Short-Term Forecasting Across a Network for the Autonomous Wind Farm

Preprint

Jennifer Annoni, Christopher Bay, Kathryn Johnson, and Paul Fleming

National Renewable Energy Laboratory

Presented at IEEE American Control Conference
Philadelphia, Pennsylvania
July 10–12, 2019

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Contract No. DE-AC36-08GO28308

Conference Paper
NREL/CP-5000-73396
July 2019

Short-Term Forecasting Across a Network for the Autonomous Wind Farm

Preprint

Jennifer Annoni, Christopher Bay, Kathryn Johnson, and Paul Fleming

National Renewable Energy Laboratory

Suggested Citation

Annoni, Jennifer, Christopher Bay, Kathryn Johnson, and Paul Fleming. 2019. *Short-Term Forecasting Across a Network for the Autonomous Wind Farm: Preprint*. Golden, CO: National Renewable Energy Laboratory. NREL/CP-5000-73396. https://www.nrel.gov/docs/fy19osti/73396.pdf.

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Contract No. DE-AC36-08GO28308

Conference Paper
NREL/CP-5000-73396
July 2019

National Renewable Energy Laboratory
15013 Denver West Parkway
Golden, CO 80401
303-275-3000 • www.nrel.gov

NOTICE

This work was authored by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Funding provided by the U.S. Department of Energy Office of Energy Efficiency and Renewable Energy Wind Energy Technologies Office. The views expressed herein do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

Cover Photos by Dennis Schroeder: (clockwise, left to right) NREL 51934, NREL 45897, NREL 42160, NREL 45891, NREL 48097, NREL 46526.

Short-Term Forecasting Across a Network for the Autonomous Wind Farm

Jennifer Annoni, Christopher Bay, Kathryn Johnson, Paul Fleming

Abstract— **In an autonomous wind farm, turbines will use information from nearby turbines to achieve wind farm-level objectives such as optimizing the overall performance of a wind farm, ensuring resiliency when other sensors fail, and adapting to changing local conditions. In this paper, the wind farm can be modeled as a network within which turbines (nodes) share information across designated communication channels, with a focus on turbines at the outside of the wind farm capturing local effects and sharing that information with downstream turbines. Understanding of varied inflow conditions can be especially important in complex terrain. This information can be used to monitor turbines, self-organize turbines into groups, and predict the power performance of a wind farm. In particular, this paper describes an autonomous wind farm that incorporates information from local sensors in real time to predict wind speed and wind direction at each turbine over a short-term horizon. Results indicate that the estimate of wind direction can be used to improve the knowledge of the wind speed and direction over the persistence method on a 10-15-minute time horizon. These short-term forecasts can also be used to facilitate advanced control methods such as feed-forward control within a wind farm.**

I. Introduction

As wind energy continues to provide more and more of the electricity on the grid, new control methods and techniques for wind plants are required. The autonomous energy grid of the future will require autonomous energy plants that can self-optimize and provide accurate forecasts of their available power, as well as other ancillary services. Control systems can enable autonomous wind plants that can self-organize turbines into groups, monitor turbine status, and control turbine performance to maximize profit and reliability of large-scale wind plants [1].

In pursuit of real-time solutions to enable this autonomous wind farm, a limited number of distributed optimization methods for wind plants have been proposed [2]–[4]. Soleimanzadeh et al. [3] presented a linearized flow field model and a distributed H_2 optimal controller to minimize turbine structural loads and provide the requested wind farm power. However, constraints on the states and inputs had to be added after the main problem formulation, resulting in a computationally complex problem that grows as the wind plant becomes larger. Others have proposed distributed methods that minimize loads and provide power reference tracking, but still require a centralized problem to be formulated to facilitate the distributed optimization [2]. We have previously proposed a distributed model predictive control solution for power tracking in wind farms [4]. The proposed model requires a linearization and does not currently leverage local measurements taken by turbine supervisory control and data acquisition (SCADA) systems.

Furthermore, short-term forecasting for wind farm power and wind conditions has been of great interest to both the wind and grid communities. Accurate near-term forecasting of wind plant power enables grid operators to better balance the grid and can alleviate the demand placed on traditional providers of ancillary services. Additionally, forecast information of the wind state across a wind farm can be leveraged by control systems to more efficiently operate turbines and their response to changing environmental conditions. Short-term forecasting can be leveraged in the electricity markets to increase revenue of a wind plant if wind plants are allowed to participate in ancillary services markets. Accurate estimates of short-term wind forecasts can be used to determine the amount of power a wind plant can generate in the near-term future, allowing wind plants to more confidently bid into the electricity market [5]. Other studies have looked at machine learning, or statistical, approaches to improve very short-term forecasting over 10-minute horizons [6], [7]. These approaches require significant amounts of data to accurately predict future events. Hybrid statistical and physics-based models have been proposed; however, it is difficult to beat persistence forecasting —that is, assuming a constant power output —for short timescales [8], [9].

The method proposed in this paper is an algorithm that continues the work of [10] and takes advantage of the topology of a wind farm while incorporating local measurements from nearby turbines to determine wind direction and wind speed at an individual turbine. This approach demonstrates short-term forecasting using a physical model of the wind farm where atmospheric information is propagated along network edges to provide a short-term prediction of the wind farm power. To provide a short-term forecast, first the

J. Annoni is with the National Renewable Energy Laboratory; email: jennifer.annoni@nrel.gov.

C. Bay is with the National Renewable Energy Laboratory; email: christopher.bay@nrel.gov.

K. Johnson has a joint appointment with the University of Colorado School of Mines and the National Renewable Energy Laboratory; email: kjohnson@mines.edu.

P. Fleming is with the National Renewable Energy Laboratory; email: paul.fleming@nrel.gov.

This work was supported in part by the U.S. Department of Energy under Contract No. DE-AC36-08GO28308 with the National Renewable Energy Laboratory. Funding for the work was provided in part by the DOE Office of Energy Efficiency and Renewable Energy, Wind Energy Technologies Office. The authors are solely responsible for any omission or errors contained herein.

wind direction needs to be determined. This is discussed in Section III-A. After the wind direction is determined, the upstream and downstream turbines can be identified and the wind speeds across the wind farm can be propagated from upstream turbines to downstream turbines at 10-15-minute time horizons (see Section III-B). Finally, we compute the power of the wind farm and compare the results to persistence forecasting (see Section IV). Conclusions and future work are presented in Section V.

II. Autonomous Wind Farm

Typically, turbines in a wind farm operate individually without taking into account aerodynamic interactions within a wind farm. It has been shown that operating individual turbines sub-optimally can improve the overall performance of a wind farm [11]. However, there are many challenges to developing and implementing real-time wind farm controllers. Wind farms are complex dynamical systems that are difficult to model with sufficient accuracy without large computational costs. In addition, wind farms have large time delays that make traditional model-based control difficult.

The autonomous wind farm was introduced in [10] and creates a foundation for implementing real-time control algorithms that have the potential to increase the performance of a wind plant. The autonomous wind farm self-organizes turbines into groups, monitors, and controls its performance in real time based on existing SCADA data. In this framework, turbines take advantage of data from nearby turbines to make more informed decisions to improve controllability, observability, and predictability. In this paper, we extend the previous work to use both wind direction and wind speed measurements from SCADA data, sharing this inflow information with nearby turbines to improve the short-term forecast of the power output of a wind plant at timescales of 10-15 minutes. If we know the short-term forecast in real time, we can use that information for control to increase energy capture, decrease structural loads, or optimize across these objectives.

A. Wind Farm as a Network

To implement real-time algorithms for an autonomous wind farm, the wind farm can be modeled as an undirected or directed network, depending on how inter-turbine communication is considered. Turbines in a wind farm can be considered the nodes, and the edges consist of established communication links between nearby turbines. Information is communicated across these edges to determine local atmospheric conditions—such as wind direction or wind speed—at a particular turbine. In the wind farm case, information is embedded in the wake of the turbine. The wake has information about the operation of the upstream turbine.

1) Undirected Network: An undirected network is a network in which information is exchanged in both directions along an edge. In this paper, this type of network is used to determine the wind direction across the wind farm. A wind farm can be modeled as an undirected graph where turbines are communicating with connected turbines and

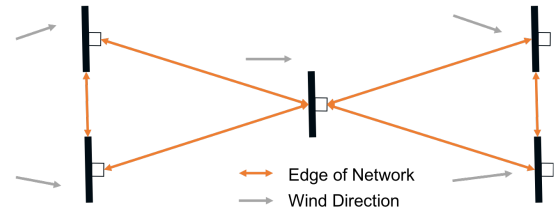

Fig. 1: Example of a network drawn to determine the wind direction consensus, looking down from above a wind farm with rotors represented by vertical black lines.

information flows both ways, rather than from one turbine to the next turbine. Modeling a wind farm as an undirected network allows for relevant spatial information to be used to determine the local atmospheric conditions.

2) Directed Network: A directed network is a network in which each edge has a direction and information flows in one direction from one node to another. A wind farm can also be modeled as a directed graph with information flowing from upstream turbines to downstream turbines. This network type will be used to determine the wind speed. The turbine interactions—that is, their wake interactions—determine the network topology. Under this paradigm, the network topology is determined by current atmospheric conditions including wind direction.

III. Short-Term Forecasting

This section details the approach to perform a NEwork-based Short-Term (NEST) forecast across an autonomous wind farm. Using SCADA data, we will determine the wind direction across the wind farm. Next, using this information, we can determine which turbines are upstream and which turbines are downstream. This information can be used in propagating wind speeds across the wind farm. Finally, knowing the wind speed across the wind farm, we can compute the predicted power 10-15-minutes into the future.

A. Wind Direction Consensus

To determine the wind direction across the wind farm, we use a consensus-based approach that uses an undirected graph to robustly determine the wind direction at every turbine considering both the turbine's measurements and those of its nearest 10 neighbors. In this subsection, we summarize the key elements of [10] that are used in our forecast; for more details on wind direction consensus, see [10]. The SCADA measurements recorded at each turbine are used to determine a robust estimate of wind direction at every turbine. This approach allows the wind direction to vary across a wind farm. It is assumed that the wind directions recorded at the turbines are with reference to true north and that the wind direction varies smoothly across the wind farm. This allows for turbines to come to an agreement on the wind direction in an "almost" consensus way; i.e., the wind directions of nearby turbines are close, but not necessarily the same.

1) Network Topology: The network topology, i.e., graph structure, is important for incorporating local information and taking advantage of the structure of the wind farm to perform real-time optimizations. In this paper, turbine communications can be defined by the nearest N turbines [4], see Fig. 1. This algorithm defines the graph structure based on the nearest 10 turbines. Alternative approaches can be used to cluster turbines to optimally exchange information such as connectivity, hierarchical, or k-means algorithms [12].

2) Node and Edge Objective Functions: For this problem, each turbine uses its own wind direction measurement as well as the wind direction measurement from the connected turbines to determine its local wind direction. The objective of the individual turbine, i, i.e., the node objective, is to minimize the error between the wind direction measurement measured at turbine i and the estimated wind direction. In addition to the node objective, the edge objective incorporates information from nearby turbines to ensure a robust measurement of the wind direction at an individual turbine. The optimization problem can be written as:

$$\underset{x_i}{\text{minimize}} \quad \underbrace{\sum_{i \in \mathcal{V}}^{N} (x_{i,measure} - x_i)^2}_{\text{node objective}} + \underbrace{\sum_{j,k \in \mathcal{E}} w_{jk} |x_j - x_k|}_{\text{edge objective}} \tag{1}$$

where N is the number of turbines, $x_{i,measure}$ is the wind direction measurement recorded at the turbine i, w_{jk} is a weight placed on the connection between turbines j and k, x_j is the estimated wind direction at turbine j, and x_k is the estimated wind direction at turbine k. In this case, the node objective function is convex and can be updated with a closed-form solution [10]. The edge objective minimizes the differences in estimated wind direction between neighboring turbines. In this paper, the weights w_{jk}, are set to 1. However, different weights can be used, including weights that vary based on distance. Future work will optimize the weighting between turbine communications to better integrate the data.

3) Alternating Direction Method of Multipliers: To solve (1), we use the Alternating Direction Method of Multipliers (ADMM) [13]. This algorithm is particularly useful in this case, as each individual turbine can solve its own optimization in parallel, communicate the solution to neighboring turbines, and iterate this process until each node within the wind farm network has converged. In this paper, each turbine determines the local wind direction at each individual turbine by only talking to its nearest neighbors as indicated in Section III-A.1. ADMM is used to solve a network optimization with connecting nodes such that:

$$\begin{aligned} \underset{x_i}{\text{minimize}} \quad & \sum_{i}^{N} (x_{i,measure} - x_i)^2 \\ & + \lambda \sum_{(j,k) \in \mathcal{E}} w_{jk} \|z_{jk} - z_{kj}\|_2 \\ \text{subject to} \quad & x_i = z_{ij}, \quad j \in N(i) \end{aligned}$$

where z_{jk} is a copy of x_j at turbine k such that the wind farm reaches an "almost" consensus of the wind direction across the wind farm.

The distributed optimization problem is solved by minimizing the augmented Lagrangian:

$$\begin{aligned} \mathcal{L}_\rho(x, z, u) = \sum_{i \in \mathcal{V}} f_i(x_i) + \sum_{(j,k) \in \mathcal{E}} \lambda w_{jk} \|z_{jk} - z_{kj}\|_2 \\ -(\rho/2) \left(\|u_{jk}\|_2^2 + \|u_{kj}\|_2^2 \right) \\ +(\rho/2) \left(\|x_j - z_{jk} + u_{jk}\|_2^2 + \|x_k - z_{kj} + u_{kj}\|_2^2 \right) \end{aligned}$$

where u is the scaled dual variable and $\rho > 0$ is the penalty parameter. The following steps are used in an iterative way to solve for x, z, and u:

$$x^{m+1} = \underset{x}{\text{argmin}} \quad \mathcal{L}_\rho(x, z^m, u^m) \tag{2}$$

$$z^{m+1} = \underset{z}{\text{argmin}} \quad \mathcal{L}_\rho(x^{m+1}, z, u^m) \tag{3}$$

$$u^{m+1} = u^m + \left(x^{m+1} - z^{m+1} \right) \tag{4}$$

There are closed-form solutions to the above equations that are shown in [10].

This setup provides an incentive for the difference between the connected nodes to be zero. For the wind farm example, this means that turbines near each other should have similar wind direction measurements. There are two penalty parameters, λ and ρ, that can be used to weigh an individual turbine's measurement against the measurements of nearby turbines.

B. Wind Speed

After the wind direction is determined across the wind farm at each individual turbine, the leading turbines can be determined and this information can be used to compute the wind speed at each turbines across the farm. Similar to wind direction, the network topology of the wind farm must first be determined.

1) Network Topology: This topology differs from the wind direction network that is shown in Fig. III-B.1. In particular, the wind speed network is a directed graph in that information only flows from upstream turbines to downstream turbines. The directed network is determined using the local wind direction established in Section III-A. Specifically, the wind direction at an individual turbine is used to assess which turbines are upstream of that specified turbine. The clusters of turbines can be different sizes than in the wind direction network. For the wind speed network, information from all the upstream turbines is used to determine the wind speed at a particular downstream turbine. Information is only exchanged in one direction, from an upstream to a downstream turbine, making this a directed graph. Wind speed is propagated along the edges of a network using nonlinear dynamics. Using the wind speed computed at each turbine, we can compute the expected power of the wind farm in the 10-15-minute time horizon.

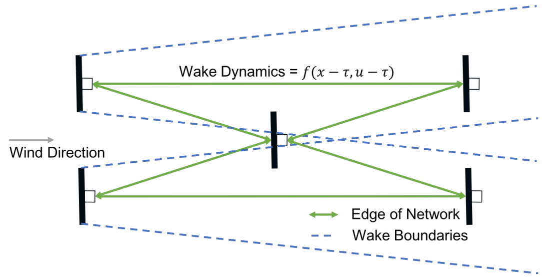

Fig. 2: Example of a network to determine wind speed across a wind farm.

2) Dynamics Across a Network: Using the directed network defined above, the wind speed is projected at a downstream turbine using a nonlinear wind turbine wake model and time delay applied across the edges of the network. Wind speeds propagate downstream along edges, determined by the wind direction, based on:

$$\dot{x}_i(t) = \frac{1}{N} \sum_j A_{ij} h_i \left(x_j(t-\tau), u_j(t-\tau) \right) \quad j \in \mathcal{N}(i) \tag{5}$$

where N is the number of turbines upstream of turbine i, j indicates the indices of the turbines upstream of turbine i, x is the wind speed at turbine i, u_j are the control actions at turbine j, h_i is a nonlinear wake model representing impacts of all of the turbines upstream of turbine i, and τ is the time it takes for the effects of the upstream turbine to reach the downstream turbine.

3) Wind Turbine Wake Model: The wind turbine wake model $h_i(x_j, u_j)$ is used to characterize the aerodynamic interactions in a wind farm along an edge in the wind speed network. The model, known as FLORIS [14], was used for this work and incorporates several models from recent papers including [15], [16]. In particular, this model uses a Gaussian wake shape to describe the velocity deficit behind a turbine and includes the effects of turbulence in the wake and local atmospheric conditions. An analytical expression for the three-dimensional velocity deficit behind the turbine in the far wake can be derived from the simplified Navier-Stokes equations:

$$\frac{\mathrm{u(x,y,z)}}{U_\infty} = 1 - Ce^{-(\mathrm{y}-\delta)^2/2\sigma_\mathrm{y}^2} e^{-(\mathrm{z}-\mathrm{z}_h-\mathrm{z}_e)^2/2\sigma_\mathrm{z}^2} \tag{6}$$

where u is the velocity in the wake, U_∞ is the free-stream velocity, x is the streamwise direction, y is the spanwise direction, δ is the wake centerline, z is the vertical direction, z_h is the hub height, z_e is the elevation of a turbine in a wind farm, σ_y is the wake expansion in the y direction, σ_z is the wake expansion in the z direction, and C is the velocity deficit at the wake center. These parameters are further defined in [16].

The wind farm investigated in this paper is in complex terrain. Terrain is taken into account by adding an elevation variable, z_e, to the wake model. This wake model incorporates some of the effects of complex terrain. Future work will include improving this model to include more complex effects of complex terrain such as wakes traversing hills and other terrain features.

Finally, a turbine model is used in this wake model to provide a realistic description of turbine interactions in a wind farm. The turbine model consists of a power coefficient, C_P, and thrust coefficient, C_T, based on wind speed. The exact turbine information was not known for this study. Rather, the information from the NREL 5 MW turbine [17] was scaled and used for this study. This approximation produces some error in the NEST forecast.

4) Time Delay: The time delay used to transport wind speed, determined by FLORIS, from an upstream turbine to a downstream turbine is computed using Taylor's frozen turbulence hypothesis [18]. In particular, the waked velocity, u, is used to advect the flow from the upstream turbine to the downstream turbine:

$$\tau = \frac{d}{\mathrm{u}} \tag{7}$$

where d is the distance between turbines. This will provide a local wind speed estimate at each turbine for future predictions.

C. Power Calculation

Once the wind direction and wind speed at each turbine have been calculated using the network theory methods described in Sections III-A-III-B, the turbine power is computed using [19]:

$$P_i = \frac{1}{2} \rho A C_P(u_i) u_i^3 \tag{8}$$

where ρ is the air density, A is the rotor swept area, C_P is the power coefficient of the turbine, which is a function of the local wind speed, u_i. For this study, only SCADA data were used and therefore the air density was set to a constant $\rho = 1.225$ kg/m^3, which introduces another source of error into the NEST forecast in addition to the C_P scaling mentioned earlier. However, the results shown in Section IV are promising, indicating that SCADA data can be used to compute the forecast of a wind farm in the near term with limited knowledge of the wind turbines in the wind farm. Further information about atmospheric conditions from meteorological towers, including air density, and information about the specific turbine can improve the NEST forecast in future work.

IV. Results

A. Simulation Setup

The NEST forecast was demonstrated on a subset of wind turbines in a wind farm used in the Wind Forecasting Improvement Project, also known as WFIP2 [20]. Some of these turbines are shown as the black dots in Figures 3 and 4. Note that the scale has been removed for proprietary reasons in both figures. SCADA data were used at 1-minute time intervals from individual turbines over approximately 8 months. The data channels of interest were the measured wind direction, wind speed, and measured power at each turbine.

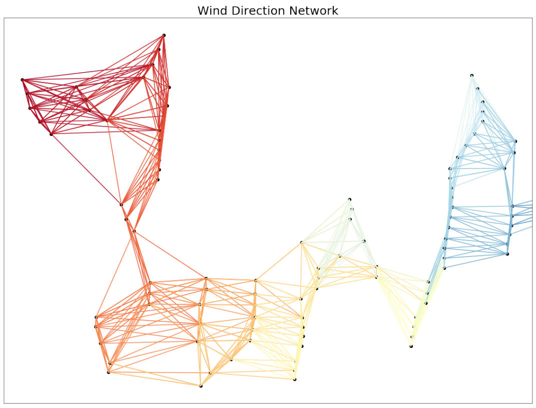

Fig. 3: Wind direction network where each turbine is connected to the nearest 10 turbines.

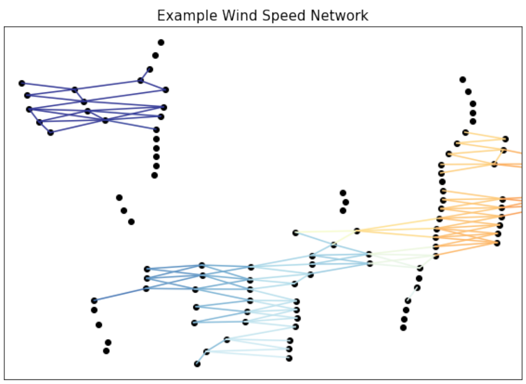

Fig. 4: Wind speed network where each turbine is connected to the turbines upstream of it based on the local wind direction.

B. Network Topology

As described in Section III-A.1 and III-B.1, the network topology was constructed to determine the wind direction across the wind farm by connecting each turbine to the nearest 10 turbines as shown in Fig. 3. The wind direction was then computed using the consensus approach discussed in Section III-A. Once the wind direction was determined for a particular time step, the wind speed network could be defined as shown in Fig. 4 where upstream turbines are connected to downstream turbines. The network looks more linear rather than the web-like appearance of the wind direction network. For this example, the wind direction was determined to be approximately 270° with 7° of variation across the wind farm. The wind speed dynamics are propagated along the edges of the wind speed network to determine the approximate wind speed at some future time.

C. Forecasting Errors

To demonstrate the NEST forecast, 830 hours of SCADA data were evaluated by this algorithm. The algorithm takes 0.75 s to analyze one time step of data. The NEST forecast will be compared to a persistence forecast. Persistence forecasting assumes the conditions have not changed over the last time interval. As the time interval becomes smaller, the persistence forecast becomes more accurate. Fig. 5 (top) shows the time series of 1400 minutes comparing the actual data, shown in black, with the estimated power output from the NEST forecast, shown in red, and the persistence forecast, shown in blue. The bottom plot shows the percent error between the estimated power output using the short-term forecasting method and persistence. Qualitatively, in the time period of 2400-3000 minutes, the NEST forecast has a consistently lower error than the persistence forecast. During the period between 1600-2100 minutes, the wind farm is experiencing curtailment. The NEST forecast as described in this paper does not have the necessary inputs to predict curtailment. However, during periods of curtailment, the NEST approach can be used to compute the possible power in the wind (1675-2100 minutes). This predictive capability might be useful for future wind plants bidding into ancillary services markets to determine the amount of power a wind farm can produce. In addition, it can be used to compensate wind farms for what they would have produced had they not been forced to curtail power.

Case	10-minute		15-minute	
	NEST	Persistence	NEST	Persistence
Down-ramp	3.6%	5.7%	5.6%	9.2%
Up-ramp	4.7%	4.8%	6.4%	7.8%
Overall	3.8%	4.3%	4.6%	6.0%

TABLE I: Short-Term Forecasting Results. The percentages represent the mean absolute error between the forecast and the actual power produced by the wind farm.

Finally, the NEST forecast was compared to a persistence forecast in up/down ramp events and overall accuracy. Up and down ramps were characterized as 1% changes over the 10- or 15-minute intervals. The overall accuracy was determined after excluding curtailment events. If a wind farm is caught in an up or down ramp, persistence forecasting is not sufficient. The results are quantified in Table I. Table I shows that the NEST forecast can outperform persistence forecasting in each of these categories. The percentages represent the mean absolute error between the forecast and the actual power produced by the wind farm. As expected, the differences between accuracies decreases as the time scale decreases from 15-minutes to 10-minutes. With a better forecast on up and down ramp events.

V. Conclusions

This paper presents a short-term forecasting algorithm to predict power in a wind farm on a network over the time horizons of 10-15 minutes. This method was demonstrated on a subset of turbines in a wind farm in complex terrain. It was shown that the NEST method out-performs the persistence algorithm especially on down ramping events.

The NEST approach relies on only SCADA data in a wind farm. Future work will include improvements to this approach, including: 1) using a more realistic wind turbine model derived from data, 2) incorporating real-time data into the FLORIS model to better account for changing conditions,

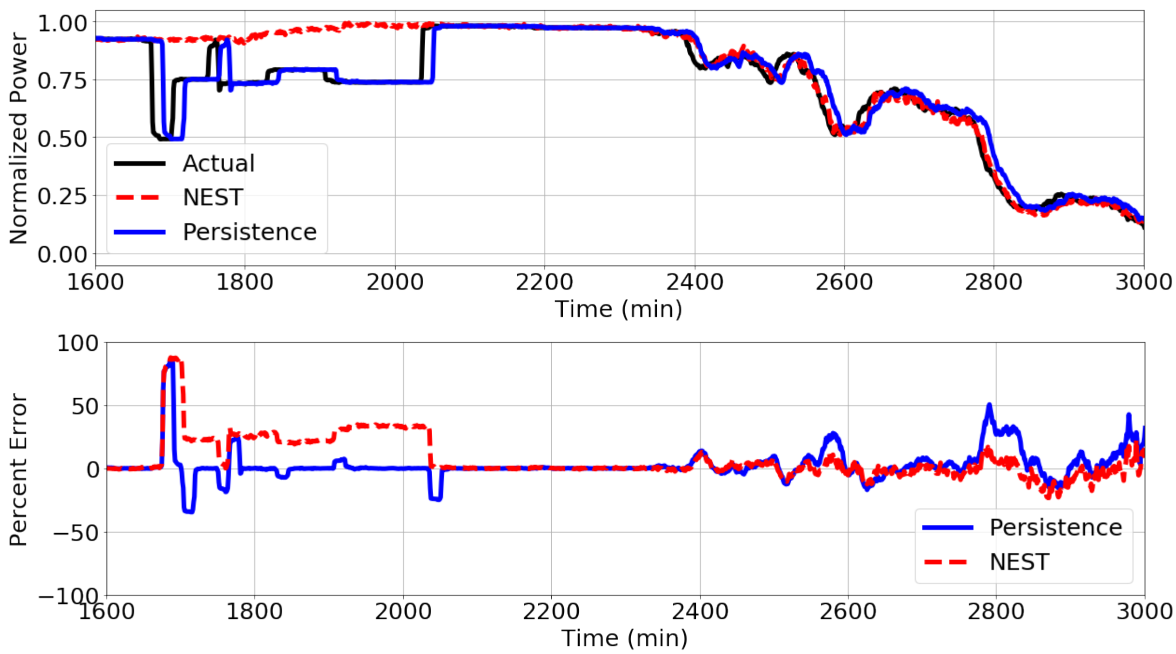

Fig. 5: (Top) This figure shows a time series of the actual power output of the wind farm, shown in black, the estimated power output from the short-term forecast, shown in red, and the persistence forecast, shown in blue, (Bottom) This figure shows the percent error computed by persistence, shown in blue, and the short-term forecasting approach outlined in this paper, shown in red.

and 3) using nearby atmospheric measurements (within a few kilometers) such as air density and turbulence intensity to better inform the NEST forecast. Lastly, this approach can be integrated with statistical approaches to provide a more robust hybrid approach to short-term forecasting.

References

[1] B. D. Kroposki, E. Dall-Anese, A. Bernstein, Y. Zhang, and B. S. Hodge, "Autonomous energy grids: Preprint," National Renewable Energy Laboratory.(NREL), Golden, CO (United States), Tech. Rep., 2017.

[2] V. Spudić, C. Conte, M. Baotić, and M. Morari, "Cooperative distributed model predictive control for wind farms," *Optimal Control Applications and Methods*, vol. 36, no. 3, pp. 333–352, 2015.

[3] M. Soleimanzadeh, R. Wisniewski, and K. Johnson, "A distributed optimization framework for wind farms," *Journal of Wind Engineering and Industrial Aerodynamics*, vol. 123, pp. 88–98, 2013.

[4] C. Bay, J. Annoni, T. Taylor, L. Pao, and K. Johnson, "Active power control for wind farms using distributed model predictive control and nearest neighbor communication," in *American Control Conference (ACC)*. IEEE, 2018.

[5] R. Barthelmie, F. Murray, and S. Pryor, "The economic benefit of short-term forecasting for wind energy in the uk electricity market," *Energy Policy*, vol. 36, no. 5, pp. 1687–1696, 2008.

[6] J. Dowell and P. Pinson, "Very-short-term probabilistic wind power forecasts by sparse vector autoregression," *IEEE Transactions on Smart Grid*, vol. 7, no. 2, pp. 763–770, 2016.

[7] P. Pinson, "Very-short-term probabilistic forecasting of wind power with generalized logit–normal distributions," *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, vol. 61, no. 4, pp. 555–576, 2012.

[8] X. Wang, P. Guo, and X. Huang, "A review of wind power forecasting models," *Energy procedia*, vol. 12, pp. 770–778, 2011.

[9] H. Holttinen, P. Meibom, C. Ensslin, L. Hofmann, J. Mccann, J. Pierik *et al.*, "Design and operation of power systems with large amounts of wind power," in *VTT Research Notes 2493*. Citeseer, 2009.

[10] J. Annoni, C. Bay, K. Johnson, E. Dall'Anese, E. Quon, T. Kemper, and P. A. Fleming, "A framework for autonomous wind farms: Wind direction consensus," *Wind Energy Science*, Submitted September 2018.

[11] S. Boersma, B. Doekemeijer, P. M. Gebraad, P. A. Fleming, J. Annoni, A. K. Scholbrock, J. Frederik, and J.-W. van Wingerden, "A tutorial on control-oriented modeling and control of wind farms," in *American Control Conference (ACC), 2017*. IEEE, 2017, pp. 1–18.

[12] R. Xu and D. Wunsch, "Survey of clustering algorithms," *IEEE Transactions on neural networks*, vol. 16, no. 3, pp. 645–678, 2005.

[13] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends® in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.

[14] J. Annoni, P. Fleming, A. Scholbrock, J. Roadman, S. Dana, C. Adcock, F. Porte-Agel, S. Raach, F. Haizmann, and D. Schlipf, "Analysis of control-oriented wake modeling tools using lidar field data," *Wind Energy Science*, Submitted February 2018.

[15] A. Niayifar and F. Porté-Agel, "A new analytical model for wind farm power prediction," in *Journal of Physics: Conference Series*, vol. 625, no. 1. IOP Publishing, 2015, p. 012039.

[16] M. Bastankhah and F. Porté-Agel, "Experimental and theoretical study of wind turbine wakes in yawed conditions," *Journal of Fluid Mechanics*, vol. 806, pp. 506–541, 2016.

[17] J. Jonkman, S. Butterfield, W. Musial, and G. Scott, "Definition of a 5-mw reference wind turbine for offshore system development," *National Renewable Energy Laboratory, Golden, CO, Technical Report No. NREL/TP-500-38060*, 2009.

[18] D. Schlipf, D. Trabucchi, O. Bischoff, M. Hofsäß, J. Mann, T. Mikkelsen, A. Rettenmeier, J. Trujillo, and M. Kühn, "Testing of frozen turbulence hypothesis for wind turbine applications with a scanning lidar system," in *15th International symposium for the advancement of boundary layer remote sensing*. ISARS, 2010.

[19] T. Burton, D. Sharpe, N. Jenkins, and E. Bossanyi, *Wind energy handbook*. John Wiley & Sons, 2001.

[20] S. Pearse, P. J. Munoz, and B. Kosovic, "Visualization of ncar's wind forecast improvement project 2," in *Proceedings of the Practice and Experience in Advanced Research Computing 2017 on Sustainability, Success and Impact*. ACM, 2017, p. 77.