

Comments on: Space-time wind speed forecasting for improved power system dispatch

Amanda S. Hering

Published online: 27 February 2014
© Sociedad de Estadística e Investigación Operativa 2014

1 Introduction

I would first like to thank the authors for a well-written paper that illustrates the benefit of collaborating across disciplines. This type of work is very important for renewable energy to become integrated into mainstream utility systems, and the authors have clearly demonstrated how advanced statistical forecasting models can be incorporated into optimization models for scheduling power while simultaneously minimizing operating costs. As the percentage of a country's total energy needs that are met with variable renewable energies, such as wind, increases, the potential to disrupt the stability of the utility to deliver uninterrupted power also increases (Marquis et al. 2011). Therefore, having good short-term wind forecasts will be even more important in the coming years, and the development of such models and testbeds as done in this paper by Zhu et al. (2014), hereafter referred to as ZGGX, will be highly relevant.

In this comment on the ZGGX paper, I would like to focus my attention on the Rotating Regime-switching Space–Time Diurnal (RRSTD) model. In the following Sect. 2, novel aspects of this model will be discussed, and then potential difficulties in fitting the model will be identified. The RRSTD is a generalization of the RSTD model developed by Gneiting et al. (2006) and is also related to the TDD model developed by Hering and Genton (2010). ZGGX show that the RRSTD is better than both persistence and autoregressive models, but they have not compared their results to the RRSTD or the TDD models. ZGGX also state that "... the RRSTD model has the potential to improve forecasts further if more information on monthly wind patterns

This comment refers to the invited paper available at: doi:[10.1007/s11749-014-0351-0](https://doi.org/10.1007/s11749-014-0351-0).

A. S. Hering (✉)
Applied Mathematics and Statistics, Colorado School of Mines, Golden, CO 80401, USA
e-mail: ahering@mines.edu

are available”, hypothesizing that their model would do even better if perhaps more data were available. Thus, in Sect. 3, I will compare the performance of the RRSTD model with the RSTD and TDD models using 3 years of 10-min wind data from the same locations used in the ZGGX paper. Section 4 concludes with a brief discussion.

2 RRSTD model

Using off-site information to make wind speed forecasts at a particular location of interest is inherently difficult since the direction that the wind is blowing changes throughout the day and by season. Thus, those “upwind” locations that may provide some information on the wind speed at the location of interest in the future as wind propagates across a region are not constant (Kretzschmar et al. 2004). By identifying regimes, or distinct wind conditions driven by prevailing weather patterns, it is possible to improve upon wind speed forecasting models that rely solely on information collected at the site of interest, as shown by Gneiting et al. (2006) and ZGGX. ZGGX have developed a novel way to objectively classify to which regime an observation belongs and select the off-site location used to switch the regimes by showing which combination of regimes and off-site location minimizes the mean absolute error (MAE) in the training set. However, it should be clarified that as it has been presented here, the RRSTD model is not a pure space–time model, or a model makes both predictions in space and forecasts in time. The RRSTD model uses spatial and temporal information in its forecasts, but it does not make spatial predictions, and this seems to conflict with their statement that the RRSTD is “... helpful when the goal is to evaluate the potential wind power production at certain locations before building wind farms there”.

The authors choose a two-regime model with equal sized regimes, and in this geographic region there is little to suggest that more than two regimes would be necessary. However, the method of choosing the regimes based on the one with the smallest MAE across each different angle becomes computationally intractable when the number of regimes is greater than 2. With two equally sized regimes, only 180 unique sets of regimes need to be explored for each of 12 months. This results in $180 \times 12 = 2,160$ sets of regimes, and $180 \times 12 \times 2 = 4,320$ models that must be fit. Here, I am assuming that the predictors in the model are already known, but in practice this will not be the case. As ZGGX have described it, the predictors must be selected for each possible regime, and this increases computing time. Nevertheless, this large number of models is still manageable.

If two regimes are chosen and are allowed to be of different sizes, then we must choose two angles of the 360° of the circle to divide the circle into two regimes. Choosing angles at either $(30^\circ, 70^\circ)$ or $(70^\circ, 30^\circ)$ would yield the same set of regimes, so the order of the chosen angles is not important. Thus, for each month, there are $\binom{360}{2} = 64,620$ unique sets of regimes, and a total of $64,620 \times 12 \times 2 = 1,550,880$ models to fit in each month and regime. Similarly, for three regimes, there are $\binom{360}{3} = 7,711,320$ regimes, and $7,711,320 \times 12 \times 3 = 277,607,520$ models to fit. Assuming

that the computing time it takes to fit one model in each regime is 0.008 s, it will require approximately 3.4 h to fit all of the two regime models, but over 25 days to fit all of the three-regime models. This clearly becomes prohibitive with four or more regimes when it takes over 8 years to fit all of the models that would be required. Allowing the number of regimes to vary from 1 month to the next and using this method to choose the number of regimes would add another level of complexity to this problem of choosing regimes.

Thus, it seems that an exhaustive search of all possible regimes is only feasible when the number of regimes is small. In cases where the number of regimes to choose is not clear or is almost certainly more than two, an algorithm to find candidate regimes could be applied first. Standard clustering methods, like k -means, have been used in other work to identify wind regimes (Burlando et al. 2008; Clifton and Lundquist 2012), but do not account for the variable or temporal dependence in winds. Kazor and Hering (2014) show how parametric and nonparametric mixture-based models perform in clustering wind speed and direction and find that unconstrained Gaussian mixture models perform relatively well with respect to correctly identifying regimes and selecting the correct number of regimes.

However, assuming that an exhaustive search of regimes is possible, ZGGX plot the MAE for each $\theta^* \in \{1^\circ, 2^\circ, \dots, 180^\circ\}$ on the training data by month at Vansycle in their Fig. 3. The authors state that the smallest MAE is achieved for $\theta = 8^\circ, 9^\circ, 175^\circ$, and 29° in the months August, September, October, and November, respectively. The first three values are still very close to $\theta = 0$, so it would be nice to see results on the testing data comparing the RSTD with the RRSTD. In addition, when $\theta = 8^\circ$ is chosen in August at Vansycle, it appears to occur in an area where the angles produce highly variable values of MAE. Could the results on the testing data have actually been better if θ is chosen not simply to minimize MAE, but where the MAE is both stable and low? For example, the range of angles from approximately 120° to 130° in August appear to produce very low MAE values and are also more consistently low. It is also not clear which off-site location, s^* , at which the regimes are defined to switch, produces the lowest value of MAE in ZGGX's Fig. 3, and this will be discussed further in the next section.

3 Model comparisons

The authors state that "... the performance of the RRSTD model "would be better if more data were available". The main purpose of this section is to flesh out this statement and to determine how well the RRSTD model performs compared to the RSTD and the TDD models on a more complete set of 10-min observations recorded in 2004, 2005, and 2006 at the same three stations. These data were used in Yoder et al. (2013), and a minimal amount of quality control was necessary. Such high-quality wind speed and direction data observed at a temporally and spatially dense network of sites has not historically been easy to obtain. However, thanks to balancing authorities such as the Bonneville Power Administration, such data are being archived and made available publicly; see <http://transmission.bpa.gov/business/operations/wind/MetData.aspx> and <http://www.mesonet.ttu.edu/wind.html> for examples.

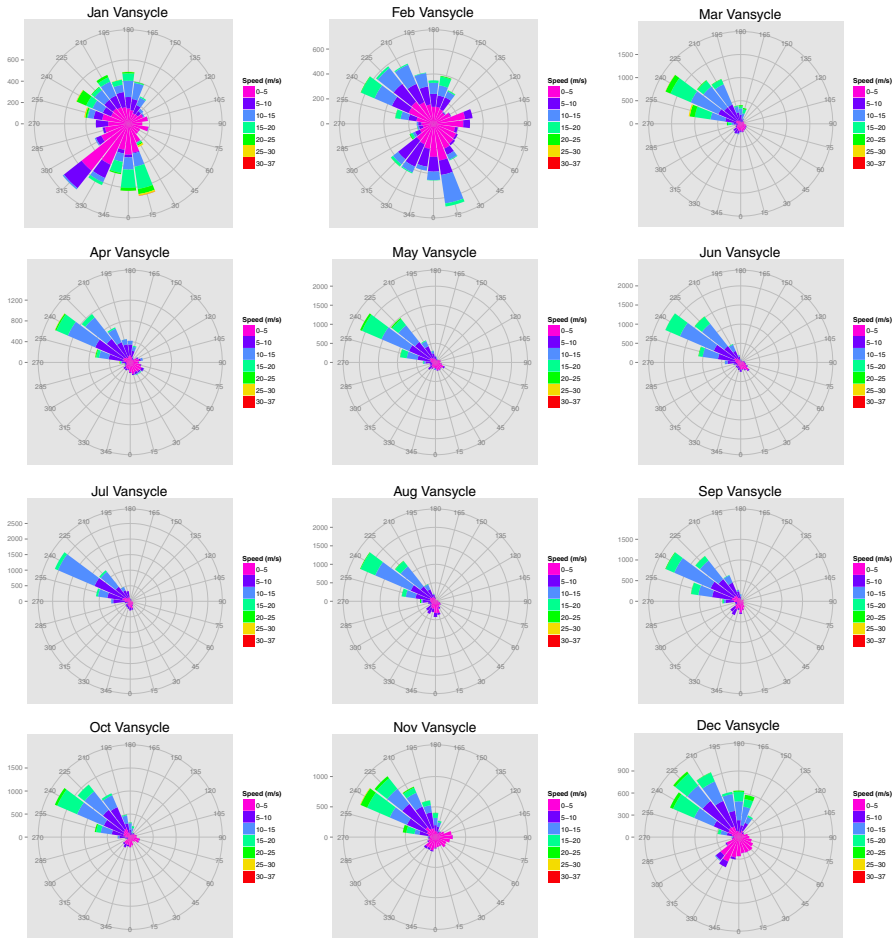


Fig. 1 Monthly wind roses for Vansycle. The monthly observations in 2004 and 2005 are combined (color figure online)

The variables for both the RSTD and TDD models are selected using the 2004 and 2005 training data to predict the 1 h ahead wind speed, μ_{t+6} , at Vansycle, Kennewick, and Goodnoe Hills. Then, the variables chosen are used in their forecasting models to predict wind speed in 2006. A 45-day window of observations is used to estimate the parameters in the models by minimizing CRPS, just as ZGGX do. The RRSTD model is also fit using the 2004 and 2005 training data. In addition to selecting variables to use in Eq. (2) of ZGGX, the wind source, s^* , and the dividing angle, θ , are chosen for each month and location.

But first, we do some exploratory analysis of this new dataset. Figures 1, 2 and 3 show the wind roses at each location where the speeds and directions for each month in 2004 and 2005 are combined. ZGGX only had the later months in 2002 upon which to investigate wind speed and direction behavior, and we also see two distinct modes in August through November (compared with their Fig. 2), but the winds begin to

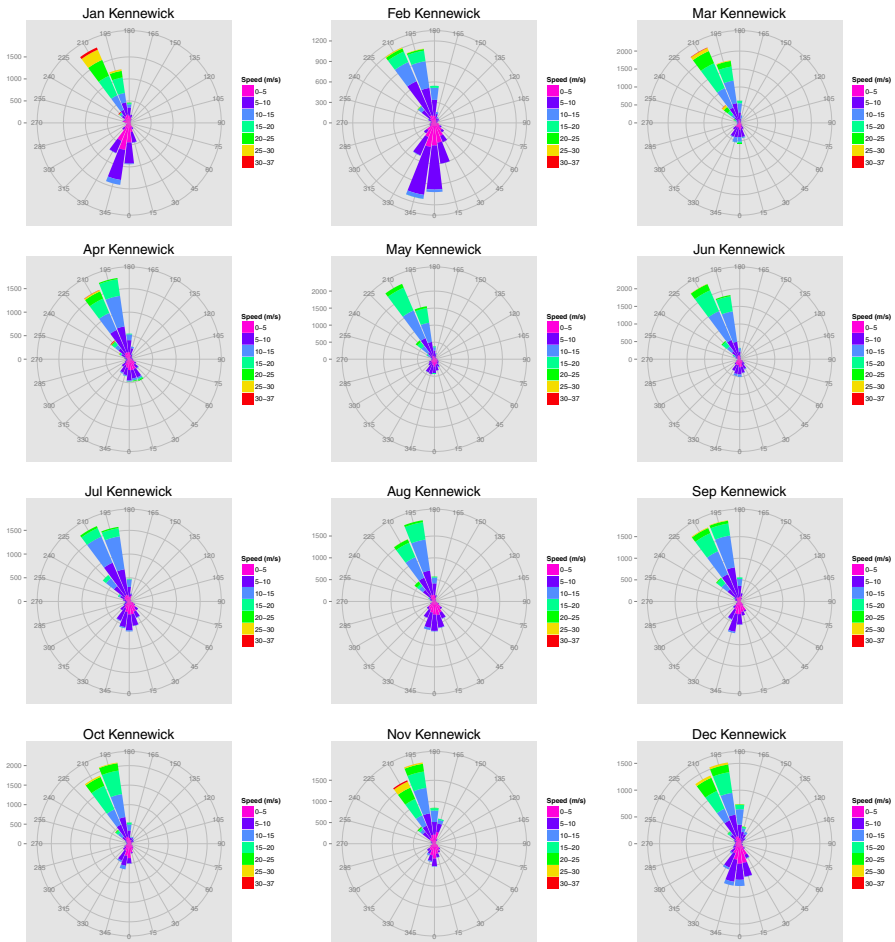


Fig. 2 Monthly wind roses for Kennewick. The monthly observations in 2004 and 2005 are combined (color figure online)

shift in the winter months of December, January, and February. Kennewick continues to display a clear dichotomy of directions, but both Vansycle and Goodnoe Hills demonstrate more variability in the winter months than any other time of the year. Thus, it may be that in these months, a different angle and wind source indicator (as compared to the RSTD model) will produce improved forecasts.

Then, in the course of implementing the RRSTD model, the following two details were not explicitly stated in the ZGX paper:

- First, to choose the wind source indicator, the reader is referred to [Gneiting et al. \(2006\)](#). In that paper, box plots of the wind speeds at Vansycle given that the winds at an off-site location were in either the easterly or westerly regime 2 h prior are constructed. After examining these pairs of box plots across multiple potential off-site locations, the location such that the box plots producing the sharpest separation between the Vansycle winds is chosen. To fully implement the RRSTD and to

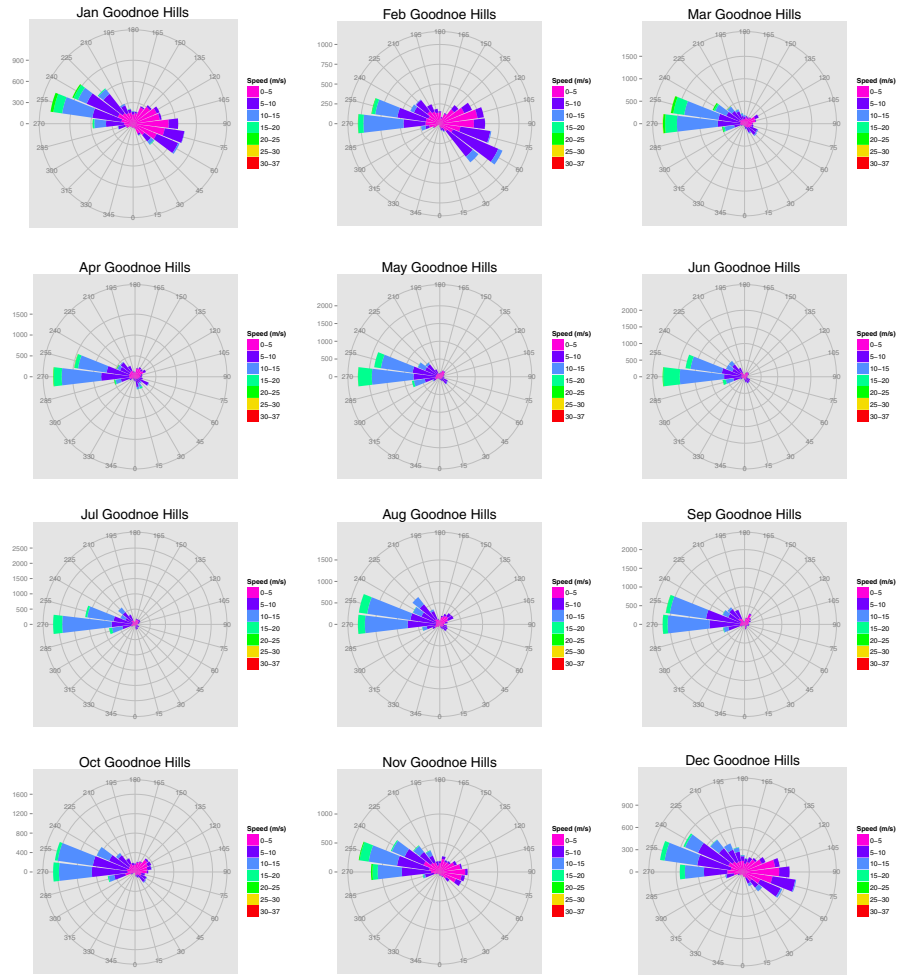


Fig. 3 Monthly wind roses for Goodnoe Hills. The monthly observations in 2004 and 2005 are combined (color figure online)

automatically choose the wind source indicator, s^* , some quantitative measure of the “sharpness” of the separation must be used. For example, the difference in the medians of the box plots could be used, and I will use this difference to choose s^* . However, a more complicated algorithm would be needed for three or more regimes.

- Secondly, the way in which the MAE values on the training data are computed is not entirely clear. I assume that once the predictors are chosen with BIC and the parameters are estimated, then these parameters are fixed and are used to make predictions. This is what I use to implement the RRSTD model.

As a minor point, the authors did not scale the winds at each location to the same height above ground level (agl). This is easily done with the power or logarithm law. Yoder et al. (2013) found very little practical difference between the two laws, so I

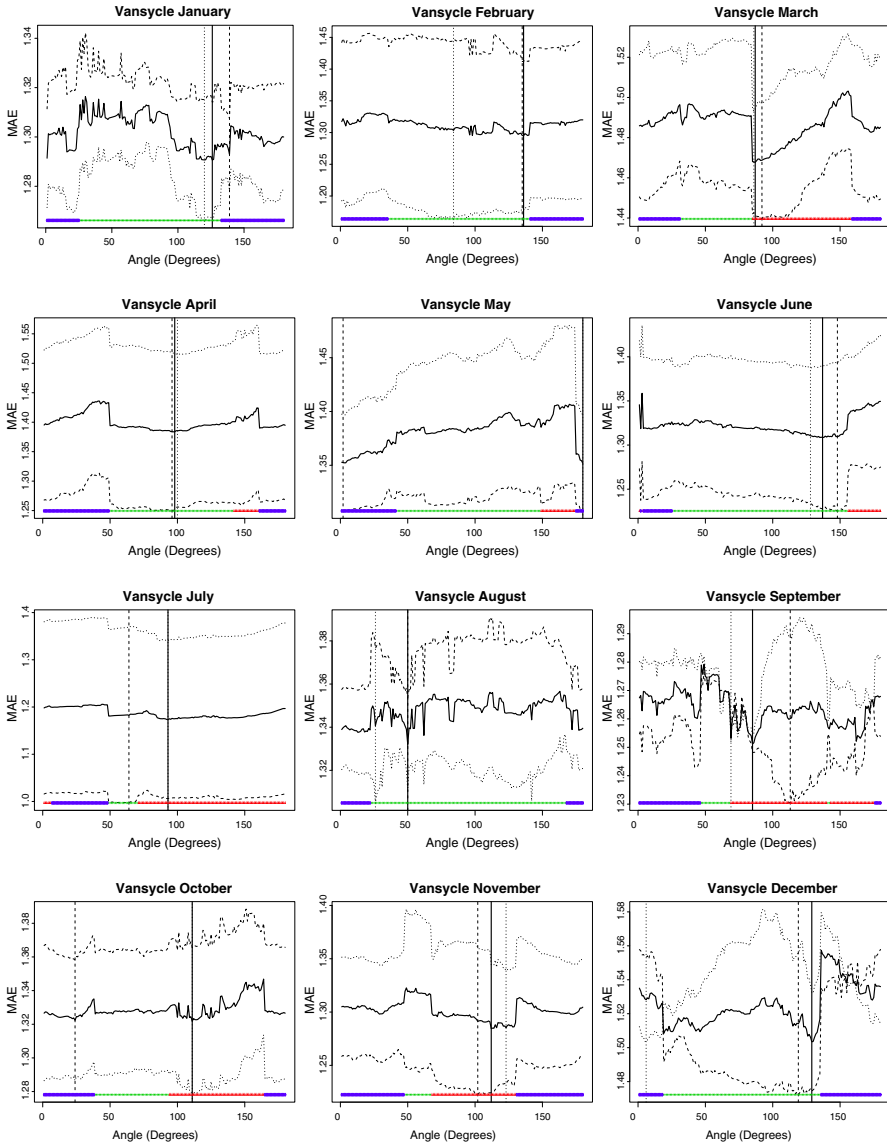


Fig. 4 Plots of 1-h ahead prediction MAE at Vansycle based on a two-regime RRSTD model. The *dashed line* is 2004 training data, the *dotted line* is the 2005 training data, and the *solid line* is the average of the MAE for 2004 and 2005. The corresponding *vertical line* types indicate the angle at which the minimum is obtained. *Red triangles, green diamonds, and blue circles* along the bottom indicate that the location chosen to switch the regimes for that particular angle are Vansycle, Kennewick, and Goodnoe Hills, respectively (color figure online)

will use the power law with coefficient of friction equal to $1/7$, and I scale all of the wind speeds to 80 m agl, a common utility-grade turbine hub height. I then used the following algorithm to build the RRSTD model:

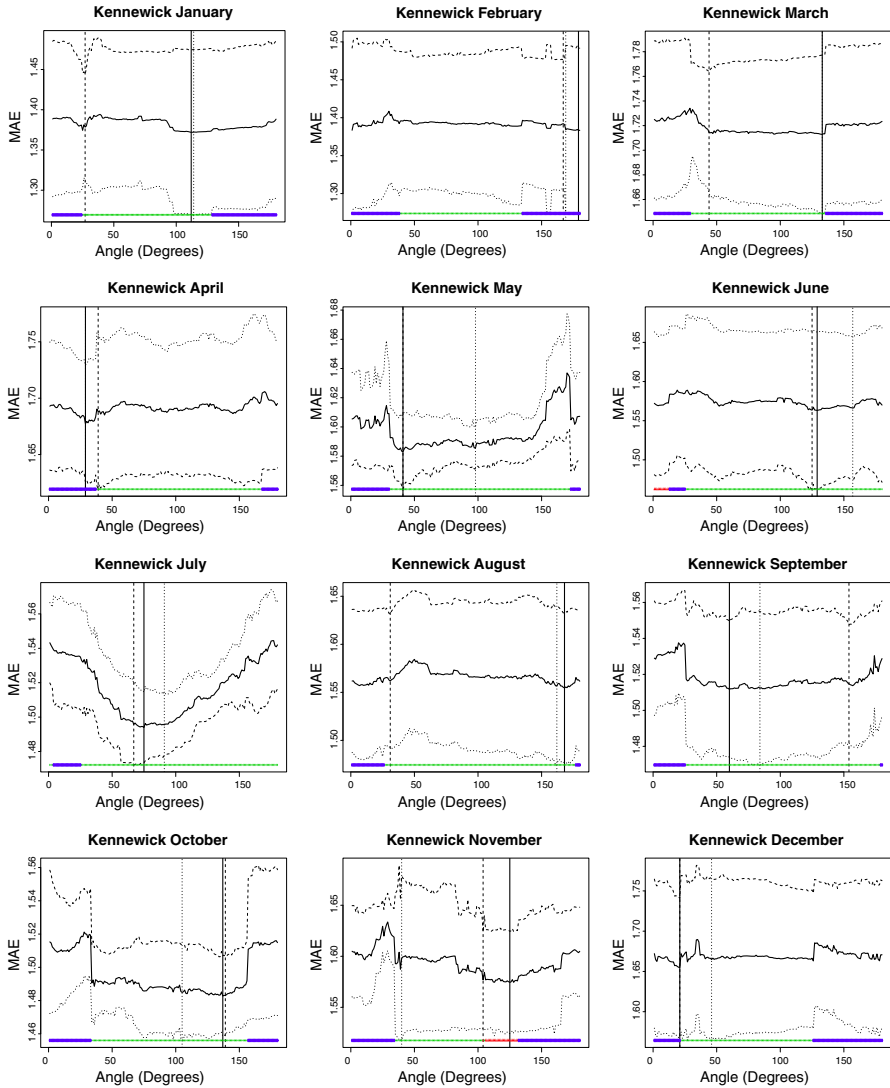


Fig. 5 Plots of 1-h ahead prediction MAE at Kennebec based on a two-regime RRSTD model. The *dashed line* is 2004 training data, the *dotted line* is the 2005 training data, and the *solid line* is the average of the MAE for 2004 and 2005. The corresponding *vertical line* types indicate the angle at which the minimum is obtained. *Red triangles, green diamonds, and blue circles* along the bottom indicate that the location chosen to switch the regimes for that particular angle are Vanscyle, Kennebec, and Goodnoe Hills, respectively (color figure online)

1. For a given month and location, choose a dividing angle, θ^* , from the set $\{1^\circ, 2^\circ, \dots, 180^\circ\}$. Note that I used the combined 2004 and 2005 data for a given month.
2. For θ^* , choose s^* , the wind source indicator by choosing the site such that the difference in median wind speeds in each regime 1 h prior is the largest.

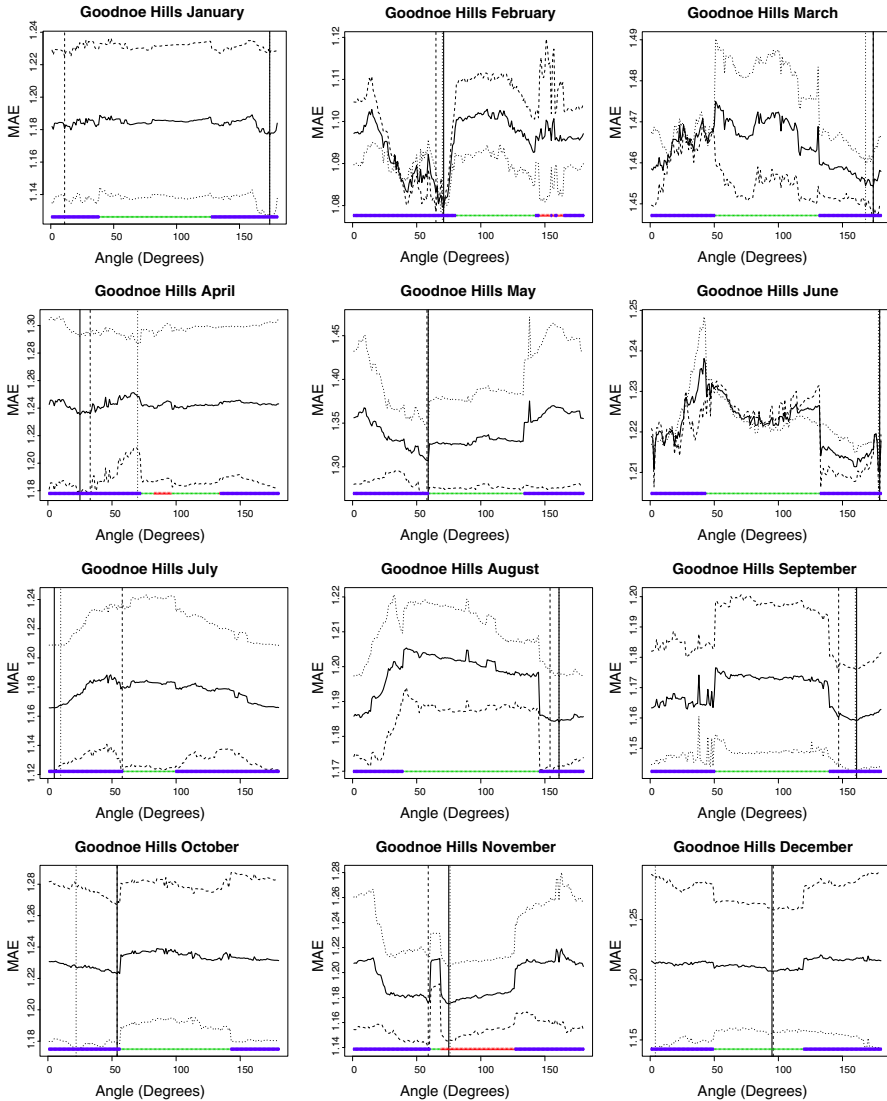


Fig. 6 Plots of 1-h ahead prediction MAE at Goodnoe Hills based on a two-regime RRSTD model. The *dashed line* is 2004 training data, the *dotted line* is the 2005 training data, and the *solid line* is the average of the MAE for 2004 and 2005. The corresponding *vertical line* types indicate the angle at which the minimum is obtained. *Red triangles*, *green diamonds*, and *blue circles* along the bottom indicate that the location chosen to switch the regimes for that particular angle are Vanscyle, Kennewick, and Goodnoe Hills, respectively (color figure online)

3. Choose the variables for the given values of θ^* and s^* in ZGGX's Eq. (2) using BIC.
4. Forecast as follows:
 - (a) Forecast the 2004 wind speeds using the parameter estimates based on the 2005 month.

- (b) Forecast the 2005 wind speeds using the parameter estimates based on the 2004 month.
- 4. Store the MAE values for both of 2004 and 2005.
- 5. Repeat steps 1–5 for each location, month, and value of θ^* .

Figures 4, 5 and 6 show the MAE by month at each of the three locations. Along the bottom of each plot, a character is plotted to indicate the wind source. If Vansycle is the wind source, a red triangle is plotted, and green diamonds and blue circles correspond to Kennewick and Goodnoe Hills, respectively. Both the MAE for 2004 (dashed line) and 2005 (dotted line) are plotted along with the average of the two (solid line). In some cases, such as Vansycle April or Goodnoe Hills December, the minimum of all three curves agree closely. In other cases, such as Kennewick November, the minimums occur at very different angles, but in most cases, at least one of the minimum yearly MAE values is close to the minimum of the averages across the 2 years. In all cases, the angle chosen to switch the regimes appears to be a reasonable one that splits the distribution of winds almost directly “in half.” For example, to forecast at Vansycle in January, the wind source is Kennewick at 126°. On comparing this with the January distribution of winds at Kennewick in Fig. 2, this angle cuts the two modes in half.

Finally, Table 1 gives the MAE of the RSTD, TDD, and RRSTD models in addition to the persistence forecast when applied to the 2006 testing data. Of these four models, the TDD does the best overall at Vansycle and ties with the RSTD model at Kennewick and Goodnoe Hills. The RRSTD model does not do as well overall as either of these two models, but for some months, such as January at Kennewick or November at Goodnoe Hills, it beats both models. In particular, it beats or ties with the RRSTD

Table 1 MAE by month and overall for persistence and each of the three forecasting models

Site	Model	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Overall
VS	PSS	1.83	1.54	1.75	1.50	1.57	1.42	1.44	1.38	1.28	1.30	1.87	1.46	1.53
	RSTD	1.76	1.52	1.68	1.42	1.48	1.27	1.33	1.27	1.19	1.21	1.81	1.38	1.44
	TDD	1.71	1.48	1.67	1.44	1.47	1.27	1.32	1.24	1.19	1.20	1.78	1.39	1.43
	RRSTD	1.75	1.51	1.69	1.44	1.48	1.27	1.33	1.27	1.19	1.22	1.82	1.41	1.45
	AVG	1.72	1.48	1.66	1.42	1.47	1.25	1.31	1.25	1.17	1.19	1.79	1.37	1.42
KW	PSS	1.95	1.69	1.67	1.62	1.77	1.58	1.67	1.57	1.42	1.49	1.99	1.48	1.66
	RSTD	1.88	1.65	1.63	1.56	1.69	1.53	1.58	1.50	1.34	1.47	1.97	1.46	1.60
	TDD	1.87	1.63	1.63	1.56	1.70	1.54	1.56	1.50	1.34	1.47	1.92	1.47	1.60
	RRSTD	1.86	1.63	1.63	1.57	1.72	1.53	1.57	1.52	1.40	1.48	1.96	1.47	1.61
	AVG	1.86	1.62	1.61	1.55	1.69	1.52	1.55	1.49	1.33	1.46	1.94	1.45	1.59
GH	PSS	1.56	1.39	1.35	1.26	1.29	1.16	1.33	1.17	1.06	1.21	1.50	1.23	1.29
	RSTD	1.50	1.36	1.32	1.23	1.26	1.11	1.28	1.14	1.03	1.17	1.45	1.19	1.25
	TDD	1.50	1.35	1.32	1.24	1.25	1.12	1.27	1.15	1.04	1.17	1.45	1.20	1.25
	RRSTD	1.50	1.36	1.32	1.23	1.28	1.12	1.27	1.15	1.04	1.20	1.44	1.19	1.26
	AVG	1.49	1.35	1.31	1.23	1.25	1.11	1.26	1.13	1.03	1.17	1.44	1.18	1.25

The “AVG” row is the forecast obtained by averaging the RSTD, TDD, and RRSTD forecasts
 The lowest value in each column is indicated in bold.

model in the more variable months of January and February at all three locations, and when tested at a location outside of the Columbia River Gorge, the benefit of using the RRSTD model may become even more pronounced.

4 Discussion

I would like to conclude by thanking the editors of *TEST* for calling attention to this important applied problem and for providing an opportunity to discuss the work of ZGGX and to add to the growing body of knowledge related to wind forecasting. One final point that I would like to make is that even though the RRSTD model did not perform as well overall as the RSTD or the TDD models on this particular dataset, I averaged the forecasts produced by the three models, indicated by the AVG row in Table 1, and this average had the best overall (and best in most months) MAE! By developing different models that may emphasize different characteristics of the atmospheric state, we may be able to do better jointly than any one model can do individually. Here, I have used equal weights for the average, but if all three models were run simultaneously, then the weights could be estimated independently on each 45-day training window, and the overall results could be even better.

References

- Burlando M, Antonelli M, Ratto CF (2008) Mesoscale wind climate analysis: identification of anemological regions and wind regimes. *Int J Climatol* 28:629–641
- Clifton A, Lundquist J (2012) Data clustering reveals climate impacts on local wind phenomena. *J Appl Meteorol Climatol* 51:1547–1557
- Gneiting T, Larson K, Westrick K, Genton MG, Aldrich E (2006) Calibrated probabilistic forecasting at the stateline wind energy center: the regime-switching space-time method. *J Am Stat Assoc* 101:968–979
- Hering AS, Genton MG (2010) Powering up with space-time wind forecasting. *J Am Stat Assoc* 105:92–104
- Kazor K, Hering AS (2014) Statistical identification of local and regional wind regimes (submitted)
- Kretschmar R, Eckert P, Cattani D (2004) Neural network classifiers for local wind prediction. *J Appl Meteorol* 43:727–738
- Marquis M, Wilczak J, Ahlstrom M, Sharp J, Stern A, Smith JC, Calvert S (2011) Forecasting the wind to reach significant penetration levels of wind energy. *Bull Am Meteorol Soc* 92:1159–1171
- Yoder M, Hering AS, Navidi WC, Larson K (2013) Short-term forecasting of categorical changes in wind power with Markov chain models. *Wind Energy* <http://onlinelibrary.wiley.com/doi/10.1002/we.1641/abstract>. Accessed 22 jan 2014
- Zhu X, Genton MG, Gu Y, Xie L (2014) Space-time wind speed forecasting for improved power system dispatch. *Test* (to appear)