predictions

The traditional deterministic view of a large number of power system operators translates to preferring single-valued forecasts. These so-called point predictions are seen as easier to appraise and handle at the time of making decisions.

When describing at time t the random variable $\mathbf{Y}_{\mathbf{s},t+\mathbf{k}}$ of a set of locations $\mathbf{s}$ and lead times $\mathbf{k}$, point forecasts comprise a summary value for each and every marginal distribution of $\mathbf{Y}_{\mathbf{s},t+\mathbf{k}}$ in time and in space. Typically, if one aims at minimizing a quadratic criterion (i.e., in a Least Squares sense), a point forecast for location s and lead time k corresponds to the conditional expectation for $Y_{s,t+k}$ given the information set available up to time t, the chosen model and estimated parameters. With respect to a predictive density $\hat{f}_{s,t+k|t}$ for location s and lead time k, that point forecast therefore corresponds to

$$(4.1) \qquad \hat{y}_{s,t+k|t} = \int_0^1 y \hat{f}_{s,t+k|t}(y)\, dy.$$

Integration is between 0 and 1 since one is dealing with power values normalized by the nominal capacity of the wind farm or group of wind farms of interest.

To issue point predictions at time t, the forecaster utilizes an information set Ω_t, a set containing measurements Ω_t^o (including observations of power and of relevant meteorological variables, the notation "o" meaning "observation") over the area covered, as well as meteorological forecasts Ω_t^f (with "f" for "forecast") for these relevant variables, $\Omega_t \subseteq \Omega_t^o \cup \Omega_t^f$. Based on this wealth of available information, different types of models of the general form

$$(4.2) \qquad Y_{s,t+k} = h(\Omega_t) + \varepsilon_{s,t+k}$$

were proposed, where $\varepsilon_{s,t+k}$ is a noise term with zero mean and finite variance.

Indeed, when focusing on a single location (a wind farm), it may be that point forecasts can be issued in an inexpensive way based on local measurements only, and in a linear time-series framework. The first proposal in the literature is that of Brown, Katz and Murphy (1984), using Auto-Regressive Moving Average (ARMA) models for wind speed observations and for lead times between a few hours and a few days. When focusing on wind power directly for very short range (say, for lead times less than 2 hours), even simpler Auto-Regressive models of order p, that is,

$$(4.3) \quad Y_{s,t+k} = \theta_0 + \sum_{i\in\mathcal{L}} \theta_i Y_{s,t-i+1} + \sigma\varepsilon_{s,t+k},$$

are difficult to outperform, possibly after data transformation (Pinson, 2012). In the above, $\mathcal{L} \subset \mathbb{N}^+$ is a set of lag values of dimension p, while $\varepsilon_{s,t+k}$ is a standard Gaussian noise term, scaled by a standard deviation value σ. In addition, $k = 1$ if the AR model is for 1-step ahead prediction only, or to be used in an iterative fashion for k-step ahead prediction, while $k > 1$ if one uses the AR model for direct k-step ahead forecasting.

These models were generalized by a few authors by accounting for off-site observations and/or by accounting for regime-switching dynamics of the time-series. A regime-switching version of the model in (4.3) assumes different dynamic behaviors in the various regimes, as expressed by

$$(4.4) \quad Y_{s,t+k} = \theta_0^{(r_t)} + \sum_{i\in\mathcal{L}} \theta_i^{(r_t)} Y_{s,t-i+1} + \sigma^{(r_t)}\varepsilon_{t+k},$$

where r_t is a realization at time t of a regime sequence defined by discrete random variables, with $r_t \in \{1, 2, \ldots, R\}$, $\forall t$, and R is the number of regimes. The number of lags and the noise variance may differ from one regime to another. The regime sequence can be defined based on an observable process, like wind direction at time t or a previous wind power measurement, yielding models of the Threshold Auto-Regressive (TAR) family, which are common in econometrics (Tong, 2011). As an example for wind speed modeling and forecasting, Reikard (2008) proposed to consider temperature as driving the regime-switching behavior in wind dynamics. In contrast, the class of Markov-Switching Auto-Regressive (MSAR) models, also popular in econometrics since the work of Hamilton (1989), assumes that the regime sequence relies on an unobservable process. MSAR models were shown to be able to mimic the observable switching in wind power dynamics, especially offshore, that cannot be explained by available meteorological measurements (Pinson and Madsen, 2012).

Incorporating off-site information in regime-switching models of the form of (4.4) was proposed by Gneiting et al. (2006) and subsequently in a more general form by Hering and Genton (2010), when focusing on a data set for the Columbia Basin of eastern Washington and Oregon in the US. The model in the regime-switching space–time (RST) approach originally proposed by Gneiting et al. (2006) can be formulated as

$$
\begin{aligned}
Y_{s,t+k} = {}& \theta_0^{(r_t)} + \sum_{i\in\mathcal{L}} \theta_i^{(r_t)} Y_{s,t-i+1} \\
& + \sum_{s_j\in\mathcal{S}} \sum_{i\in\mathcal{L}_j} \nu_{j,i}^{(r_t)} Y_{s_j,t-i+1} + \sigma^{(r_t)}\varepsilon_{t+k},
\end{aligned}
$$

$$(4.5)$$

that is, in the form of a TARX model (TAR with exogenous variables), where a set of terms is added to the regime-switching model of (4.4), for observations at off-site locations $s_j \in \mathcal{S}_X$ and for a set of lagged values $i \in \mathcal{L}_j$ at this location. Such models allow considering advection and diffusion of upstream information, but require extensive expert knowledge for optimizing the model structure.

Conditional parametric AR (CP–AR) models are another natural generalization of regime-switching models,

$$(4.6) \quad \begin{aligned} Y_{s,t+k} = \theta_0(\mathbf{x}_t) + \sum_{i \in \mathcal{L}} \theta_i(\mathbf{x}_t) Y_{s,t-i+1} \\ + \sigma(\mathbf{x}_t)\varepsilon_{s,t+k}, \end{aligned}$$

where instead of considering various regimes with their own dynamics, the AR coefficients are replaced by smooth functions of a vector (of low dimension, say, less than 3) of an exogenous variable $\mathbf{x}$, for instance, wind direction only in Pinson (2012). The noise variance can be seen as a function of $\mathbf{x}$, or as a constant, for simplicity. CP–AR models are relevant when switches between dynamic behaviors are not that clear. Meanwhile, they also require fairly large data sets for estimation, which are more and more available today. Their use is motivated by empirical investigations at various wind farms, where it was observed that specific meteorological variables (e.g., wind direction, atmospheric stability) can substantially impact power generation dynamics and predictability in a smooth manner.

Other forms of conditional parametric models were proposed for further lead times, also requiring additional meteorological forecasts as input. As an example, a simplified version of the CP–ARX model (CP–AR with exogenous variables) of Nielsen (2002) writes

$$(4.7) \quad \begin{aligned} Y_{s,t+k} \\ = \theta_0^c(\mathbf{x}_t) \cos\left(\frac{2\pi h_{t+k}}{24}\right) + \theta_0^s(\mathbf{x}_t) \sin\left(\frac{2\pi h_{t+k}}{24}\right) \\ + \theta_1(\mathbf{x}_t) Y_{s,t} + \theta_2(\mathbf{x}_t) g(\hat{u}_{t+k|t}, \hat{v}_{t+k|t}, k) \\ + \sigma \varepsilon_{s,t+k}, \end{aligned}$$

where $\hat{u}_{t+k|t}$ and $\hat{v}_{t+k|t}$ are forecasts of the wind components (defining wind speed and direction) at the level of the wind farm of interest. The vector $\mathbf{x}_t$ includes wind direction and lead time. In addition, g is used for a nonlinear conversion of the information provided by meteorological forecasts to power generation, for instance, modeled with nonparametric nonlinear regression (local polynomial or spline-based). The

model in (4.7) finally includes diurnal Fourier harmonics for the correction of periodic effects that may not be present in the meteorological forecasts, with h_{t+k} the hour of the day at lead time k.

Besides (4.7), a number of alternative approaches were introduced in the past few years for predicting wind power generation up to 2–3 days ahead based on both measurements and meteorological forecasts. Notably, neural networks and other machine learning approaches became popular after the original proposal of Kariniotakis, Stavrakakis and Nogaret (1996) and more recently with the representative work of Sideratos and Hatziargyriou (2007). For all of these models, parameters are commonly estimated with Least Squares (LS) and Maximum Likelihood (ML) approaches (and a Gaussian assumption for the residuals $\varepsilon_{s,t+k}$), potentially made adaptive and recursive so as to allow for smooth changes in the model parameters (accepting some form of nonstationarity), while reducing computational costs. It was recently argued that employing entropy-based criteria for parameter estimation may be beneficial, as in Bessa, Miranda and Gama (2009), since they do not rely on any assumption for the residual distributions. A more extensive review of alternatives statistical approaches to point prediction of wind speed and power can be found in Zhu and Genton (2012).

As an illustration, Figure 5 depicts example point forecasts for the 5 aggregated zones of Western Denmark, issued on 16 March 2007 at 06 UTC based on the methodology described by Nielsen (2002). These have an hourly resolution up to 43 hours ahead, in line with operational decision-making requirements. The well-captured pattern for the first lead times originates from the combination of the trend given by meteorological forecasts with the autoregressive component based on local observational data. For the further lead times, the dynamic wind power generation pattern is mainly driven by the meteorological forecasts, though nonlinearly converted to power and recalibrated to the specific conditions at these various aggregated zones.

In contrast with the introductory part of this section, where it was mentioned that point forecasts corresponded to conditional expectation estimates, Gneiting (2010, and references therein) discussed the more general case of quantiles being optimal point forecasts in a decision-theoretic framework. Indeed, in view of the operational decision-making problems described in Section 3, it is the case that if one accounts for the utility function of the decision-makers at the time of issuing predictions, such forecasts would then become
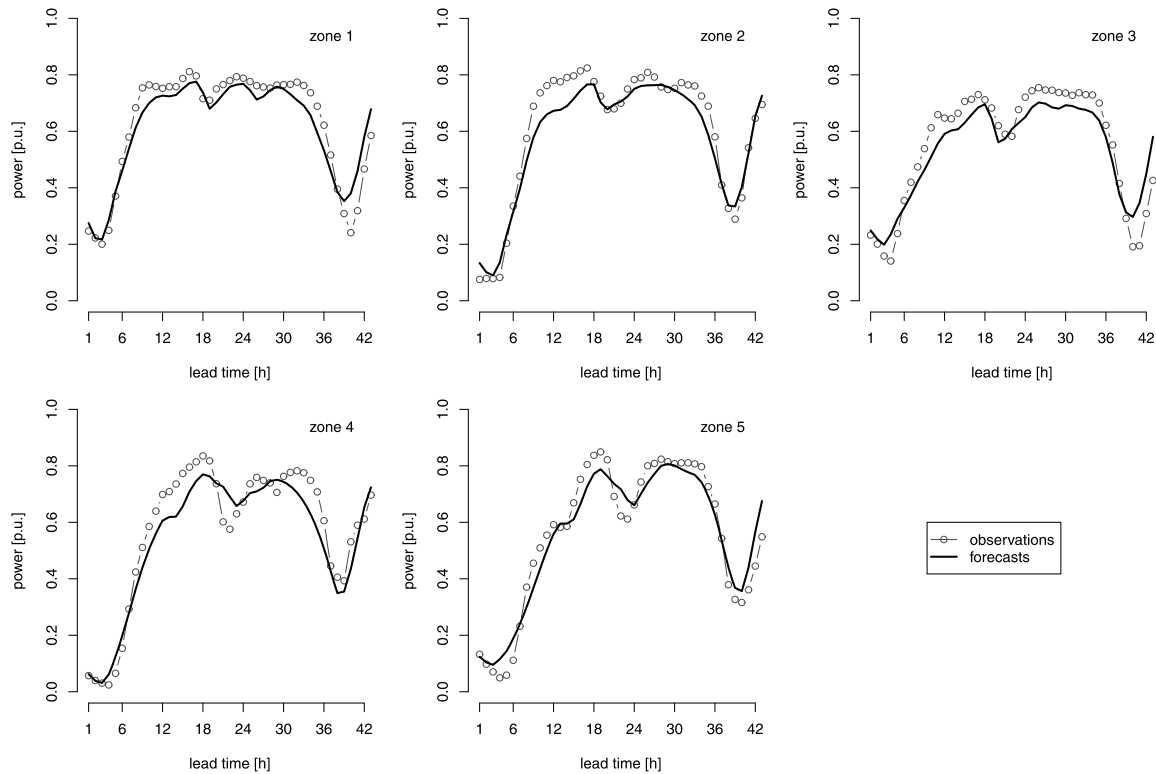
FIG. 5. *Example episode with point forecasts for the* 5 *aggregated zones of Western Denmark, as issued on* 16 *March* 2007 *at* 06 *UTC. Corresponding power measurements, obtained a posteriori, are also shown.*

specific quantiles,

$$\hat{y}_{s,t+k|t} = \hat{F}^{-1}_{s,t+k|t}(\alpha), \quad (4.8)$$

whose nominal level α is determined from the utility function and the structure of the problem itself. The information set and models to be used for issuing quantile forecasts are similar in essence to those for point predictions in the form of conditional expectations. The estimation of model parameters is then performed based on the check function criterion of Koenker and Bassett (1978) or any general scoring rules for quantiles (Gneiting and Raftery, 2007), instead of quadratic and likelihood-based criteria. An example approach to point forecasting of wind power generation where point forecasts actually are quantiles of predictive densities is that of Møller, Nielsen and Madsen (2008), based on time-adaptive quantile regression.

4.2 Predictive Marginal Densities

Point forecasts in the form of conditional expectations are somewhat "just the mean of whatever may happen." These are not optimal inputs to a large class of decision problems. Since the nominal level of quantile forecasts to be used instead may vary in time

while depending upon the problem itself, or might be even unknown, issuing predictive densities certainly is more relevant. Given the random variable $\mathbf{Y_{s,t+k}}$ whose characteristics are to be predicted, these actually are predictive marginal densities $\hat{f}_{s,t+k|t}$ for all locations and lead times involved, individually, with $\hat{F}_{s,t+k|t}$ the corresponding predictive c.d.f.s.

Today such a type of wind power forecast is issued in both parametric and nonparametric frameworks. In the former case, based on an assumption for the shape of predictive marginal densities (for instance, motivated by an empirical investigation), one has

$$\hat{f}_{s,t+k|t} = f(y_{s,t+k}; \hat{\boldsymbol{\theta}}_{s,t+k|t}), \quad (4.9)$$

where f is the density function for power to be generated at location s and time $t + k$, for the chosen probability distribution, for example, truncated/censored Gaussian or Beta. In (4.9), $\hat{\boldsymbol{\theta}}_{s,t+k|t}$ is the predicted value for the vector of parameters fully characterizing that distribution, for instance, a vector of parameters consisting of location and scale parameters for the truncated/censored Gaussian and Beta distributions. For these classes of distributions characterized by such limited sets of parameters only, point forecasts as condi-

tional expectations, complemented by a variance estimator, for example, using exponential smoothing or based on an ARCH/GARCH process, permit to directly obtain location and scale parameters of predictive marginal densities. This reliance on a limited number of parameters may be seen as desirable since it eases subsequent estimation and related computational cost.

Models for the density parameters take a general form similar to that in (4.2) (and subsequent models in Section 4.1), that is, based on linear or nonlinear models with input a subset Ω_t from the information set at time t. Example parametric approaches include the RST approach of Gneiting et al. (2006) for predicting wind speed with truncated Gaussian distributions and that of Pinson (2012) using Generalized Logit–Normal distributions for wind power, also compared with censored Gaussian and Beta assumptions. Similarly, Lau and McSharry (2010) proposed employing Logit–Normal distributions for aggregated wind power generation for the whole Republic of Ireland.

In contrast, nonparametric approaches, since they do not rely on any assumption for the shape of predictive densities, translate to focusing on sets of quantile forecasts defining predictive c.d.f.s. These are conveniently summarized by such sets of quantile forecasts,

$$
\hat{F}_{s,t+k|t} = \{\hat{q}_{s,t+k|t}^{(\alpha_i)}; 0 \le \alpha_1 < \cdots < \alpha_i < \cdots
$$
$$
(4.10) \qquad\qquad\qquad < \alpha_l \le 1\},
$$

with nominal levels α_i spread over the unit interval, though, in practice, $\hat{F}_{s,t+k|t}$ is obtained by interpolation through these sets of quantile forecasts. Actually, nonparametric approaches to quantile forecasts may suffer from a limited number of relevant observations for the very low and high nominal levels α, say, $\alpha, 1 - \alpha < 0.05$, therefore potentially compromising the quality of resulting forecasts. This was observed by Manganelli and Engle (2004) when focusing on risk quantification approaches in finance, and more particularly on dynamic quantile regression models for very low and high levels. Even though the application of interest here is different, the numerical aspects of estimating models for quantiles of wind power generation for very low and high levels are similar. It may therefore be advantageous under certain conditions to define nonparametric predictive densities for their most central part, say, $\alpha, 1 - \alpha > 0.05$, while parametric assumptions could be employed for the tails.

A number of approaches for issuing nonparametric probabilistic forecasts of wind power were proposed and benchmarked over the last decade. In the most standard case, these are obtained from already generated point predictions and, potentially, associated meteorological forecasts. Maybe the most well-documented and widely applied methods are the simple approach of Pinson and Kariniotakis (2010) consisting in dressing the available point forecasts with predictive densities of forecast errors made in similar conditions, as well as the local quantile regression of Bremnes (2004) and the time-adaptive quantile regression of Møller, Nielsen and Madsen (2008), to be used for each of the defining quantile forecasts. The approach of Møller, Nielsen and Madsen (2008) comprises an upgraded version of the original proposal of Nielsen, Madsen and Nielsen (2006), where quantile forecasts of wind power generation are conditional to previously issued point forecasts and to input wind direction forecasts. As for point predictions, neural network and machine learning techniques became increasingly popular over the last few years for generating nonparametric probabilistic predictions based on a set of quantiles (Sideratos and Hatziargyriou, 2012). In contrast to these methods using single-valued forecasts of wind power and meteorological variables as input, a relevant alternative relies on meteorological ensemble predictions, that is, sets of multivariate space–time trajectories for meteorological variables as issued by meteorological institutes [see Leutbecher and Palmer (2008) and the references therein], which are then transformed to the wind power space. Ensemble forecasts attempt at dynamically representing uncertainties in meteorological forecasts (as well as spatial, temporal and inter-variable dependencies), by jointly accounting for misestimation in the initial state of the Atmosphere and for parameter uncertainty in the model dynamics. To obtain probabilistic forecasts of wind power generation from such meteorological ensembles, conventional approaches combine nonlinear regression and kernel dressing of the ensemble trajectories, as in the alternative proposals of Roulston et al. (2003); Taylor, McSharry and Buizza (1999) and Pinson and Madsen (2009). In a similar vain, a general method for the conversion of probabilistic forecasts of wind speed to power based on stochastic power curves, thus accounting for additional uncertainties in the wind-to-power conversion process in a Bayesian framework, was recently described by Jeon and Taylor (2012).

Example nonparametric forecasts are shown in Figure 6 for the same period as in Figure 5, as obtained by applying the method of Pinson and Kariniotakis (2010) to the already issued point predictions and their input
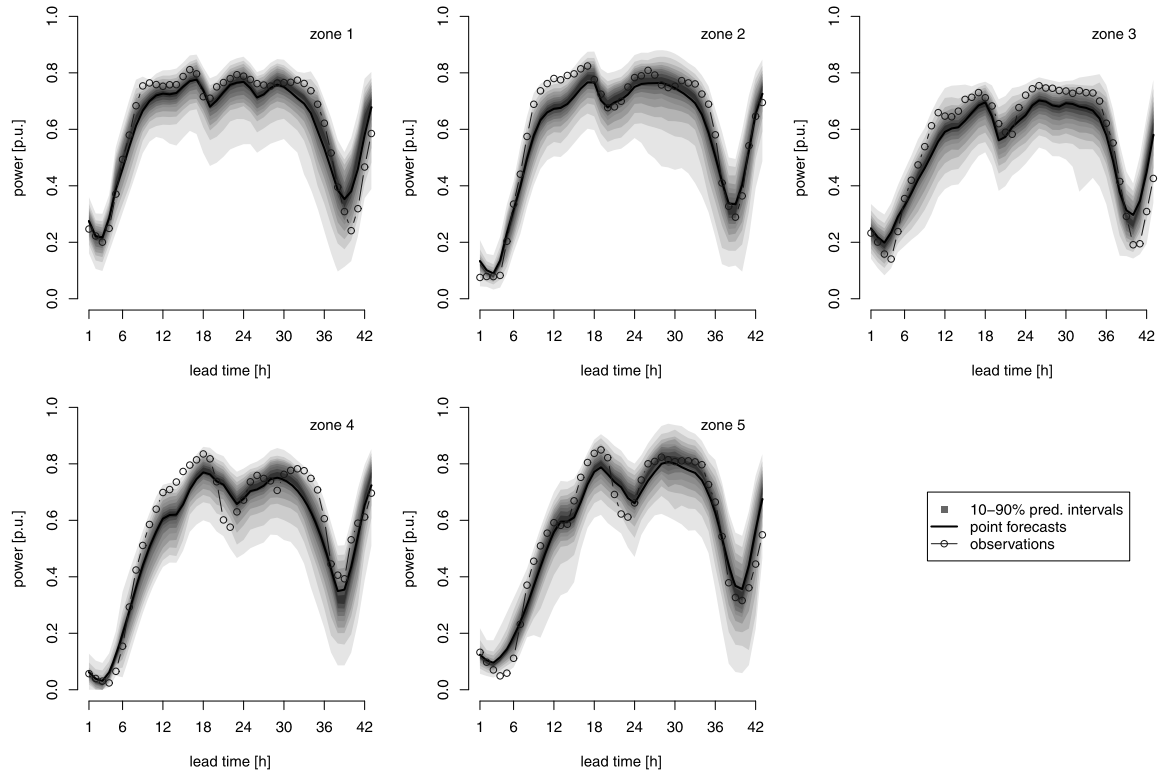
FIG. 6. *Example episode with probabilistic forecasts for the* 5 *aggregated zones of Western Denmark* (*and corresponding measurements obtained a posteriori*), *as issued on* 16 *March* 2007 *at* 06 *UTC. They take the form of so-called river-of-blood fan charts* [*termed coined after Wallis* (2003)], *represented by a set of central prediction intervals with increasing nominal coverage rates* (*from* 10% *to* 90%).

meteorological forecast information. The characteristics of these predictive densities smoothly evolve as a function of a number of factors, for example, lead time, geographical location, time of the year and level of power generation (since nonlinear and bounded power curves shape forecast uncertainty). By construction, and through adaptive estimation, these predictive densities are probabilistically calibrated, meaning that observed levels for the defining quantile forecasts correspond to the nominal ones. This is a crucial property of probabilistic forecasts to be used as input to decision problems such as those of Section 3, since a probabilistic bias in the forecasts would yield suboptimality of resulting operational decisions. Actually, in addition, probabilistic calibration is also a prerequisite for application of the methods described in the following in order to generate trajectories.

4.3 Spatio-Temporal Trajectories

Both point forecasts and predictive densities are suboptimal inputs to decision-making when spatial and temporal dependencies are involved. It is then required to fully describe the density of the spatio-temporal

process $\mathbf{Y}_{\mathbf{s},t+\mathbf{k}}$. Following a proposal by Pinson et al. (2009) for wind power and, more recently, by Möller, Lenkoski and Thorarinsdottir (2013) for multiple meteorological variables, the probabilistic forecast $\hat{F}_{\mathbf{s},t+\mathbf{k}|t}$ can be fully characterized under a Gaussian copula by the predictive marginal c.d.f.s $\hat{F}_{s,t+k|t}$, $\forall s, k$, and by a space–time covariance matrix $\hat{\mathbf{C}}_t$ linking all locations and lead times. In that case, using notation similar to that of Möller, Lenkoski and Thorarinsdottir (2013),

$$
(4.11) \quad \begin{aligned}
&\hat{F}_{\mathbf{s},t+\mathbf{k}|t}(\mathbf{y}_{\mathbf{s},t+\mathbf{k}}|\hat{\mathbf{C}}_t) \\
&= \Phi_{mn}\big(\big\{\Phi^{-1}\big(\hat{F}_{s,t+k|t}(y)\big)\big\}_{s,k}|\hat{\mathbf{C}}_t\big),
\end{aligned}
$$

where $\mathbf{y}_{\mathbf{s},t+\mathbf{k}}$ was defined in (2.3) and Φ is the c.d.f. of a standard Gaussian variable, while Φ_{mn} is that for a multivariate Gaussian of dimension $m \times n$. Going beyond the Gaussian copula simplification, one could envisage employing more refined copulas, though at the expense of additional complexity. The interested reader may find an extensive introduction to copulas in Nelsen (2006).

This type of construction of multivariate probabilistic forecasts for wind power generation in space and

in time has clear advantages. Indeed, given that all predictive densities forming the marginal densities are calibrated, it may be assumed that one deals with a latent space–time Gaussian process consisting of successive multivariate random variables $\mathbf{Z}_t$ (each of dimension $m \times n$) with realizations given by

$$(4.12) \quad \begin{aligned} \mathbf{z}_t = \{ &\Phi^{-1}(\hat{F}_{s,t+k|t}(y_{s,t+k})); \\ &s = s_1, s_2, \ldots, s_m, k = 1, 2, \ldots, n \}. \end{aligned}$$

Consequently, this latent Gaussian process can be used for identifying and estimating a suitable parametric space–time structure or, alternatively, if $m \times n$ is low and the sample size large, for the tracking of the non-parametric (sample) covariance structure, for instance, using exponential smoothing.

Similarly, one of the advantages of this construction of multivariate probabilistic forecasts based on a Gaussian copula is that it is fairly straightforward to issue space–time trajectories. Remember that these are the prime input to a large class of stochastic optimization approaches, such as the advanced version of the problems presented in Sections 3.1 and 3.2, where representation of space–time interdependencies is required. Such trajectories also are a convenient way to visualize the complex information conveyed by these multivariate probabilistic forecasts, as hinted by Jordà and Marcellino (2010), among others. Let us define by

$$(4.13) \quad \begin{aligned} \hat{\mathbf{y}}_{\mathbf{s},t+\mathbf{k}|t}^{(j)} = \{ &\hat{y}_{s,t+k|t}^{(j)}; s = s_1, s_2, \ldots, s_m, \\ &k = 1, 2, \ldots, n \}, \\ &j = 1, 2, \ldots, J, \end{aligned}$$

a set of J space–time trajectories issued at time t. As an illustrative example, Figure 7 gathers a set of $J = 12$ space–time trajectories of wind power generation for the same episode as in Figures 5 and 6. The covariance structure $\hat{\mathbf{C}}_t$ used to fully specify the space–time interdependence structure is obtained by exponential smoothing of the sample covariance of the latent Gaussian process. The trajectories are then obtained by first randomly sampling from a multivariate Gaussian variable with the most up-to-date estimate of the space–time covariance structure. Denote by $\mathbf{z}_t^{(j)}$ the jth sample, whose components $z_{s,t+k}^{(j)}$ will directly relate to a location s and a lead time k in the following. These multivariate Gaussian samples are converted to wind power generation by a transformation which is the inverse of that in (4.12). This yields

$$(4.14) \quad \hat{y}_{s,t+k}^{(j)} = \hat{F}_{s,t+k|t}^{-1}(\Phi(z_{s,t+k}^{(j)})) \quad \forall s, k, j.$$

This type of visualization allows to appraise the temporal correlation in wind power generation and potential forecast errors through predictive densities, giving an extra level of information if compared to the probabilistic forecasts of Figure 6. There are obvious limitations stemming from the dimensionality of the random variable of interest. For instance, here, the spatial interdependence structure, though serving to link these trajectories, is nearly impossible to appreciate.

5. DISCUSSION: UPCOMING CHALLENGES

Three decades of research in modeling and forecasting of power generation from the wind have led to a solid understanding of the whole chain from taking advantage of available meteorological and power measurements, as well as meteorological forecasts, all the way to using forecasts as input to decision-making. Today, methodologies are further developed in a probabilistic framework, even though forecast users may still prefer to be provided with single-valued predictions. Some important challenges are currently under investigation or identified as particularly relevant for the short to medium term. These are presented and discussed below, with emphasis placed on new and better forecasts, and forecast verification, as well as bridging the gap between forecast quality and value.

5.1 Improved Wind Power Forecasts: Extracting More out of the Data

Improving the quality of wind power forecasts is a constant challenge, with strong expectations linked to the increased commitment of the meteorological community to issue better forecasts of relevant weather variables, mainly surface wind components. This will come, among other things, from a better description of the physical phenomena involved, especially in the boundary layer, as well as from an increased resolution of the numerical schemes used to solve the systems of partial differential equations.

Meanwhile, for statisticians, there are paths toward forecast improvement that involve a better utilization of available measurement data, combining measurements available on site and additional observations spatially distributed around that site. Wind forecasts used to issue power forecasts over a region seldom capture fully the spatio-temporal dynamics of power generation owing to, for example, a too coarse resolution (spatial and temporal) and timing errors with respect to passages of weather fronts. However, all distributed meteorological stations and wind turbines may serve as
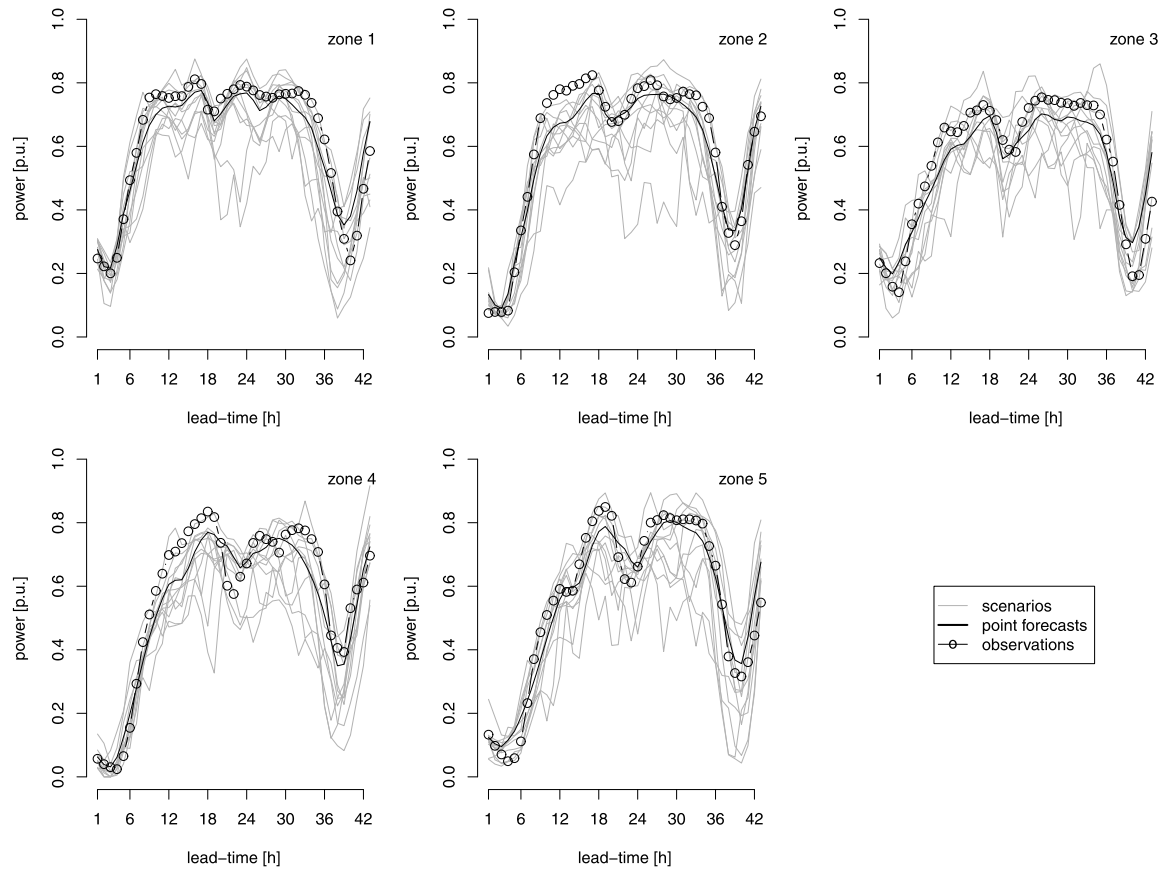
FIG. 7.   *Set of* 12 *space–time trajectories of wind power generation for the* 5 *aggregated zones of Western Denmark*, *issued on the* 16 *March* 2007 *at* 06 *UTC*.

sensors in order to palliate for these deficiencies. For the example of the Western Denmark data set, Girard and Allard (2013) explored the spatio-temporal characteristics of residuals after capturing local dynamics at all individual sites, hinting at the role of prevailing weather conditions on the space–time structure. For the same data set, Lau (2011) investigated an anisotropic space–time covariance model based on a Lagrangian approach, conditional to prevailing wind direction over the region. Based on such analysis, it is required to propose nonlinear and nonstationary spatio-temporal models for wind power generation, for instance, using covariance structures conditional to prevailing weather conditions, in the spirit of Huang and Hsu (2004). An advantage will be that, instead of having to identify and estimate models for every single site of interest (more than 400 for the Western Denmark data set), and at various spatial and temporal resolutions of relevance to forecast users, a single model would fit all purposes at once. Even though more complex and potentially more costly in terms of parameter estimation, they could lead

to a substantial overall reduction in computational time and expert knowledge necessary to set up and maintain all individual models. Alternatively, approaches relying on stochastic partial differential equations ought to be considered owing to appealing features and recent advances in their linkage to spatio-temporal covariance structures, as well as improved computational solving (Lindgren, Rue and Lindström, 2011). Challenges there, however, relate to the complexity of the stochastic processes involved, requiring to account for the state-dependent diffusion part, and also for changes in the very dynamics of wind components, as induced by a number of weather phenomena. It is not clear how all these aspects could be accounted for in a compact set of stochastic differential equations, which could be solved with existing numerical integration schemes.

The increasing availability of high-dimensional data sets, with a large number of relevant meteorological and power systems variables, possibly at high spatial and temporal resolutions, gives rise to a number of challenges and opportunities related to data aggregation. These challenges have already been identified

in other fields, for example, econometrics, where aggregation has shown its interest and potential limitations. A relevant example work is that of Giacomini and Granger (2004), which looks at the problem of pooling and forecasting spatially correlated data sets. On the one hand, considering different levels of aggregation for the wind power forecasting problem can permit to ease the modeling task, by identifying groups of turbines with similar dynamic behavior which could be modeled jointly. On the other hand, this would lower the computational burden by reducing model size and complexity. Proposals related to aggregation should, however, fully consider the meteorological aspects at different temporal and spatial scales, which may dynamically condition how aggregate models would be representative of geographically distributed wind farms. One could build on the classical Space–Time Auto-Regressive (STAR) model of Cliff and Ord (1984) by enhancing it to having dynamic and conditional space–time covariance structures. In a similar vain, dynamic models for spatio-temporal data such as those introduced by Stroud, Müller and Sansó (2001) and follow-up papers are appealing, since they provide an alternative approach to data aggregation by seeing the overall spatial processes as a linear combination of a limited number of local (polynomial) spatial processes in the neighborhood of appropriately chosen locations. Overall, various relevant directions to space–time modeling could be explored, based on the substantial literature existing for other processes and in other fields.

5.2 New Forecast Methodologies and Forecast Products

As a result of these efforts, new types of forecasts will be available to decision-makers in the form of continuous surfaces and trajectories, from which predictions with any spatial and temporal resolution could be dynamically extracted. Similar to the development of meteorological forecasting, the need for larger computational facilities might call for centralizing efforts in generating and issuing wind power predictions. Actually, in the opposite direction, a share of practitioners request predictions of lower complexity that could be better appraised by a broader audience and more easily integrated into existing operational processes. For instance, since accommodating the variability of power fluctuations with successive periods of fast-increasing and fast-decreasing power generation is seen as an issue by some system operators in the US and in Australia, methodologies were proposed for the prediction of so-called ramp events, where the definition

of these "ramp events" is based on the very need of the decision-maker (Bossavy, Girard and Kariniotakis, 2013; Gallego et al., 2013).

Besides, even though alternative parametric assumptions for predictive marginal densities have been analyzed and benchmarked, for example, Beta (Bludszuweit, Domínguez-Navarro and Llombart, 2008), truncated and censored Gaussian, and Generalized Logit–Normal (Pinson, 2012), there is no clear superiority of one over the others, for all potential lead times, level of aggregation and wind dynamics themselves. This certainly originates from the nonlinear and bounded curves representing the conversion of wind to power, known to shape predictive densities. Such curves may additionally be time-varying, uncertain and conditional on various external factors. This is why future work should consider these curves as stochastic power curves, also described by multivariate distributions, as a generalization of the proposal of Jeon and Taylor (2012). Their impact on the shape of predictive densities ought to be better understood. Then, combined with probabilistic forecasts of relevant explanatory variables, for instance, from recalibrated meteorological ensembles, stochastic power curves would naturally yield probabilistic predictions of wind power generation, in a Bayesian framework. This is since stochastic power curves comprise models of the joint density of meteorological variables and of corresponding wind power generation. Predictive densities of wind power generation would then be obtained by applying Bayes rule, that is, by passing probabilistic forecasts of meteorological variables through such stochastic power curve models.

To broaden up, and since operational decision-making problems are based on interdependent variables (power generation from different renewable energy sources, electric load and potentially market variables), multivariate probabilistic forecasts for relevant pairing, or for all of them together, should be issued in the future, with the weather as the common driver. Similar to the proposal of Möller, Lenkoski and Thorarinsdottir (2013) for multivariate probabilistic forecasts of meteorological variables, one could generalize the space–time trajectories of Section 4.3 to a multivariate setup. Alternatives should be thought of, allowing to directly obtain such spatio-temporal and multivariate predictions, instead of having to go through predictive marginal densities first.

5.3 Verifying Probabilistic Forecasts of Ever-Increasing Dimensionality

Forecast verification is a subtle exercise already for the most simple case of dealing with point forecasts, to be based on the joint distribution of forecasts and observations (Murphy and Winkler, 1987). Focus is to-day on verifying forecasts in a probabilistic framework, for instance, following the paradigm of Gneiting, Balabdaoui and Raftery (2007) originally introduced for the univariate case, based on calibration and sharpness of predictive marginal densities. The nonlinear and double-bounded nature of the wind power stochastic process (possibly also a discrete-continuous mixture) renders the verification of probabilistic forecasts more complex, especially for their calibration. It generally calls for an extensive reliability assessment conditional on variables known to impact the shape of predictive densities: level of power, wind direction, etc. In addition, the benchmarking and comparison of forecasting methods ought to account for sample size and correlation issues, since evaluation sets often are of limited size (though of increasing length now that some wind farms have been operating for a long time), while correlation in forecast errors and other criteria (skill score values, probability integral transform) is necessarily present for forecasts with lead times further than one step ahead. Verifying high-dimensional forecasts, like space–time trajectories in the most extreme case, based on small samples will necessarily yield score values that may not fully reflect actual forecast quality even though the score used is proper. Indeed, the deviations from the expected score value, which could be estimated better with larger samples, would be substantial. Correlation issues may only magnify this problem, since they somewhat reduce the effective sample size for estimation. An illustration of the combined effects of sampling and correlation on the verification of probabilistic forecasts can be found in Pinson, McSharry and Madsen (2010).

Going from univariate to multivariate aspects, Gneiting et al. (2008) explained how the previously introduced paradigm can be readily generalized for multivariate probabilistic predictions, yielding an evaluation framework including skill scores and diagnostic tools. An application to the verification of temporal trajectories of wind power generation in Pinson and Girard (2012) illustrated its potential limitations stemming from the high-dimensionality (there, $n = 43$ lead times) of the underlying random variables. Following the discussion in Section 5.1, it is clear that new views

on forecast verification ought to be introduced and evaluated as dimensionality increases. For instance, recent work by Hering and Genton (2011) showed how to compare spatial predictions in a framework inspired by the Diebold–Mariano test and with limited assumptions on the spatial processes themselves, thus permitting to deal with high-dimensional predictions by focusing on their spatial structure.

5.4 Bridging the Gap Between Forecast Quality and Value

Murphy (1993) introduced 3 types of goodness for weather forecasts, also valid and relevant for other types of predictions like for wind power. Out of these 3, quality and value play a particular role: (i) quality relates to the objective assessment of how well forecasts describe the stochastic process of interest (and its realizations), regardless of how the forecasts may be used subsequently, while (ii) value corresponds to the economic/operational gain from considering forecasts at the decision-making stage. Through the introduction of representative operational decision problems in Section 3, it was shown that optimal forecasts as input to decision-making in a stochastic optimization framework take the form of quantiles, predictive marginal densities or, finally, trajectories describing the full spatio-temporal process. However, it is not clear today how improving the quality of these forecasts, for instance, in terms of reduced skill score values or increased probabilistic calibration, may lead to added value for the decision-makers, especially when they might use these forecasts sub-optimally. In practice, this will call for more analytic work in a decision-theoretic framework, by better linking skill scores of the forecasters and utility of the decision-makers, as well as for a number of simulation studies in order to simulate the usage of forecasts of varying quality as input to a wide range of relevant operational problems. Full benefits from integrating wind power generation into existing power systems and through electricity markets will only be obtained by optimally integrating forecasts in decision-making.

ACKNOWLEDGMENTS

REFERENCES

ACKERMANN, T. (2012). *Wind Power in Power Systems*, 2nd ed. Wiley, Chichester.

BESSA, R. J., MIRANDA, V. and GAMA, J. (2009). Entropy and correntropy against minimum square error in offline and online three-day ahead wind power forecasting. *IEEE T. Power Syst.* **24** 1657–1666.

BILLINTON, R. and ALLAN, R. N. (1984). *Reliability Evaluation of Power Systems*. Plenum Press, London.

BLUDSZUWEIT, H., DOMÍNGUEZ-NAVARRO, J. A. and LLOM-BART, A. (2008). Statistical analysis of wind power forecast er-rors. *IEEE T. Power Syst.* **23** 983–991.

BOSSAVY, A., GIRARD, R. and KARINIOTAKIS, G. (2013). Forecasting ramps of wind power production with numerical weather prediction ensembles. *Wind Energy* **16** 51–63.

BOTTERUD, A., WANG, J., MIRANDA, V. and BESSA, R. J. (2010). Wind power forecasting in US electricity markets. *The Electricity Journal* **23** 71–82.

BREMNES, J. B. (2004). Probabilistic wind power forecasts using local quantile regression. *Wind Energy* **7** 47–54.

BROWN, B. G., KATZ, R. W. and MURPHY, A. M. (1984). Time series models to simulate and forecast wind speed and wind power. *J. Appl. Meteor.* **23** 1184–1195.

CLIFF, A. D. and ORD, J. K. (1984). Space–time modeling with applications to regional forecasting. *Trans. Inst. Br. Geogr.* **66** 119–128.

CONEJO, A., CARRION, M. and MORALES, J. M. (2010). *Decision-Making under Uncertainty in Electricity Markets*. Springer, New York.

CONRADSEN, K., NIELSEN, L. B. and PRAHM, L. P. (1984). Re-view of Weibull statistics for estimation of wind speed distribu-tions. *J. Appl. Meteor.* **23** 1173–1183.

DOHERTY, R. and O'MALLEY, M. (2005). A new approach to quantify reserve demand in systems with significant installed wind capacity. *IEEE T. Power Syst.* **20** 587–595.

GALLEGO, C., COSTA, A., CUERVA, A., LANDBERG, L., GREAVES, B. and COLLINS, J. (2013). A wavelet-based ap-proach for large wind power ramp characterisation. *Wind En-ergy* **16** 257–278.

GIACOMINI, R. and GRANGER, C. W. J. (2004). Aggregation of space–time processes. *J. Econometrics* **118** 7–26. MR2030964

GIEBEL, G., BROWNSWORD, R., KARINIOTAKIS, G., DEN-HARD, M. and DRAXL, C. (2011). The state-of-the-art in short-term prediction of wind power: A literature overview, 2nd ed. Technical report. Available at orbit.dtu.dk.

GIRARD, R. and ALLARD, D. (2013). Spatio-temporal propaga-tion of wind power forecast errors. *Wind Energy*. To appear.

GNEITING, T. (2010). Quantiles as optimal point forecasts. *Int. J. Forecasting* **27** 197–207.

GNEITING, T., BALABDAOUI, F. and RAFTERY, A. E. (2007). Probabilistic forecasts, calibration and sharpness. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **69** 243–268. MR2325275

GNEITING, T. and RAFTERY, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *J. Amer. Statist. Assoc.* **102** 359–378. MR2345548

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 space–time method. *J. Amer. Statist. Assoc.* **101** 968–979. MR2324108

GNEITING, T., STANBERRY, L. I., GRIMIT, E. P., HELD, L. and JOHNSON, N. A. (2008). Assessing probabilistic forecasts of multivariate quantities, with an application to ensemble predic-tions of surface winds. *TEST* **17** 211–235. MR2434318

HAHN, H., MEYER-NIEBERG, S. and PICKL, S. (2009). Electric load forecasting methods: Tools for decision making. *European J. Oper. Res.* **199** 902–907.

HAMILTON, J. D. (1989). A new approach to the economic analy-sis of nonstationary time series and the business cycle. *Econo-metrica* **57** 357–384. MR0996941

HASLETT, J. and RAFTERY, A. E. (1989). Space–time modelling with long-memory dependence: Assessing Ireland's wind power resource (with discussion). *J. R. Stat. Soc. Ser. C Appl. Stat.* **38** 1–50.

HERING, A. S. and GENTON, M. G. (2010). Powering up with space–time wind forecasting. *J. Amer. Statist. Assoc.* **105** 92–104. MR2757195

HERING, A. S. and GENTON, M. G. (2011). Comparing spatial predictions. *Technometrics* **53** 414–425. MR2850473

HUANG, H. C. and HSU, N. J. (2004). Modeling transport ef-fects on ground-level ozone using a non-stationary space–time model. *Environmetrics* **15** 251–268.

JEON, J. and TAYLOR, J. W. (2012). Using conditional kernel density estimation for wind power density forecasting. *J. Amer. Statist. Assoc.* **107** 66–79. MR2949342

JORDÀ, Ò. and MARCELLINO, M. (2010). Path forecast evalua-tion. *J. Appl. Econometrics* **25** 635–662. MR2758077

JOSKOW, P. L. (2011). Comparing the costs of intermittent and dis-patchable electricity generating technologies. *Amer. Econ. Rev.* **100** 238–241.

KARINIOTAKIS, G. N., STAVRAKAKIS, G. S. and NOGA-RET, E. F. (1996). Wind power forecasting using advanced neu-ral networks models. *IEEE T. Energy Conver.* **11** 762–767.

KOENKER, R. and BASSETT, G. JR. (1978). Regression quantiles. *Econometrica* **46** 33–50. MR0474644

LANDBERG, L. (1999). Short-term prediction of the power produc-tion from wind farms. *J. Wind Eng. Ind. Aerodyn.* **80** 207–220.

LANDBERG, L. and WATSON, S. J. (1994). Short-term prediction of local wind conditions. *Bound.-Layer Meteor.* **70** 171–195.

LANGE, M. (2005). On the uncertainty of wind power predictions—Analysis of the forecast accuracy and statistical distribution of errors. *J. Solar Energ.-T. ASME* **127** 177–184.

LANGE, M. and FOCKEN, U. (2006). *Physical Approach to Short-Term Wind Power Prediction*. Springer, Berlin.

LAU, A. (2011). Probabilistic wind power forecasts: From aggregated approach to spatiotemporal models. Ph.D. thesis, Univ. Oxford, Oxford, UK.

LAU, A. and MCSHARRY, P. (2010). Approaches for multi-step density forecasts with application to aggregated wind power. *Ann. Appl. Stat.* **4** 1311–1341. MR2758330

LETCHER, T. M. (2008). *Future Energy: Improved, Sustainable and Clean Options for Our Planet*. Elsevier, Amsterdam.

LEUTBECHER, M. and PALMER, T. N. (2008). Ensemble forecasting. *J. Comput. Phys.* **227** 3515–3539. MR2400226

LINDGREN, F., RUE, H. and LINDSTRÖM, J. (2011). An explicit link between Gaussian fields and Gaussian Markov random fields: The stochastic partial differential equation approach. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **73** 423–498. MR2853727

MANGANELLI, S. and ENGLE, R. F. (2004). A comparison of value-at-risk models in finance. In *Risk Measures for the 21st Century* (G. Szegö, ed.) 123–143. Wiley, Chichester.

MATOS, M. A. and BESSA, R. J. (2010). Setting the operating reserve using probabilistic wind power forecasts. *IEEE T. Power Syst.* **26** 594–603.

MÖLLER, A., LENKOSKI, A. and THORARINSDOTTIR, T. L. (2013). Multivariate probabilistic forecasting using ensemble Bayesian model averaging and copulas. *Q. J. Royal Met. Soc.* **139** 982–991.

MØLLER, J. K., NIELSEN, H. A. and MADSEN, H. (2008). Time-adaptive quantile regression. *Comput. Statist. Data Anal.* **52** 1292–1303. MR2418566

MORTHORST, P. E. (2003). Wind power and the conditions at a liberalized power market. *Wind Energy* **6** 297–308.

MUR AMADA, J. and BAYOD RÚJULA, A. (2010). Variability of wind and wind power. In *Wind Power* (S. M. Muyeen Vukovar, ed.) 289–320. Intech Open, Rijeka, Croatia.

MURPHY, A. H. (1993). What is a good forecast? An essay on the nature of goodness in weather forecasting. *Wea. Forecasting* **8** 281–293.

MURPHY, A. H. and WINKLER, R. L. (1987). A general framework for forecast verification. *Mon. Wea. Rev.* **115** 1330–1338.

NELSEN, R. B. (2006). *An Introduction to Copulas*, 2nd ed. Springer, New York. MR2197664

NIELSEN, T. S. (2002). Online prediction and control in nonlinear stochastic systems. Ph.D. thesis, Dept. Informatics and Mathematical Modelling, Technical Univ. Denmark.

NIELSEN, H. A., MADSEN, H. and NIELSEN, T. S. (2006). Using quantile regression to extend an existing wind power forecasting system with probabilistic forecasts. *Wind Energy* **9** 95–108.

ORTEGA-VAZQUEZ, M. A. and KIRSCHEN, D. S. (2009). Estimating the spinning reserve requirements in systems with significant wind power generation penetration. *IEEE T. Power Syst.* **24** 114–124.

PALMER, T. N. (2012). Towards the probabilistic Earth-system simulator: A vision for the future of climate and weather prediction. *Quart. J. Royal Met. Soc.* **665** 841–861.

PAPAVASILIOU, A. and OREN, S. S. (2013). Multiarea stochastic unit commitment for high wind penetration in a transmission constrained network. *Oper. Res.* **61** 578–592. MR3079732

PEÑA DIAZ, A., GRYNING, S. E. and MANN, J. (2010). On the length-scale of the wind profile. *Quart. J. Royal Met. Soc.* **136** 2119–2131.

PINSON, P. (2012). Very-short-term probabilistic forecasting of wind power with generalized logit–normal distributions. *J. R. Stat. Soc. Ser. C Appl. Stat.* **61** 555–576. MR2960738

PINSON, P., CHEVALLIER, C. and KARINIOTAKIS, G. (2007). Trading wind generation from short-term probabilistic forecasts of wind power. *IEEE T. Power Syst.* **22** 1148–1156.

PINSON, P. and GIRARD, R. (2012). Evaluating the quality of scenarios of short-term wind power generation. *Appl. Energ.* **96** 12–20.

PINSON, P. and KARINIOTAKIS, G. (2010). Conditional prediction intervals of wind power generation. *IEEE T. Power Syst.* **25** 1845–1856.

PINSON, P. and MADSEN, H. (2009). Ensemble-based probabilistic forecasting at Horns Rev. *Wind Energy* **12** 137–155.

PINSON, P. and MADSEN, H. (2012). Adaptive modelling and forecasting of offshore wind power fluctuations with Markov-switching autoregressive models. *J. Forecast.* **31** 281–313. MR2924797

PINSON, P., MCSHARRY, P. E. and MADSEN, H. (2010). Reliability diagrams for nonparametric density forecasts of continuous variables: Accounting for serial correlation. *Quart. J. Royal Met. Soc.* **136** 77–90.

PINSON, P., NIELSEN, H. A., MADSEN, H., PAPAEFTHYMIOU, G. and KLÖCKL, B. (2009). From probabilistic forecasts to statistical scenarios of short-term wind power production. *Wind Energy* **12** 51–62.

RAIFFA, H. and SCHAIFER, R. (1961). *Applied Statistical Decision Theory*. Division of Research, Harvard Business School, Boston.

REIKARD, G. (2008). Using temperature and state transitions to forecast wind speed. *Wind Energy* **11** 431–443.

ROULSTON, M. S., KAPLAN, D. T., HARDENBERG, J. and SMITH, L. A. (2003). Using medium-range weather forecasts to improve the value of wind energy production. *Renew. Energ.* **28** 585–602.

SIDERATOS, G. and HATZIARGYRIOU, N. D. (2007). An advanced statistical method for wind power forecasting. *IEEE T. Power Syst.* **22** 258–265.

SIDERATOS, G. and HATZIARGYRIOU, N. D. (2012). Probabilistic wind power forecasting using radial basis function neural networks. *IEEE T. Power Syst.* **27** 1788–1796.

SKYTTE, K. (1999). The regulating power market on the Nordic power exchange Nord Pool: An econometric analysis. *Energ. Econ.* **21** 295–308.

STROUD, J. R., MÜLLER, P. and SANSÓ, B. (2001). Dynamic models for spatiotemporal data. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **63** 673–689. MR1872059

TAYLOR, J. W., MCSHARRY, P. E. and BUIZZA, R. (1999). Wind power density forecasting using ensemble predictions and time series models. *IEEE T. Energy Conver.* **24** 775–782.

TOL, R. S. J. (1997). Autoregressive conditional heteroscedasticity in daily wind speed measurements. *Theo. Appl. Clim.* **56** 113–122.

TONG, H. (2011). Threshold models in time series analysis—30 years on. *Stat. Interface* **4** 107–118. MR2812802

VAUTARD, R., CATTIAUX, J., YIOU, P., THÉPAUT, J. N. and CIAIS, P. (2010). Northern Hemisphere atmospheric stilling

partly attributed to an increase in surface roughness. *Nat. Geosci.* **3** 756–761.

VINCENT, C. L., GIEBEL, G., PINSON, P. and MADSEN, H. (2010). Resolving nonstationary spectral information in wind speed time series using the Hilbert–Huang transform. *J. Appl. Meteor. Clim.* **49** 253–267.

WALLIS, K. F. (2003). Chi-squared tests of interval and density forecasts, and the bank of England's fan charts. *Int. J. Forecasting* **19** 165–175.

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

WOLAK, F. A. (2013). Regulating competition in wholesale electricity supply. In *Economic Regulation and Its Reform: What Have We Learned*? (N. L. Rose, ed.). Univ. Chicago Press. To appear.

WWEA (WORLD WIND ENERGY ASSOCIATION) (2012). 2012—Half-year report. Technical report. Available at http://www.wwindea.org/webimages/Half-year_report_2012.pdf.

ZHU, X. and GENTON, M. G. (2012). Short-term wind speed forecasting for power system operations. *Int. Stat. Rev.* **80** 2–23. MR2990340

ZUGNO, M., JÓNSSON, T. and PINSON, P. (2013). Trading wind energy based on probabilistic forecasts of both wind generation and market quantities. *Wind Energy*. **16** 909–926.

ZUGNO, M., MORALES, J. M., PINSON, P. and MADSEN, H. (2013). Pool strategy of a price-maker wind power producer. *IEEE T. Power Syst.* **28** 3440–3450.