The Winds of Tomorrow

Veer Renewables Embraces GPU-Powered Weather Modeling

Prepared by

Mike Optis, President, Veer Renewables

2023-10-30

Executive Summary

At Veer Renewables, we continue to push the boundaries of numerical weather prediction (NWP) modeling and data analytics to provide advanced products and services to our wind energy stakeholders. In line with this commitment to innovation, we are excited to announce a transformative shift in our NWP computational capabilities. After meticulous evaluation, we are transitioning away from the standard Weather Research and Forecasting (WRF) software – grounded in conventional central processing units (CPUs) – to the cutting-edge AceCAST software by TempoQuest. This WRF-equivalent platform harnesses the raw speed and efficiency of graphical processing units (GPUs), offering greatly accelerated and cost-efficient WRF simulations – a direct advantage we are eager to extend to our growing list of clients.

AceCAST is the product of over five years of rigorous R&D at TempoQuest and empowers WRF users to achieve striking performance optimizations using the massive parallelism of GPU hardware versus traditional CPUs. With a broad set of refactored WRF physics, dynamic modules, and namelist options ingrained with NVIDIA CUDA or OpenACC GPU programming methodologies, AceCAST offers an effortless drop-in alternative to the conventional WRF workflow.

Our decision to pivot from the industry standard WRF software to AceCAST was not taken lightly. Rather, the decision comes after months of evaluation and comparisons between the two software products. In this report, we share validation results from a recent WakeMap analysis centered on the offshore U.S. Mid-Atlantic. Here we conduct a year-long side-by-side assessment of WakeMap simulations with both WRF and AceCAST and assess modeled atmospheric variables including wind speed, turbulent kinetic energy, and power production.

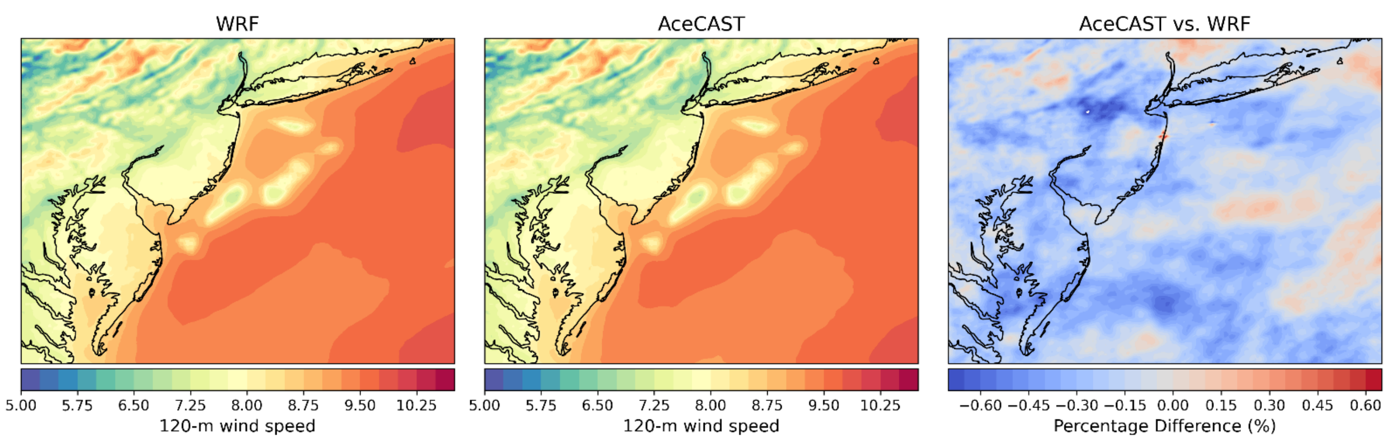

Figure E1: Mean annual wind speeds estimated in the U.S. Mid-Atlantic offshore region using WRF 4.4.2 (left) and AceCAST (built equivalent to WRF 4.4.2). Differences in the mean winds are shown in the rightmost panel. Not the offshore deficits caused by the modeled wind turbines.

The validation findings are unequivocal. We observed only minimal variances between WRF and AceCAST, as illustrated in Figure E1. These variances are of similar magnitude to those one might encounter when running WRF with a different compiler or on different CPU architectures.

Given such consistent performance, Veer Renewables has confidently embraced AceCAST as our go-to NWP platform for WakeMap and other WRF-based products. In doing so, we are not just advancing our technological edge; we are committed to delivering these cost and time savings directly to our clients, reaffirming our position as an industry leader in innovation, speed, and affordability.

Introduction

Veer Renewables and WRF Modeling

Veer Renewables stands at the forefront of WRF modeling innovation in the wind energy industry. In February 2023, we proudly launched WakeMap – a first-of-its-kind WRF-based product that leverages the wind farm parameterization in WRF to model wind farm wake impacts on the local and regional wind resource. Described in detail in our white paper, WakeMap is now used by a growing list of wind energy developers to accurately quantify the historically underestimated impacts of neighboring wind farms on their current or planned projects.

In early 2024, Veer Renewables will begin expanding its WRF-based product offerings, including wind maps and timeseries data, with options to include wind turbines or not. Users will be able to request simulations, as well as download and visualize results, all through an interactive web interface.

NWP Acceleration through GPUs

In recent years, Graphics Processing Units (GPUs) have emerged as powerful computational tools, offering dramatic advancements in parallel processing capabilities. Originally designed for rendering graphics, GPUs are now increasingly recognized for their aptitude in handling data-intensive tasks, making them an indispensable asset in a range of scientific computations. Especially in the realm of NWP modeling, GPUs hold considerable promise. Compared to the traditional Central Processing Units (CPUs), GPUs can process thousands of operations simultaneously, which is particularly advantageous for NWP modeling where the spatial domain is typically segmented into multiple gridded regions that are processed concurrently. By pivoting to GPU-based computations, not only can there be a significant reduction in simulation time, but also substantial cost savings in terms of energy consumption and infrastructure investments.

TempoQuest and AceCAST

TempoQuest's flagship software, AceCAST, represents a leap in NWP modeling by significantly accelerating the WRF model. Harnessing the power of GPUs, AceCAST—born from half a decade of rigorous research—offers unparalleled performance enhancements compared to traditional CPU-based approaches. This GPU optimization not only makes it the world's fastest and most detailed weather forecasting model but also ensures meteorologists and end-users access more precise insights into localized weather phenomena at a fraction of the cost. The transformational capability of AceCAST, embedded with NVIDIA CUDA and OpenACC programming techniques, serves as a seamless substitute for established WRF configurations, typically boosting modeling processing speeds by a factor of five.

Intent of Study

As the winds of change sweep across the NWP computational landscape, the emergence of AceCAST introduces both promise and a need for validation. For the wind energy industry in particular, the accuracy and consistency of NWP models are paramount. While WRF has long served as the industry benchmark, the purported benefits of AceCAST – greatly enhanced speed and cost-efficiency – beckon for a comprehensive comparative analysis. Thus, in this report we summarize a meticulous examination of the two software products, comparing their performance in modeling the regional wind resource, the power output from turbines, and other atmospheric parameters within the nascent mid-Atlantic U.S. offshore wind areas.

Domain and Model Setup

The spatial domain for this analysis (Figure 1) is centered on the mid-Atlantic offshore region around Delaware to New York. A large 3-km domain is nested by a larger 9-km domain (not shown). Modeled wind farms are shown in green and are all currently in various stages of planning and construction.

We selected this offshore region given the expected impacts of "long wakes". These are conditions where, especially under stable atmospheric conditions, wakes can extend far beyond conventional expectations, often exceeding 50 kilometers in length. The U.S. mid-Atlantic offshore wind region is particularly susceptible to long wakes. The vast stretches of upwind fetch in the dominant southwestern wind direction, coupled with the induced stable stratification when warm air from inland areas like Delaware, Virginia, and North Carolina flows over the colder ocean, provides ideal conditions for long wakes to occur.

For these reasons, there is growing concern that engineering-based wake models are drastically underestimating wake losses in this region. In response, wind energy stakeholders are increasingly turning to WakeMap for a better understanding and more accurate quantification of long wake impacts in this unique weather region.

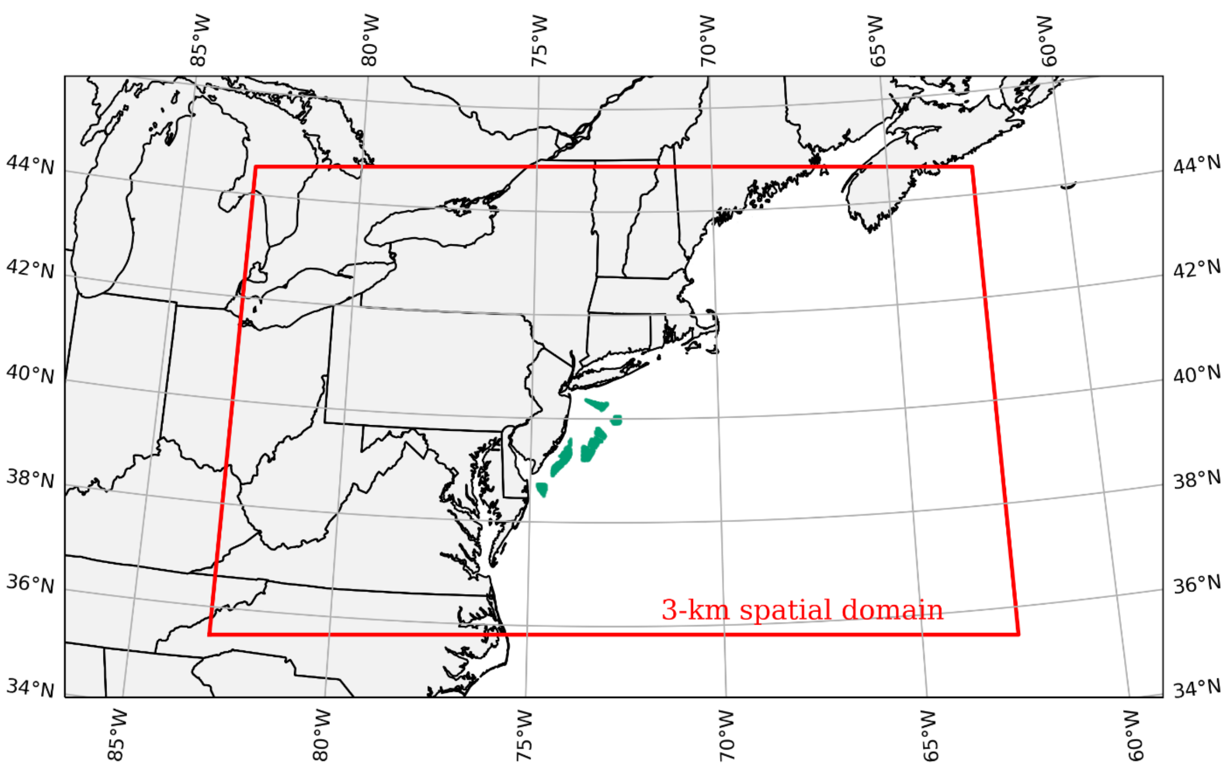

Figure 1: Spatial domain used in this analysis. The inner 3-km domain is shown as a red box and modeled turbines are shown in green. The 9-km outer grid is not shown.

In Figure 2 we show more detailed layouts of the modeled wind farms. Wind farm layouts and turbine technical specifications have been provided by an industry partner. Due to confidentiality considerations, we are unable to share specific details in this report. Instead, we provide approximate number of turbines per Lease Area in Table 1. We further note that wind turbine capacities modeled represent a range of capacities currently proposed by the wind turbine manufacturers for the US East Coast market..

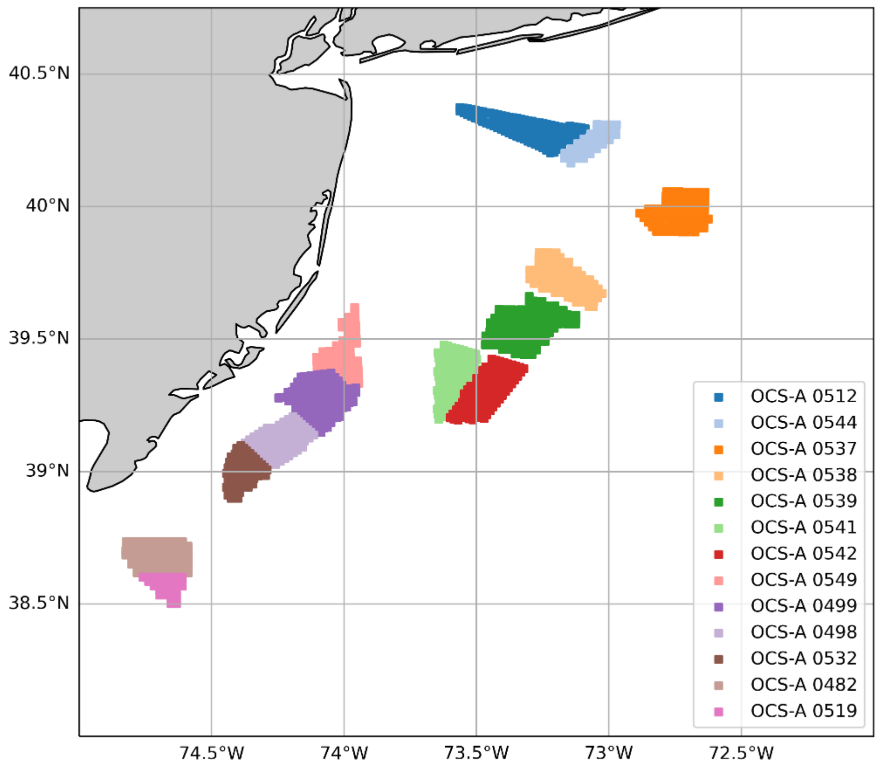

Figure 2: Lease Areas modeled in this validation study.

Lease Area	Number of Turbines (rounded to nearest 20)
OCS-A 0512	180
OCS-A 0544	40
OCS-A 0537	120
OCS-A 0538	120
OCS-A 0539	160
OCS-A 0541	120
OCS-A 0542	120
OCS-A 0549	80
OCS-A 0499	200
OCS-A 0498	100
OCS-A 0532	80
OCS-A 0482	80
OCS-A 0519	40

Table 1: Approximate number of turbines modeled for each offshore wind Lease Area.

We launch a full year of atmospheric simulations representative of long-term conditions by constructing a Typical Meteorological Year (TMY). Specifically, we examine monthly statistics between 2000 and 2022 from the ERA5 reanalysis product, extracted at the center of the model domain. For each calendar month, we examine distributions of 100-m wind speed, 100-m wind direction, and 2-m temperature, and identify which year most closely resembles the long-term 2000-2022 period. This year is then selected in the TMY for that calendar month. The process continues until the optimal 12 calendar months are selected. Based on this approach, we arrive at the 12-month TMY for this analysis, summarized in Table 2.

Month	Year
January	2019
February	2015
March	2020
April	2006
May	2015
June	2001
July	2012
August	2020
September	2018
October	2011
November	2017
December	2006

Table 2: Constructed TMY used in this validation study.

The CPU-powered WRF and GPU-powered AceCAST simulations share an identical model setup, summarized in Table 3. We selected WRF Version 4.4.2 and TempoQuest provided an equivalent AceCAST distribution. Boundary forcings are provided by the ERA5 reanalysis product. We run simulations over 2-day periods with a 12-hour spin-up. These shorter simulation lengths are used to minimize WRF model drift (i.e., inner domain simulation diverging from the ERA5 boundary forcing), which is common for simulations longer than a few days. To construct a full year of WakeMap simulations, we simply concatenate the separate 2- or 3-day simulations into a single timeseries. Data is output every 10 minutes (instantaneous, not averaged) and post-processed to calculate atmospheric variables at multiple heights above ground level.

Parameter	Value
WRF Version	4.4.2
Boundary forcing	ERA5 reanalysis
Spatial resolution	3 km
Nesting	9km – 3km
Vertical resolution	20 m up to 200 m height, then increasing
Timestep	30 seconds
Data output intervals	10 minutes
Simulation length	2-3 days with 12-hour spin-up
Planetary boundary layer scheme	Mellor-Yamada-Nakanishi-Niino
WFP scheme	Fitch
Wind farm turbulent kinetic energy generation factor	0.25

Table 3: Key setup parameters for both WRF and AceCAST.

Validation Results

Mean Statistics

We begin by examining mean maps of 120-m wind speed, power production from the wind farms, and TKE, as shown in Figure 3. Here we see very little divergence between the mean WRF and AceCAST results. Deviations between wind speeds are less than 1%, as are the modeled power estimates. Variances in TKE are slightly higher, peaking at +/- 1.5%.

These differences are of similar magnitude to those observed when WRF is launched on different CPU architectures or built using different compilers. Overall, AceCAST is doing an exceptional job of reproducing mean annual atmospheric variables over the entire domain.

However, we do note two anomalies in Figure 3, both of which relate to very minor differences in the grid cell locations between the two models:

1. In the top-right plot, we see a single anomaly in wind speed difference along the coast of New Jersey. Upon examination, we found that different compilations of the WRF Pre-processing System (WPS) resulted in a single grid cell being characterized as water by the AceCAST WPS compilation and land by the standard WRF WPS compilation. As a result, we see slightly higher wind speeds in AceCAST due to the lower surface roughness of water compared to land.
2. In the mid-right plot, we see a noticeable anomaly in mean power comparisons at the Empire Wind Lease Area. Upon investigation, we determined that this anomaly was again caused by slight differences in the WPS compilations. Specifically, different number of decimal places preserved in the latitude and longitude coordinates resulted in a single turbine being placed in a lower grid cell in the AceCAST run compared to the WRF run (see Figure 4)

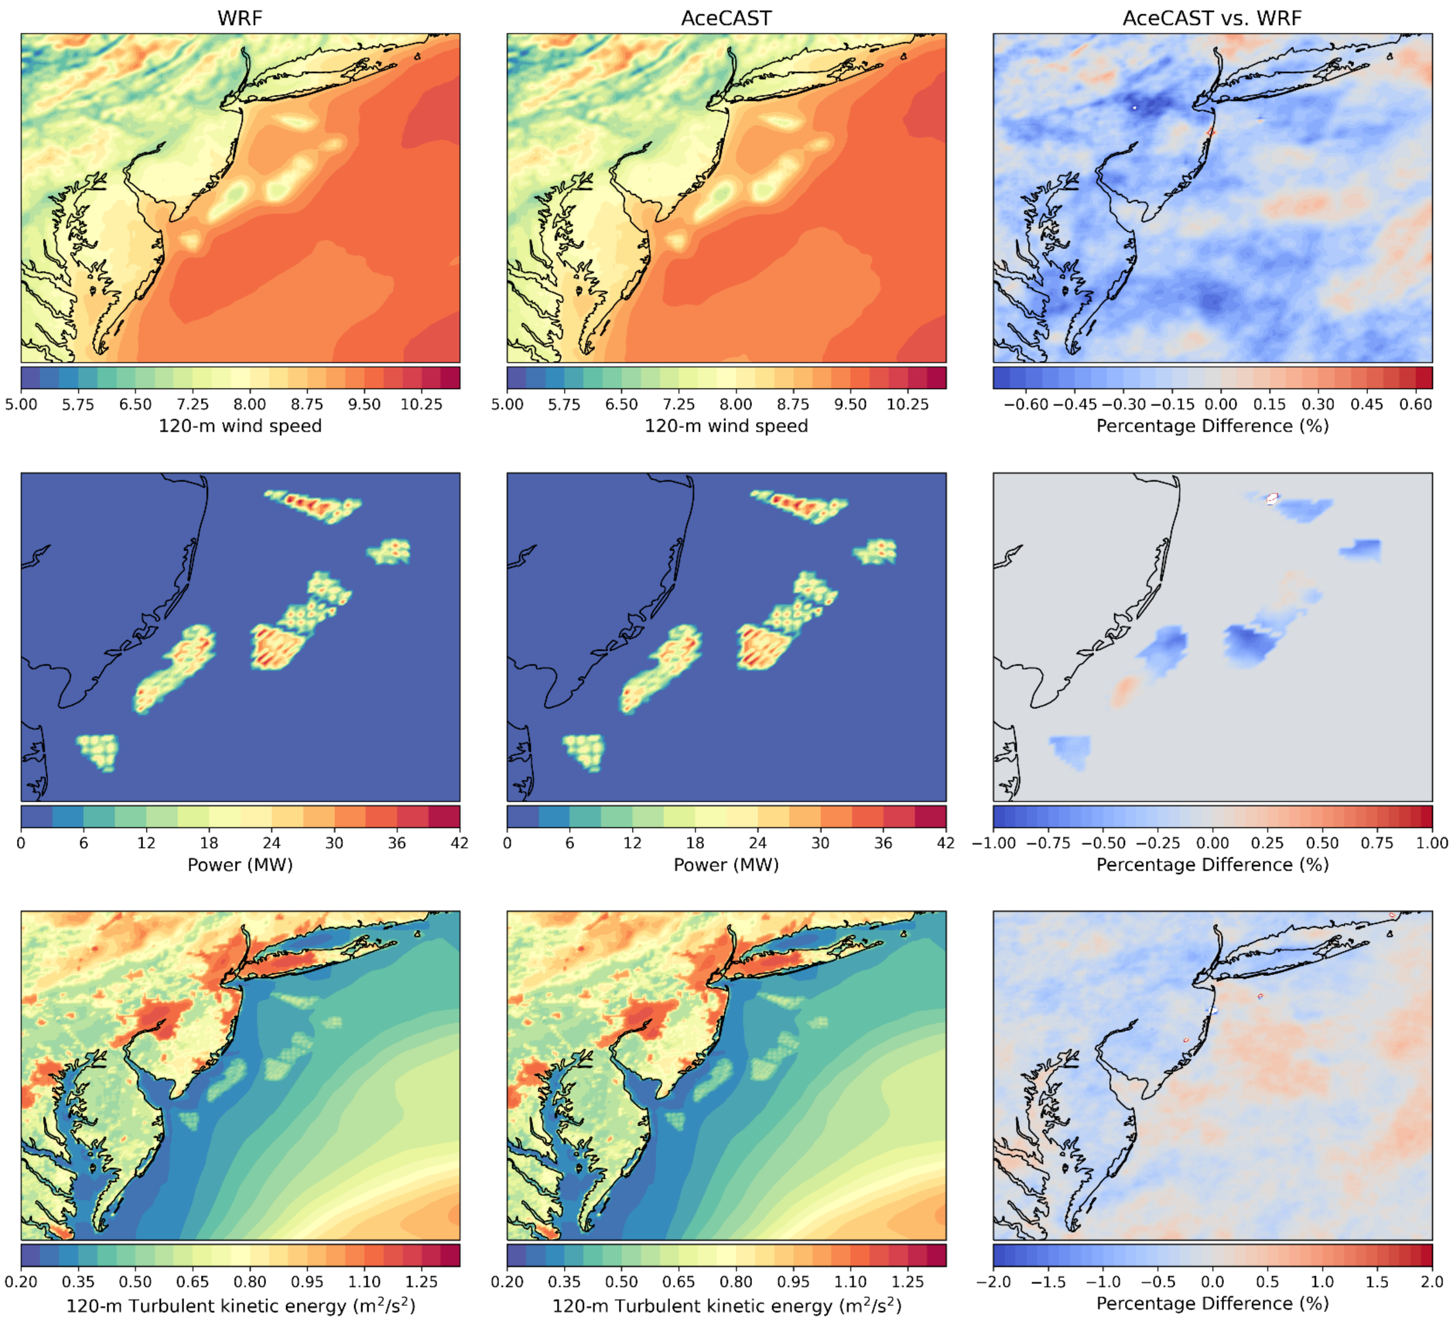

Figure 3: Mean maps of 120-m wind speed (top row), power (middle row) and 120-m turbulent kinetic energy (bottom row) modeled using WRF (left column) and AceCAST (center column). Percentage differences between the WRF and AceCAST results are shown in the right column

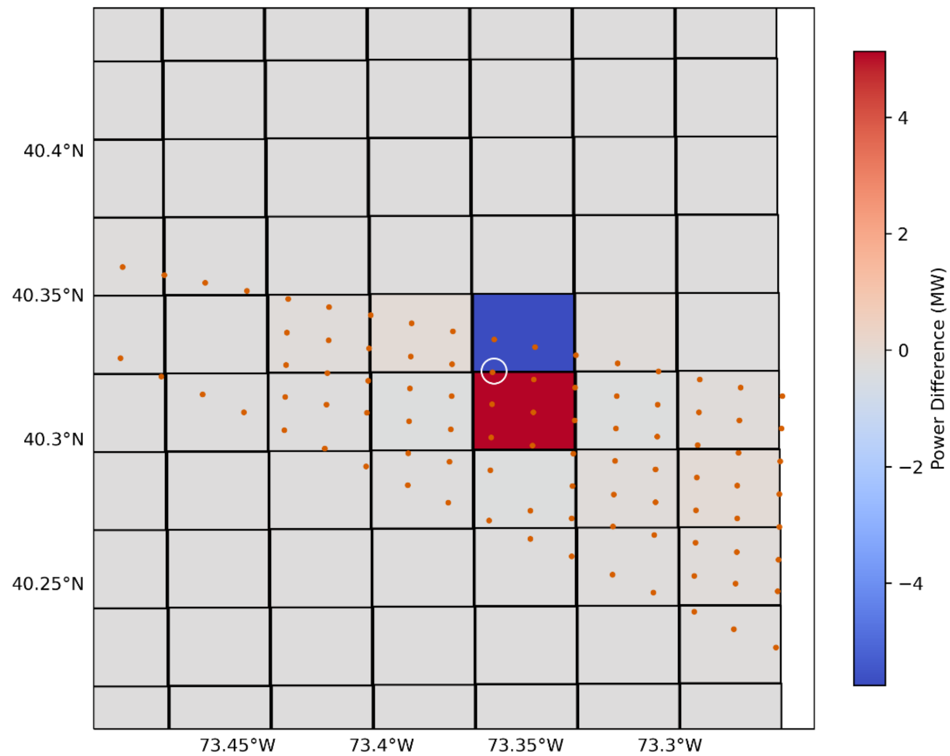

Figure 4: Percentage difference in mean power between AceCAST and WRF centered at the Empire Wind Lease Area (turbines shown in orange). The observed anomaly is caused by different grid cell placements of a single wind turbine (circled in white).

In Table 4, we compare mean capacity factor (CF) estimates across the different Lease Areas, which have been anonymized at the request of our industry partner. First, we acknowledge much lower CF values than generally predicted by industry standard methods, which we attribute to the extreme atmospheric stability observed in this region and the inability of standard wake models to fully capture these expected long wake impacts. Next, we observe very little difference in CF estimates between WRF and AceCAST. Overall, AceCAST tends to estimate slightly lower wind speeds than WRF within the Lease Areas, resulting in slight decreases in the CF estimates.

Lease Area	Modeled Capacity Factor (%)		Difference (%)
	AceCAST	WRF	
Lease Area A	36.37	36.58	-0.21
Lease Area B	36.08	36.23	-0.15
Lease Area C	33.21	33.38	-0.17
Lease Area D	35.50	35.72	-0.22
Lease Area E	41.16	41.16	0.00
Lease Area F	37.21	37.37	-0.16
Lease Area G	42.60	42.76	-0.16
Lease Area H	43.36	43.56	-0.20
Lease Area I	40.63	40.66	-0.03
Lease Area J	38.22	38.31	-0.09
Lease Area K	34.99	35.17	-0.18
Lease Area L	41.11	41.30	-0.19
Lease Area M	36.71	36.71	0.00
Lease Area N	35.38	35.35	0.03
Lease Area O	44.98	45.08	-0.10
		Mean:	**-0.12**

Table 4: Capacity factor estimates across the Lease Areas, which have been anonymized at the request of our industry partner.

Timeseries Analysis

Next, we perform timeseries analyses of the AceCAST and WRF simulations. In Figure 5, we provide a snapshot of key atmospheric variables over January 2019. As shown in the figure, agreement between the WRF and AceCAST timeseries is exceptional, with only small deviations in atmospheric parameters.

In Table 5, we examine timeseries statistics of wind speeds, spatially averaged across each Lease Area. The table shows very high correlations between the WRF and AceCAST timeseries, with R^2 coefficients equal to 0.989 or greater. Also in the table, we fit the wind speed data to a Weibull distribution and report the *A* and *k* parameters. Again, we see very strong agreement between WRF and AceCAST. For visual validation purposes, we plot distributions of 120-m wind speeds at Lease Area A in Figure 6. As we can see, the AceCAST and WRF wind speed distributions are nearly identical.

More extensive timeseries validation was conducted in this study, including plots for different Lease Areas and additional statistical metrics such as bias and root-mean-squared-error for the different atmospheric parameters. However, for the sake of brevity, these results are not presented in this report. Across all spatial regions, performance metrics, and atmospheric variables, we found that AceCAST provides a near-perfect replication of the WRF model.

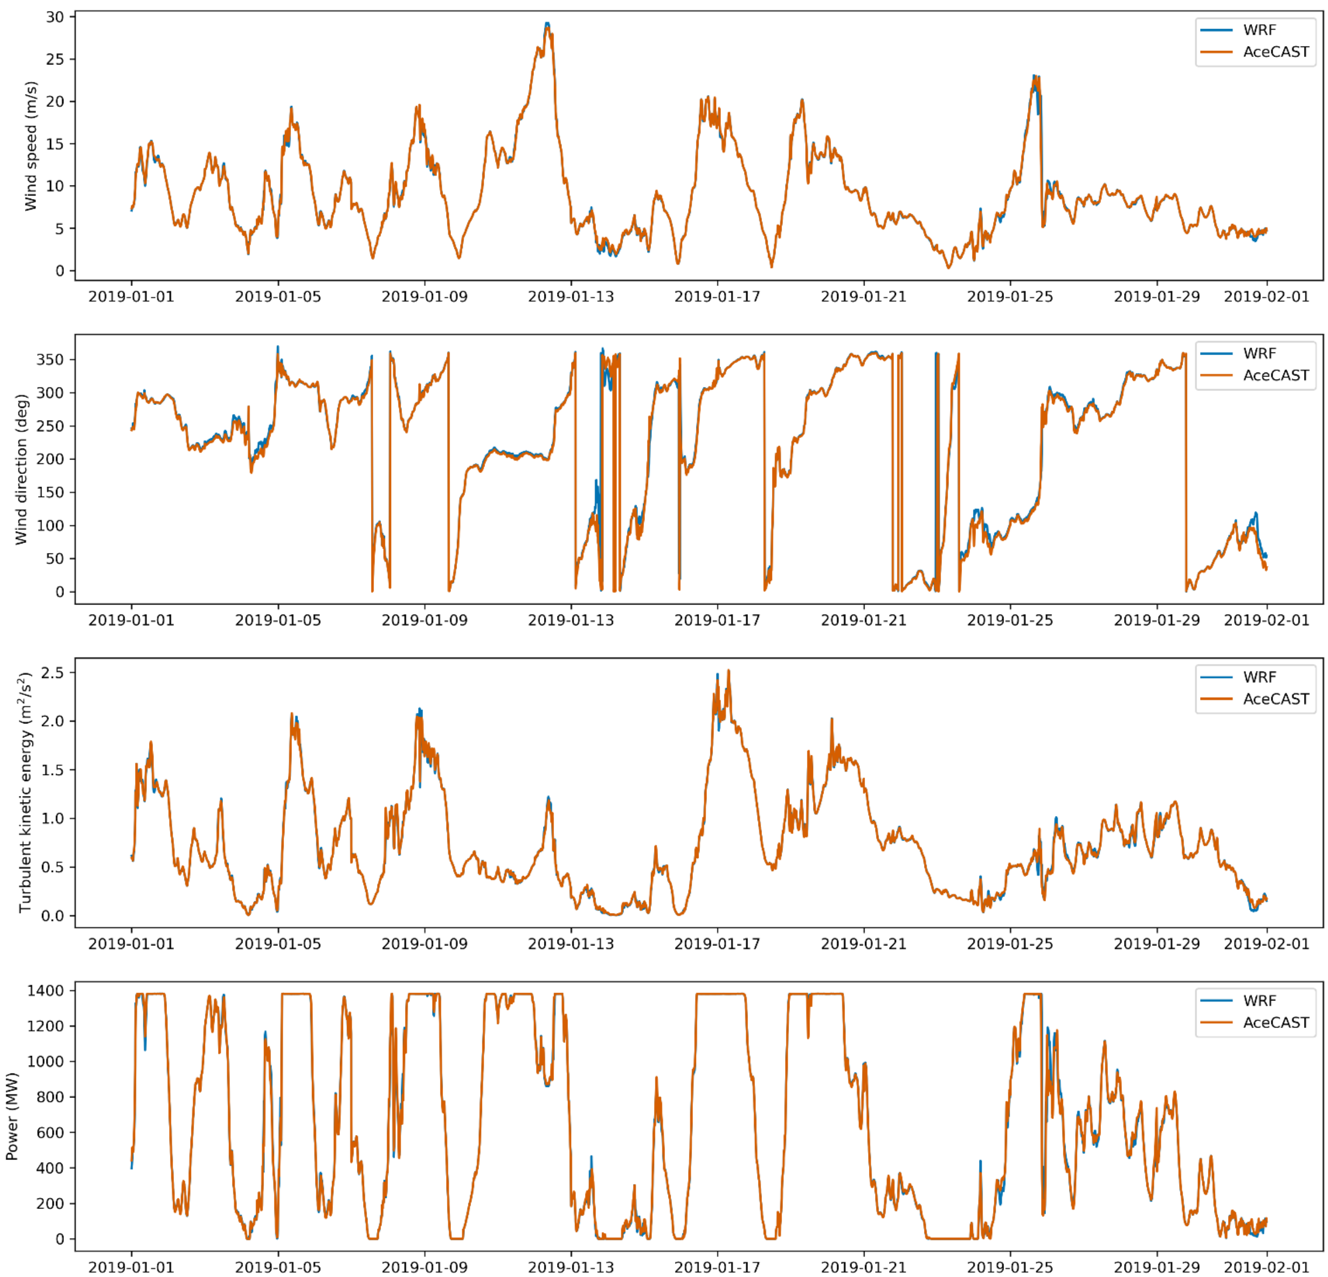

Figure 5: Timeseries visualization of key atmospheric parameters in January 2019 within Lease Area A.

Lease Area	Correlation Coefficient (AceCAST-WRF)	Weibull Parameter A		Weibull Parameter k	
		WRF	AceCast	WRF	AceCAST
Lease Area A	0.992	9.20	9.18	1.79	1.79
Lease Area B	0.991	8.70	8.68	1.90	1.90
Lease Area C	0.990	8.37	8.35	1.84	1.84
Lease Area D	0.989	8.96	8.93	1.92	1.92
Lease Area E	0.989	9.36	9.35	1.87	1.86
Lease Area F	0.988	8.97	8.94	1.81	1.81
Lease Area G	0.992	9.61	9.59	1.97	1.97
Lease Area H	0.990	9.99	9.96	1.85	1.85
Lease Area I	0.992	9.56	9.55	1.86	1.86
Lease Area J	0.992	9.31	9.30	1.83	1.84
Lease Area K	0.992	8.99	8.97	1.77	1.77
Lease Area L	0.990	9.73	9.70	1.85	1.85
Lease Area M	0.991	8.95	8.93	1.93	1.93
Lease Area N	0.991	9.49	9.48	1.94	1.94
Lease Area O	0.990	9.95	9.92	1.96	1.95

Table 5: Wind speed timeseries statistics between AceCAST and WRF

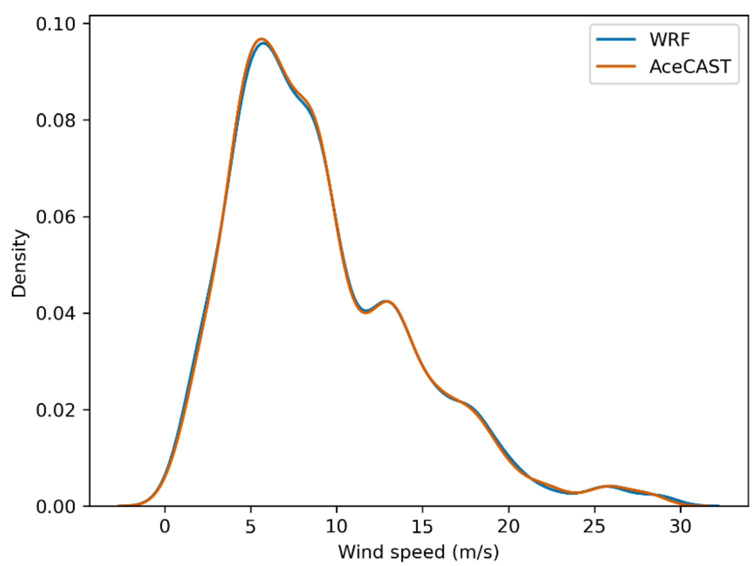

Figure 6: Distributions of mean 120-m wind speeds across Lease Area A.

Conclusions

The results from this validation study are unequivocal: AceCAST is a near-perfect and greatly accelerated replication of the standard CPU-based WRF model. Whether looking at mean atmospheric parameters modeled across the entire domain or timeseries analysis at specific coordinates, the agreement between AceCAST and WRF is exceptional.

Given such consistent performance, Veer Renewables has confidently embraced AceCAST as our go-to NWP platform for WakeMap and other WRF-based products. In doing so, we are not just advancing our technological edge; we are committed to delivering these cost and time savings directly to our clients, reaffirming our position as an industry leader in innovation, speed, and affordability.

About Veer Renewables

Veer Renewables was founded by Dr. Mike Optis in 2022 and provides advanced R&D solutions for wind energy stakeholders, with a special focus on wake modeling, wind resource assessment, and operational performance analysis. Dr. Optis is a world-leading expert in the intersection of meteorology, data science, and wind energy. Throughout his 15-year career, he has held roles with consultants, developers, and research institutes, including four years as a senior scientist at the National Renewable Energy Laboratory.

Document History

Version 1.0	2023-10-30	