A Fetch-Based Statistical Method to Bias Correct and Downscale Wind Speed over Unresolved Water Bodies

BRYAN P. HOLMAN AND STEVEN M. LAZARUS

Department of Ocean Engineering and Sciences, Florida Institute of Technology, Melbourne, Florida

MICHAEL E. SPLITT

College of Aeronautics, Florida Institute of Technology, Melbourne, Florida

(Manuscript received 16 February 2017, in final form 7 July 2017)

ABSTRACT

This paper presents a method to bias correct and downscale wind speed over water bodies that are unresolved by numerical weather prediction (NWP) models and analyses. The dependency of wind speeds over water bodies to fetch length is investigated as a predictor of model wind speed error. Because model bias is found to be related to the forecast wind direction, a statistical method that uses the forecast fetch to remove wind speed bias is developed and tested. The method estimates wind speed bias using recent forecast errors from similar stations (i.e., those with comparable fetch lengths). As a result, the bias correction is not tied to local observations but instead to locations with similar land–water characteristics. Thus, it can also be used to downscale wind fields over inland and coastal water bodies. The fetch method is compared to four reference bias correction methods using one year's worth of wind speed output from three NWP analyses in Florida. The fetch method yields a bias error near zero and results in a reduction of the mean absolute error that is comparable to the reference methods. The fetch method is then used to bias correct and downscale a coarse analysis to 500-m grid spacing over a coastal estuary in central Florida.

1. Introduction

Numerical weather prediction (NWP) models are essential tools for weather forecasting. Yet, current operational models are prone to systematic errors (biases) due to coarse model resolution and imperfections in model physics, initial conditions, and boundary conditions (Mass et al. 2002). One effective method of removing biases is statistical postprocessing (hereafter referred to as postprocessing) of model output. In general, NWP models cannot realize their full potential without postprocessing model output (Gneiting and Raftery 2005).

This study focuses on postprocessing methods for 10-m wind speed forecasts. Popular methods for this purpose include running-mean bias correction (e.g., Stensrud and Yussouf 2005; Cheng and Steenburgh 2007), Kalman filtering (e.g., Louka et al. 2008; Cassola and Burlando 2012), machine learning methods (e.g., Salcedo-Sanz et al. 2009), and model output statistics (MOS; Glahn and Lowry 1972; Carter 1975). Comparison studies have shown that MOS is the most successful method for reducing wind speed forecast error (Cheng and Steenburgh 2007; Müller 2011). However, MOS requires an extensive training dataset and a static model, motivating the development of more adaptable postprocessing techniques.

Traditional postprocessing methods, such as those just mentioned, are tied to locations with verifying observations. Yet bias-corrected forecasts are needed at locations without verifying observations (Mass et al. 2008). Yussouf and Stensrud (2006) used a Cressman weighting scheme (Cressman 1959) to interpolate wind speed biases onto a grid over Oklahoma. While distance-weighted interpolation schemes may be appropriate for regions with little topographical variability and marginal differences in land-use characteristics, it is problematic in regions of complex terrain. Because of complex interactions with topography, Engel and Ebert (2007) suggest that wind postprocessing methods need a combined physical and statistical approach.

De Rooy and Kok (2004) presented such an approach and applied it to downscale model wind speed over the

Corresponding author: Bryan P. Holman, bholman2013@my.fit.edu

DOI: 10.1175/WAF-D-17-0016.1

© 2017 American Meteorological Society. For information regarding reuse of this content and general copyright information, consult the AMS Copyright Policy (www.ametsoc.org/PUBSReuseLicenses).

Unauthenticated | Downloaded 12/28/22 07:05 PM UTC

Netherlands. They first identified and removed the systematic error due to the difference in surface roughness between the model grid cell and verifying observation locations. The residual bias was then attributed to large-scale model error. Using roughness length estimates derived from a 1-km land-use grid, they applied their physical/statistical method to wind speed forecasts. However, their approach did not address flow-dependent (anisotropic) bias near water boundaries, and thus it was not applied in the coastal zone.

Representative, high-resolution wind information in nearshore and coastal areas is essential for applications such as forcing hydrodynamic models (e.g., Weaver et al. 2016) and wind energy resource assessments (e.g., Al-Yahyai et al. 2010; Charabi et al. 2011). However, even dynamically downscaled winds are error prone at coastal sites and in regions of complex terrain (Mortensen et al. 2006; Gastón et al. 2008). Acquiring sufficient data for wind energy assessments through dynamical downscaling is time and resource intensive, motivating the development of statistical methods (Zhang et al. 2015).

Here, we introduce an inexpensive, physically based statistical method to bias correct model wind speed output at observation-free locations near and over water bodies. As such, the main advantage of our method is that it is ideally suited for use within a downscaling framework. It is developed based on the work of Mass et al. (2008), who estimated surface temperature biases at grid points using errors from physically similar observing locations (i.e., locations with similar elevation and land use). Here, in the context of wind speed, we show that forecast error over unresolved water bodies is, in part, fetch dependent (defined here as the distance wind blows over water in a constant direction without obstruction) and thus coupled to the forecast wind direction. We outline an approach that uses errors from stations with similar forecast fetch lengths to bias correct wind speed, and compare the results against popular local bias correction methods.

The remainder of this paper is outlined as follows. Section 2 describes the area of study and the model and verification data used. Section 3 outlines the fetch method and details four other postprocessing methods used for comparison. Section 4 presents the results, including a downscaling case study, and section 5 concludes with a discussion of the results and provides suggestions for future work.

2. Data

a. Study area

The Indian River Lagoon (IRL) in Florida is approximately 250 km long, has an average width of 2–4 km, and varies on the order of 1–3 m in depth. It is an estuarine system comprised of the Indian River, Banana River, and Mosquito Lagoon. The northern half of the IRL and nearby Cape Canaveral are shown in Fig. 1e. As a result of the IRL's narrowness and complex land/water mask, it is either poorly resolved or completely unresolved by current operational NWP models and analyses. This can result in modeled surface winds that are unrepresentative of winds observed over the IRL.

Wind stress is the primary forcing for circulation throughout most of the IRL (Smith 1990). Hydrodynamic models, such as the Advanced Circulation model (ADCIRC; Luettich and Westerink 2004), rely on accurate wind forcing from NWP models. In a recent IRL circulation study, Weaver et al. (2016) demonstrate the value in subkilometer wind forcing. In that study the authors downscaled GFS wind output with the Weather Research and Forecasting (WRF; Skamarock et al. 2008) Model. However, dynamical downscaling is not feasible for most operational users, especially within the ensemble framework.

b. Observations

High quality wind observations over the IRL, as in many other estuaries, are sparse. WeatherFlow Inc. (http://www.weatherflow.com) provides the most extensive network of quality controlled observations over the IRL. WeatherFlow operates three stations with open (i.e., over water) fetch greater than 1 km in all directions (XPAR, XCCB, and XJEN) and two additional stations on the IRL shoreline (XMER and XRPT).

To ensure significance in the training dataset, we expand our area of focus beyond the IRL to incorporate nine additional sites in four water basins in Florida that are appropriate for this study. The basins include Tampa Bay, Lake Okeechobee, the St. John's River, and Lake Tohopekaliga. Two stations are located in Tampa Bay, one from WeatherFlow Inc. (XTAM), and one from the National Ocean Service (NOS; MTBF1). Four stations, all from the South Florida Water Management District (SFWMD) are on Lake Okeechobee (L001, L005, L006, and LZ40). Two stations, from the NOS network, are located near the western shore of the St. John's River (GCVF1 and BKBF1). The SFWMD station S61W is located along the southern shore of Lake Tohopekaliga. The locations of the 14 stations and their respective water basins are shown in Fig. 1.

In cases where the anemometer height is not at 10 m, a 10-m wind speed equivalent U_{10} was calculated following Hsu et al. (1994):

$$U_{10} = U\left(\frac{10\,\mathrm{m}}{z}\right)^{0.11}, \tag{1}$$

Unauthenticated | Downloaded 12/28/22 07:05 PM UTC

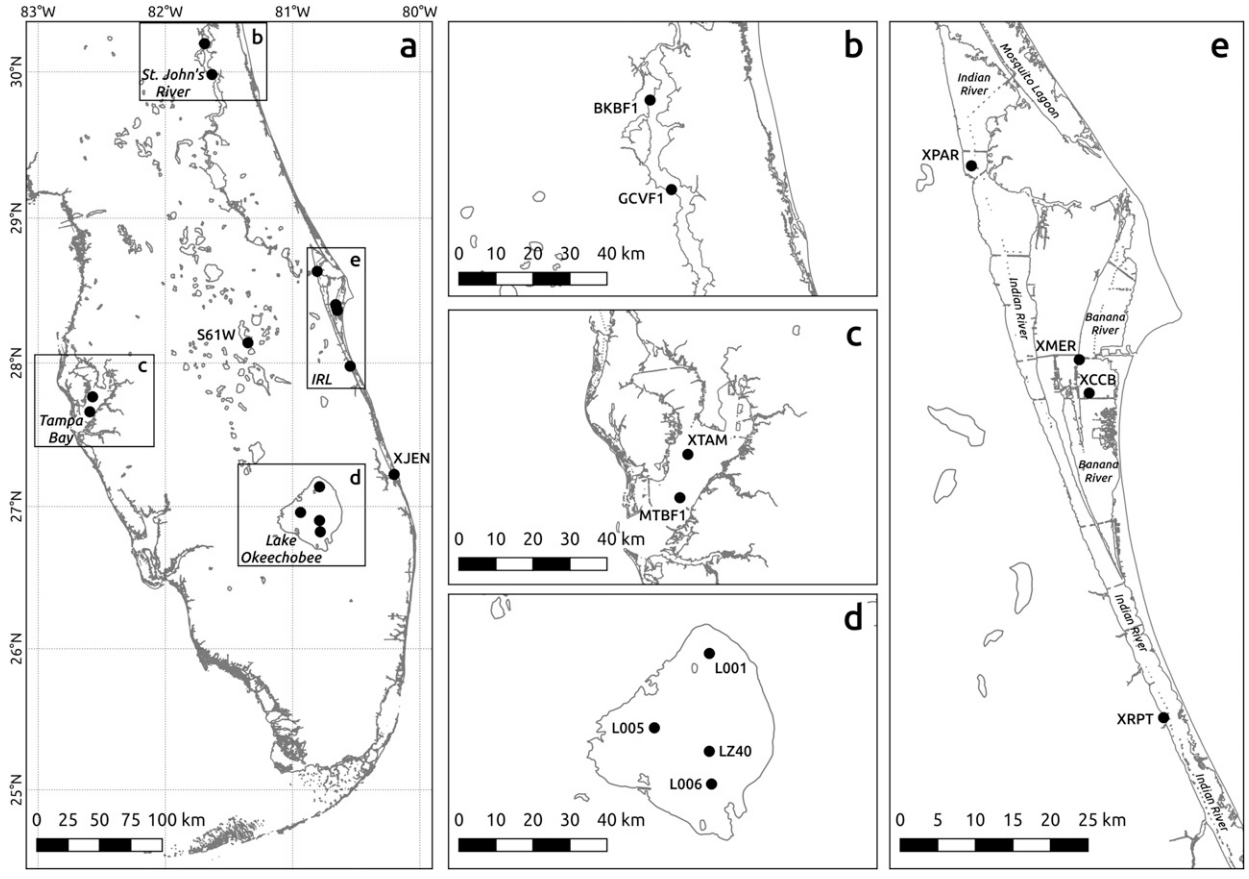

FIG. 1. The locations of the 14 surface stations (Table 1) used in this study. (a) A regional view of the FL peninsula, with enlargements of (b) the St. John's River, (c) Tampa Bay, (d) Lake Okeechobee, and (e) the IRL. Note that XJEN, which is south of the region shown in (e), is considered to be part of the IRL basin.

where U is the observed wind speed at the anemometer height z. The exponent in Eq. (1) was determined empirically by Hsu et al. (1994) to be a good fit for near-neutral conditions at sea. The near-neutral assumption is not appropriate for the surface layer over Florida for much of the year, but provides a sufficient zero-order adjustment to normalize the heights of the observations. Table 1 shows the anemometer heights for each station, as well as the location, network, basin, and data coverage during the study period (10 October 2014–10 October 2015). Most wind data provided by these stations are reported in intervals ranging from 5 to 15 min. Only the observed wind speeds closest in time (within a 10-min window) to the 3-hourly NWP model output were used.

The heights in Table 1 represent the height of the anemometer above water (or above land for stations located along the shoreline). Water level changes throughout the temporal coverage of this study were negligible except for the four stations over Lake Okeechobee (marked in Table 1). While the heights listed in Table 1 are the average heights throughout the study period, the daily water level at each Lake Okeechobee station was used to determine the anemometer height above water when calculating the 10-m equivalent wind speeds. For example, the average anemometer height for station L001 was 7.0 m, with a range from 6.3 to 7.8 m.

c. *Model output*

Three surface (10 m) wind datasets from the National Centers for Environmental Prediction (NCEP) were used in this study including model output from the Global Ensemble Forecast System (GEFS) and Short-Range Ensemble Forecast (SREF) Advanced Research version of WRF (ARW) control members, and the Real-Time Mesoscale Analysis (RTMA; De Pondeca et al. 2011). The GEFS produces global output every 6 h (0000, 0600, 1200, and 1800 UTC) on a 1.0° latitude–longitude grid. The SREF-ARW also generates output every 6 h, offset from the GEFS by 3 h (0300, 0900, 1500, and 2100 UTC). The SREF-ARW uses a Lambert conic conformal (LCC) grid that covers a majority of North America and the surrounding ocean with horizontal grid

Unauthenticated | Downloaded 12/28/22 07:05 PM UTC

TABLE 1. Metadata for the 14 stations used in this study, grouped by basin.

Station ID	Station name	Basin	Network	Location	Height (m)	Data coverage (%)
XPAR	Parrish Park North	IRL	WeatherFlow	28.63°N, 80.81°W	5.5	94.0
XMER	Banana River 528	IRL	WeatherFlow	28.40°N, 80.66°W	9.1	67.2
XRPT	Rocky Point	IRL	WeatherFlow	27.98°N, 80.55°W	6.1	93.9
XCCB	Banana River 520	IRL	WeatherFlow	38.36°N, 80.65°W	4.9	93.9
XJEN	Jensen Beach	IRL	WeatherFlow	27.22°N, 80.20°W	4.9	94.0
L001	Lake Okeechobee north end	Okeechobee	SFWMD	27.14°N, 80.79°W	7.0[a]	97.0
L005	Lake Okeechobee west end	Okeechobee	SFWMD	26.96°N, 80.94°W	7.1[a]	96.6
L006	Lake Okeechobee south end	Okeechobee	SFWMD	26.82°N, 80.78°W	7.1[a]	96.8
LZ40	Lake Okeechobee center	Okeechobee	SFWMD	26.90°N, 80.79°W	7.0[a]	96.8
GCVF1	Red Bay Point	St. John's	NOS	29.98°N, 81.63°W	8.5	97.8
BKBF1	I-295 bridge	St. John's	NOS	30.19°N, 81.69°W	9.7	98.2
MTBF1	Middle Tampa Bay	Tampa Bay	NOS	27.66°N, 82.59°W	6.7	97.9
XTAM	Tampa Bay Cut J	Tampa Bay	WeatherFlow	27.77°N, 82.57°W	14.6	94.0
S61W	Lake Tohopekaliga	Tohopekaliga	SFWMD	28.14°N, 81.35°W	15.8	96.4

[a] The height of the stations in Lake Okeechobee indicate an average anemometer height above water during the study period.

spacing near 40 km. The RTMA is an hourly analysis product with 2.5-km grid spacing over the continental United States. While only the control members of these ensemble systems are used here, in future work the fetch method will be used to downscale the full ensemble suite of the GEFS and SREF.

This study compares wind observations against both analyses and model forecasts with an emphasis on the former. We focus on analyses for two reasons. The first is to mitigate the impact of forecast errors and investigate the effect that NWP grid spacing has on the fetch method. To accomplish this, we apply the fetch method to the 2.5-km RTMA as well as the coarser SREF-ARW and GEFS analyses. Second, NWP model analyses have many applications, including model verification and initial (IC) and boundary condition (BC) forcing, etc., for dynamical downscaling (e.g., Davis et al. 2008; Carvalho et al. 2014; Li et al. 2012). For example, hindcasting studies of waves and circulation often use wind forcing from downscaled analyses (e.g., Weaver et al. 2016; Chen and Curcic 2016). In particular, we want to determine if the fetch method can successfully downscale coarse analysis wind speed fields to provide inexpensive yet representative ensemble wind forcing for future hindcasting studies as well as input to simple wind–wave parameterizations.

A major source of error in 10-m wind speed forecasts is related to roughness differences between model grid cells and observation stations (Strassberg et al. 2008). This is particularly true when a model resolves the location of an overwater observation station as land, or vice versa. Figures 2a–c show the land masks of the GEFS, SREF-ARW, and RTMA, respectively. Despite its size (~50 km wide), Lake Okeechobee is resolved as land by both the GEFS and SREF-ARW. The St. John's River is also fully resolved as land by both the GEFS and SREF-ARW, while the RTMA resolves only the widest sections as water. The station locations in Tampa Bay are represented as water by both the GEFS and SREF-ARW, although portions of the bay itself are resolved as land. The northern IRL is resolved as land by the GEFS, while the southern portion of the IRL is resolved as water. The SREF-ARW resolves only one of the IRL stations as land, but all the IRL stations are on the edge of the SREF-ARW land mask. The RTMA provides the most representative land mask for all of the water basins studied here, but the shorelines of the Tampa Bay, St. John's River, and the IRL are still too complex for the 2.5-km grid spacing. Figure 2c shows a close up of the RTMA land mask over Cape Canaveral in the northern IRL. XPAR, while sited over water and more than a kilometer away from the IRL shoreline, is on the edge of the RTMA land mask. Additionally, XMER, while on the IRL shoreline, is resolved approximately 700 m inland on the RTMA grid.

The standard practice when postprocessing model output is to interpolate to the location of the verifying observations. The impact that interpolation has on postprocessing over unresolved water bodies is unclear. Because of this, postprocessing was performed using output from the closest grid cell to each verifying station as well as that interpolated from the four adjacent grid cells using inverse distance weighting. This results in six distinct wind fields: both gridded and interpolated winds from two models and an analysis.

3. Methods

Our hypothesis is that one source of systematic wind speed error over small bodies of water is that they are poorly resolved by NWP models. The fetch method applied here attempts to determine systematic wind

Unauthenticated | Downloaded 12/28/22 07:05 PM UTC

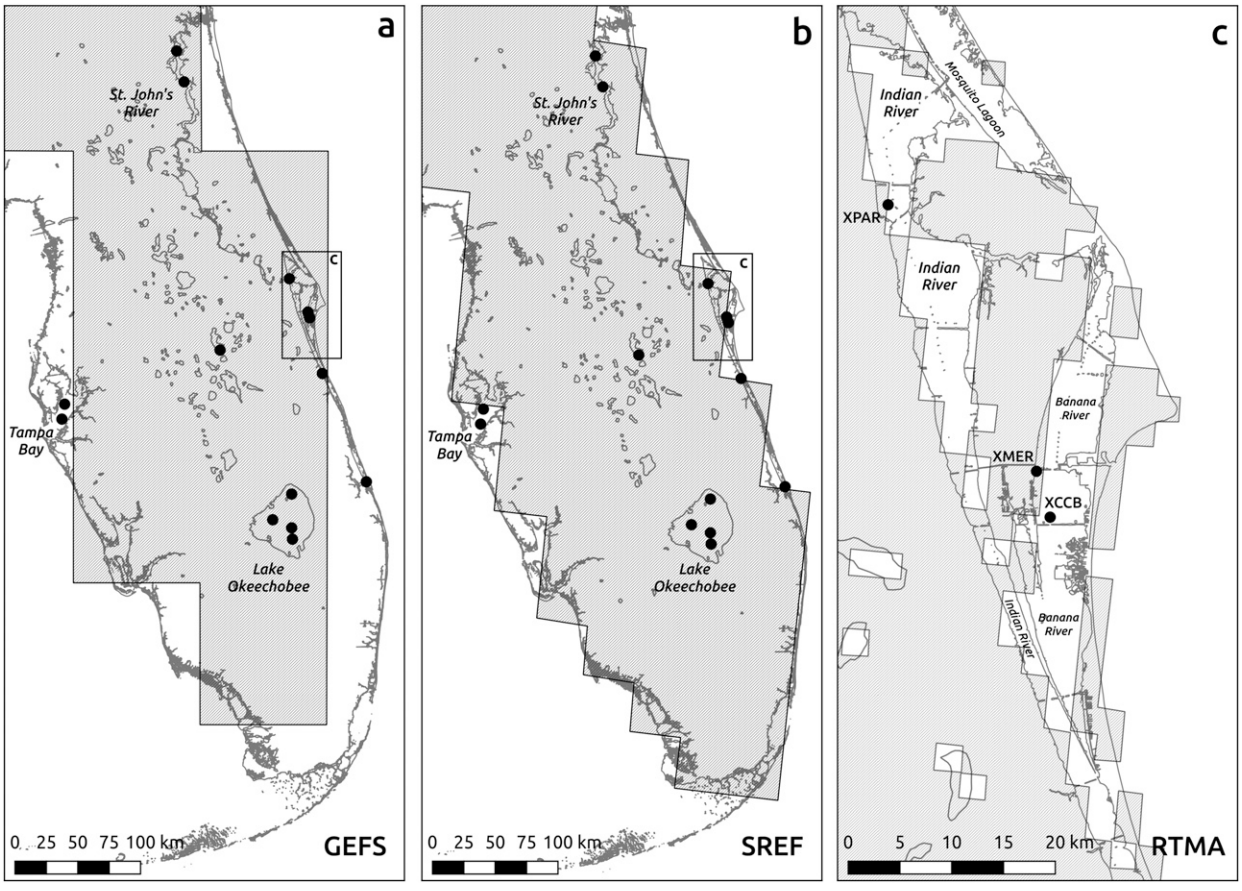

FIG. 2. The land masks for the (a) GEFS, (b) SREF-ARW, and (c) RTMA. The gray shading delineates land for each of the respective models. The 2.5-km RTMA land mask is much more representative of the FL peninsula as a whole, so only a close-up of Cape Canaveral is shown.

speed error in NWP model output using error estimates from similar stations. If successful, the technique could be used to provide downscaled, gridded, bias-corrected wind speeds over coastal estuaries, lakes, etc. In this section we describe the fetch approach. We also describe four known bias correction methods for wind speed, to which the fetch method will be compared.

Systematic error (bias) is determined by comparing NWP output to verifying observations within a predefined period, hereinafter referred to as the training data. The training dataset consists of analysis/observation winds valid at the time of each analysis cycle. The GEFS and SREF-ARW models, which both cycle four times per day, are offset by 3 h. To be consistent with the GEFS and SREF-ARW, the hourly RTMA winds were only verified every 3 h, providing eight analysis/observation pairs per day.

When implementing bias correction methods, it is important to determine the optimal number of analysis/observation pairs to include in the training dataset. One way to optimize the process is by varying the number of days included in the training dataset (i.e., the window length). When window lengths are large, more data are available to establish robust models of systematic error. Conversely, window lengths should be short enough to respond to model and weather regime changes. For example, model biases change with the seasons (Cheng and Steenburgh 2007). In this study, we tested window lengths from 10 to 60 days in increments of 5 days throughout the year-long study period. We found that a window length of 30 days resulted in the lowest mean absolute error (MAE) and, thus, was used for each bias correction method.

In the following section, we discuss the bias correction methods used in this study. Each method is applied to 3-hourly model output for a 1-yr period beginning 10 October 2014. The biases for each method are determined, independently, for each model cycle.

a. The fetch method

The objective of this method is to provide gridded, downscaled, bias-corrected wind speeds over unresolved

Unauthenticated | Downloaded 12/28/22 07:05 PM UTC

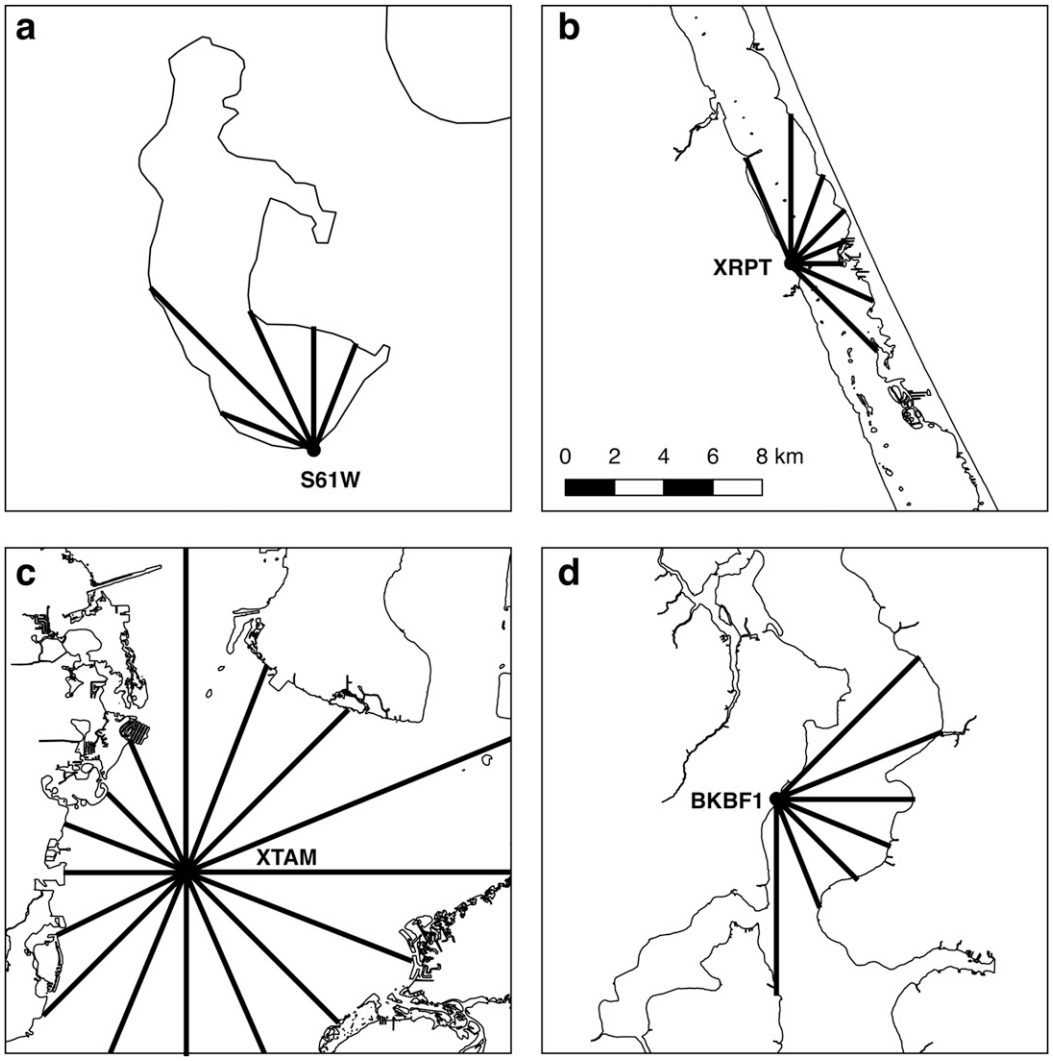

FIG. 3. Fetch rays for the (a) S61W, (b) XRPT, (c) XTAM, and (d) BKBF1 stations. The scale is the same for all four panels. Although stations XRPT and BKBF1 are "shoreline" locations, they are both located at the ends of piers, and thus have nominal downwind nearshore fetches on the order of 100 m.

water bodies. Because most grid cells lack observations, a traditional local bias correction approach (e.g., MOS) is not possible. In contrast, our method builds training datasets using recent errors obtained from similar observing stations—that is, regional observations that are also near or over unresolved water bodies.

Before applying the method, the fetch at each station was determined using QGIS, an open source desktop geospatial software package (QGIS Development Team 2016), for 16 directions (every 22.5°), starting at due north (see Fig. 3 for examples). Short fetches (less than 1 km) are accurate to 25 m while longer fetches were determined to 100-m accuracy. Locations at the ends of piers (such as XRPT) were considered to be open water and thus an offshore wind would have a fetch equal to the length of the pier.

The steps of the fetch method are relatively straightforward. For a given over-water location, we begin by constructing a training dataset composed of recent errors (within 30 days) from similar observing stations (the 14 stations presented in section 2b). The analysis/observation pairs are stratified by analysis cycle. For each of these pairs, we determine the fetch using the analysis wind direction (in 22.5° bins) and corresponding fetch lengths. For the stations and directions used here, 51% of fetches are below 5 km, 30% are between 5 and 15 km, and the remaining 19% are greater than 15 km (Fig. 4). Shorter fetches (<5 km) occur mainly at the IRL and St. John's stations, while

Unauthenticated | Downloaded 12/28/22 07:05 PM UTC

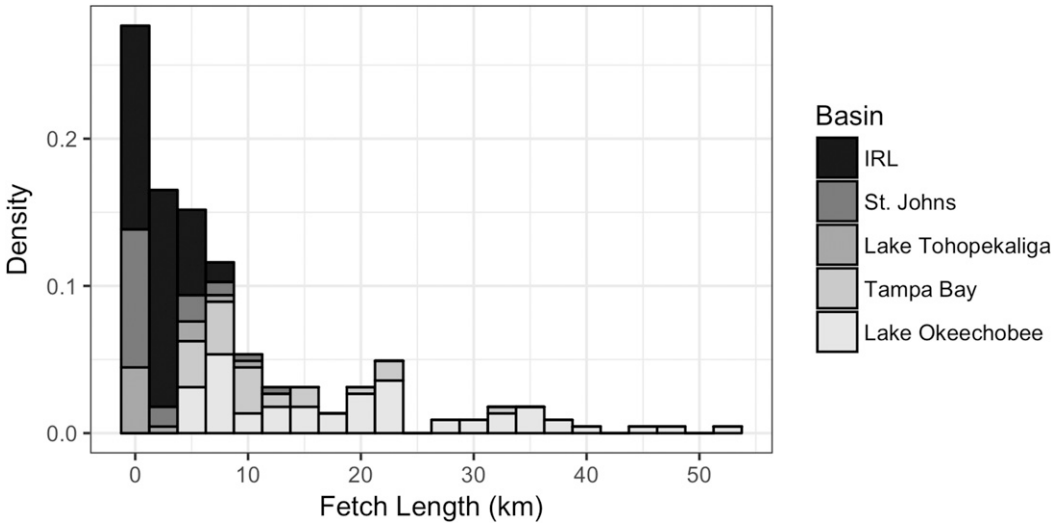

FIG. 4. Histogram of the 224 possible fetch lengths (14 stations, 16 directions), binned every 2.5 km and colored by basin.

longer fetches (>15 km) occur at the Tampa Bay and Lake Okeechobee stations.

To estimate a representative bias at a location within the analysis grid, the training data are trimmed so that they contain only relevant pairs with fetch lengths consistent with the given analysis fetch at our point of interest. This is done by sorting the training data by fetch length and retaining 20% of the pairs that are nearest to the given analysis fetch. However, for fetch lengths near zero (50 km), only 10% of the pairs with the smallest (largest) fetches are preserved in the final training dataset. With 14 available stations and a window length of 30 days this refined dataset contained, on average, 72 pairs. Preserving a percentage, as opposed to all pairs within a given fetch length window (e.g., ±5 km), is advantageous since shorter fetches (<10 km) are much more common than longer fetches at these stations (Fig. 4). We also tested retaining 10% and 30% of the available training data, but found retaining 20% to reduce MAE the most (not shown).

The average wind speed error (bias), for the given analysis fetch length, is then calculated assuming that each pair of the refined dataset has equal weight. The bias is then removed from the analysis wind speed by subtraction. Thus, in contrast to traditional data assimilation, interpolation, and statistical downscaling approaches, a station's proximity does not directly impact the weight given to its observations here. Instead, fetch length, rather than distance, is a more meaningful predictor of wind speed error for these water locations.

As an example, the fetch method is applied at station S61W using the 30 July 2015 interpolated GEFS analyses. The 30-day (prior) training dataset for each analysis cycle is shown in Fig. 5. The scatterplots were generated using interpolated GEFS analysis/observation pairs from 13 stations (S61W was not included in the training data). The exclusion of S61W's data is critical here, since the fetch approach is designed for use at observation-free locations. The black lines depict the bias estimate as a function of analysis fetch using the trimmed training set methodology described earlier. For the 1200 UTC analysis (Fig. 5b), the wind direction at S61W is out of the SSW, which corresponds to a fetch of 0 km and a bias estimate of $+1.12\,\mathrm{m\,s^{-1}}$. Subtracting this from the analysis wind speed ($1.99\,\mathrm{m\,s^{-1}}$) yields a bias-corrected value of $0.87\,\mathrm{m\,s^{-1}}$. For this and the other methods used in this study, the bias-corrected wind speed was set to zero if the correction produced a negative wind speed. This occurred 0.35% of the time.

Figure 5 indicates that there is a relationship, albeit noisy, between wind speed error and analysis fetch. Negative (positive) wind speed error indicates that the analysis wind speed is less (greater) than the observed. As the fetch approaches zero, the analysis wind speeds are biased high (positive wind speed error) for the 0000, 0600, and 1200 UTC cycles. As the fetch increases, the wind speed error quickly becomes negative, while asymptotically approaching values on the order of $-2\,\mathrm{m\,s^{-1}}$. For the most part, the positive bias is absent in the 1800 UTC training dataset (Fig. 5c), and the wind speed error is relatively constant (near $-2\,\mathrm{m\,s^{-1}}$) for all fetch lengths. The same data are shown in Figs. 4e–h but the x axis has been transformed to a natural log-based scale. In these plots the bias estimates (black lines) resemble negatively oriented trend lines, with the exception of the 1800 UTC analysis cycle, which has a near-zero slope. Plots of wind

Unauthenticated | Downloaded 12/28/22 07:05 PM UTC

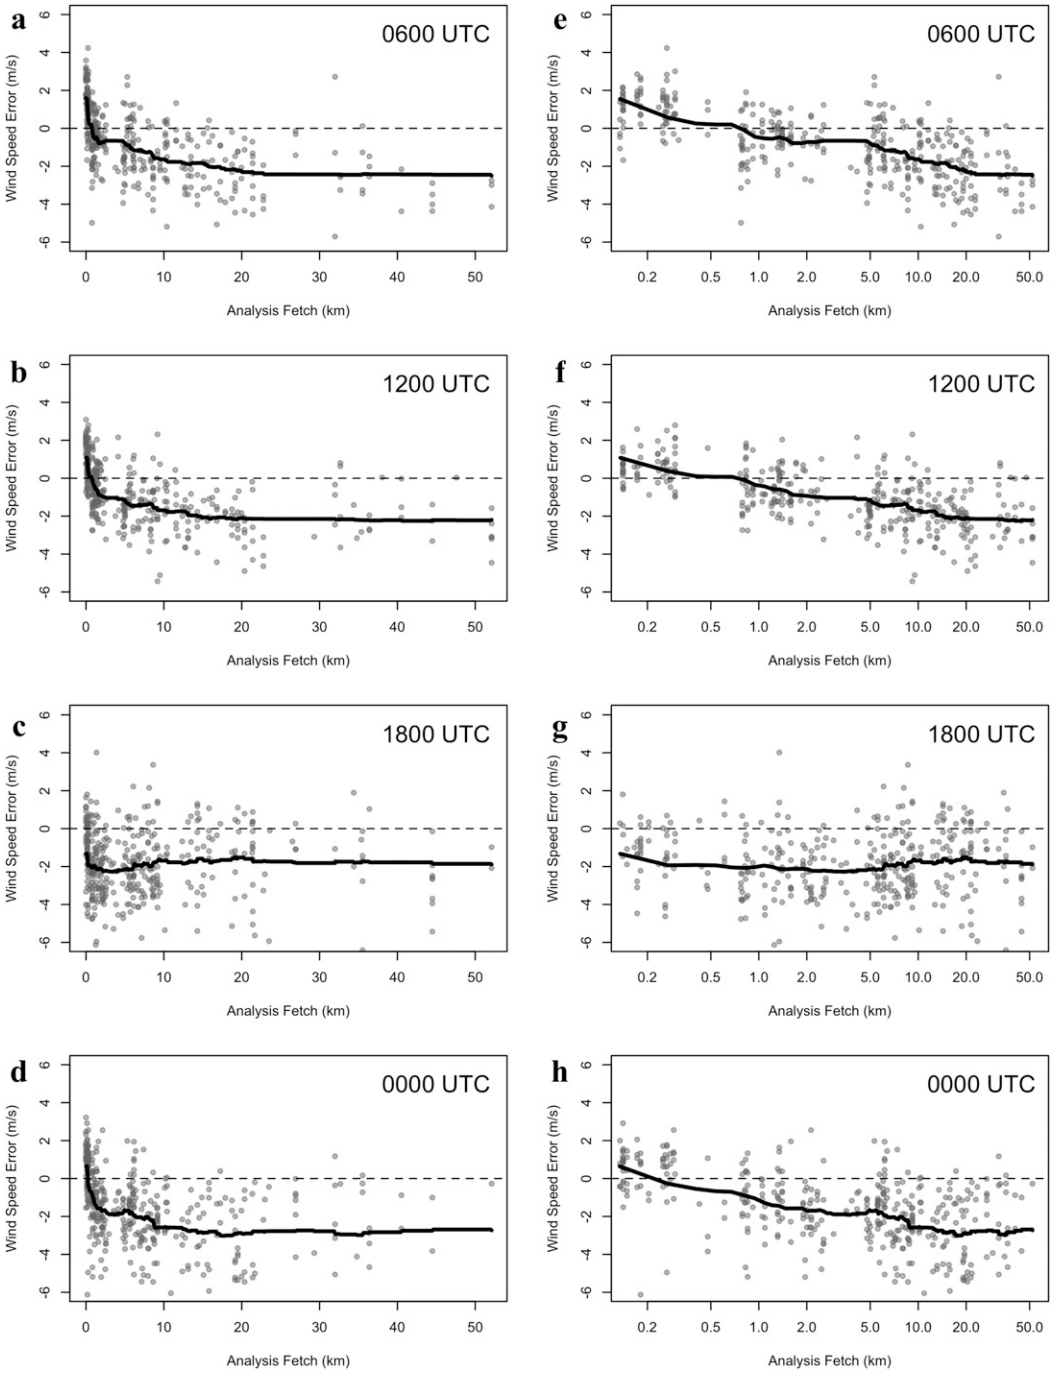

FIG. 5. Training dataset (analysis/observation pairs) composited from a 30-day prior window (to 30 Jul 2015) and 13 stations (S61W was not used). The data are stratified by the GEFS forecast cycles at (a) 0600, (b) 1200, (c) 1800, and (d) 0000 UTC. (e)–(h) As in (a)–(d), but with a natural log-scale x axis. The plots are arranged (from top to bottom) to correspond with the local diurnal cycle, i.e., middle of the night (0600 UTC), morning (1200 UTC), early afternoon (1800 UTC), and early evening (0000 UTC). The solid black line represents the bias estimate given the fetch length (see text for details).

Unauthenticated | Downloaded 12/28/22 07:05 PM UTC

speed error versus fetch length for each month (not shown) show that these diurnal error trends are largely present year-round. As such, the 1800 UTC anomaly is discussed in more detail in section 5.

A total of eight different fetch estimation methods were tested using both the analysis and observed wind directions. Because there are bridges in both the IRL and along the western edge of Tampa Bay, as well as spoil islands in the former, we investigate whether or not the fetch should be limited by these features. To test this, we determined fetch lengths that both stopped and continued at bridges and spoil islands. Also, the effective fetch (see U.S. Army Coastal Engineering Research Center 1977; Keddy 1982), which uses weighted averaging of nearby fetches to account for complex shorelines, was also calculated. Thus, the fetch method was tested using both effective and direct fetch, with and without bridges and spoil islands, and for both observed and analysis winds. While MAE differences were small (i.e., less than $0.03\,\mathrm{m\,s^{-1}}$ for gridded GEFS output) between the eight scenarios, the obstacle-limited direct fetch, calculated using the analysis wind direction, resulted in the lowest MAE. As a result, this fetch estimate was used for the remainder of the paper.

b. Comparative bias correction methods

To evaluate the fetch method, we apply it to the three analyses at the 14 station locations (Table 1), using a standard data withdrawal methodology that exempts the verifying location from the training dataset. Using the observed wind speed, we then compare the fetch approach, at each station, against four popular bias correction methods.

1) RUNNING MEAN

The running mean bias correction method calculates the mean wind speed error over the previous 30 days. Given m analysis/observation pairs within the 30-day window, the running mean bias is defined as

$$\mathrm{bias}_{\mathrm{RM}} = \frac{1}{m}\sum_{i=1}^{m}(a_i - o_i), \tag{2}$$

where a_i and o_i are the analysis and observation, respectively. The bias-corrected analysis is calculated by subtracting the average bias from the current analysis:

$$a_{\mathrm{RM}} = a - \mathrm{bias}_{\mathrm{RM}}. \tag{3}$$

2) HOURLY RUNNING MEAN

Past studies have shown that wind speed error varies with the time of day (e.g., Cheng and Steenburgh 2007; Sweeney et al. 2013). This is true of the stations used in this study. as illustrated in Fig. 6a, which shows the observed versus analysis wind speed error at station XRPT for the interpolated GEFS (circles) and SREF-ARW (squares) as a function of forecast cycle, as well as the RTMA (triangles). The error is averaged over the study period. The analysis wind speeds are biased low with the exception of the 0600 and 1200 UTC GEFS. The magnitude of the wind speed error is largest during the afternoon and evening. To account for this diurnal-based error, an additional running mean bias correction method [i.e., Eq. (3)] is introduced whereby training datasets are separated by forecast cycle; for example, 0000 UTC training data are only used for bias correcting 0000 UTC analyses.

3) LINEAR REGRESSION

The running mean methods can only identify and remove additive errors. However, it has been shown (e.g., Engel and Ebert 2007) that NWP models tend to under- (over-) forecast strong (weak) winds. In general, we found this to be the case for the stations used in this study. For example, a scatterplot (Fig. 6b) of observed versus interpolated GEFS analysis wind speeds at XRPT (for the study period) and corresponding linear regression indicates that slope error is present and thus should be taken into account.

The method is developed by regressing NWP model output (predictor) versus observed wind speeds (predictand). The linear model analysis wind speed a_{LM} is given by

$$a_{\mathrm{LM}} = \beta_0 + \beta_1 a, \tag{4}$$

where a represents the current analysis wind speed and β_0 and β_1 are regression coefficients found by minimizing the sum-squared error between the model output and observations in the training dataset (Wilks 2011). As in the running mean-hour method, separate training datasets are constructed for each analysis cycle.

4) WIND DIRECTION

The observation stations used in this study are either directly over water or along the shoreline. Three of the water basins are narrow (except Lake Okeechobee), and thus stations located in these regions will be highly sensitive to the upwind surface roughness, resulting in directionally varying wind speed biases. Figure 6c shows the average wind speed error versus wind direction (at XRPT) for the GEFS, SREF-ARW, and RTMA analyses. The analysis wind speed error is negative (biased low) for flow ranging (clockwise) from northwesterly to southerly, where water is upwind from XRPT (see Fig. 3b). Conversely, the bias is

Unauthenticated | Downloaded 12/28/22 07:05 PM UTC

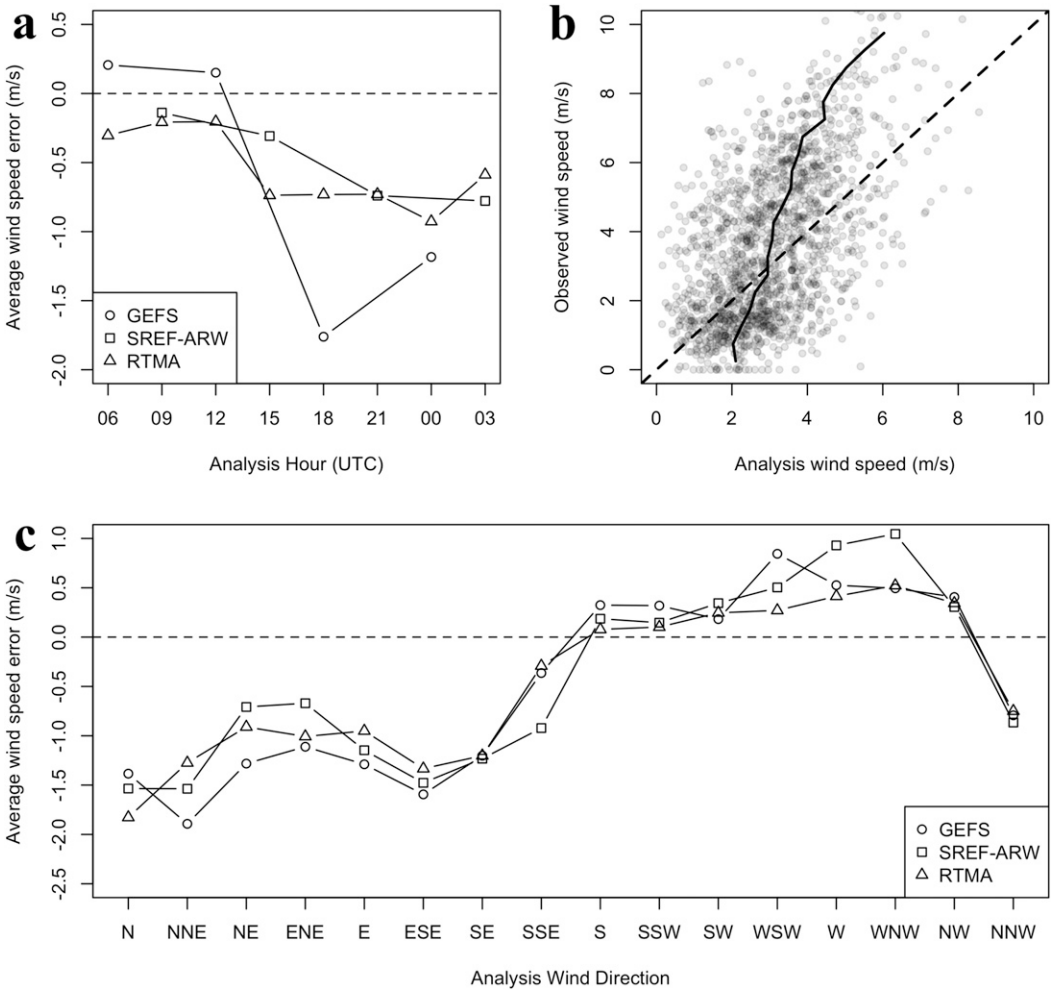

FIG. 6. Analysis bias at station XRPT for the 1-yr study period (see text). (a) Average bias vs analysis cycle for the GEFS, SREF-ARW, and RTMA. (b) Scatterplot of observations vs interpolated GEFS analysis wind speeds (the dashed line represents a perfect fit while the solid line is the average analysis wind speed given the observed wind speed). (c) As in (a), but as a function of analysis wind direction.

positive for southwesterly-to-northwesterly flow, with land upstream. Thus, a direction-based approach could isolate and correct for these anisotropic wind speed errors.

Sweeney et al. (2013) proposed a method that limits training pairs to those that ''match'' the current analysis wind direction. Here, we define a matching pair if wind direction error is less than 30°. Tests were performed using a range of matching wind directions from 20° to 45°, but we found 30° to reduce MAE the most. In the case of an infrequent wind direction, there could be insufficient analysis/observation pairs to identify a meaningful direction-based bias. In this case, the direction-based and running mean methods are combined and weighted according to a user-specified number of matching pairs in the training dataset (set to 20 here). Thus, if there were 0, 10, and 20 analysis/observations pairs in the training dataset that matched the current analysis wind direction, the running mean method would receive full, half, and zero weight, respectively, and the direction-based method would receive zero, half, and full weight, respectively. We tested three additional observation thresholds (5, 10, and 15) and found that 20 resulted in the lowest MAE for the method.

We also tested restricting the analysis/observation training pairs to those where both the analysis and observed wind direction matched the current analysis wind direction. This method was less effective at reducing MAE, probably because of an insufficient number of matching analysis/observation pairs. It has been noted (e.g., Cheng and Steenburgh 2005; Bao et al. 2010) that wind direction observations are not reliable for light winds. To account for this, an implementation of

Unauthenticated | Downloaded 12/28/22 07:05 PM UTC

the wind direction method was applied where the training dataset was restricted to pairs with observed winds greater than 2.5 m s^{-1}. However, this resulted in slightly higher MAE values, so this restriction was not imposed.

c. Verification

The bias correction methods are evaluated using the mean absolute error [Wilks (2011), Eq. (8.29)]:

$$\mathrm{MAE} = \frac{1}{n}\sum_{i=1}^{n}\left|a_i - o_i\right|, \tag{5}$$

where n is the product of all stations and all run cycles. The MAE is therefore the average absolute error of the analysis wind speed at the 14 verifying stations. The MAE was calculated for each station, over the 1-yr period, and then averaged across all stations for each method. Similarly, the root-mean-square error (RMSE) was also used to evaluate the bias correction methods, but those results were nearly identical to the MAE analysis and are not included here.

The mean bias error [BE; Wilks (2011), Eq. (8.32)] was similarly calculated:

$$\mathrm{BE} = \frac{1}{n}\sum_{i=1}^{n}(a_i - o_i). \tag{6}$$

The mean bias error also reflects the average wind speed bias across the entire dataset. For a bias correction method to be effective, its output should have a BE near zero.

4. Results

a. Results with analysis wind fields

A Cleveland dot plot (Cleveland 1984) of the MAE and BE associated with the raw model and bias-corrected wind speed for each of the six model/analysis outputs discussed in section 2c is shown in Fig. 7. For each of the analysis outputs, the bias correction methods are sorted in order of decreasing average MAE. Error bars represent 95% confidence bounds estimated by bootstrapping the set of wind speed errors for each bias correction method and analysis product. For the raw data, the largest error, using the gridded (i.e., nearest neighbor) output, is in the coarsest of the three products, the GEFS (MAE of 1.84 m s^{-1}). Interpolation only marginally reduces the error (1.76 m s^{-1}). The opposite is true of the SREF-ARW where the MAE of the interpolated output is higher than the nearest neighbor. The MAEs for the two RTMA outputs are almost identical. Thus, interpolation does not necessarily reduce the MAE, and its impact is mitigated as the horizontal grid spacing decreases (see section 5). The bias-corrected output behaves similarly (e.g., linear model applied to the GEFS), with lower MAE when applied to interpolated output (1.20 m s^{-1}) compared to the nearest neighbor (1.25 m s^{-1}). For the RTMA, interpolation results in little or no impact on the MAE for any of the bias correction methods.

Whether applied to the interpolated or nearest-neighbor analysis, each of the bias correction methods reduces the MAE by at least 18%. In terms of MAE reduction, the most effective local methods are the linear model for the GEFS (32.0%) and SREF-ARW (25.7%) and the wind direction method for the RTMA (25.2%). Regardless of which bias correction method is applied to the interpolated GEFS output, the resulting MAE is lower than the MAE from the raw interpolated RTMA output. This finding is consistent with results from other studies that show that postprocessing is, in general, more effective at reducing model error than increasing model resolution (e.g., Louka et al. 2008; Müller 2011).

The MAE dot plot (Fig. 7, left) also shows that the fetch method is comparable to the other methods. While this method has higher error than the other bias correction methods when using the nearest-neighbor approach, it yields lower MAEs when compared with the interpolated GEFS running mean (all hours) method, and lower MAEs than both the running mean (all hours) and wind direction methods applied to the interpolated SREF-ARW. For the RTMA, the fetch approach results in the highest MAEs of the bias correction methods, although the decrease in MAEs is comparable with the results of the other methods. The shortcomings of the fetch method here are likely due to the RTMA's improved ability (compared with the GEFS and SREF-ARW) to resolve these inland water bodies (see Fig. 2).

The right panels in Fig. 7 show that the average bias error (across all stations and times) for each raw output is negative. This does not hold true for all stations. For example, the shoreline station XMER, which is sited between a building and trees, had a positive average bias throughout the study period (not shown). The bias associated with raw gridded GEFS output is the most negative, near −1.5 m s^{-1}. Interpolation alone decreases the bias to −1.24 m s^{-1}, while interpolating the SREF-ARW output results in a more negative bias. Interpolation has very little impact on the RTMA bias. More importantly, each of the bias correction methods reduces the error to near zero in all versions of the analyses.

Other studies (e.g., Cheng and Steenburgh 2007) have shown that the efficacy of bias correction can vary seasonally. To investigate this, a 60-day center-average

Unauthenticated | Downloaded 12/28/22 07:05 PM UTC

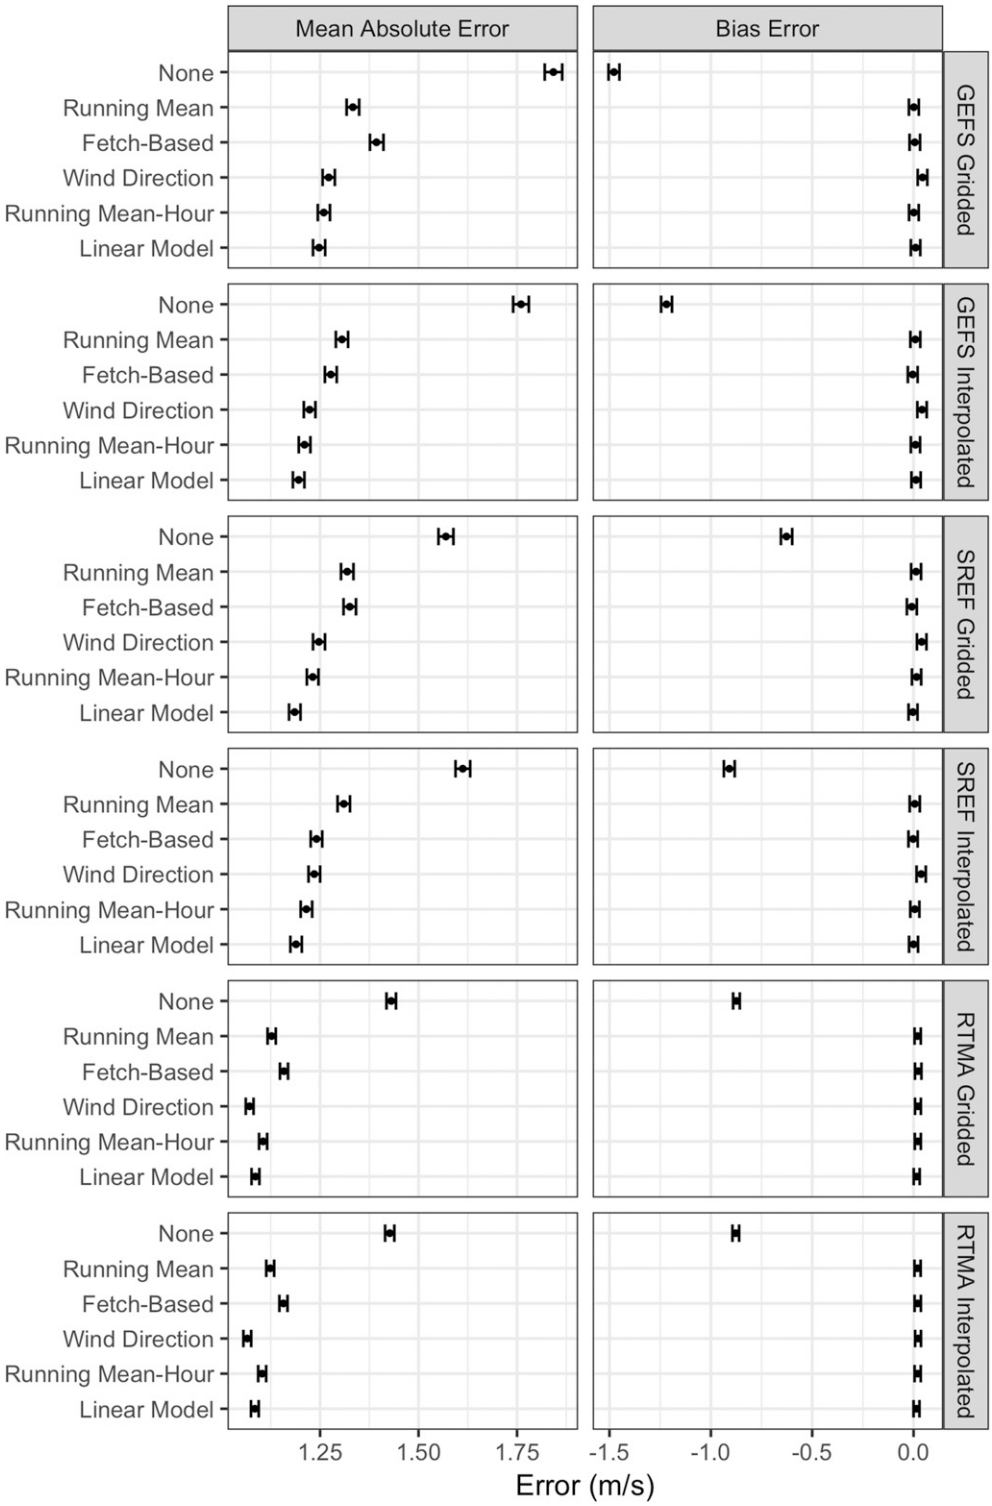

FIG. 7. Cleveland dot plot of the 14-station average (left) MAE and (right) BE during the 1-yr study period. Results are shown for each bias correction method and analysis (gridded and interpolated). Error bars represent 95% confidence intervals of respective group means determined by bootstrapping. See text for details.

Unauthenticated | Downloaded 12/28/22 07:05 PM UTC

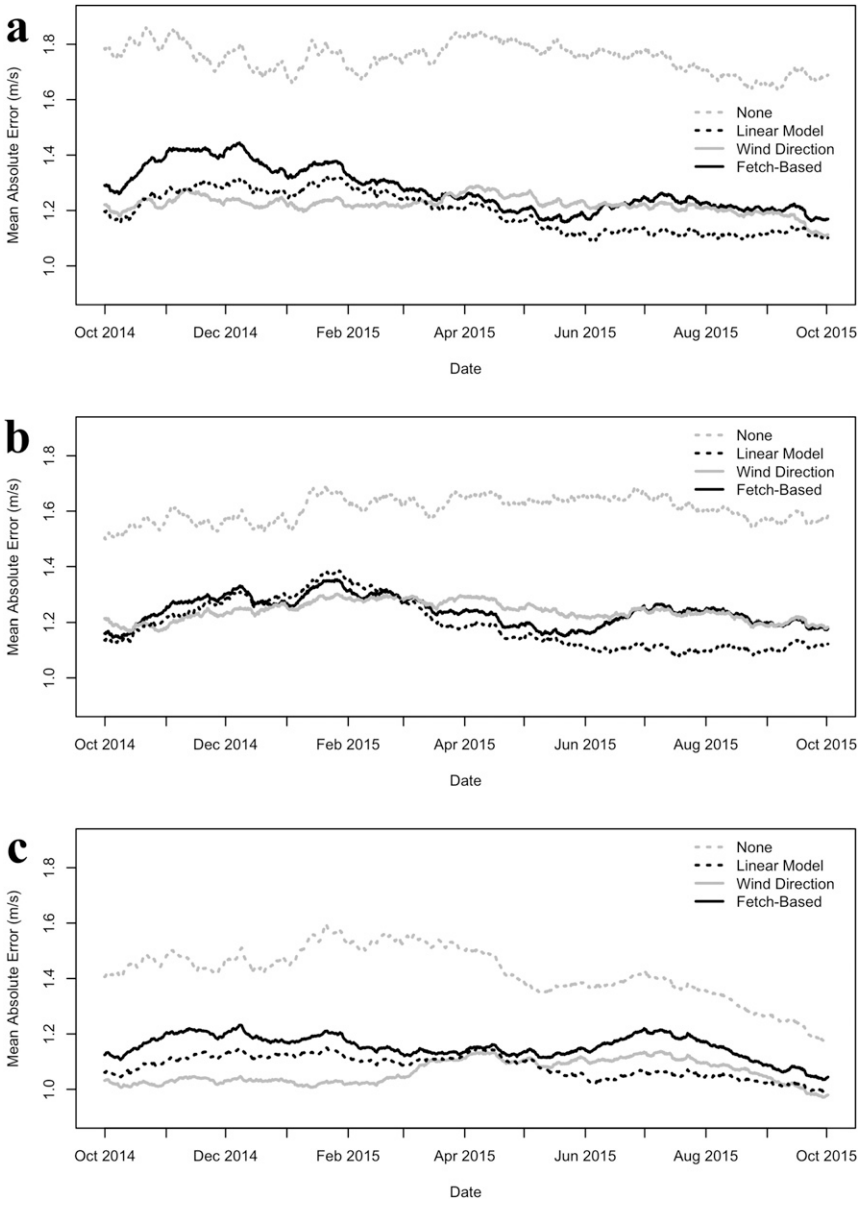

FIG. 8. Time series of the 60-day centered-mean average MAE from the raw interpolated analyses [i.e., no bias correction (none); dashed gray lines], and from the bias-corrected interpolated analyses using the linear model (dashed black line), wind direction method (solid gray line), and fetch method (solid black line). Results are shown for (a) GEFS, (b) SREF-ARW, and (c) RTMA.

MAE for each bias correction method and each analysis was calculated. The data window was extended to include 60 additional days, from 10 September 2014 through 10 November 2015, so that a full year was still available after the centered averages were calculated. Figures 8a–c show the 60-day center-averaged MAEs for the interpolated GEFS, SREF-ARW, and RTMA, respectively. For brevity, the raw analysis data (dashed gray lines) along with the output from just three of the bias correction methods are shown, including the linear (dashed black lines), wind direction (solid gray lines), and fetch (solid black lines) models. The linear method was selected because it resulted in the lowest MAEs for both the GEFS and SREF-ARW while the wind direction approach produced the lowest MAE for the RTMA.

These time series provide additional insight into the fetch method's performance. As is the case with the other two bias correction methods shown, the fetch-based

Unauthenticated | Downloaded 12/28/22 07:05 PM UTC

60-day center-averaged MAE (hereafter referred to as average MAE) is always lower than the results of the raw analyses. For the GEFS, the fetch approach is comparable to the direction method from April to October but exhibits relatively large differences during the winter. While the direction method displays little temporal variability, the fetch approach has a December peak, decreasing through May (from 1.45 to 1.2 m s^{-1}). For the SREF-ARW, the fetch method is similar to the other methods from late December through early March. Thereafter, the three methods diverge until early summer where the fetch and wind direction methods are nearly identical. For the RTMA, the fetch method always has the highest MAE. Ultimately, the fetch approach is competitive with the other techniques, but the degree to which it does so varies during portions of the year, suggesting that other predictors might be used to improve its performance.

Additionally, the time series of average MAEs shows that the wind direction (linear) method is generally superior to the other approaches during the cool (warm) season, when winds are predominantly driven by synoptic-scale (mesoscale) patterns. The implications of this result warrant further investigation, which is beyond the scope of this paper. Also, the downward trend in the average MAE for the raw RTMA beginning in April 2015 is the result of an upgrade on 14 April 2015 (McClung 2015). After this date, the bias correction methods remove a smaller percentage of the error, suggesting that the upgrade has decreased the ratio of systematic error to random error. A time series of 60-day centered-averaged bias error results confirms this, as the magnitude of the average bias error decreases from near −1.1 to −0.6 m s^{-1} following the upgrade (not shown).

b. *Extending the method to forecasts*

The fetch method is applied to model forecasts to investigate its utility within a forecast framework. The 1° GEFS produces forecasts every 6 h out to 384 h (16 days) and the SREF-ARW every 3 h out to 87 h (approximately 3.5 days). The fetch method is applied to 24- and 48-h forecast wind fields for each GEFS and SREF-ARW run cycle and results are verified against station observations during the same 1-yr period used for the analyses. For comparison purposes, the suite of bias correction methods with a training window of 30 days is applied to the forecast wind fields. The training dataset includes all but the most recent (i.e., 24 and 48 h) forecast/observation pairs since the verifying observations occurred after the associated forecast run cycle. Here, we only consider the interpolated GEFS and SREF-ARW control output.

Figure 9 shows the MAEs associated with the raw output and bias correction methods for both the GEFS and SREF-ARW analyses and forecasts (24 and 48 h). Error bars represent 95% confidence intervals determined by bootstrapping. The MAE represents the error for all stations for the study period. The raw GEFS MAE is actually lower for the forecast wind speed (1.60 and 1.62 m s^{-1} for the 24- and 48-h forecasts, respectively) than the analysis wind speed (1.76 m s^{-1}). This is not the case however for the SREF-ARW, where the forecast MAE is slightly larger than the analysis (1.61 versus 1.63 and 1.72 m s^{-1}, respectively). The MAE from all bias correction methods increases with forecast hour, peaking at 48 h whether applied to the GEFS or SREF-ARW. Also, the bias correction methods' ability to reduce MAE decreases with increasing forecast hour. This is due to the magnitude of the bias associated with the raw models decreasing as the forecast hour increases (not shown). Despite this, each bias correction method produces a statistically distinct MAE, reducing it by 10%–20% compared to that of the raw models. This reduction is in line with results from other studies (e.g., Cheng and Steenburgh 2007).

For both the GEFS and SREF-ARW forecasts, the linear model is the most effective of the bias correction methods. The largest differences (i.e., the largest reduction in MAE) are on the order of 6% between the linear model and the next best method (running mean hour) for both the 24- and 48-h SREF-ARW. The performance of these two methods suggests that 1) the bias in the forecasts is increasingly tied to a diurnal component and 2) accounting for both additive and multiplicative errors to model the bias is more effective. The ability of the fetch method to reduce the MAE also decreases with forecast hour, and falls between (slightly better than) the running mean and wind direction methods for the GEFS (SREF-ARW) forecasts. More importantly, the fetch-based MAE for the 48-h GEFS forecasts is only slightly larger (0.04 m s^{-1}) than the raw RTMA, a state-of-the-art, high-resolution (2.5-km horizontal grid spacing) analysis.

c. *Case study: Easterly flow, 5 May 2015*

The previous results indicate that the fetch method is effective at reducing wind speed error over unresolved water bodies. An advantage of the method is that it only requires observations from similar stations rather than proximity locations. As a result, it can be used to estimate and remove biases in any region with characteristics analogous to those of the training data. Here, the method is applied to both downscale and bias-corrected interpolated GEFS analysis wind speeds onto a 500-m grid over the IRL. For this case, the flow is easterly across the region, as indicated by the 0000 UTC 5 May 2015 surface analysis (Fig. 10). The 500-m grid covers

Unauthenticated | Downloaded 12/28/22 07:05 PM UTC

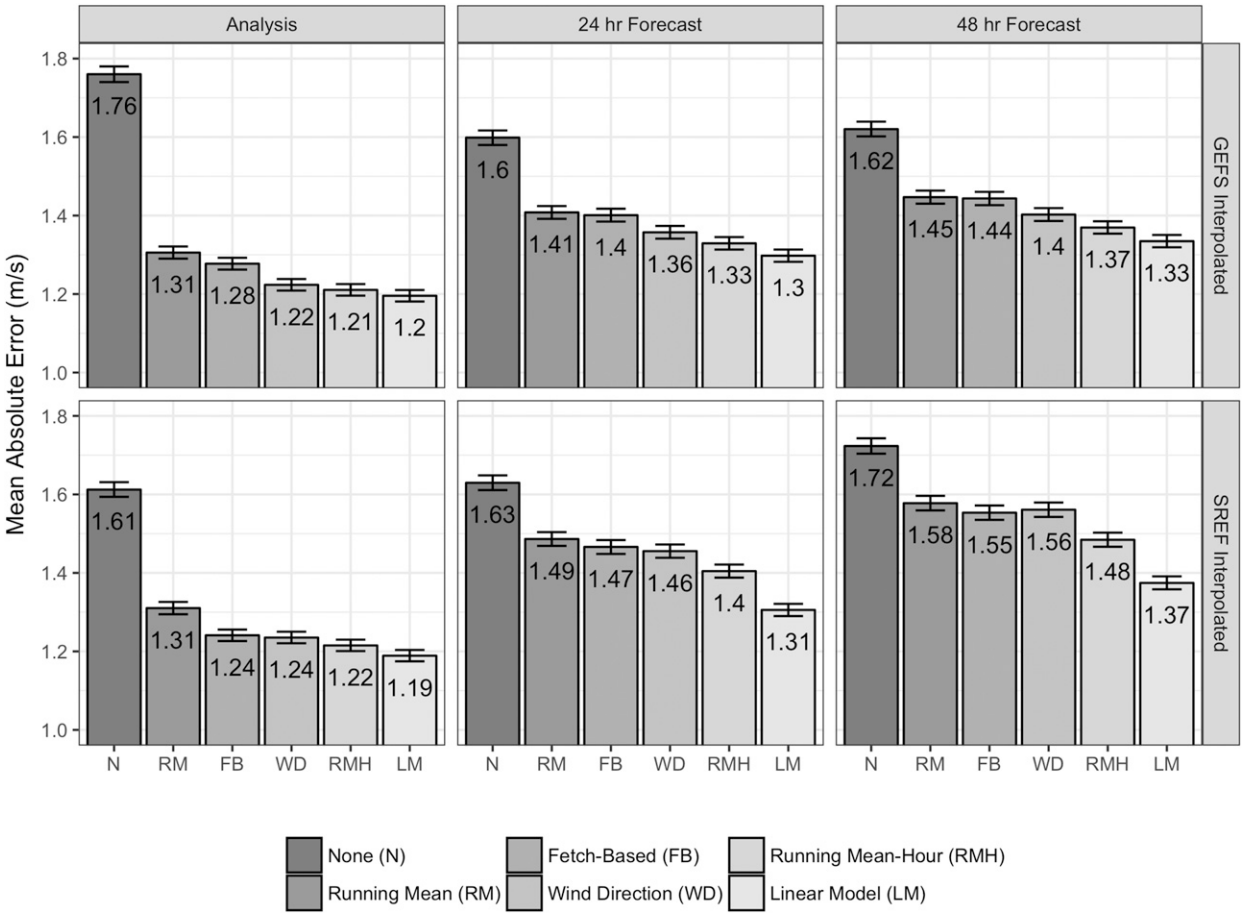

FIG. 9. MAE ($m\,s^{-1}$) averaged over a year's worth of data (see text for details) from (left) the raw and bias-corrected analyses and the (center) 24- and (right) 48-h forecasts for the interpolated (top) GEFS and (bottom) SREF-ARW. Error bars represent 95% confidence intervals for MAE values.

the portion of the IRL near Cape Canaveral, the same area shown in Fig. 2c.

The 0000 UTC 5 May 2015 interpolated raw GEFS analysis is shown in Fig. 11a. High pressure to the north results in easterly winds over the region that gradually decrease in strength from east to west (ocean to land). We present this specific case because postprocessing methods are most useful during quiescent regimes where model bias does not change much from day to day (Cheng and Steenburgh 2007). Also, this particular synoptic setting occurs frequently over the region. Thus, it provides a common scenario where the fetch method performs at its best. Analysis of other case studies allowed us to identify situations where the method did not perform well, and these are discussed in section 5.

Using the linear method described in section 3b(3), the GEFS analysis winds are bias corrected and then spread using a simple Cressman weighting scheme (Fig. 11b), as described by Yussouf and Stensrud (2006). The radius of influence, which is set to 50 km, results in a smooth "bias corrected" analysis. The corrected analysis is influenced by the biases at the XPAR, XMER, XCCB, and XRPT stations (just south of the image domain). The linear model produces negative wind speed biases at XPAR ($-1.37\,m\,s^{-1}$), XCCB ($-1.94\,m\,s^{-1}$), and XRPT ($-0.65\,m\,s^{-1}$), as well as a positive bias at XMER ($+1.64\,m\,s^{-1}$). The linear approach increases the wind speed over the entire domain compared with the raw GEFS analysis. However, the positive bias at XMER reduces the impact on the grid cells near the center of the image (Fig. 11c). Although bias information is used, the corrections are spread isotropically with no distinction drawn between land- and water-based grid cells. As a result, the IRL is still not resolved.

In contrast, the fetch-based analysis is shown in Fig. 12. For this case study, the fetches for 13 846 grid cells were estimated, in 10-m increments, from the center of the grid cell in the direction upwind of the raw interpolated GEFS analysis wind direction using the

Unauthenticated | Downloaded 12/28/22 07:05 PM UTC

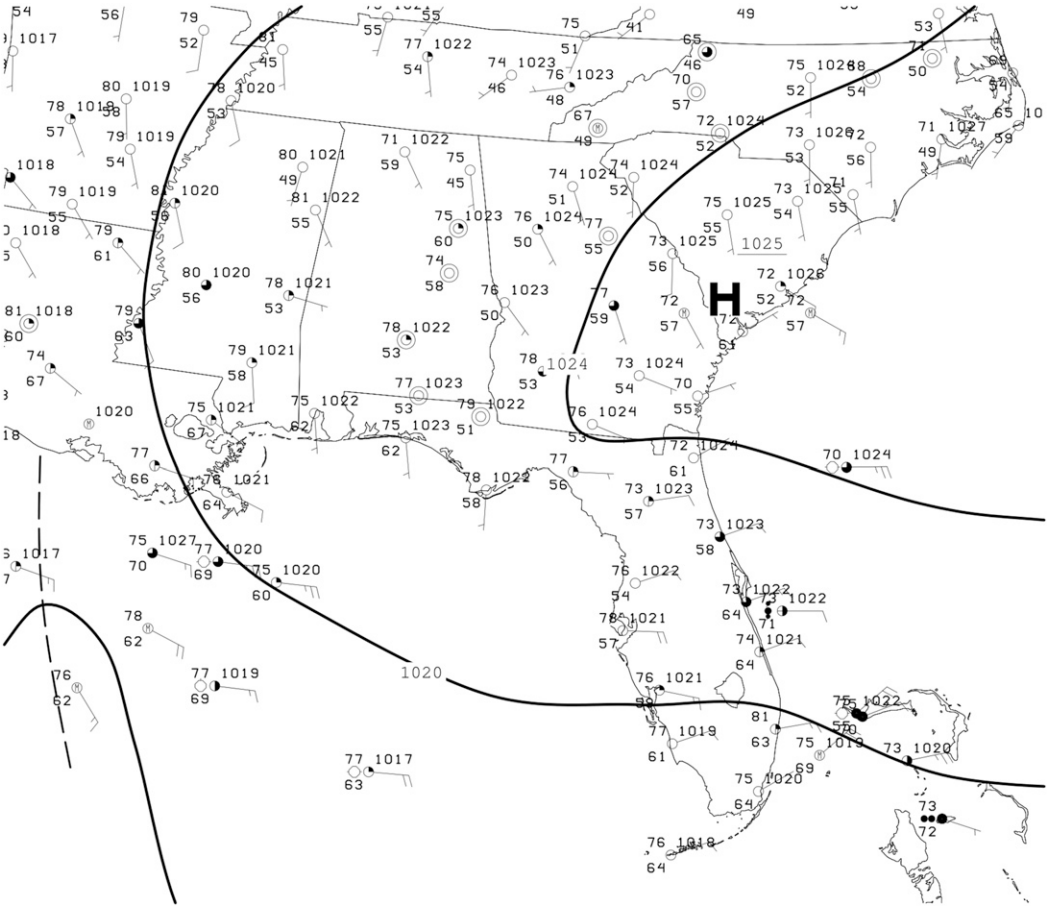

FIG. 10. Surface analysis valid at 0000 UTC 5 May 2015.

2011 National Land Cover Database (NLCD) land-use dataset (Homer et al. 2015). From these fetch lengths, the corresponding bias was estimated by fitting the training dataset comprising 0000 UTC analysis/observation pairs from the 14 stations over a 30-day period ending 4 May 2015 (Fig. 13). The fit indicates increasingly negative bias as the fetch increases. The raw analysis winds are first interpolated onto the 500-m grid and then corrected using the corresponding fetch-dependent bias estimate. For illustration purposes, the ocean is masked here. Over land, where the fetch is zero, the fetch method increases wind speeds by $0.19\,\mathrm{m\,s^{-1}}$, corresponding with the zero-fetch bias estimate from the training data shown in Fig. 13. Over the IRL, wind speeds increase with increasing (easterly) fetch, with differences (compared to the raw GEFS) exceeding $3\,\mathrm{m\,s^{-1}}$ on the downwind side (i.e., western shorelines) of the IRL. The magnitude of the largest adjustments is consistent with GEFS-estimated fetch lengths on the order of 10 km (Fig. 13). A closer look at the IRL near Cape Canaveral, with the shorelines intentionally removed, is shown in Fig. 12b. The IRL is clearly identifiable, manifest as the systematic (and physically consistent) east-to-west wind speed increase (light-to-dark gray shading) across the lagoon.

We examine a cross section (white line in Fig. 12b) that intersects station XCCB following the wind (Fig. 14). The wind speeds from the raw GEFS (dotted line), linear model (dashed line), and fetch method (solid line) are interpolated every 100 m along the cross section. Gray rectangles signify where the cross section traverses either land or a bridge, and the verifying observation for XCCB ($9.29\,\mathrm{m\,s^{-1}}$) is shown as the black dot at its location. The wind speeds produced by the raw and linear models are smooth and nondescript, with wind speeds slowly decreasing moving WSW along the cross section. In general, the fetch-generated wind speeds increase along the water portions of the transect, and relax toward the GEFS over land. Upwind of land, the transition is gradual as a result of interpolation artifacts. Ultimately, the fetch-based wind speeds are more representative of XCCB.

Unauthenticated | Downloaded 12/28/22 07:05 PM UTC

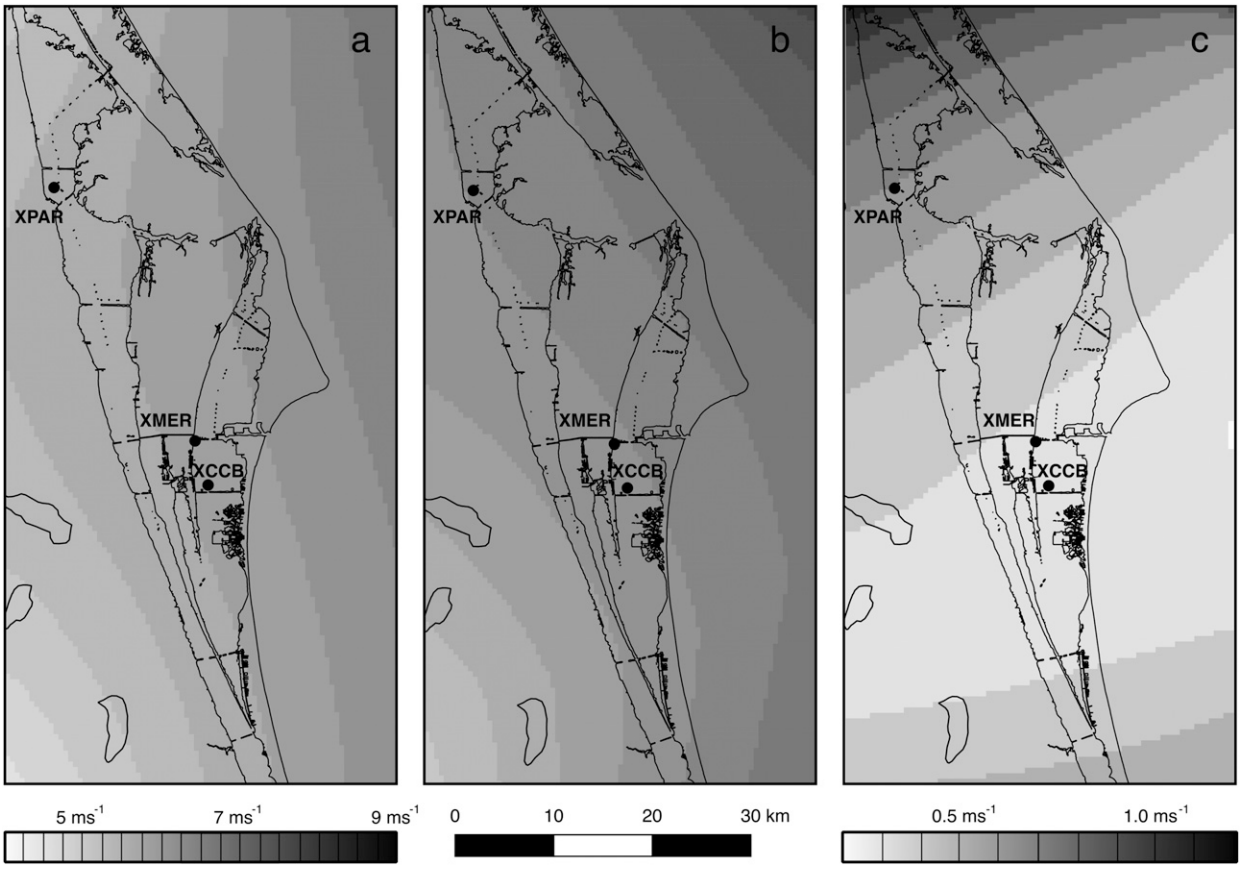

FIG. 11. Gridded (500 m) wind speed analyses for the Cape Canaveral region valid 0000 UTC 5 May 2015. Shown are the (a) raw GEFS, (b) bias-corrected (via linear model) GEFS interpolated using a Cressman weighting scheme, and (c) the difference between the two analyses [(b) − (a)]. Shading varies across the panels with contour intervals and wind speed ranges of 0.25 and 4.0–9.0 m s^{-1}, respectively, in (a) and (b) and 0.1 and 0.2–1.2 m s^{-1}, respectively, in (c).

The observed wind speeds at XPAR, XMER, and XCCB are compared with the GEFS analysis and the bias-corrected winds from both the linear and fetch methods in Table 2. In this case, the interpolated GEFS wind speeds are too low at all three stations. The linear model (Fig. 11b) yields slightly higher wind speeds at all three stations. While it is a close match at XMER, it underestimates wind speed at the ''open water'' stations XPAR and XCCB. The fetch method over- (under-) estimates the wind speed at XPAR (XCCB), but is more representative of the wind speeds at the three stations, with an associated MAE of 0.25 m s^{-1}, the smallest of the three.

5. Summary and discussion

In this study, a statistical method is introduced that bias corrects and downscales wind speed over unresolved water bodies. The fetch method is designed to correct for wind speed biases at locations without observations by estimating the bias using recent wind speed errors from similar stations. The method was applied, along with four additional ''local'' bias correction techniques, to one year's worth of both gridded and interpolated analysis wind fields from the GEFS control, SREF-ARW control, and RTMA simulations. All bias correction methods produced near-zero bias error and reduced MAE by at least 18.8% compared with the gridded or interpolated GEFS, SREF-ARW, and RTMA analyses. The fetch approach's reduction in MAE was slightly better than the running mean method, and the results are comparable to those of the other bias correction methods. The fetch method is capable of reducing the analysis wind speed error in a coarse model (e.g., 1° GEFS) to levels below that of a high-resolution analysis (e.g., RTMA). The fetch approach was also successfully applied to forecasts. Similar to the other bias correction methods, its effectiveness at reducing model error decreases with increasing forecast hour.

A case study was presented for which the fetch method was used to downscale (from ~100-km grid

Unauthenticated | Downloaded 12/28/22 07:05 PM UTC

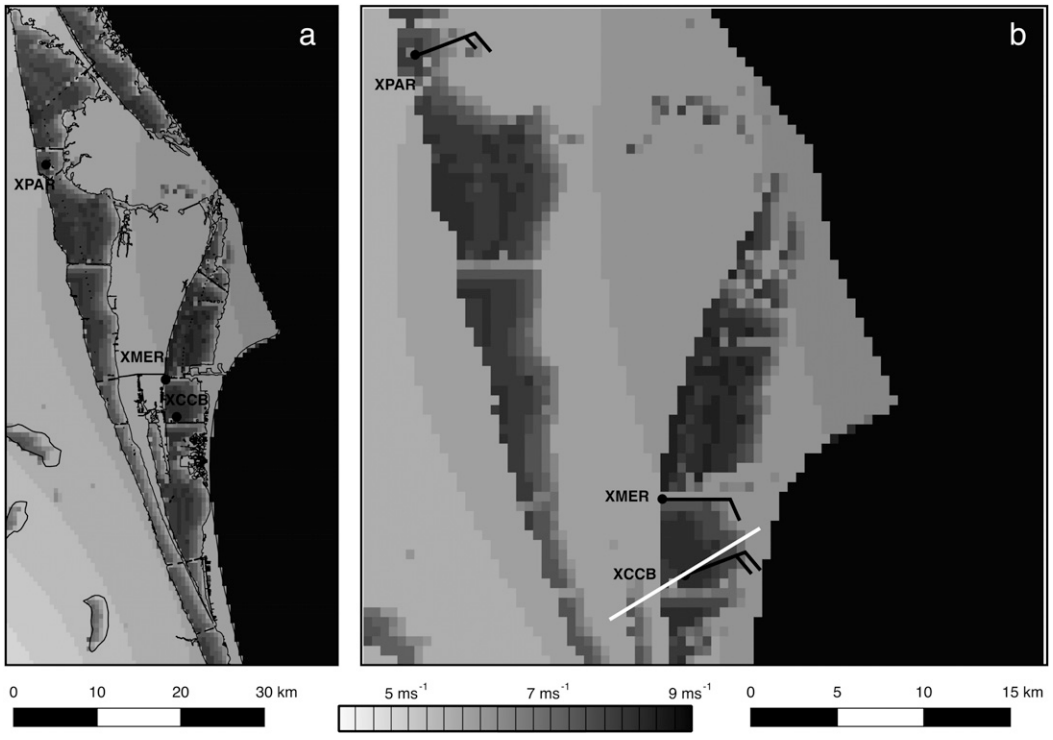

FIG. 12. (a) The fetch-based bias-corrected wind speed analysis valid 0000 UTC 5 May 2015 and (b) an enlarged view of the Cape Canaveral region (displayed without shorelines) and observed winds (half barb and full barb are 2.5 and 5 m s^{-1}, respectively) at XPAR, XMER, and XCCB. For each (500 m) grid cell, the fetch lengths were calculated using the upwind direction from the raw interpolated GEFS analysis. The corresponding bias was estimated by fitting training data composed of 0000 UTC analysis/observation pairs from 14 stations over the 30-day period prior to the analysis. The white line at the bottom of (b) is the location of the cross section analyzed in Fig. 14.

spacing to 500 m) an analysis wind speed field over the Indian River Lagoon (IRL). The downscaled wind speeds resolve the previously unresolved IRL and reduced the average error at three observing stations by approximately 1.5 m s^{-1}. The fetch-based method produces downscaled wind speeds that are more consistent with the intricate land/water mask of the coastal estuary—in contrast to the local approaches where bias estimates are spread (interpolated) using inverse distance-based weights.

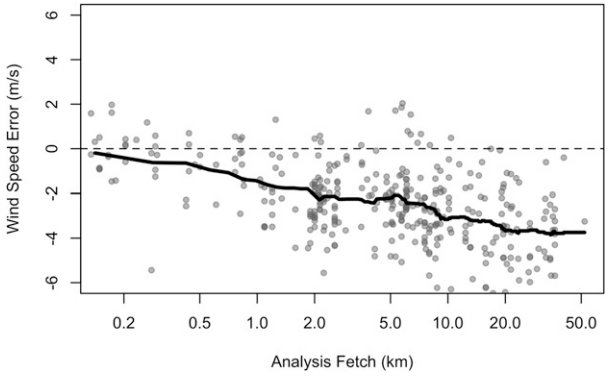

FIG. 13. Training data (analysis/observation pairs) generated using the wind speeds from 14 stations over a 30-day window (5 Apr–4 May 2015) and the interpolated 0000 UTC GEFS analysis cycle. Note that the x axis is log scale.

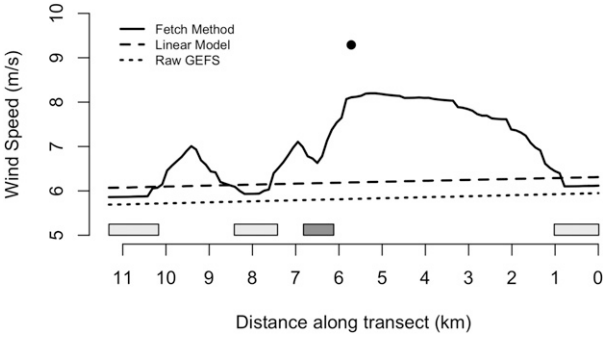

FIG. 14. Wind speed (m s^{-1}) from the raw GEFS (dotted line), linear model (dashed line), and fetch method (solid line) along the transect shown in Fig. 12b. The data were interpolated every 100 m. The verifying observation, XCCB, is shown (black dot) along with the individual land segments and bridges (light and dark gray rectangles, respectively). The distance decreases eastward in the upwind direction.

Unauthenticated | Downloaded 12/28/22 07:05 PM UTC

TABLE 2. Analysis wind speeds ($m\,s^{-1}$) from left to right including the raw interpolated GEFS, as well as the bias-corrected linear ($WSPD_{LM}$) and fetch ($WSPD_{FETCH}$) models, and the verifying wind speeds (subscript OBS) for stations XPAR, XMER, and XCCB for the 5 May 2015 case study (Figs. 11 and 12). The bottom row lists the MAE ($m\,s^{-1}$) associated with each method.

Station	$WSPD_{GEFS}$	$WSPD_{LM}$	$WSPD_{FETCH}$	$WSPD_{OBS}$
XPAR	5.57	6.26	7.47	7.02
XMER	5.83	6.21	6.28	6.32
XCCB	5.82	6.19	8.12	9.29
MAE	1.80	1.32	0.25	—

Training datasets were constructed using both nearest-neighbor and interpolated model output to create observation/analysis pairs. Compared to the nearest neighbor, the interpolation results were mixed, with a lower (higher) MAE for the GEFS (SREF-ARW) analysis winds and little impact on the RTMA winds. This is an artifact of the location of the observations with respect to the model grids. Assigning a value of 1 for land and 0 for water, the aggregate land value for the 14 stations, using the nearest-neighbor approach, is 11 for the GEFS and 8 for the SREF-ARW (see Fig. 2). After interpolation, the distance-weighted values decrease (increase) to 8.89 (10.25) for the GEFS (SREF-ARW). Thus, for the GEFS (SREF-ARW), interpolation incorporates more (less) water-friendly wind information. In general, one might expect that the nearest-neighbor approach would produce higher MAE results for coarser-resolution analyses.

During the cool season, from October through March, the wind direction method is more effective at reducing average MAEs than is either the linear or the fetch model (this is especially true for both the GEFS and RTMA). In the summer, this pattern reverses, with the linear model the clear winner in all three analyses. Since the linear model is stratified by cycle (and the wind direction method is not), this suggests that the analysis error is driven more by wind direction during the synoptic season, while the warm season exhibits more of a diurnal component, where sea, lake, and land breezes dominate the central Florida wind climate. A best approach for ''local'' bias correction at water-friendly stations might incorporate both the linear and wind direction methods. This could be done, for example, by optimally combining the two methods (e.g., Sweeney et al. 2013), or by using machine learning methods such as artificial neural networks (e.g., Salcedo-Sanz et al. 2009).

The error associated with the fetch method exhibits some temporal variability that is model (analysis) dependent. The largest variations in MAE occur in connection with the GEFS, where it decreases from a December peak to a June minimum. While the fetch method has the highest MAE throughout the year (for the RTMA), it is still substantially less than the raw analyses. Ultimately, the fetch approach is competitive, but might be improved by taking into account additional predictors.

In developing the fetch approach, we assume that wind speed error over unresolved water bodies is, to a large extent, fetch dependent. The method performs at its worst when this assumption is weak. For example, the training dataset for the 1800 UTC 30 July 2015 analysis (Fig. 5c) indicates that wind speed error is less dependent on fetch at 1800 UTC than the other cycles (i.e., a relatively flat fit line). This cycle exhibited similar behavior during the spring, autumn, and winter (not shown). Preliminary analysis suggests that this may, in part, be related to sea, lake, and/or land breezes. At 1800 UTC, stations with short fetches, in the IRL and St. John's River, are typically impacted by the sea breeze, which is not well resolved by the coarse-resolution models used here. This may result in higher observed wind speeds (compared to the analyses), driving the negative bias at short fetches. Possible future research might focus on integrating mesoscale features such as sea, lake, and land breezes into the fetch method. However, doing so would likely require the use of advanced statistical (e.g., data assimilation) and/or dynamical downscaling approaches, significantly increasing computational expense.

The fetch approach also performs poorly when stability effects significantly impact overwater wind speeds. For example, during the winter, prefrontal southwest flow can advect warm air over a much cooler (by 10°C or more) Lake Okeechobee. Lake fetches are long enough to allow the development of a stable internal boundary layer, which can result in relatively low winds at anemometer height. We found many cases like this in the training data (not shown). Under these circumstances, the wind speed error can actually become more positive with increasing fetch, opposite to what is normally observed. This suggests that additional improvements might be realized by incorporating stability effects into the fetch method.

In this work we estimated the bias as the average wind speed error of the trimmed dataset. Given the overall good performance of the linear model, the fetch method might also be improved by incorporating a relationship between wind speed error and fetch length within the trimmed dataset using a linear regression framework. Restricting the training dataset to pairs with closely matching analysis and observed wind direction could also improve results. Also, the postprocessing methods used here estimated bias using the most recent wind speed errors within a specified training window. Results

Unauthenticated | Downloaded 12/28/22 07:05 PM UTC

could be improved by expanding the temporal extent of the training window and only using training pairs that serve as analogs to the upcoming analysis/forecast (Delle Monache et al. 2011).

Acknowledgments. This work was supported with funds provided by NOAA/NWS CSTAR Award NA14NWS4680014. The authors thank WeatherFlow Inc. for providing observation data used in this research. We are very appreciative of the comments made by three anonymous reviewers, whose insights and suggestions greatly improved this manuscript.

REFERENCES

Al-Yahyai, S., Y. Charabi, and A. Gastli, 2010: Review of the use of numerical weather prediction (NWP) models for wind energy assessment. *Renewable Sustainable Energy Rev.*, **14**, 3192–3198, doi:10.1016/j.rser.2010.07.001.

Bao, L., T. Gneiting, E. P. Grimit, P. Guttorp, and A. E. Raftery, 2010: Bias correction and Bayesian model averaging for ensemble forecasts of surface wind direction. *Mon. Wea. Rev.*, **138**, 1811–1821, doi:10.1175/2009MWR3138.1.

Carter, G. M., 1975: Automated prediction of surface wind from numerical model output. *Mon. Wea. Rev.*, **103**, 866–873, doi:10.1175/1520-0493(1975)103<0866:APOSWF>2.0.CO;2.

Carvalho, D., A. Rocha, M. Gómez-Gesteira, and C. S. Santos, 2014: WRF wind simulation and wind energy production estimates forced by different reanalyses: Comparison with observed data for Portugal. *Appl. Energy*, **117**, 116–126, doi:10.1016/j.apenergy.2013.12.001.

Cassola, F., and M. Burlando, 2012: Wind speed and wind energy forecast through Kalman filtering of numerical weather prediction model output. *Appl. Energy*, **99**, 154–166, doi:10.1016/j.apenergy.2012.03.054.

Charabi, Y., S. Al-Yahyai, and A. Gastli, 2011: Evaluation of NWP performance for wind energy resource assessment in Oman. *Renewable Sustainable Energy Rev.*, **15**, 1545–1555, doi:10.1016/j.rser.2010.11.055.

Chen, S. S., and M. Curcic, 2016: Ocean surface waves in Hurricane Ike (2008) and Superstorm Sandy (2012): Coupled model predictions and observations. *Ocean Modell.*, **103**, 161–176, doi:10.1016/j.ocemod.2015.08.005.

Cheng, W. Y., and W. J. Steenburgh, 2005: Evaluation of surface sensible weather forecasts by the WRF and the Eta Models over the western United States. *Wea. Forecasting*, **20**, 812–821, doi:10.1175/WAF885.1.

——, and ——, 2007: Strengths and weaknesses of MOS, running-mean bias removal, and Kalman filter techniques for improving model forecasts over the western United States. *Wea. Forecasting*, **22**, 1304–1318, doi:10.1175/2007WAF2006084.1.

Cleveland, W. S., 1984: Graphical methods for data presentation: Full scale breaks, dot charts, and multibased logging. *Amer. Stat.*, **38**, 270–280, doi:10.2307/2683401.

Cressman, G. P., 1959: An operational objective analysis system. *Mon. Wea. Rev.*, **87**, 367–374, doi:10.1175/1520-0493(1959)087<0367:AOOAS>2.0.CO;2.

Davis, C., and Coauthors, 2008: Prediction of landfalling hurricanes with the Advanced Hurricane WRF Model. *Mon. Wea. Rev.*, **136**, 1990–2005, doi:10.1175/2007MWR2085.1.

Delle Monache, L., T. Nipen, Y. Liu, G. Roux, and R. Stull, 2011: Kalman filter and analog schemes to postprocess numerical weather predictions. *Mon. Wea. Rev.*, **139**, 3554–3570, doi:10.1175/2011MWR3653.1.

De Pondeca, M. S. , and Coauthors, 2011: The real-time mesoscale analysis at NOAA's National Centers for Environmental Prediction: Current status and development. *Wea. Forecasting*, **26**, 593–612, doi:10.1175/WAF-D-10-05037.1.

De Rooy, W. C., and K. Kok, 2004: A combined physical–statistical approach for the downscaling of model wind speed. *Wea. Forecasting*, **19**, 485–495, doi:10.1175/1520-0434(2004)019<0485:ACPAFT>2.0.CO;2.

Engel, C., and E. Ebert, 2007: Performance of hourly operational consensus forecasts (OCFs) in the Australian region. *Wea. Forecasting*, **22**, 1345–1359, doi:10.1175/2007WAF2006104.1.

Gastón, M., E. Pascal, L. Frías, I. Martí, U. Irigoyen, E. Cantero, S. Lozano, and L. Loureiro, 2008: Wind resources map of Spain at mesoscale. Methodology and validation. *European Wind Energy Conf.*, Brussels, Belgium, European Wind Energy Association. [Available online at http://secure.cener.com/documentos/wind-resources-map-mesoscale-PaperEwec08.pdf.]

Glahn, H. R., and D. A. Lowry, 1972: The use of model output statistics (MOS) in objective weather forecasting. *J. Appl. Meteor.*, **11**, 1203–1211, doi:10.1175/1520-0450(1972)011<1203:TUOMOS>2.0.CO;2.

Gneiting, T., and A. E. Raftery, 2005: Weather forecasting with ensemble methods. *Science*, **310**, 248–249, doi:10.1126/science.1115255.

Homer, C. G., and Coauthors, 2015: Completion of the 2011 National Land Cover Database for the conterminous United States—Representing a decade of land cover change information. *Photogramm. Eng. Remote Sens.*, **81**, 345–354.

Hsu, S. A., E. A. Meindl, and D. B. Gilhousen, 1994: Determining the power-law wind-profile exponent under near-neutral stability conditions at sea. *J. Appl. Meteor.*, **33**, 757–765, doi:10.1175/1520-0450(1994)033<0757:DTPLWP>2.0.CO;2.

Keddy, P. A., 1982: Quantifying within-lake gradients of wave energy: Interrelationships of wave energy, substrate particle size and shoreline plants in Axe Lake, Ontario. *Aquat. Bot.*, **14**, 41–58, doi:10.1016/0304-3770(82)90085-7.

Li, Y., X. Wang, and M. Xue, 2012: Assimilation of radar radial velocity data with the WRF hybrid ensemble–3DVAR system for the prediction of Hurricane Ike (2008). *Mon. Wea. Rev.*, **140**, 3507–3524, doi:10.1175/MWR-D-12-00043.1.

Louka, P., G. Galanis, N. Siebert, G. Kariniotakis, P. Katsafados, G. Kallos, and I. Pytharoulis, 2008: Improvements in wind speed forecasts for wind power prediction purposes using Kalman filtering. *J. Wind Eng. Ind. Aerodyn.*, **96**, 2348–2362, doi:10.1016/j.jweia.2008.03.013.

Luettich, R. A., and J. J. Westerink, 2004: Formulation and numerical implementation of the 2D/3D ADCIRC finite element model version 44.XX. Institute of Marine Sciences, University of North Carolina, 74 pp. [Available online at http://www.unc.edu/ims/adcirc/adcirc_theory_2004_12_08.pdf.]

Mass, C. F., D. Ovens, K. Westrick, and B. A. Colle, 2002: Does increasing horizontal resolution produce more skillful forecasts? The results of two years of real-time numerical weather prediction over the Pacific Northwest. *Bull. Amer. Meteor. Soc.*, **83**, 407–430, doi:10.1175/1520-0477(2002)083<0407:DIHRPM>2.3.CO;2.

——, J. Baars, G. Wedam, E. Grimit, and R. Steed, 2008: Removal of systematic model bias on a model grid. *Wea. Forecasting*, **23**, 438–459, doi:10.1175/2007WAF2006117.1.

Unauthenticated | Downloaded 12/28/22 07:05 PM UTC

McClung, T., 2015: Amended: RTMA and URMA upgrade: Effective April 14, 2015. NWS Tech. Implementation Notice TIN 15-08, 4 pp. [Available online at http://www.nws.noaa.gov/os/notification/tin15-08rtma_urma_aaa.htm.]

Mortensen, N. G., and Coauthors, 2006: Wind atlas for Egypt: Measurements, micro- and mesoscale modelling. *European Wind Energy Conf.*, Athens, Greece, European Wind Energy Association. [Available online at http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.486.2414&rep=rep1&type=pdf.]

Müller, M. D., 2011: Effects of model resolution and statistical postprocessing on shelter temperature and wind forecasts. *J. Appl. Meteor. Climatol.*, **50**, 1627–1636, doi:10.1175/2011JAMC2615.1.

QGIS Development Team, 2016: QGIS: A free and open source Geographic Information System. Open Source Geospatial Foundation Project. [Available online at http://www.qgis.org/.]

Salcedo-Sanz, S., A. M. Perez-Bellido, E. G. Ortiz-García, A. Portilla-Figueras, L. Prieto, and D. Paredes, 2009: Hybridizing the fifth generation mesoscale model with artificial neural networks for short-term wind speed prediction. *Renewable Energy*, **34**, 1451–1457, doi:10.1016/j.renene.2008.10.017.

Skamarock, W. C., and Coauthors, 2008: A description of the Advanced Research WRF version 3. NCAR Tech. Note NCAR/TN-475+STR, 113 pp., doi:10.5065/D68S4MVH.

Smith, N. P., 1990: Computer simulation of tide-induced residual transport in a coastal lagoon. *J. Geophys. Res.*, **95**, 18 205–18 211, doi:10.1029/JC095iC10p18205.

Stensrud, D. J., and N. Yussouf, 2005: Bias-corrected short-range ensemble forecasts of near surface variables. *Meteor. Appl.*, **12**, 217–230, doi:10.1017/S135048270500174X.

Strassberg, D., M. A. LeMone, T. T. Warner, and J. G. Alfieri, 2008: Comparison of observed 10-m wind speeds to those based on Monin–Obukhov similarity theory using IHOP_2002 aircraft and surface data. *Mon. Wea. Rev.*, **136**, 964–972, doi:10.1175/2007MWR2203.1.

Sweeney, C. P., P. Lynch, and P. Nolan, 2013: Reducing errors of wind speed forecasts by an optimal combination of post-processing methods. *Meteor. Appl.*, **20**, 32–40, doi:10.1002/met.294.

U.S. Army Coastal Engineering Research Center, 1977: *Shore Protection Manual.* Vol. 1. 3rd ed. U.S. Government Printing Office, 514 pp.

Weaver, R. J., P. Taeb, S. M. Lazarus, M. E. Splitt, B. P. Holman, and J. A. Colvin, 2016: Sensitivity of modeled estuarine circulation to spatial and temporal resolution of input meteorological forcing of a cold frontal passage. *Estuarine Coastal Shelf Sci.*, **183**, 28–40, doi:10.1016/j.ecss.2016.10.014.

Wilks, D. S., 2011: *Statistical Methods in the Atmospheric Sciences.* 3rd ed. Elsevier, 676 pp.

Yussouf, N., and D. J. Stensrud, 2006: Prediction of near-surface variables at independent locations from a bias-corrected ensemble forecasting system. *Mon. Wea. Rev.*, **134**, 3415–3424, doi:10.1175/MWR3258.1.

Zhang, J., C. Draxl, T. Hopson, L. Delle Monache, E. Vanvyve, and B. M. Hodge, 2015: Comparison of numerical weather prediction based deterministic and probabilistic wind resource assessment methods. *Appl. Energy*, **156**, 528–541, doi:10.1016/j.apenergy.2015.07.059.

Unauthenticated | Downloaded 12/28/22 07:05 PM UTC