

How accurate is a machine learning-based wind speed extrapolation under a round-robin approach?

Nicola Bodini¹, Mike Optis¹

¹National Renewable Energy Laboratory, Golden, Colorado, USA

E-mail: nicola.bodini@nrel.gov

Abstract. As the size of commercial wind turbines keeps increasing, having accurate ways to vertically extrapolate wind speed is essential to obtain a precise characterization of the wind resource for wind energy production. Recently, machine learning has been proposed and applied to extrapolate wind speed to hub heights. However, previous studies trained and tested the machine learning methods at the same site, giving them an unfair advantage over the conventional extrapolation techniques, which are instead more universal. Here, we use data from four sites in Oklahoma to test a round-robin validation approach for machine learning, under which we train a random forest at a site, and test it at a different site, where the model has no prior knowledge of the wind resource. We quantify how the accuracy of this technique varies with distance from the training site, and we find that it outperforms conventional techniques for wind extrapolation at all the considered spatial separations. We then assess how the accuracy of the machine-learning based approach varies when it is used to predict wind speed in a wind farm far wake. Finally, we explore as case study the performance of the random forest in extrapolating winds during a low-level jet event.

1. Introduction

An accurate assessment of the wind resource at hub height is necessary for an efficient and bankable wind farm project. However, direct measurement of wind speed at the constantly increasing height of the hub of commercial wind turbines is oftentimes challenging and expensive, so that it is common practice to vertically extrapolate the wind resource from lower and more easily accessible levels [1]. Conventional techniques for wind speed vertical extrapolation include the use of a power law and a logarithmic profile. The power law is a simple empirical relationship, that requires the knowledge of wind speed at two levels to calculate a shear coefficient, which is then applied to extrapolate the wind resource at other heights [2]. On the other hand, the logarithmic profile is based on the Monin-Obukhov Similarity Theory [3]. While simple, the limits in accuracy of these two conventional methods have been shown in various studies [4, 5]. Notably, low-level jets cannot be represented by these simple profiles [6], and wind profiles observed offshore often deviate from these conventional laws [7].

Recently, machine learning has been proposed as a new method to vertically extrapolate winds. While promising, all the published studies on the topic [8, 9, 10] assess the performance of machine learning techniques in vertically extrapolating the wind resource at the same location where the algorithm has been trained, and where it already has knowledge of the wind resource. Also, in real-world applications, the wind resource is measured at the instrument location, but it then needs to be extrapolated at hub height at the location of the wind turbines within the



find farm. Moreover, the ability of machine learning techniques in extrapolating wind speed data under waked conditions has not been assessed to date.

To be able to fully recommend the use of machine learning techniques over the simple power law and logarithmic law, the spatial variability of the performance improvements of the machine learning approaches should first be assessed. Here, we propose a round-robin validation of a machine learning-based method for wind speed extrapolation. We consider a random forest, and use data collected at four different sites in central United States. We test the performance of the machine learning-based approach when used to predict wind speed aloft at a location different from the training site, and contrast results with the predictions from conventional techniques, which instead have a universal nature, as they can be applied at any site. Moreover, we assess how the extrapolation accuracy of the proposed technique is impacted when new wind farms are built in the vicinity of the site of interest, to understand whether training data observed when the wind farm is operating are needed to obtain accurate predictions of wind speed aloft. Finally, we show how accurate the proposed technique is when it is applied to extrapolate the wind resource during low-level jets.

2. Data

We use observations at four locations spanning a 100 km wide region at the Southern Great Plains (SGP) atmospheric observatory, managed by the Atmospheric Radiation Measurement (ARM) Research Facility, in north-central Oklahoma (Figure 1). The topography in the region is relatively simple, with less than 200 m of elevation difference across the different instrument sites. Several wind farms are located in the region, as shown by the clusters of dots in the map in Figure 1.

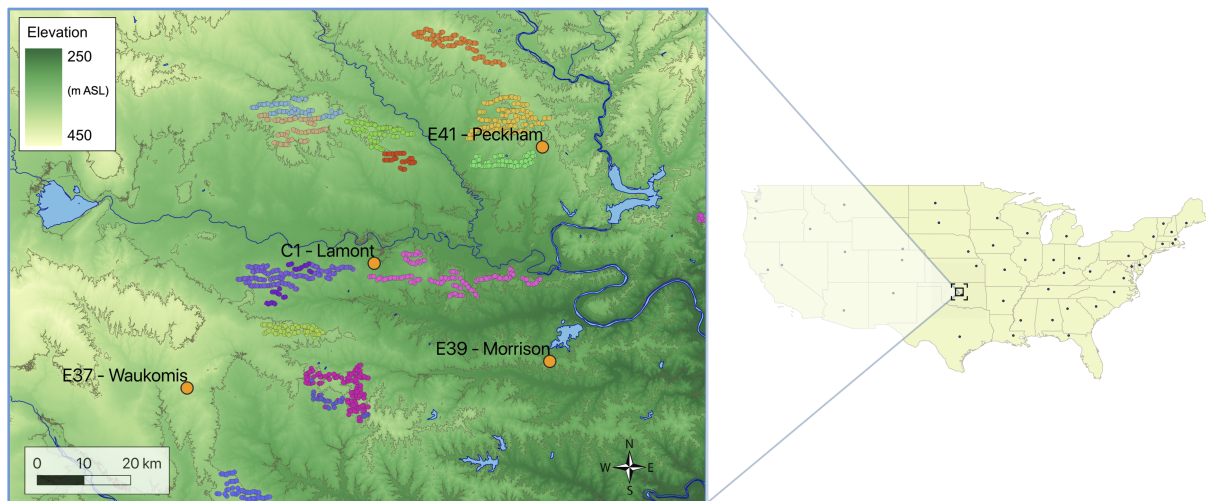


Figure 1. Map of the instrument locations at the SGP site used in this study. The clusters of smaller dots represent the wind farms in the area.

At each of the four sites, we use data from a Halo Streamline scanning lidar [11]. We retrieve [12] horizontal wind speed from the conical sector scans performed by the instrument once every 15 minutes. Each scan had a duration of about 1 minute to complete. The retrieval assumes horizontal homogeneity of the wind field over the scanning volume [13], an assumption that can be considered acceptable given the simple topography at the considered site. As quality controls, we remove periods of precipitation, and we apply a filter with a SNR threshold of -21

dB. From the lidar, 30-minute average data every 26 m, from 65 to 195 m AGL, are used for the vertical extrapolation analysis.

In addition, at each site we use surface measurements at 4 m AGL, recorded by sonic anemometers at a 10-Hz resolution. We use 30-minute average observations of wind speed, turbulent kinetic energy (TKE) and Obukhov length. Following standard approaches in the boundary layer community, a 30-minute period is used for the Reynolds decomposition for turbulent fluxes [14, 15]. Precipitation periods are excluded from the analysis.

3. Methods

3.1. Conventional extrapolation techniques

For the power law, we extrapolate wind speed U measured at $z_1 = 65$ m AGL up to $z_2 = 143$ m AGL using instantaneous shear α calculated from data at 4 and 65 m AGL:

$$U(z_2) = U(z_1) \left(\frac{z_2}{z_1} \right)^\alpha \quad (1)$$

For the logarithmic law, we extrapolate wind speed from $z_1 = 65$ m AGL up to $z_2 = 143$ m AGL as:

$$U(z_2) = U(z_1) + \frac{u_*}{\kappa} \left[\ln \left(\frac{z_2}{z_1} \right) - \Psi_m \left(\frac{z_2}{L}, \frac{z_1}{L} \right) \right] \quad (2)$$

where u_* is friction velocity, $\kappa = 0.41$ is the von Kármán constant, L is the Obukhov length, and Ψ_m is a stability correction. We classify atmospheric stability as a function of the Obukhov length: we consider stable conditions when $L > 0$ m, unstable conditions when $L \leq 0$ m. Following some of the most widely chosen functions for wind energy applications, we use the stability correction Ψ_m proposed by [16] in stable conditions, and the expression suggested in [17] under unstable cases.

3.2. Proposed machine learning approach

At each location, we train a random forest to extrapolate 30-min average wind speed at 143 m AGL. Given the proof-of-concept nature of this analysis, we choose the random forest as a standard but powerful machine learning approach. While we acknowledge that an exhaustive comparison and selection of different machine learning algorithms could lead to optimized results, we defer such a complete comparison to a later study. We use as input features the lidar wind speed at 65 m AGL, lidar wind direction at 65 m AGL (expressed in terms of its sine and cosine, to preserve the cyclical nature of this variable), time of day (its sine and cosine), sonic anemometer wind speed at 4 m AGL, sonic anemometer wind direction at 4 m AGL (its sine and cosine), turbulent kinetic energy, and Obukhov length. We use the `RandomForestRegressor` module in Python's Scikit-learn to implement our model. We consider the hyperparameters in Table 1 to fine-tune the random forest and optimize its accuracy. We use a five-fold cross validation, and sample 30 sets of hyperparameters. To minimize the effects of temporal autocorrelation in the observations, we do not shuffle the data, and use 80% of the whole data set for training, and the remaining 20% for testing.

4. Results

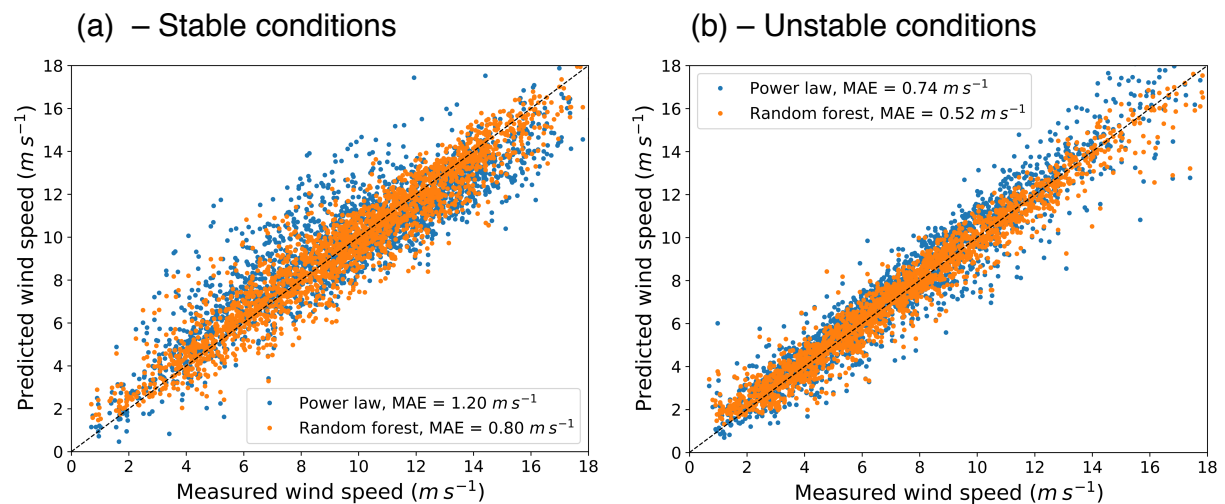
4.1. Performance assessment of the machine learning-based wind speed extrapolation

To assess the performance of the proposed machine learning model in extrapolating wind speed, we use 20 months of data, collected from 13 November 2017 to 23 July 2019. Throughout this period, lidar and surface observations were available at all the four considered sites, and all the wind farms in the vicinity of the instruments had already been built, with the final layout shown in Figure 1.

Table 1. Algorithm Hyperparameters Considered for the Random Forest and Their Considered Values in the Cross Validation.

Hyperparameter	Possible Values
Number of estimators	10–800
Maximum depth	4–40
Maximum number of features	1–10
Minimum number of samples to split	2–11
Minimum number of samples for a leaf	1–15

First, we perform a same-site comparison of the performance of the proposed random forest against the conventional techniques for wind speed extrapolation. Figure 2 compares observed with machine learning- and power law-predicted 30-min average wind speed at 143 m AGL for one of the four considered sites, for different stability regimes. Both the considered techniques

**Figure 2.** Predictive performance of the random forest compared to the use of power law for wind speed vertical extrapolation, for stable and unstable conditions. Comparison between measured 30-min average wind speed at 143 m AGL and prediction by both methods, for the site E39.

can extrapolate winds more accurately under unstable conditions, with the mean absolute error (MAE) in stable conditions about two times larger than what found in unstable cases. The random forest outperforms the power law in vertically extrapolating wind speed in both the considered stability conditions, with a 33% reduction in MAE for stable conditions, and a 30% reduction in unstable conditions. Similar results are found at the three other sites considered in this analysis.

Similarly, Figure 3 compares observed winds at 143 m AGL with the predictions from the logarithmic profile and the random forest. As already found for the power law, the average MAE for the logarithmic profile is larger in stable conditions. Also, the random forest provides more accurate predictions of extrapolated winds compared to the log profile, especially in stable conditions, with a reduction in MAE of 35%. In unstable conditions, the reduction in MAE decreases to about -17% . The same-site comparison between machine-learning based

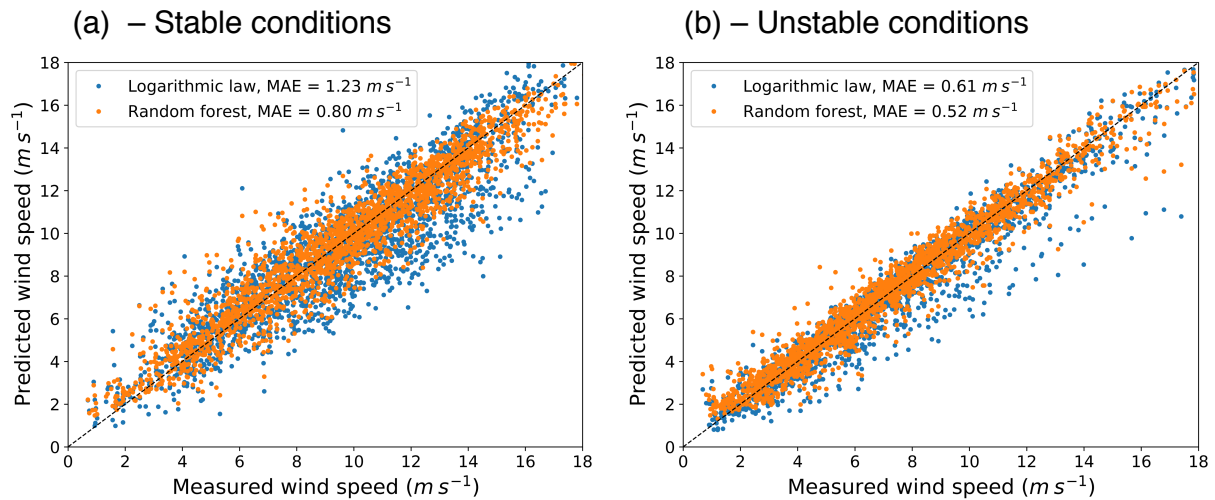


Figure 3. Predictive performance of the random forest compared to the use of logarithmic law for wind speed vertical extrapolation, for stable and unstable conditions. Comparison between measured 30-min average wind speed at 143 m AGL and prediction by both methods, for the site E39.

extrapolation and conventional techniques clearly shows that the random forest outperforms both the power law and logarithmic profile.

However, the proposed machine learning based approach has an inherent advantage under the same-site comparison performed so far, as the algorithm is trained at the same site where is then used for predictions. On the other hand, the power law and logarithmic profile have a more universal nature, and can be applied at any site. Therefore, to obtain a fairer comparison of the performance of the various techniques, we perform a round-robin validation, to use the random forest trained at each site to extrapolate wind speed at the remaining three sites. We summarize the performance metrics of this validation in Figure 4, where we show MAE (normalized, at each site, by the value from the same-site test of the random forest) as a function of the separation between training and testing site, for the considered approaches. We find that the performance of the random forest approach degrades when the algorithm is tested at a site different than the training one with, on average, errors $\sim 15\%$ larger. However, even under those circumstances, the machine learning-based approach outperforms the conventional techniques for wind speed extrapolation, at all the four considered sites.

4.2. Impact of new wind farms on the proposed extrapolation technique

As we do not see a strong trend in the relationship between MAE and separation between training and testing sites, we can speculate that other factors might be involved in the different performance found at the various sites. Here, we explore whether the presence of wind farms in the vicinity of the considered sites has an impact on the results we found. We consider the measurement site C1 - Lamont, and its relationship with the closest wind farm (pink dots in Figure 1), which was built in 2017 and whose turbines are as close as 3.5 km from the instruments at the site C1. While this distance is significant, we can still expect the site to be affected by wind farm wake effects, which have been observed to have an impact on the wind flow at comparable distances [18].

To explore the impact of the new wind farm on the performance of the machine learning algorithm, we consider and contrast observations from two different periods: from 2 November

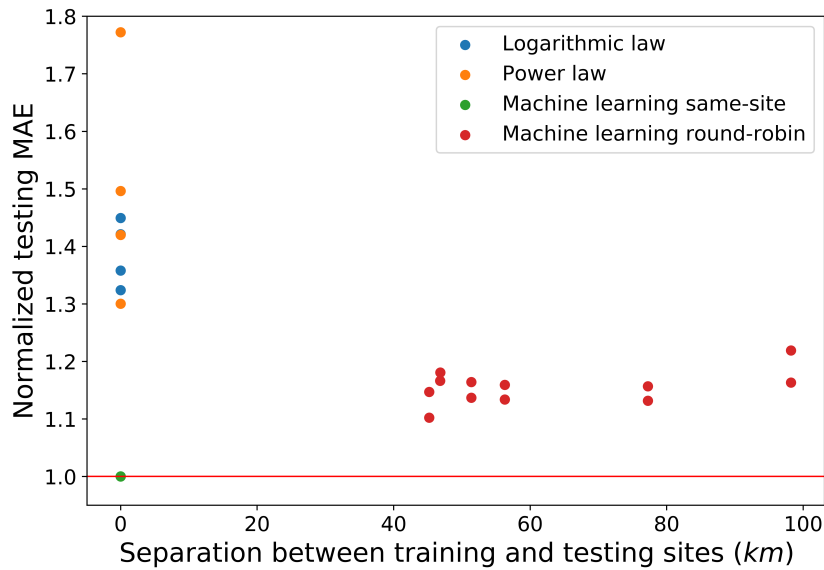


Figure 4. Performance of round-robin approach to predict 30-min average wind speed at 143 m AGL. Normalized testing MAE as a function of the spatial separation between training and testing sites, for all considered extrapolation methods.

Table 2. Comparison of MAE in extrapolating 30-minute average wind speed at 143 m AGL for different combinations of training and testing periods at the C1 - Lamont site.

Training period	Testing period	MAE (m s^{-1})
Pre-wind farm	Pre-wind farm	0.66
Pre-wind farm	Post-wind farm	0.74

2015 to 31 December 2016 (i.e. pre-wind farm) and from 1 January 2018 to 14 January 2020 (i.e. post-wind farm). We train a random forest using the pre-wind farm portion of the data, and test it using both the pre- and post-wind farm data sets. Table 2 compares the MAE values for the relevant different combinations of training and testing periods. We find that when the random forest is trained on the pre-wind farm observations, and then tested on the post-wind farm data, the MAE is about 10% larger than when the testing period is also prior to the construction of the wind farm. Therefore, to maximize the prediction accuracy when using the proposed approach in practical wind energy applications, a slight reduction in accuracy should be taken into account, if observations affected by the presence of the wind farm of interest are not included in the training set.

4.3. Case study: extrapolation performance during Low-Level Jet (LLJ) events

The encouraging results that the machine-learning based extrapolation approach achieves can have significant consequences for a more accurate assessment of the wind resource at sites relevant for wind energy production. To this regard, it is important to consider how the proposed approach performs under specific wind flow conditions that have a large potential for wind energy generation. Low-level jets, which are frequently recorded at the SGP site [19], represent a perfect candidate for the purpose: while they can determine large energy production, they are hard to represent and quantify with the conventional extrapolation techniques [6].

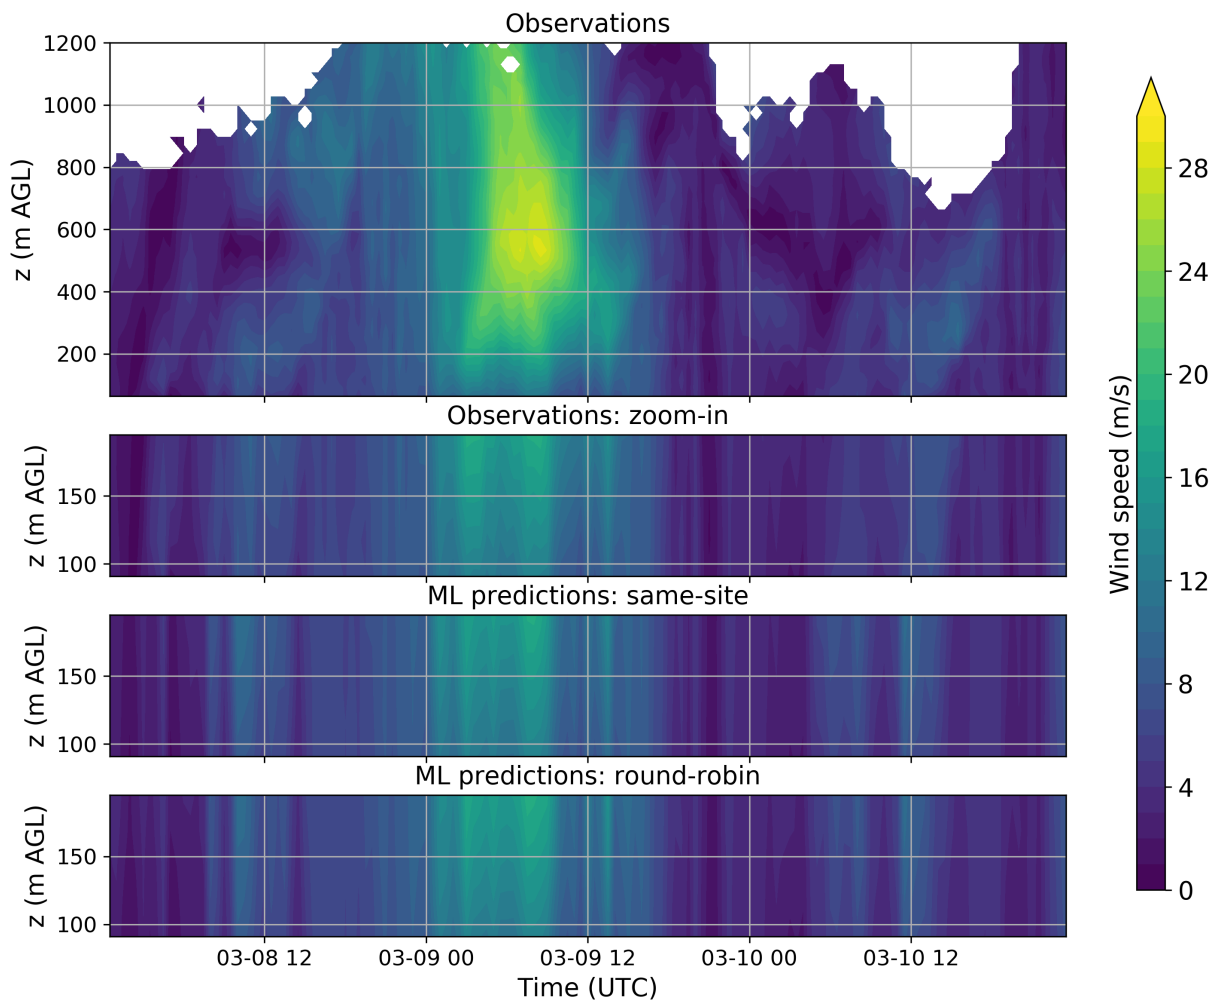


Figure 5. Time-height cross section of observed and machine-learning extrapolated wind speed during a low-level jet event on 9 March 2018.

We select as case study a strong LLJ event that was recorded over the night between March 8-9 2018, with wind speeds over 25 m s^{-1} at $\sim 500 \text{ m AGL}$. We test the performance of random forests in extrapolating the wind resource at the heights of interest for wind energy production. We use as inputs the same set of features used for the main analysis already described, and train various random forests to extrapolate winds at 91, 117, 143, 169, and 195 m AGL. We train and test the random forests at site C1, and then replicate the analysis under the round-robin validation approach, using the site E37 for training, and site C1 for testing. The observed values and machine learning predictions of wind speed are shown in Figure 5. The random forest is capable of accurately extrapolating the wind resource aloft throughout the LLJ event, both under the same-site case and the round-robin one. The MAE in extrapolated winds increases with height from 0.30 m s^{-1} at 91 m AGL to 1.04 m s^{-1} at 195 m AGL for the same-site approach, with an approximately linear trend. As found for the main analysis, the MAE values for the round-robin case are about 5-10% larger at each height.

5. Conclusions

With the constant increase of the size of commercial wind turbines, especially offshore, the need for practical but accurate techniques to extrapolate the wind resource at heights relevant for wind energy production becomes a primary need for a bankable wind energy project. The simplicity of the conventional extrapolation techniques, which include the use of a power law and a logarithmic profile, comes together with their inherent limitations in accuracy.

Machine learning-based techniques for wind speed vertical extrapolation can provide accurate predictions at wind speed aloft. Our results show that although the random forest best performs when used to extrapolate the wind resource at the site where it has been trained, it outperforms the conventional extrapolation techniques of power law and logarithmic profile even when used to extrapolate wind speed at a different site, at the spatial scales sampled in our analysis (up to ~ 100 km of horizontal separation). Under these circumstances, we find, on average, a reduction in mean absolute error between 15 and 20% over the conventional methods, with the largest benefits obtained under stable conditions. The proposed approach is also capable of accurately extrapolating the wind resource during low-level jet events, which have a primary importance for wind energy production in the region tested for the analysis. We have also found that the performance of the machine learning extrapolation approach can be negatively affected if important modifications of the wind flow, for example caused by wind farm wakes, are not included in the data set used to train the model.

Future work can extend this analysis to complex terrain, where the effect of topography on the round-robin validation would likely lead to a faster degradation of the performance of the machine-learning based approach as the distance increases from the training site, after the terrain-induced wind flow features become more and more different from what observed at the training site. Similarly, the same analysis could be replicated offshore, where the heights of interest for wind energy production are even larger, and the installation of multiple remote sensing instruments is more expensive and prohibitive. Moreover, it is worth exploring whether combining observations from multiple sites and terrains could lead to a more universal pre-trained machine learning model, which could then be applied to extrapolate the wind resource at any other location. Finally, with regard to the assessment of the impact of wind farm wakes on the extrapolation performance of the proposed machine learning approach, further analysis can be done using observational sites closer to wind farms.

Acknowledgments

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 in the article 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.

References

- [1] Brower M 2012 *Wind resource assessment: a practical guide to developing a wind project* (Hoboken, New Jersey: John Wiley & Sons)
- [2] Peterson E W and Hennessey Jr J P 1978 *Journal of Applied Meteorology* **17** 390–394
- [3] Monin A S and Obukhov A M 1954 *Contrib. Geophys. Inst. Acad. Sci. USSR* **151**
- [4] Lubitz W D 2009 *Wind Engineering* **33** 259–271
- [5] Optis M, Monahan A and Bosveld F C 2016 *Wind Energy* **19** 1053–

- 1072 (Preprint <https://onlinelibrary.wiley.com/doi/pdf/10.1002/we.1883>) URL
<https://onlinelibrary.wiley.com/doi/abs/10.1002/we.1883>
- [6] Sisterson D, Hicks B, Coulter R and Wesely M 1983 *Solar Energy* **31** 201–204
 - [7] Högström U, Smedman A S and Bergström H 2006 *Wind Engineering* **30** 269–286
 - [8] Türkan Y S, Aydoğmuş H Y and Erdal H 2016 *An International Journal of Optimization and Control: Theories & Applications (IJOCTA)* **6** 179–187
 - [9] Mohandes M A and Rehman S 2018 *IEEE Access* **6** 77634–77642
 - [10] Vassallo D, Krishnamurthy R and Fernando H J in review *Wind Energy Science Discussion*
 - [11] Newsom R 2012 Doppler lidar (dl) handbook Tech. rep. DOE Office of Science Atmospheric Radiation Measurement (ARM) Program
 - [12] Frehlich R, Meillier Y, Jensen M L, Balsley B and Sharman R 2006 *Journal of applied meteorology and climatology* **45** 821–837
 - [13] Browning K and Wexler R 1968 *Journal of Applied Meteorology* **7** 105–113
 - [14] De Franceschi M and Zardi D 2003 *Boundary-Layer Meteorology* **108** 289–303
 - [15] Babić K, Bencetić Klaić Z and Večenaj Ž 2012 *Geofizika* **29** 35–51
 - [16] Beljaars A and Holtslag A 1991 *Journal of Applied Meteorology* **30** 327–341
 - [17] Dyer A and Hicks B 1970 *Quarterly Journal of the Royal Meteorological Society* **96** 715–721
 - [18] Lundquist J, DuVivier K, Kaffine D and Tomaszewski J 2019 *Nature Energy* **4** 26–34
 - [19] Klein P, Bonin T, Newman J, Turner D, Chilson P, Wainwright C, Blumberg W, Mishra S, Carney M, Jacobsen E *et al.* 2015 *Bulletin of the American Meteorological Society* **96** 1743–1764