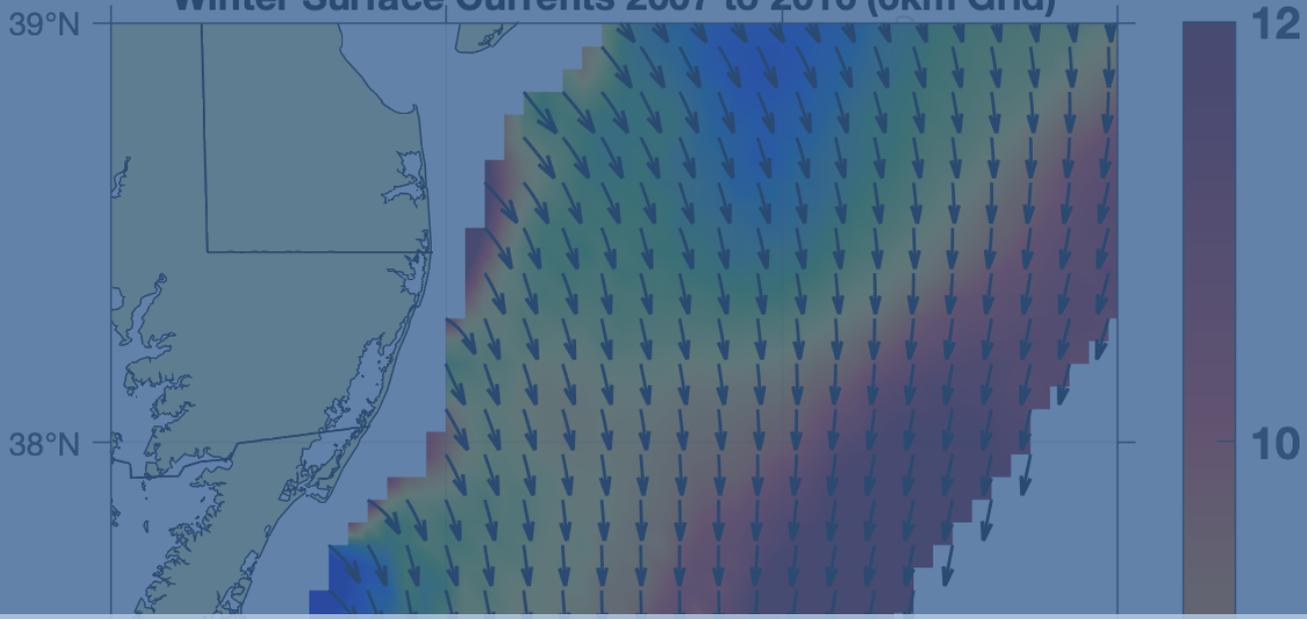


Winter Surface Currents 2007 to 2016 (6km Grid)



Oceanographic HF Radar Data Preservation in Wind Turbine

Final Report



CODAR



ODU



ECU



**NATIONAL
OFFSHORE WIND**
RESEARCH & DEVELOPMENT CONSORTIUM

Final Report

Final Report

Prepared for:

National Offshore Wind Research and Development Consortium

Prepared by:

CODAR Ocean Sensors LTD.¹

Mountain View, California

Authors: Josh Trockel¹, Dale Trockel¹, Chad Whelan¹,
Teresa Updyke², Hugh Roarty³, Patterson Taylor⁴, Mike Muglia⁴

Additional Affiliations:

Old Dominion University²

Rutgers University³

East Carolina University⁴

Title of the project:

Oceanographic HF Radar Data Preservation in

Wind Turbine Interference Mitigation

Agreement number: 111

Notice

This report was prepared by CODAR Ocean Sensors LTD. in the course of performing work contracted for and sponsored by the New York State Energy Research and Development Authority (hereafter “NYSERDA”). The opinions expressed in this report do not necessarily reflect those of NYSERDA or the State of New York, and reference to any specific product, service, process, or method does not constitute an implied or expressed recommendation or endorsement of it. Further, NYSERDA, the State of New York, and the contractor make no warranties or representations, expressed or implied, as to the fitness for particular purpose or merchantability of any product, apparatus, or service, or the usefulness, completeness, or accuracy of any processes, methods, or other information contained, described, disclosed, or referred to in this report. NYSERDA, the State of New York, and the contractor make no representation that the use of any product, apparatus, process, method, or other information will not infringe privately owned rights and will assume no liability for any loss, injury, or damage resulting from, or occurring in connection with, the use of information contained, described, disclosed, or referred to in this report.

NYSERDA makes every effort to provide accurate information about copyright owners and related matters in the reports we publish. Contractors are responsible for determining and satisfying copyright or other use restrictions regarding the content of reports that they write, in compliance with NYSERDA’s policies and federal law. If you are the copyright owner and believe a NYSERDA report has not properly attributed your work to you or has used it without permission, please email print@nyserda.ny.gov

Information contained in this document, such as web page addresses, are current at the time of publication.

Table of Contents

Notice	2
List of Figures	6
List of Tables	8
Acronyms and Abbreviations	10
1 Executive Summary	12
2 Background	13
3 Key Findings	15
4 Recommendations	18
5 Next Steps	19
6 Acknowledgements	20
Appendix A. Data Collection	21
A.1 2022 SeaSonde Data Collection	21
A.2 Historical Data Collection	23
A.3 SCADA Data Request.....	24
A.4 Usability of the Dataset	24
Appendix B. Wind Turbine Simulations	25
B.1 Introduction	25
B.2 WTI Simulation.....	25
B.2.1 Multi-Static WTI Simulation.....	25
B.2.2 Variability in Rotation Wind Turbine Rates.....	28
B.3 Final Dataset	31
B.3.1 WTI Parameters.....	32
Appendix C. Effect of Flagging on Sea Surface Measurements	34
C.1 Background	34
C.2 Impact of WTI on HF Surface Wave Radars	34
C.3 Experimental Design.....	37
C.3.1 Study Area	38
C.3.2 Kinds of Data	39
C.3.2.1 Monostatic and Radials.....	39
C.3.2.2 Bistatic and Ellipticals	40
C.3.2.3 Totals.....	41
C.3.3 Data Flagging.....	42

C.3.4	Processing Range-Doppler Spectra	43
C.4	Results	45
C.4.1	Radials.....	45
C.4.2	Totals	45
C.5	Final Remarks	46
Appendix D. Machine Learning Wind Turbine Interference Estimation		47
D.1	Background	47
D.2	WTI Estimation Method.....	50
D.2.1	Classification and RPM, RPM Variability, and Yaw Angle Estimation	50
D.2.3	Training and Testing Datasets	51
D.2.4	WTI Peak Amplitude Estimation.....	51
D.2.4.1	Regression and K-Nearest Neighbors	51
D.2.5	Testing.....	52
D.3	Results	52
D.3.1	Classification:.....	52
D.3.2	Rotation rate, Yaw, Rotation Rate Variation Estimation:	53
D.3.3	Amplitude Estimation:	53
D.4	Discussion	54
D.4.1	Limitations.....	54
D.4.2	Improved Flagging.....	54
D.4.3	Effect of Rotation Rate Variability	55
D.5	Final Remarks	56
Appendix E. Wind Turbine Interference Flagging Real-Time Software.....		57
E.1	Software Overview	57
E.2	WTI Mitigation Software Toolchain.....	57
Appendix F. Effects of Machine Learning Wind Turbine Estimation and Multisite Operation on Sea Surface Measurements		60
F.1	Methods	60
F.1.1	Simulation.....	60
F.1.2	Flagging.....	63
F.1.3	Surface Current Processing	66
F.2	Results	66
F.2.1	ML Assignment:.....	66
F.2.2	Radial Results	67
F.2.3	Totals Results.....	72

F.3	Discussion	74
Appendix G. Field Test		76
G.1	Field Test Preparation	76
G.2	Field Test.....	79
G.3	Discussion	81
References		82

List of Figures

Figure 1: Elliptical Current Vector Coverage LISL	23
Figure 2: An example of bistatic range-Doppler spectra. Range bins 0-40 contain the monostatic signal and 90 to 120 contain the bistatic signal.	26
Figure 3: A diagram of the bistatic transmission between two radar sites.	27
Figure 4: Bistatic cross spectra with simulated WTI added in both the monostatic and bistatic echo. Peaks can be seen in range bin 10 as well as range bin 100.	28
Figure 5: A histogram of the change in rotation rate over 10 minutes for the Block Island Wind farm. ..	29
Figure 6: Spectra from BLCK when the turbines had rotation rates listed in Table 4. WTI Peaks from the wind turbine can be seen in range bins 0-2. Simulated WTI Peaks can be seen in range bins 11-13.....	30
Figure 7: A range slice from range bin 1 of the range-Doppler spectra from BLCK showing the spreading of the actual wind turbine interference.	30
Figure 8: The difference between an oscillatory wind turbine rotation rate variation compared to a linearly increase rotation rate variation.	31
Figure 9: DUCK, LISL, and CEDR radar coverage map.....	32
Figure 10: SeaSonde range-Doppler spectra from Block Island with wind turbine interference in the first range cell (top) and flagged WTI (bottom). The flagged range-Doppler cells are blacked out in the spectra plot on the bottom. The magenta lines indicate the boundaries of the Bragg regions containing the sea echo.....	35
Figure 11: An example of a radar radial current measurement. The Blue and green vectors show individual radial vector measurements, and the red vector indicates the resulting averaged radial vector. The green vector indicates one that was eliminated because of WTI. The left plot shows the average when the flagged vector is removed, and the right plot shows what the averaged radial vector would be without the flagged vector removed.....	36
Figure 12: Diagram showing how multiple radars will infill the loss of data caused by WTI mitigation.	37
Figure 13: Coverage area of LISL, DUCK, and CEDR.....	39
Figure 14: Radial Vector Map LISL	40
Figure 15: Elliptical Coverage Map of Signals Transmitted at DUCK and Received at LISL	41
Figure 16: Total Vector Map Comprising Data from LISL, CEDR, and DUCK.....	42
Figure 17: Range-Doppler Spectra. Unflagged spectra are shown in the top panel, and flagged spectra are shown in the bottom panel of the top figure. The bistatic (bottom left) and monostatic (bottom right) flagged portions are cropped and shown below.	44
Figure 18: The Doppler frequencies of the first four positive and negative harmonic wind turbine interference peaks as the rotation rate varies. The top shows the impact of rotation rates for radars using a 1 Hz sweep rate, and the bottom plot shows the impacts on radars with a 4 Hz sweep rate. The gray rectangular regions indicate the typical Doppler frequencies containing the sea echo. The blue line shows the fourth positive harmonic and how aliasing loops around in Doppler. ...	48
Figure 19: The influence of yaw angle on the amplitude of the first four positive and negative WTI harmonic peaks. The Nacelle (yaw) angel represents the angle of the nacelle relative to the radar. Each harmonic, indicated by m, is plotted in a different color. A Nacelle angle of 0 is when the turbine blades are parallel to the radar.....	48

Figure 20: The spreading of simulated wind turbine interference peaks resulting from short time variability in a turbine rotation rate. The amplitude of the WTI has been increased to show the peaks clearly. 49

Figure 21: BLCK simulated spectra showing the effects of different kinds of variability in rotation rate on the characteristics of the WTI. The leftmost figure shows an oscillating increase in rotation from 6rpm to 7rpm over an integration period. The rightmost figure shows a linear increase in rotation rate from 6rpm to 7rpm over an integration period. The peaks are wider and lower in amplitude as the variance in rotation rate increases. The way the rotation changes determines the characteristics of the WTI harmonic peaks..... 55

Figure 22: Software implementation Diagram. 58

Figure 23: Map of the study area showing the location of the Cedar Island, VA (CEDR), Little Island, VA (LISL) and Duck, NC (DUCK) SeaSonde stations. 61

Figure 24: Radial vector count for each SeaSonde station CEDR (upper left), LISL (upper right), DUCK (lower left) and LIDU (lower right). Each panel contains the data from each test case no WTI (blue), WTI inserted into spectra (red) and WTI flagged (yellow)..... 62

Figure 25: Flagging using different methods. Top shows flagging using only ML to estimate the rotation rate. The red squares indicate the location of the flagged range-Doppler bins. Bottom shows flagging using ML to estimate rotation rate, then corrected using analytical methods. The red bounding boxes match the peaks more accurately in Right. This results in different sections of the Bragg being flagged. 65

Figure 26: Statistics of the radial data for each site. Top left: CEDR, Top Right: LISL, Bottom Left: DUCK, Bottom Right: LIDU. Within each image, top panel: Correlation between the no WTI data and the three test cases, WTI inserted into spectra (blue), WTI flagged (red), WTI flagged with machine learning and analytical techniques (yellow). Middle panel: root mean square error (cm/s) between no WTI and the three cases. Bottom panel: Percentage of vectors in the affected range cells that were altered..... 68

Figure 27: Average radial vector map for the CEDR station. Each map contains the data from each test case no WTI (top), WTI inserted into spectra (middle) and WTI flagged (bottom). The mean velocity is shown as the colorbar and the standard deviation is the size of the dot. 69

Figure 28: Scatter plot of radial velocity data matching in range and bearing between the no WTI case (x axis) and the WTI flagged case (y axis) for the CEDR station at January 02, 2020 05:00 GMT. 70

Figure 29: Root mean square error for the CEDR station for each test case WTI inserted into spectra (left), WTI flagged (middle), WTI flagged with machine learning and analytical techniques (right). The locations of the two wind turbines are shown as black circles. 71

Figure 30: Same caption as Figure 29 for the LISL station..... 71

Figure 31: Same caption as Figure 29 for the DUCK station. 72

Figure 32: Same caption as Figure 29 for the LIDU station..... 72

Figure 33: Total vector map with change detected by processing using method 2 and 3. From left to right the figure shows totals processed using method 1, method 2, and method 3..... 74

Figure 34: Range slice showing normalized signal to noise ratio of loop three with no noise added (blue) and noise added (red). 77

Figure 35: An example cross spectra file at LISL during the field test. There are no observable WTI peaks. 79

Figure 36: Average radial velocities at LISL during the field test. 80

List of Tables

Table 1: Dataset Configuration Settings.....	22
Table 2: Dataset Configuration Settings – Station Alignments	22
Table 3: Percentage of Spectra Files Available in January 2020	23
Table 4: Operating parameters of the turbines during the Doppler FFT integration period for the spectra plots in the two Figures above.....	30
Table 5: The range of wind turbine parameters used for the WTI simulations.....	33
Table 6: Statistical Results of Radial Vector Comparison Between Data With no Flags and Data which has been Flagged.	45
Table 7: Statistical Results of Total Vector Comparison Between Unflagged and Flagged Data Sets	46
Table 8: Accuracy of classification of how many turbines were spinning.	52
Table 9: Accuracy of estimation of rotation rate, yaw angle, and variation. Mean error in prediction and standard deviation of the error shown.	53
Table 10: The RMSE and Pearson Correlation of prediction of the amplitude of the first positive and first negative harmonic peaks using polynomial regression.....	53
Table 11: The RMSE and Pearson Correlation of prediction of the amplitude of the first positive and first negative harmonic peaks using KNN regression.	54
Table 12: Location of the two turbines used in the study.	63
Table 13: Range cells where the wind turbines are located relative to each radar station and range cells where WTI is expected.....	63
Table 14: Summary of four methods of processing radials.....	66
Table 15: Performance metrics of ML model in predicting rotation rates.	67
Table 16: Radial statistics. Radials from CEDR, DUCK, and LISL were processed with four methods: 1. with no WTI added, 2. with WTI added and no mitigation performed, 3. with WTI added and WTI mitigation performed using only ML, and 4. with WTI added and mitigation performed using ML and analytical methods. Methods 2-4 were compared to method 1 for each site. The results were averaged together and average # vectors lost, # vectors added, # vectors changed, and RMSE in velocity among the changed vectors are reported here. The average total number of vectors compared was 45,070.	67
Table 17: Totals stats using two sites. The same comparisons shown in Table 16 were repeated at the totals level for totals processed using radials from CEDR and DUCK. The total number of vectors in method 1 is 117,829. This same grid is used for comparisons in the case of three sites as well as three sites and ellipticals.....	73
Table 18: Totals stats using three sites. The same comparisons shown in Table 16 were repeated at the totals level for totals processed using radials from CEDR, DUCK, and LISL.....	73
Table 19: Totals stats using three sites and ellipticals. The same comparisons shown in Table 16 were repeated at the totals level for totals processed using radials from CEDR, DUCK, LISL, as well as ellipticals received at LISL from DUCK.	73
Table 20: results of different machine learning models. “Separate Models” represents two different prediction models: one for when the two turbines are rotating at similar rates, another for when	

they are different. “Single Prediction Model” represents a single model used for both scenarios.
Classification accuracy is the percentage of correct classifications..... 78

Acronyms and Abbreviations

BLCK	Block Island, Rhode Island radar
CEDR	Cedar Island, VA radar
CVOW	Coastal Virginia Offshore Wind
DUCK	Duck, NC radar
FFT	Fast Fourier Transform
FMCW	Frequency Modulated Continuous Wave
GPS	Global Positioning System
HF	High Frequency
HFR	High Frequency Radar
Hz	Hertz
ITU	International Telecommunication Union
KNN	K-Nearest Neighbors
LIDU	Little Island Park to Duck multistatic radar system
LISL	Little Island Park, Virginia Beach, VA radar
NEC	Numerical Electronic Code
MHz	Megahertz
ML	Machine Learning
MRE	Marine Renewable Energy
MUSIC	Multiple Signal Classification
NOAA	National Oceanic and Atmospheric Administration
OTT	Ocean Technology Transition
RCS	Radar-Cross-Section
RMSD	Root Mean Squared Difference
RMSE	Root Mean Squared Error
RPM	Rotations per minute

SCADA	Supervisory control and data acquisition
SNR	Signal to Noise Ratio
WTI	Wind Turbine Interference

1 Executive Summary

This report summarizes the key findings of this National Offshore Wind Research and Development Consortium-funded project to mitigate offshore wind turbine interference (WTI) in oceanographic radars. Evidence and analysis of each of the findings can be found in the accompanying Appendices.

This WTI mitigation research and development has shed new light on the effects of WTI on high frequency (HF) radar networks, specifically the long-range systems operating in the 4.4 – 5.3 MHz band, and new mitigation methods have been developed that overcome shortcomings of previous methodologies. Furthermore, the funded mitigation efforts have improved CODAR’s WTI simulation tool through extensive testing and comparison with WTI found in radar data at Block Island.

The key findings of this project include:

- Increasing the geometric redundancy in surface current measurements (i.e. increasing the number of observations from different directions) is the most effective way tested to mitigate WTI. Designing radar networks or adding to existing ones to increase the amount of overlapping coverage from multiple radar sites shows the greatest reduction of the effect of WTI of any WTI mitigation method tested to date.
- Machine learning (ML) is effective at estimating rotation rate, yaw angle, and variation in rotation rate from the WTI peaks in HF Doppler spectra for a small number of turbines.
- The characteristics of the WTI in HF Doppler spectra are extremely sensitive to changes in rotation rate of the wind turbines. The sensitivity of WTI characteristics to the variability of a turbine's rotation rate has blocked efforts to separate WTI from the sea echo in the frequency domain.
- Improved WTI simulations realistically simulate changing rotation rates within the spectral integration period.
- Improved WTI simulations can be used to augment datasets to train ML models for flagging.
- A more robust and dependable WTI flagging method is achieved by combining ML model techniques with previously developed analytical techniques (Trockel et al. 2021) to estimate turbine rotation rates. The combination of methods outperforms either method in isolation.
- Real-time software which incorporates the ML rotation rate estimate technique has been developed and tested for up to two turbines.
- Real-time software which suppresses range-Doppler bins flagged with WTI from current processing has been developed and tested. The software can run as part of the real-time processing at each HFR site.

2 Background

Oceanographic high frequency radar (HFR) is a remote sensing tool that allows for the measurement and monitoring of ocean surface currents, waves, and winds in near real-time. This technology utilizes electromagnetic waves in the HF band (3-30 MHz) to measure the speed and direction of the ocean surface, providing maps of the dynamics of the ocean at a regional scale. HFR has become an important tool for oceanographers due to its ability to provide continuous, high-resolution data on ocean circulation and coastal processes, making it particularly useful for applications such as coastal hazard management, search and rescue, oil spill response (Roarty et al. 2019, Roarty et al. 2020), and marine renewable energy (MRE) applications (Muglia et al. 2020). Oceanographic HFR is poised to play an increasingly important role in understanding the complex and ever-changing ocean environment.

Networks of HFR stations have operated along the Mid Atlantic coast of the United States for the past thirty years. Ocean currents drive coastal surface waters that have profound societal impacts for the 50% of the nation's citizens that live within 50 miles of the coast (Wilson & Fischetti, 2010). The National Oceanic and Atmospheric Administration (NOAA) has 19 program requirements for surface current observations. High Frequency radar (HFR) has been recognized nationally as a cost-effective solution for providing these surface current measurements on an operational basis (Harlan, 2009).

Offshore wind power is rapidly developing off the East Coast of the United States, with several large-scale projects underway. As of early 2023 there are two operational offshore wind farms in the United States (Block Island Wind Farm – 5 turbines and Coastal Virginia Offshore Wind – 2 turbines), but more than 20 other projects are in various stages of development. These projects have the potential to generate significant amounts of clean energy, and they are expected to create thousands of jobs and spur economic growth in coastal communities. The Biden administration has set a goal of developing 30 gigawatts of offshore wind energy by 2030.

When offshore wind turbines are in the measurement area of a coastal HFR (typically up to 180 km from the coast), they act as reflectors and can cause an interference signal that can be mixed with the sea surface echo from which oceanographic data is derived. When mixed together, the wind turbine interference (WTI) causes errors when determining the direction of arrival of the sea echo resulting in velocity errors and gaps in the oceanographic data (Teague & Barrick, 2012; Trockel, Rodriguez-Alegre, Barrick, & Whelan, 2018; Wyatt et al., 2011). These distortions reduce the accuracy with which HF radars can measure oceanographic data such as sea-surface currents. To mitigate these errors, CODAR Ocean Sensors Ltd. and Rutgers

University have been at the forefront of the efforts to characterize and mitigate the WTI impacts on coastal HFR measurements (Trockel, Rodriguez-Alegre, Barrick, & Whelan, 2018; Trockel et al., 2021).

The mitigation techniques developed prior to this project were designed to detect the Doppler bins in sea echo that were impacted by the WTI and remove them from processing for currents, which can result in the loss of meaningful oceanographic data. They were also limited to systems operating with a sweep rate of 4 Hz or higher, which is a problematic waveform parameter for large HFR networks sharing the limited frequency bands available to HFR transmissions. The purpose of this project is to develop and test a mitigation method that can remove the WTI while preserving most, if not all, of the oceanographic data at any operational sweep rate.

3 Key Findings

A significant finding from this project is that sites with overlapping coverage create redundant surface current observations, which allow for the backfill and correction of data that has been changed or lost because of WTI. Designing radar networks to have increasing levels of overlapping coverage was found to be the best method of mitigation of any tested to date. Increased overlapping coverage decreased both the amount of data lost from WTI as well as reduced the error in measurements which were changed by WTI signals. This was demonstrated using three radar sites with overlapping coverage, one of which was operating with bistatic capabilities (i.e., the radar received and processed sea echo from its own transmitted signal as well as transmissions from other sites). Two tests were conducted using the location of two turbines in the overlapping coverage area of the three radar sites: 1) a set of range-Doppler bins representative of the turbines spinning at a consistent rate were flagged in the spectra at each of the three sites, and processed using the software toolchain developed during this project which suppresses flagged range-Doppler bins from sea surface current measurements. The location in Doppler was selected so that it would consistently be in the sea-surface current data. 2) WTI from the turbines was simulated and added to cross spectra from each of the three sites. WTI from the two turbines was simulated with transient wind conditions so that the rotation rate of the turbines was different between successive cross spectra. The rotation rates were once again selected to interfere with sea-surface current measurements, but unlike test 1, the Doppler bins were not consistent across successive cross spectra. The first method was used to infer the effect of a set of turbines with consistent rotation rates that placed WTI within the Bragg region. This test helps us infer the effect of WTI from larger wind turbines where WTI could be placed consistently in the same Doppler frequencies. The second test simulates the effect of transient wind conditions placing WTI into the Bragg. This test helps us infer the effect of WTI from turbines like those current in operation in the United States, where their rotation rate power-curves places WTI into the Bragg only during times of changing winds. Both tests showed that the effects of both WTI and removed data from flagging are mitigated by using more sites with overlapping coverage to produce total surface current velocities. The redundancy in multi-site data provides a safety net against both signals mixed with WTI and data loss from WTI mitigation. This is an encouraging finding because it provides a simple and effective method of improving data reliability across a wide range of WTI impacts.

This project also explored the viability of using machine learning (ML) models to separate WTI from the sea echo in SeaSonde Doppler spectra as well as estimate key turbine attributes such as rotation rate, yaw angle, and variance in rotation rate. We found that ML can successfully be used to identify wind turbine

interference as well as estimate rotation rate, yaw angle, and variance in rotation rate. ML models have also proven to be robust with low sensitivity to different noise conditions and radar operating parameters.

Through exploring the effect of rotation rate and rotation rate variability on the characteristics of WTI harmonic peaks, it was found that rotation rate variability dramatically influences the shape, width, and amplitude of harmonic peaks. A change or variation in rotation rate over the integration period of the cross spectra leads to a spread in Doppler of the WTI harmonic peaks as well as a different characteristic shape. Rotation rate changes are the results of the variability of the wind and the proprietary power curves of the turbines. While we can reliably predict the total amount of variation in rotation rate over the course of an integration period, the details of how that variation took form within the integration period are needed to resolve the amplitude and location of the peaks. The variation on this granular scale is directly related to instantaneous changes in wind conditions as well as the turbine power curve, which dictates how turbines respond to changes in wind conditions second by second. Due to the unpredictable nature of these conditions, we are unable to predict the variation in rotation rate with enough detail to effectively separate the WTI from the sea echo despite the ability to get reliable rotation rate, yaw, and rotation total variation in rotation rate across an integration period estimates from ML models.

Additionally, during this project, the WTI simulation tool developed by CODAR Ocean Sensors Ltd. (Trockel, Rodriguez-Alegre, Barrick, Whelan, et al., 2018) was updated to include different rotation rate variability models. The new models have led to simulated WTI peaks visually consistent with what is observed in SeaSonde Doppler spectra obtained at BLCK when the turbines were known to have a changing rotation rate. Allowing the rotation rate to change non-linearly over an integration period better mimics the effects of changing wind conditions during a turbine's rotation rate ramp-up and ramp-down periods. This tool can be used to help train ML models to classify WTI as well as predict rotation rates. Simulation can be used to augment real-world data where Supervisory Control And Data Acquisition (SCADA) data is available and provide datasets for training when no SCADA data can be acquired.

The rotation rate estimates provided by ML models can be used to significantly improve flagging schemes. Current flagging methods (Trockel et al., 2021) are limited to situations where the WTI harmonic peaks preserve a symmetrical nature in the cross spectra. This limits the operation of HF radars to 4 Hz. Using ML, rotation rates can be estimated regardless of the operating sweep rate of the radar. The ML predicted rotation rates are most effective when used to narrow the search windows for the analytical approach developed by Trockel et al. (2021). This hybrid approach of estimating rotation rates provides an accurate method for identifying and flagging WTI peaks in cross spectra.

Analysis of SCADA Data from the Block Island wind farm indicates that WTI is only mixed with the oceanographic data when the rotation rates of the turbines are highly sensitive to wind speed leading to wide WTI Peaks that shift in frequency from one Doppler spectra to the next. The effects of these transient peaks are reduced during the averaging used by the radars to produce radial and total sea surface current observations. Due to the effectiveness of averaging, the impacts of the transient peaks are similar when flagged or unflagged. The losses and errors in the data due to the WTI are best mitigated by increasing the reliability of these averages. This can be achieved by increasing the redundancy in oceanographic measurements by overlapping the coverage of the radars and isolating the outliers caused by WTI. The placement of WTI outside of the Bragg is due to the current size of wind turbines deployed in the field. However, as turbines increase in size, their optimal rotation rate will slow. The optimal rotation rate is where turbines are designed to spend most of their time. If the new optimal rotation rate places interference within the Bragg, the flagging methods developed in this project are expected to provide an added benefit.

4 Recommendations

Below is a summary of the recommendations for HF radar operators and offshore wind manufacturers to mitigate the adverse effects of WTI on HF radar data:

- When operating sites outside of a network, using a sweep rate of 4 Hz will limit the amount of time WTI is mixed with sea echo. *
- When operating a network, it is not always possible to operate radars at 4 Hz. In these cases, the effects of WTI can be best mitigated by adding more radars to the network. *
- As turbine sizes increase and more time is spent at rotation rates that cause interference for sea echo, flagging schemes can be implemented to real-time software to suppress the range-Doppler bins from processing that are contaminated with WTI.

*The best practice recommendations for operating HFR networks to minimize the impact of WTI are currently limited by the small number of frequency bands allotted to HFR by the International Telecommunication Union (ITU). To fully implement these methods and better mitigate WTI, the currently allotted ITU bands would need to be expanded to allow for more stations and sweep rates of 4+ Hz within a HFR network.

5 Next Steps

The work done over the course of this project will be continued through a grant funded by the National Oceanic and Atmospheric Administration's (NOAA) Ocean Technology Transition (OTT) program. Through this project, we have shown that ML can be used to flag two simulated turbines in the field of view of the radar and explored the effects on surface current measurements processed with and without WTI and flags. Through the OTT effort, the use of ML and analytical flagging on real-world data with more turbines will be investigated. If rotation rate data are not provided, cameras will be used to record and extract the rotation rates during operation. Further validation of mitigation schemes explored in this project will be performed with drifter deployments to test the downstream effects of mitigation schemes on radar surface current observation and WTI mitigation. Additionally, toward the end of this project, Ørsted has provided historical SCADA data, which will be used during the OTT-funded project to further test the methods developed during this project with real-world data.

SeaSonde software is now able to process and suppress range-Doppler segments of cross spectra that have been flagged with WTI. This software enables the use of current mitigation schemes in the field as well as facilitates further research and experimentation done in the field by both CODAR and other organizations interested in mitigating WTI. The findings from this project will be disseminated to both radar operators and other stakeholders so that the radar systems can continue to operate with minimal impact from planned wind farms.

6 Acknowledgements

The operation of the radar network during this project was funded by NOAA Award Number NA16NOS0120020 “Mid-Atlantic Regional Association Coastal Ocean Observing System (MARACOOS): Powering Understanding and Prediction of the Mid-Atlantic Ocean, Coast, and Estuaries”. Sponsor: National Ocean Service (NOS), National Oceanic and Atmospheric Administration (NOAA) NOAA-NOS-IOOS-2021-2006475, Integrated Ocean Observing System Topic Area 1: Implementation and Development of Regional Coastal Ocean Observing Systems.

The operation of the DUCK HFR was funded by the Southeast Coastal Ocean Observing Regional Association under Award No. NA16NOS0120028. SECOORA “Operating and Maintaining High Frequency Radars” Sponsor: National Ocean Service (NOS), National Oceanic and Atmospheric Administration (NOAA) NOAA-NOS-IOOS-2021-2006475, Integrated Ocean Observing System Topic Area 1: Implementation and Development of Regional Coastal Ocean Observing Systems.

Appendix A. Data Collection

A.1 2022 SeaSonde Data Collection

A month of data was collected at the LISL and DUCK stations from March 22 to April 22, 2022. Range series files (i.e. files containing data after the range FFT of the SeaSonde received signal) were collected from each site. LISL range series are available for the entire data collection period. DUCK range series are available from April 5 to April 22 and half hourly cross spectra (CSS files) will be used for the analysis from Mar 22 to April 5. The CEDR station no longer has access to power from the electric grid and this presented challenges to including this station in the dataset. Only two alternate locations near CEDR have power and would be suitable for a radar system: Wallops Island and the south end of Assateague Island. After failed attempts to obtain permission at those alternate sites, the focus shifted to running a short test of the system on Cedar Island using generator power. The equipment was re-installed on the island in two field trips on March 8 and Mar 21. On March 22, a few hours of data were collected with all three SeaSonde systems (DUCK, LISL and CEDR). CEDR collected data while a Honda generator (eu2000i) supplied power for the electronics. Several factors contributed to the limited amount of time available for this initial collection period. The generator could not be run without supervision and the island is only accessible for a window of time around high tide. However, during this time, the team was able to set up all three stations to allow for the best opportunity to clearly identify WTI in the Doppler spectra and was also able to determine appropriate alignments for multi-static operation. Table 1 lists the configuration settings for the dataset that are common to all stations, while Table 2 provides alignment settings for each station. The alignments are small precise offsets in timing for the transmit signals. The timing offsets are necessary for multiple sites operating at the same frequency so that the sea echo reflections of transmissions originating from each station will arrive at the receiving station at different times, keeping the signals separated in range in the recorded Doppler spectra.

Table 1: Dataset Configuration Settings

	Mar 22 - Apr 14	Apr 14 - Apr 18	Apr 18 - Apr 22
Frequency (MHz)	4.537183	4.537183	4.537183
Sweep Rate (Hz)	4	4	1
Blanking Period (μs)	972	1945.6	1945.6
Doppler Cells	4096	4096	1024

Table 2: Dataset Configuration Settings – Station Alignments

Alignment (μs)	DUCK	LISL	CEDR
Mar 22	5000	11807	16815
Mar 22 - Apr 14	5000	11807	N/A
Apr 14-Apr 18	8000	1000	N/A
Apr 18- Apr 22	5000	2000	N/A

The Elliptical Suite R2 software license was obtained and installed at the LISL station. Since April 1, 2022, the LISL site has collected and processed bi-static echo from the DUCK transmitter in real-time to produce elliptical surface current data in the area where the wind turbines are located (Figure 1).

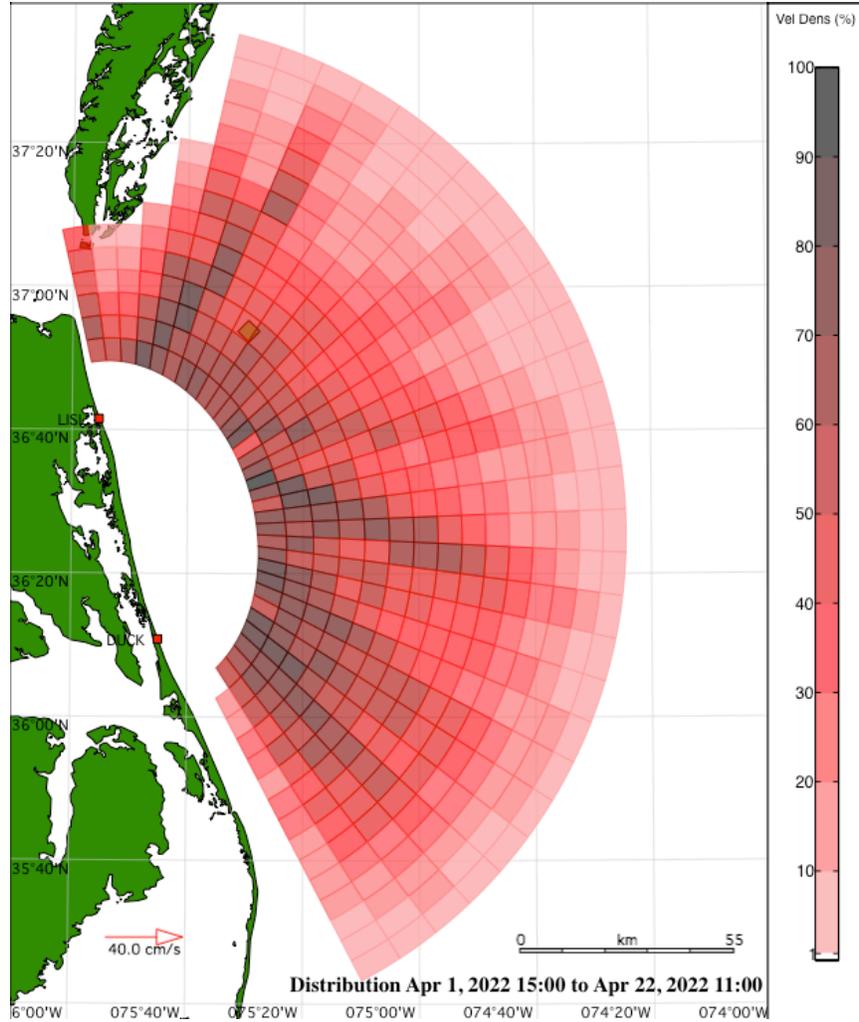


Figure 1: Elliptical Current Vector Coverage LISL

A.2 Historical Data Collection

ODU provided CODAR with a data set of SeaSonde® cross spectra from LISL, DUCK and CEDR archived prior to the rotation of any turbines in the CVOW wind farm. This historical data set spans the month of January 2020. Spectra files were recorded every half hour.

Table 3: Percentage of Spectra Files Available in January 2020

SeaSonde Site	LISL	CEDR	DUCK
Percentage of coverage in Jan 2020	98.3%	94.9%	99.4%

A.3 SCADA Data Request.

To proceed with tasks four and five of the project it is necessary to have a matched data set containing data from both the radar and the SCADA systems of the wind turbines in the field of view of the radars. A request was sent to Ørsted for a month of SCADA data from the Block Island wind turbines. Several follow-up requests were sent, but were not answered. Similar problems were encountered with all efforts to obtain SCADA data from the CVOW wind farm. We finally did receive a two-year data set of rotation data from Ørsted on April 3, 2023 but this project was concluded by then. Hopefully the rotation data can be utilized in future efforts.

Despite this setback, SCADA Data and radar data from the Block Island wind lease area obtained under a NOAA OTT project were made available for this project. The data set contains ten-minute outputs from the Block Island wind turbines, including rotation rates, and yaw angles. The data from the 5 MHz SeaSonde were supplied by Rutgers University. The combined data set includes data from March 1, 2021, to April 30, 2021.

A.4 Usability of the Dataset

The data collected, while not complete, will be sufficient for developing and testing the mitigation strategies outlined in the contract. The historic radar data set along with the radar data collected during the field test will allow the testing of backfilling gaps in the SeaSonde cross spectra caused by flagging and excluding range-Doppler bins containing WTI. Additionally, the 2022 field data set will be used in conjunction with the wind turbine interference tool developed by Trockel and Whelan (2021) to train and test the machine learning models proposed for the latter stages of this project. The simulated data will be augmented with the Block Island data set to provide a real-world example of the effectiveness of the method.

Appendix B. Wind Turbine Simulations

B.1 Introduction

One of the difficulties of developing and testing coastal HF radar wind turbine interference (WTI) mitigation tools is the limited number of radars near offshore wind turbines. Thus, WTI simulation tools are necessary so WTI mitigation strategies are in place before offshore wind turbines become more prolific. Funded by the Bureau of Ocean Energy management, Trockel, Rodriguez-Alegre, Barrick and Whelan (2018) developed a WTI simulation tool to add WTI from a single Wind turbine to a SeaSonde range-Doppler spectra file. Later Trockel et al. (2021) extended the tool to add interference from an arbitrary number of wind turbines in the coverage area of an HF Radar. This NYSERDA report describes how the WTI simulation tool was extended to simulate WTI in both the monostatic signal (i.e., the signal transmitted and received by the same radar) as well as the multi-static signal (i.e. the signal transmitted and received by different radars in the network). The simulation tool was then calibrated and used to add a simulated WTI to a network of radars in North Carolina and Virginia.

B.2 WTI Simulation

The WTI simulation tool previously developed by CODAR Ocean Sensors Ltd. was extended to enable the simulation of bistatic WTI. This tool simulates the WTI of a configurable set of wind turbines (position in latitude-longitude coordinates, as well as turbine-specific parameters such as hub height, blade length, and maximum rpm). The simulation tool uses Numerical Electronic Code (NEC) (Burke et al., 1979) to simulate the radar cross-section (RCS) of individual turbines with their blades at various rotation angles and yaw angles. The set of RCSs are then interpolated and scaled to create a simulated time series of the reflected signal from each of the turbines received by the radar. The received signal is then processed with CODAR's Dual FFT FMCW (Barrick et al., 1994) method to obtain the WTI range-Doppler spectral components which are then added to the input radar range-Doppler spectra files. Previously the simulation tool was only used for a site operating monostatically.

B.2.1 Multi-Static WTI Simulation

Using GPS pulse per second and a high-precision oscillator to time the sweeping and pulsing of each transmitter in an HF radar network, all SeaSonde stations can operate simultaneously in the same band without interference. Furthermore, the timing of the start of each station's sweep can be set such that each can receive the sea echo from the transmission of one or more the others (Lipa et al., 2009). For

oceanographic HFR employing a linear FMCW sweep, range is determined by frequency difference due to time delay of the radar echoes. By increasing the delay in start time of the sweep at a second synchronized station, the bistatic sea echo from the second station's transmission can be artificially separated in range from the monostatic sea echo. For example, Figure 2 shows an example of a bistatic range-Doppler spectra from a radar receiving sea echo from its own transmitted signal as well as that from another site. The first set of range bins (0 to 40) shows the sea echo from the signal transmitted from the receiving radar site. The second set of range bins (90 to 120) shows sea echo from the signal transmitted from a second radar. Due to the high precision sweep start delay set by the operator plus the time delay from station separation, the echo from the second site's signal is offset by 90-range bins.

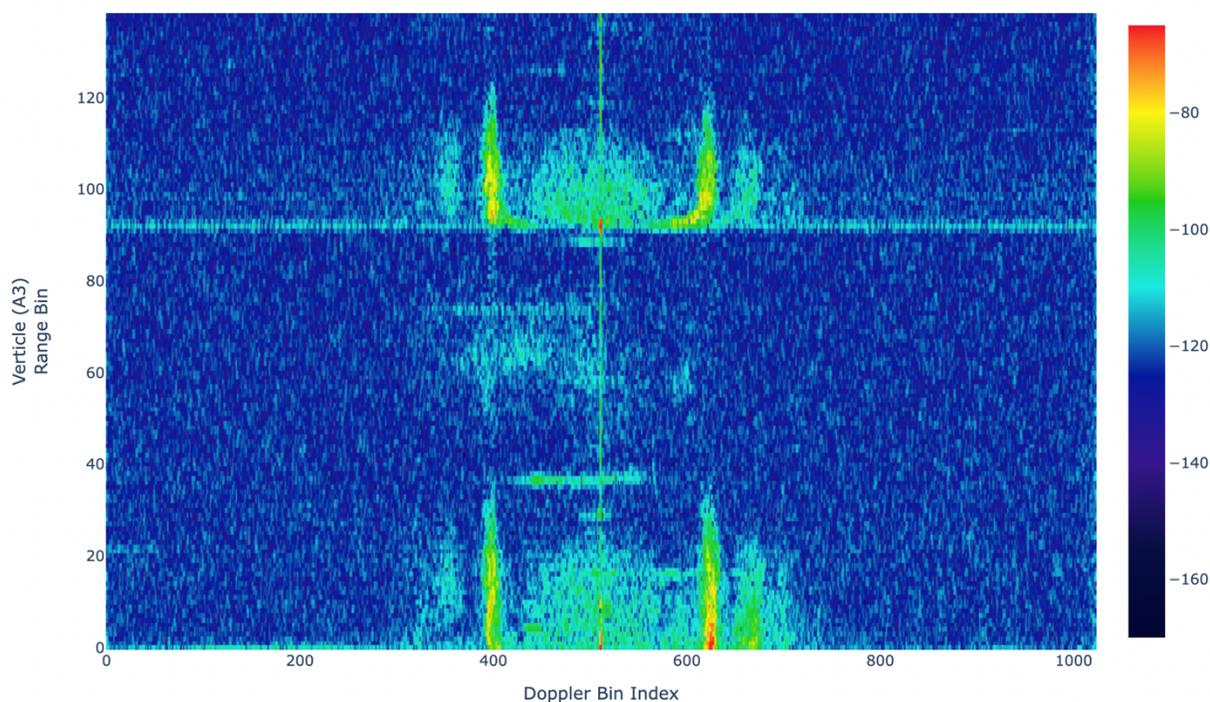


Figure 2: An example of bistatic range-Doppler spectra. Range bins 0-40 contain the monostatic signal and 90 to 120 contain the bistatic signal.

When simulating WTI in the monostatic case, the range of each turbine from the radar is used to adjust the timing of the simulated signal at the receiver which places the strong interference peaks in the correct range bin. For the multi-static case, the signal path is different and the range timings must be adjusted to account for this. Figure 3 uses a representative two-radar network operating as a bistatic pair to demonstrate how the multi-static WTI range factor is calculated. The bistatic range, R_{wti} , of the WTI transmitted from Radar 2 and received at Radar 1 is given by:

$$R_{wti} = R^* + \frac{r_1 + r_2}{2},$$

where R^* is the range offset where the bistatic signal starts (in the case of monostatic signals the offset is 0), r_1 is the distance from the transmitting radar at site one to the turbine, and r_2 is the distance from the turbine to the receiver radar at site two. Using R_{wti} as the range of the turbine when simulating the WTI signal the bistatic WTI will be placed at the correct range in range-Doppler spectra. This allows the tool to add the simulated WTI to the monostatic portion of the spectra as well as place it in the correct range and Doppler bin for each multi-static signal in the input spectra.

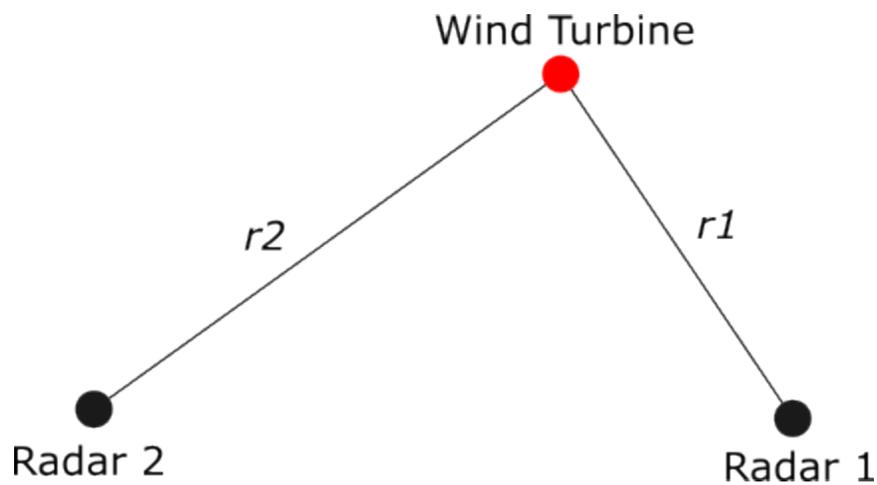


Figure 3: A diagram of the bistatic transmission between two radar sites.

An example of the bistatic range-Doppler spectra with added WTI in both the monostatic and bistatic portion of the range-Doppler spectra is shown in Figure 4.

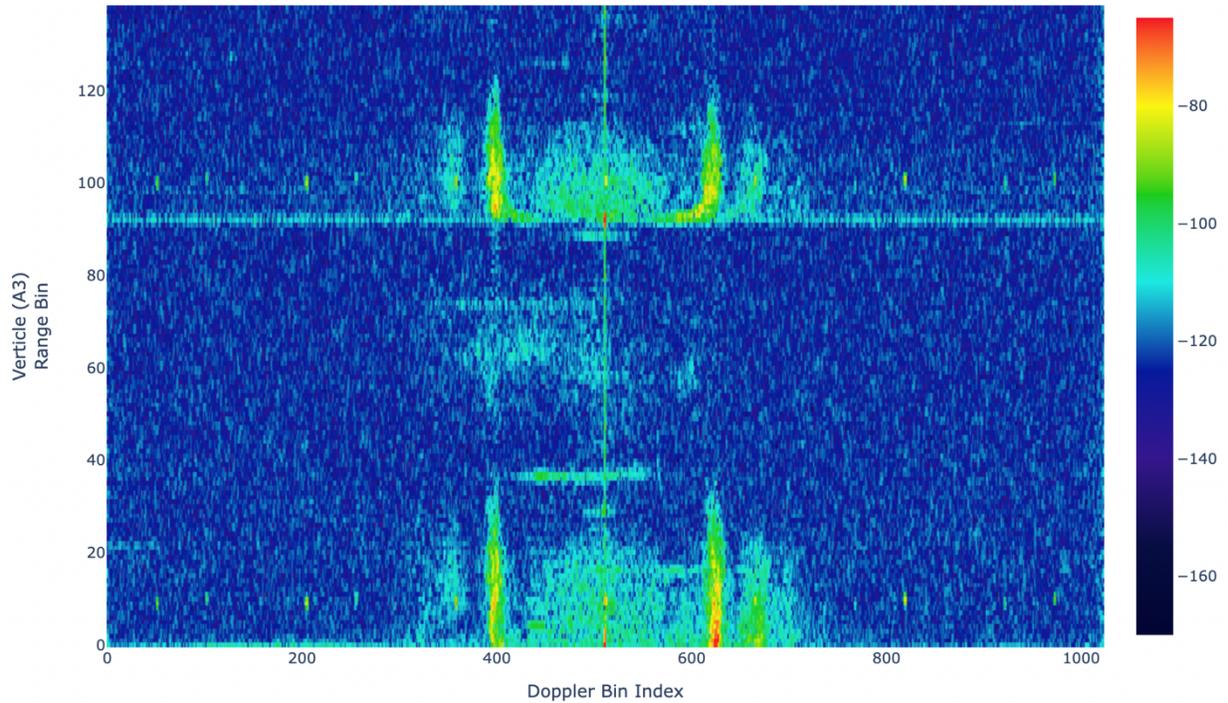


Figure 4: Bistatic cross spectra with simulated WTI added in both the monostatic and bistatic echo. Peaks can be seen in range bin 10 as well as range bin 100.

B.2.2 Variability in Rotation Wind Turbine Rates

Analysis of the wind turbine supervisory control and data acquisition (SCADA) data and SeaSonde range-Doppler spectra collected in the field of view of wind turbines, indicates that the rotation rates of the turbines are not steady, but rather are variable during the Doppler integration period of the radar. Figure 5 shows a histogram of the change in rotation rate (labeled drpm) in rotations per minute (RPM) for ten-minute time periods over one month at the Block Island wind farm. The majority of the time the rotation rate is steady, but there were regularly times when the variation over ten minutes is up to 1 RPM.

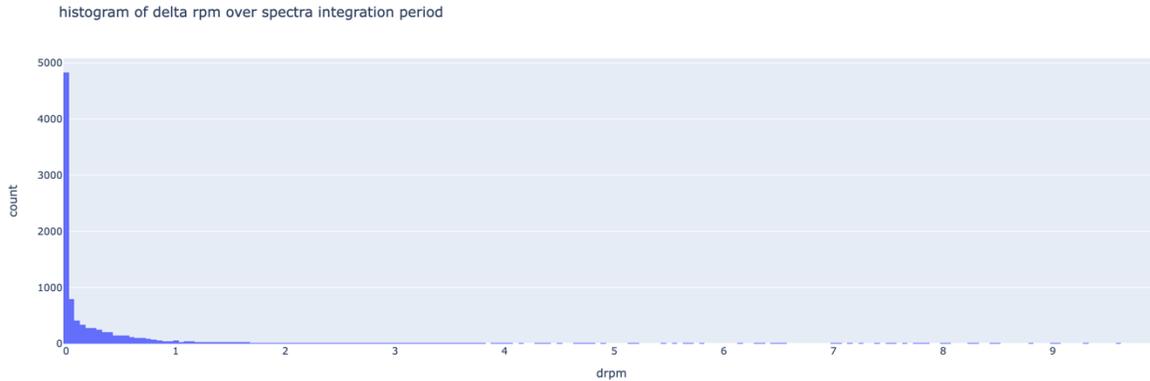


Figure 5: A histogram of the change in rotation rate over 10 minutes for the Block Island Wind farm.

The variable rotation rates cause the WTI peaks to spread out in Doppler. This can be seen in Figure 6 and Figure 7. The WTI peaks in the first range bin are much wider than the narrow two Doppler-bin wide peaks seen when the rotation rates are stable across the Doppler integration period. When the rotation rate varies, the peaks appear to be more like mounds than sharp peaks as can be seen in Figure 7 around Doppler cells 1500 and 2500.

Furthermore, the nature of the rotation rate variation over the Doppler integration period is as important as the maximum range of variation. Figure 8 shows two different changes to the rotation rate: a linear change and an oscillating change. The linear change produces smooth and wide peaks while the oscillating change produces noisy, wider peaks. The effects of the oscillating variability more closely match what is observed in SeaSonde range-Doppler spectra during the times of turbine rotation rate variability. For the simulations using a variable rotation rate, the simulation tool has been updated to have the option to use the oscillating variable rotation rates.

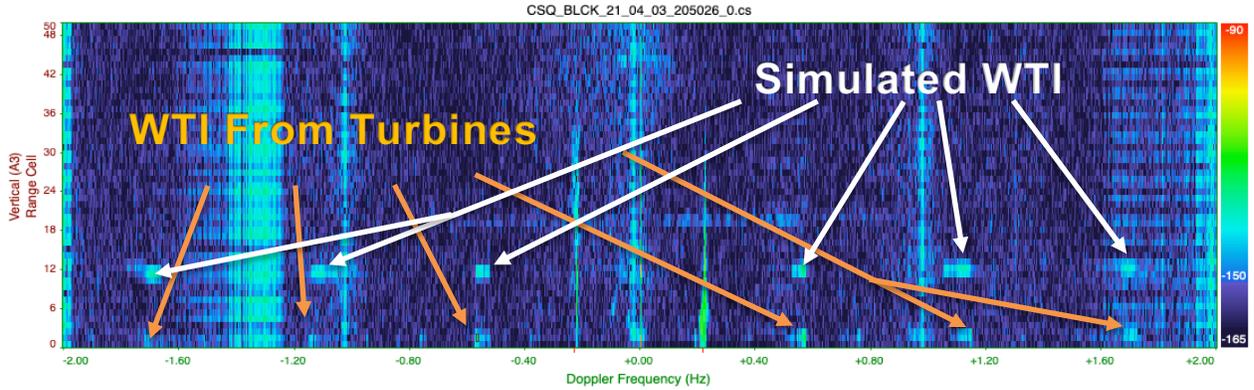


Figure 6: Spectra from BLCK when the turbines had rotation rates listed in Table 4. WTI Peaks from the wind turbine can be seen in range bins 0-2. Simulated WTI Peaks can be seen in range bins 11-13.

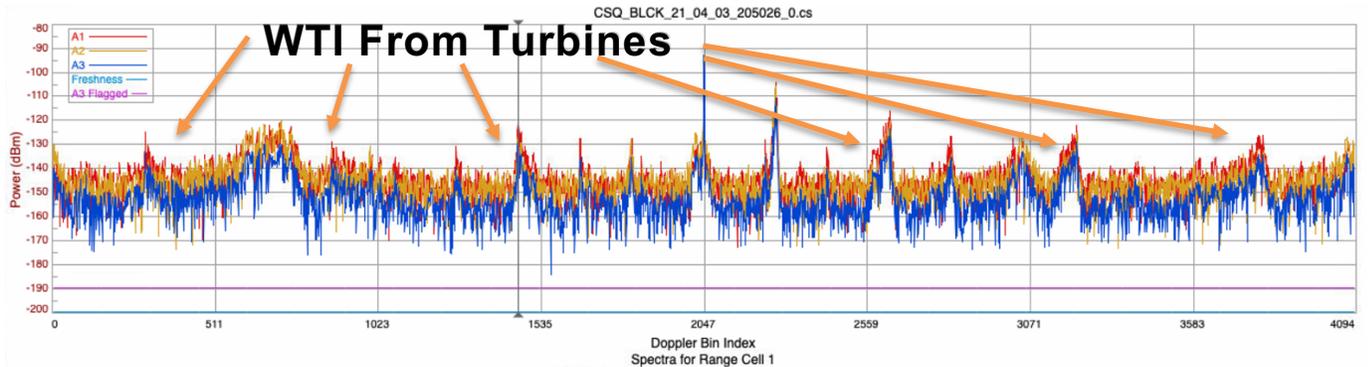


Figure 7: A range slice from range bin 1 of the range-Doppler spectra from BLCK showing the spreading of the actual wind turbine interference.

Table 4: Operating parameters of the turbines during the Doppler FFT integration period for the spectra plots in the two Figures above.

<i>Turbine</i>	<i>RPM</i>	<i>YAW</i>	<i>Scale</i>	<i>Variance In RPM</i>
<i>BIW01</i>	0.02	26°	-10	0.02
<i>BIW02</i>	11.275	45°	-10	0.17
<i>BIW03</i>	11.345	68°	-10	0.13
<i>BIW04</i>	11.32	72°	-10	0.04
<i>BIW05</i>	11.32	36°	-10	0.28

The WTI simulator is also capable of simulating various turbine interference configurations (rotation rate, yaw angle, amplitude of interference peak, and variance in rotation rate). Each of these parameters were explored in creating the test data set (Table 4). The wind turbine configurations used were the two turbines present along the coast of Virginia.

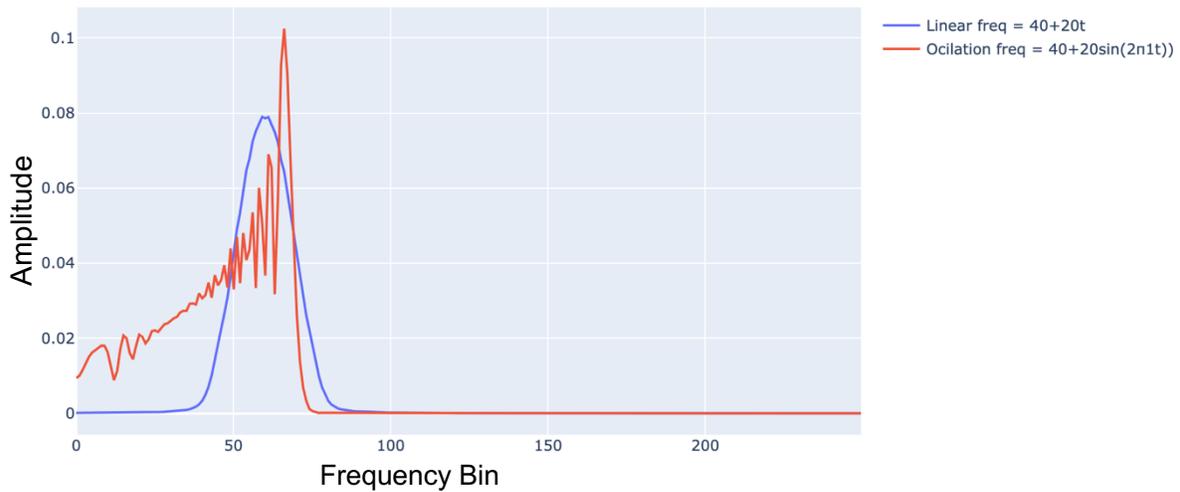


Figure 8: The difference between an oscillatory wind turbine rotation rate variation compared to a linearly increase rotation rate variation.

B.3 Final Dataset

Data from three sites off the coast of North Carolina and Virginia with overlapping coverage were used as input to the simulator. The three radar sites are located at Little Island Park, Virginia Beach, VA (LISL), Duck, NC (DUCK), and Cedar Island, VA (CEDR). A plot of the coverage of the three radars is shown in Figure 9. The input SeaSonde range-Doppler spectra cover the time period from January 01 through January 31, 2020. During this time the radars at LISL and DUCK were operating as a bistatic pair. WTI was added to each input spectra from the three sites with the WTI simulation tool. Each input spectra had WTI interference peaks placed into the monostatic and bistatic echo range bins. The output spectra covered a variety of turbine configurations.

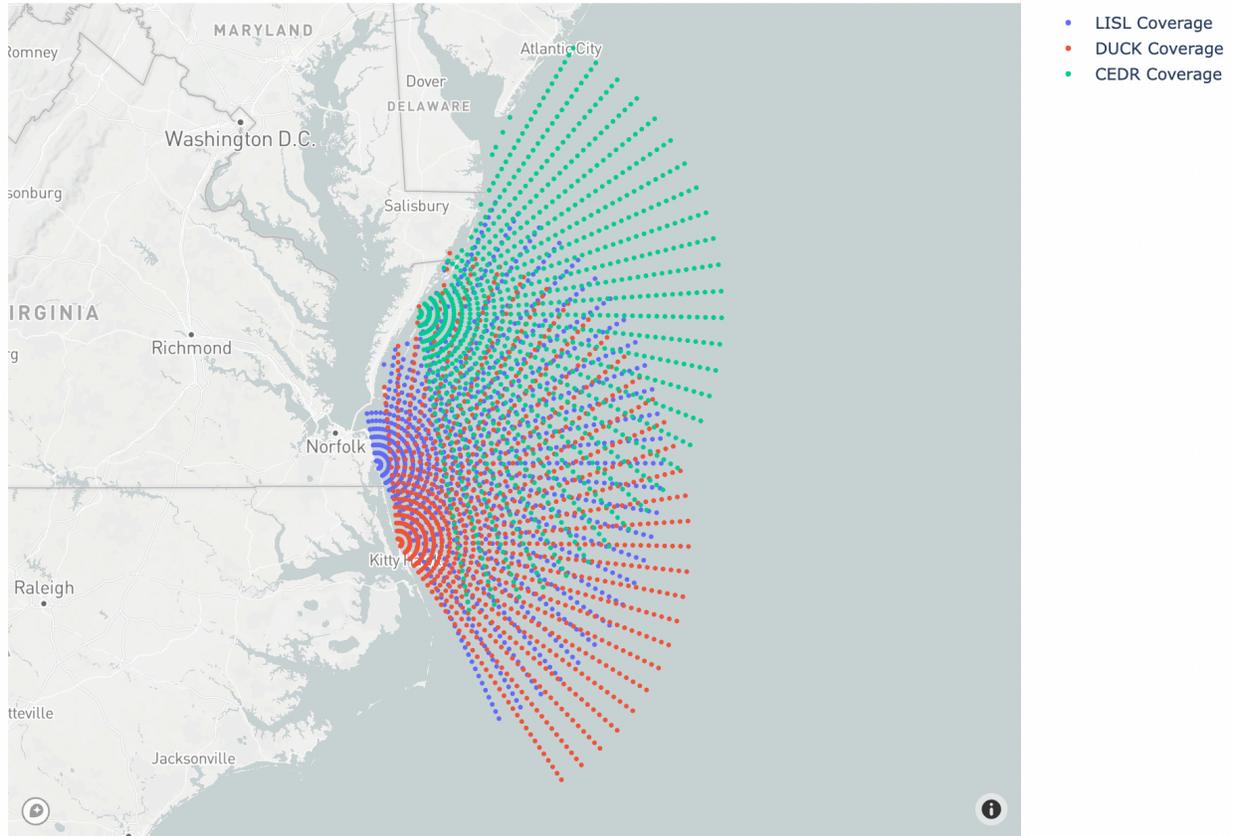


Figure 9: DUCK, LISL, and CEDR radar coverage map.

B.3.1 WTI Parameters

For each input spectra, WTI was simulated with a random selection from the parameters in Table 5. The range of parameters was selected to cover and exceed the range of possibilities that will occur in real-time on the study radar network and nearby wind turbines.

The amplitude parameter scales the WTI peaks relative to the Bragg. Therefore, an amplitude of -10 dB would scale the maximum WTI to be 10 dB below the maximum Doppler peak in the Bragg region. The range of amplitudes selected were from -20 dB below the Bragg to 10 dB Bragg level. The lower the amplitude of the WTI the more difficult it is to identify during WTI mitigation.

The location of the interference peak in Doppler is determined by the rotation rate of the wind turbine blades. The location of the peaks can therefore be specified. The locations of the WTI interference selected were in the Bragg, on the edge of the Bragg, and at the edge of the spectra.

Table 5: The range of wind turbine parameters used for the WTI simulations.

Parameter	Range
WTI Max Amplitude	-20 dB to 10 dB
Rotation Rate	In Bragg, edge of Bragg, and edge of the spectra
Rotation Rate Variability	0-2 rpm
Yaw Angle	0° to 90°

The Rotation Rate Variability represents the change in rotation rate (in RPM) over the integration period of the cross spectra. A rotation rate of 0 means there is no change, 2 means that the rotation rate goes from, for example, 2 RPM to 4 RPM over the integration period of the Doppler spectra. The variabilities included in this study range from 0 to 2 RPM and were randomly selected.

Using the parameters in Table 5, simulated WTI was added to each range-Doppler spectra collected from the DUCK, LISL, and CEDR during the month of January 2021. In the case of DUCK and LISL, the WTI was added to the bistatic range bins as well.

Appendix C. Effect of Flagging on Sea Surface Measurements

C.1 Background

Coastal oceanographic high frequency radars (HFR) measure the current of the ocean at the surface. The most common coastal HFR in the U.S. is the SeaSonde® compact cross loop system produced by CODAR Ocean Sensors Ltd. SeaSonde® HFR signals propagate in what is known as groundwave mode, following the curvature of the ocean surface beyond the horizon, giving them the ability to make observations over large areas. The sea surface current measurements from the national network of coastal oceanographic HFR are vital to United States Coast Guard for search and rescue efforts, aid in oil spill response, provide resource characterization for MRE ocean current turbine development, and current maps to ocean mariners and researchers. Secondary uses for coastal HF radar include wave measurements, wind measurements, and vessel detection. Wind turbines are in the field of view of coastal HFR act as reflectors and cause an interference signal that mixes with the oceanographic data. The wind turbine interference (WTI) causes errors or gaps in oceanographic data output.

Existing mitigation techniques (Troekel & Whelan, 2021; Wyatt et al., 2011) are designed to detect and remove oceanographic data impacted by the WTI, which can result in the loss of meaningful oceanographic data. The purpose of this section is to summarize the benefit of using multiple radars with overlapping coverage to reduce the loss of oceanographic data caused by WTI mitigation. The data used during this investigation includes data from three sites off the coast of Virginia and North Carolina with overlapping coverage. Multiple radars are used to prevent the downstream effect of removing range-Doppler bins from the SeaSonde® Doppler spectra processing. Additionally, the radars are configured to operate in a multi-static mode in which HFR receives bistatic sea echo from multiple other HFR transmissions, filling in gaps caused by WTI.

C.2 Impact of WTI on HF Surface Wave Radars

Due to the difficulty of separating WTI from the sea echo, existing state-of-the-art HFR WTI mitigation methods (Troekel & Whelan, 2021) identify range-Doppler cells containing interference and flag them to exclude them from further processing. An example of a SeaSonde range-Doppler spectra with WTI and WTI flags is shown in Figure 10. When the turbine rotation rates are such that a portion of the flagged WTI is contained within the Bragg region (the portion of the range-Doppler spectra containing the sea echo), crucial sea surface current observations may be lost or have increased error.

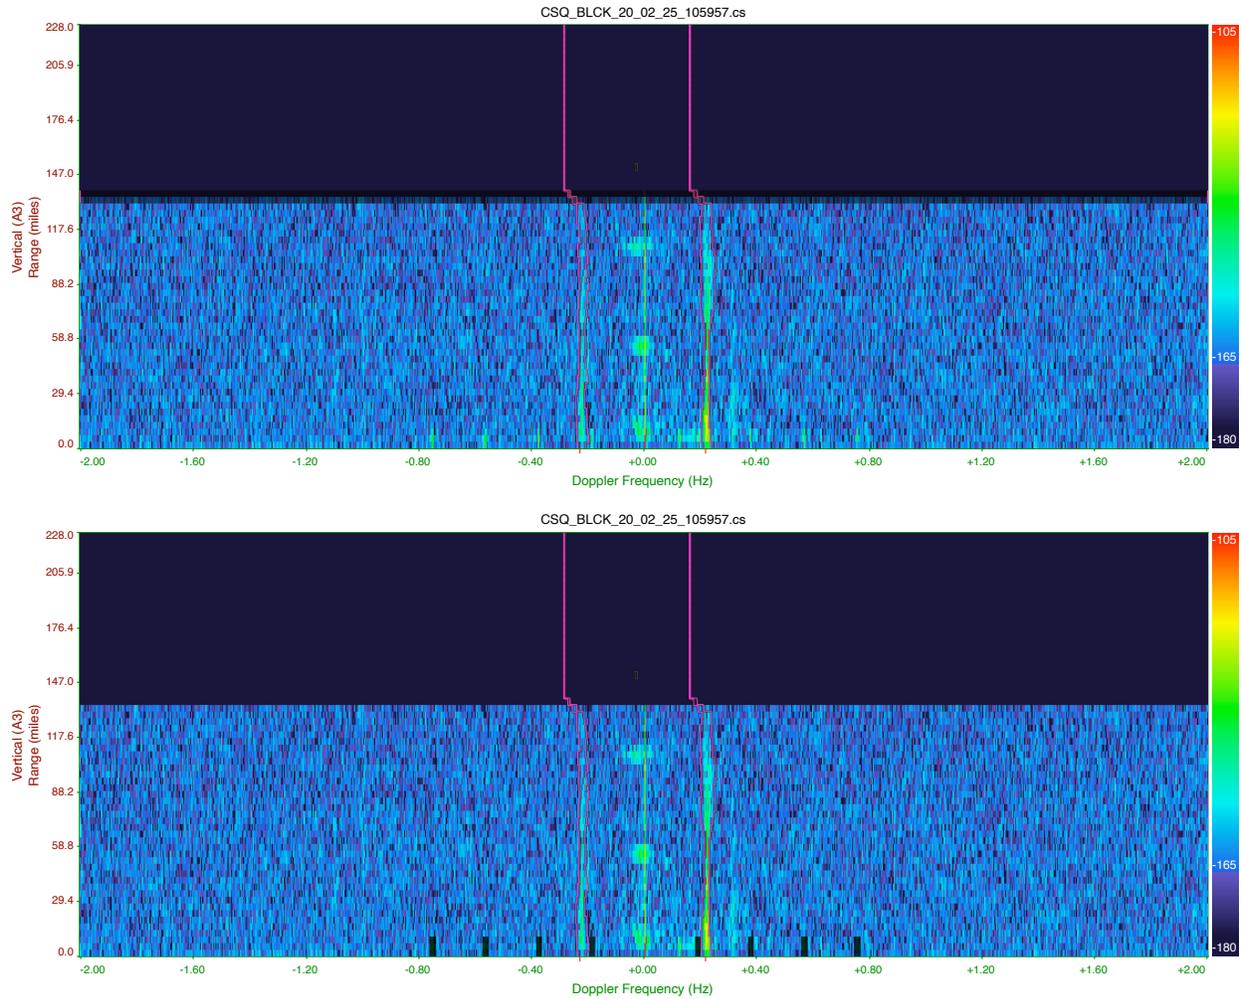


Figure 10: SeaSonde range-Doppler spectra from Block Island with wind turbine interference in the first range cell (top) and flagged WTI (bottom). The flagged range-Doppler cells are blacked out in the spectra plot on the bottom. The magenta lines indicate the boundaries of the Bragg regions containing the sea echo.

When operating in monostatic mode, where each radar is only receiving and processing sea echo from its own transmission, only the component of the sea surface current moving towards or away from the radar is measured, henceforth referred to as the radial current. The radial current measurements from two or more radars can be combined to produce the total current vector.

The SeaSonde receiver consists of three orthogonal antennas: a vertical monopole, and two horizontal cross loop antennas. Once the received signal is separated into its range-Doppler components as shown in the top of Figure 10, the ratios of the three antennas are used with a direction-finding algorithm such as Multiple Signal Classification (MUSIC) to determine the angle of arrival of the signals in each range-Doppler cell. The direction-finding algorithm places each Doppler velocity in a range-bearing bin. After a configurable

collection time, usually, half an hour or an hour, all the radial velocities in each range bearing bin are averaged to give the final radial velocity observation of the radar. When a range-Doppler bin is flagged by the WTI mitigation software, the radial current observation from the cell is eliminated from the averages. An example of the resulting impact of omitting vectors from the average is shown in Figure 11. When estimating the total current vector for a range-bearing bin from multiple sites, the error in the radial vectors from either site will propagate to the total vector calculation.

The removal of impacted data presents a challenge when it comes to WTI mitigation. If the WTI is not flagged and removed the direction-finding algorithm can place the radial vector in the wrong range bearing bin. On the other hand, if the radial vector is excluded from the range-bearing bin in which it belongs, the averaged radial vector uncertainty is increased.

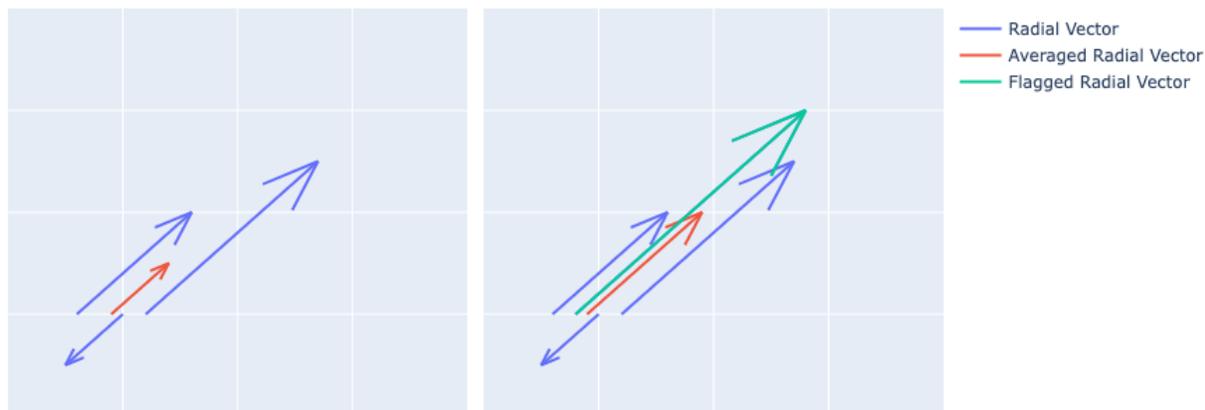


Figure 11: An example of a radar radial current measurement. The Blue and green vectors show individual radial vector measurements, and the red vector indicates the resulting averaged radial vector. The green vector indicates one that was eliminated because of WTI. The left plot shows the average when the flagged vector is removed, and the right plot shows what the averaged radial vector would be without the flagged vector removed.

It is expected that the operation of multiple radars operating multistatically, whereby each radar can receive the sea echo from the transmission of one or more other radars, will reduce the impact of omitting the flagged data. For this data infill method, the fact that the location of the WTI in range-Doppler space is a function of the rotation rates of the turbines is exploited. Figure 12 shows a simplified example of how this works. If three radar sites are at different locations and a wind turbine places a WTI peak in the Bragg region, the WTI mitigation will flag/remove the signal from the same Doppler bin at each site, though at possibly different ranges. However, since the radars are at different locations, the radial velocity component of a given patch of the sea surface is different at each radar. Since the same patch of ocean will likely have

different radial velocities at each radar, flagging the Doppler cells impacted by the WTI removes surface current measurements from different patches of the ocean.

Similarly, if two sites are set up to operate as a bistatic pair, with one receiving the sea echo from the other’s transmission producing elliptical current vectors (as opposed to radial current vectors), there is again redundancy in the velocity measurements which can only partially be impacted by the removal of a single range-Doppler bin at each site. In each case, having three radar sites operating monostatically or having two sites operating bistatically, there can be up to three measurements of the sea surface at each location, only one of which would be removed with the WTI mitigation. Since only two velocity components are needed to resolve the total velocity at each location, the loss and error in data resulting from the removal of flagged WTI should be reduced.

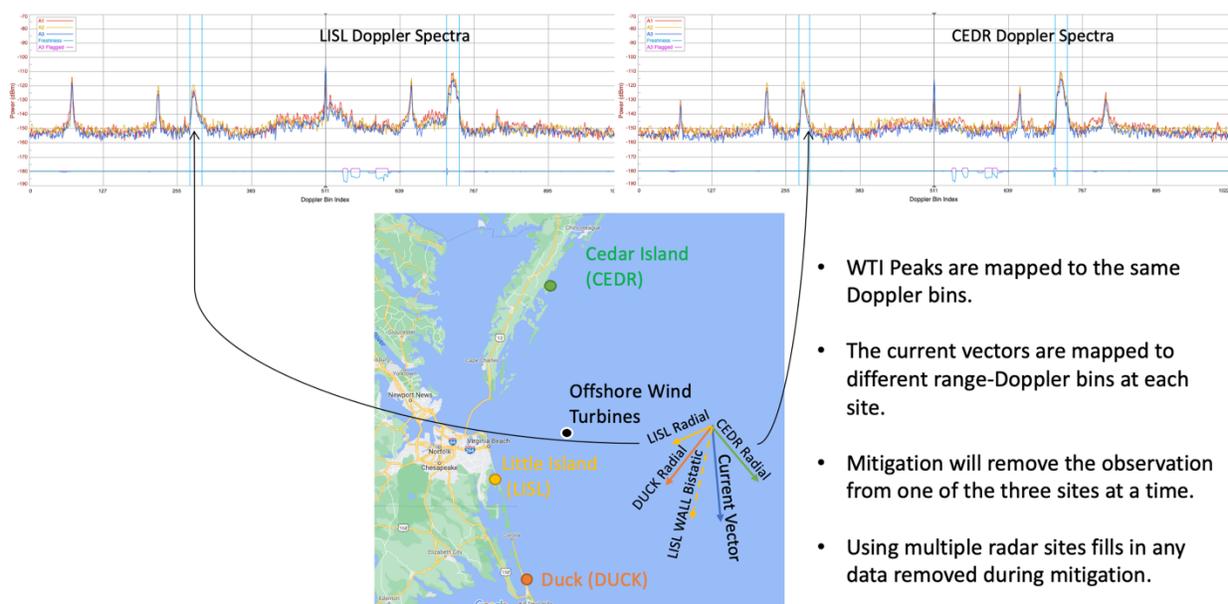


Figure 12: Diagram showing how multiple radars will infill the loss of data caused by WTI mitigation.

C.3 Experimental Design

The effectiveness of using multiple radars to backfill data removed from WTI mitigation was assessed using three radar sites with overlapping coverage—two of them operating as a bistatic pair. A latitude-longitude location in the coverage area of the three sites was chosen as the location of a representative wind turbine. To simulate the loss of data from the WTI mitigation, a Doppler velocity in the Bragg region was removed from the range bin corresponding to the representative wind turbine. The resulting data set has a loss of

data similar to what would be expected at a radar site near a single wind turbine running the WTI mitigation software.

The impact of the missing data was measured at each site by looking at the change in the resulting radial vectors from each site. Additionally, total vectors were calculated using the range-Doppler data with and without the removed Doppler cells. Three experiments were run using totals processed with data from only two sites operating monostatically, three sites operating monostatically, and three sites with two of the three operating bistatically. The difference in the total vectors of each of the above was compared with and without the deleted Doppler cells.

C.3.1 Study Area

The data used during this investigation includes data from three sites off the coast of North Carolina and Virginia with overlapping coverage. The three radar sites are located at Little Island Park, Virginia Beach, VA (LISL), Duck, NC (DUCK), and Cedar Island, VA (CEDR). The coverage area of the three radar sites is shown in Figure 13. The data set contains a month of data from LISL, DUCK, and CEDR. The data spans January 01 through January 31, 2020. During this time the radars at CEDR and DUCK were operating monostatically. The radar at LISL was operating bistatically with DUCK, site code LIDU. This data was collected prior to the installation of any wind turbines within the coverage area.

A site operating bistatically means that it receives and processes signals not only from itself but also from a neighboring site. The monostatic portion of the range-Doppler spectra refers to the portion of the range-Doppler spectra that contains information from the radar itself. In other words, LISL's monostatic range-Doppler spectra contain signals transmitted and received by LISL. The bistatic portion of LISL contains information from signals transmitted from DUCK and received at LISL. Thus, LISL site can receive, and process signals transmitted both by itself and the neighboring DUCK radar. This provides more information when processing the files into totals as objects (such as the sea surface) can be observed from two different radar signals.

The monostatic portion of the spectra for the three sites also has considerable overlap, increasing the coverage area for which there are redundant measurements of the total vectors. Data coverage maps of the three radars can be seen in Figure 13. These coverages are based on coastline angles and a maximum range of 204 kilometers.

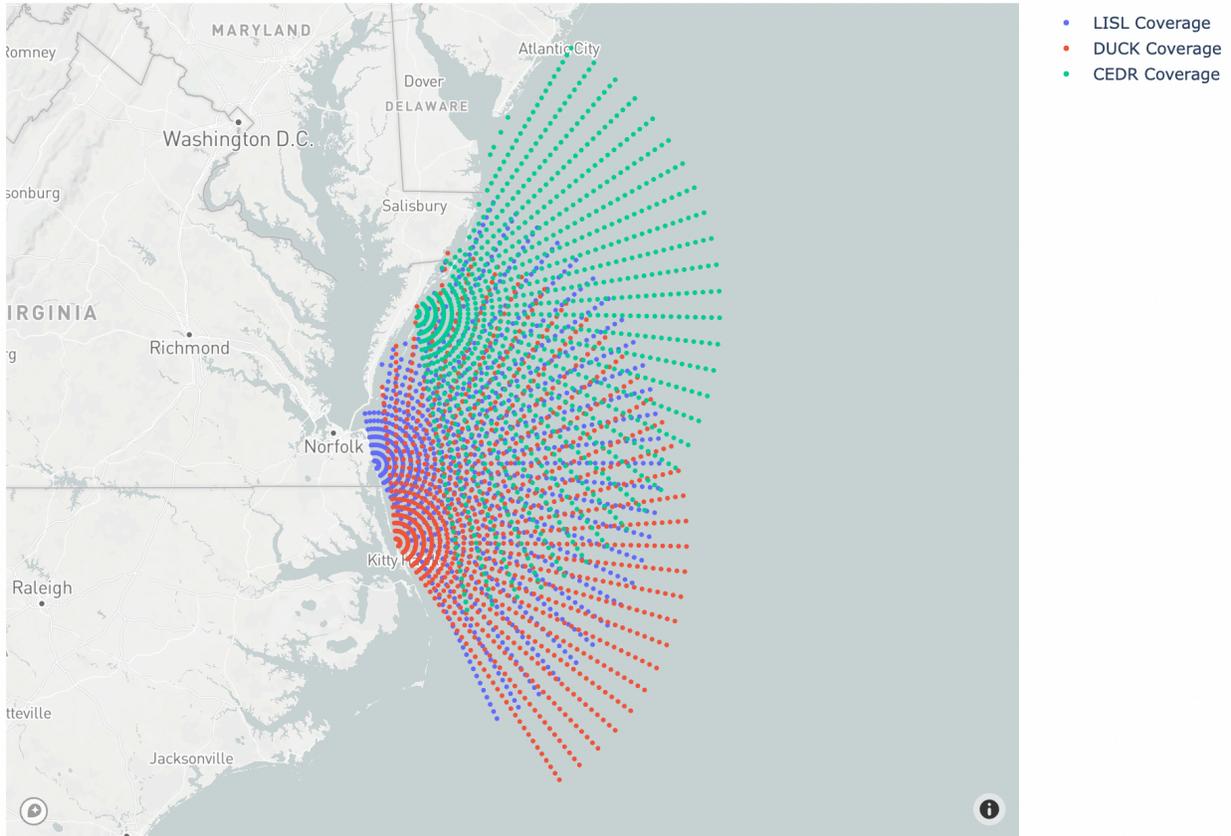


Figure 13: Coverage area of LISL, DUCK, and CEDR.

C.3.2 Kinds of Data

The monostatic and bistatic spectra are processed into radials and ellipticals, respectively. These radials and ellipticals can then be combined to form totals.

C.3.2.1 Monostatic and Radials

The monostatic signal is stored in radial files. Radial files also contain a temporal average. The radial files used in this analysis contain one hour of data on either side of the file date that have been averaged together to form the radial file. This means these radial files represent a moving 180-minute average of data. Radials were processed at each site for the duration of the field experiments. An example of a radial map can be seen in Figure 14.

Radials contain the range and bearing of the current observation as well as the radial velocity (the velocity toward or away from the radar). The range is half the distance traveled by the signal. The bearing is calculated using MUSIC from the antenna loop ratios. The radial velocity is calculated based on the Doppler

shift of the radar’s transmission frequency. The radial velocities across the coverage area can provide information as to the currents at that time.

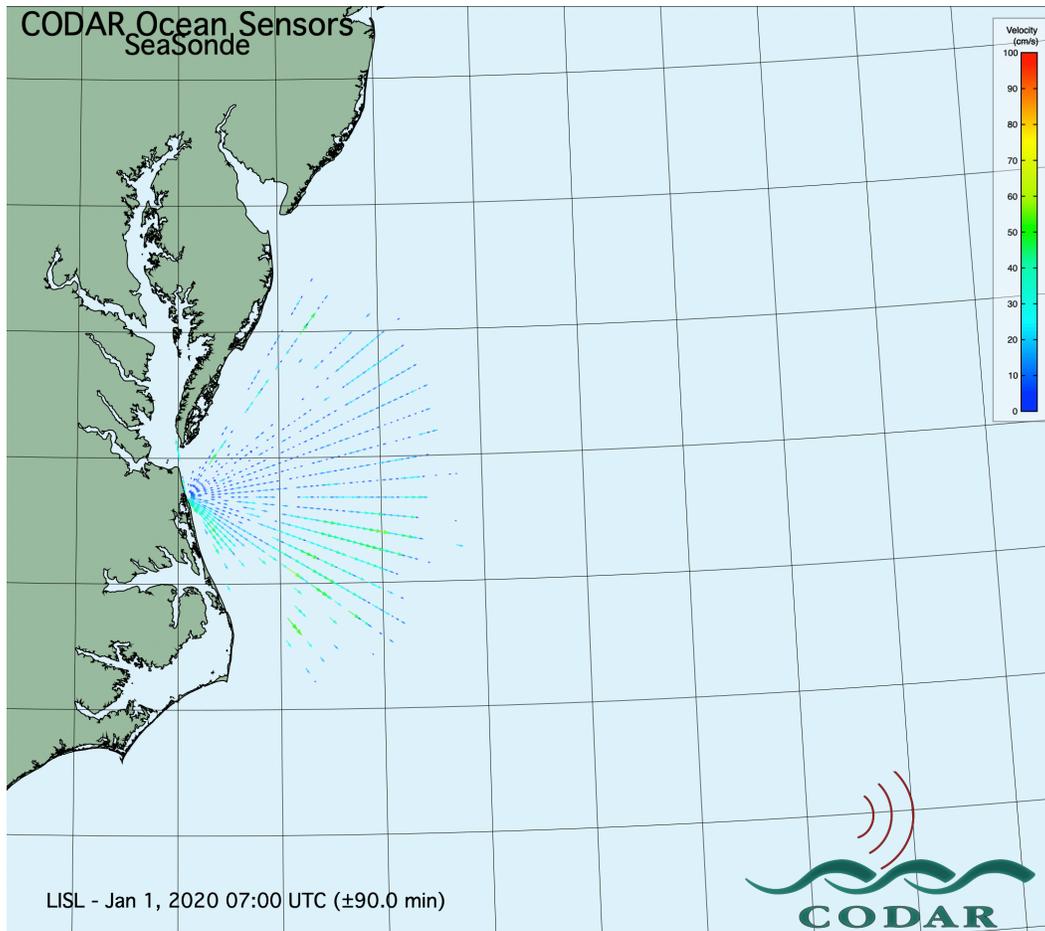


Figure 14: Radial Vector Map LISL

C.3.2.2 Bistatic and Ellipticals

The bistatic portion of each range-Doppler spectra file is processed into elliptical files. These elliptical files are like radials in that they are a temporal moving average of 2 hours of data. The elliptical files relate to the signal transmitted from a neighboring radar and received at the radar creating the files. In the case of LISL, the elliptical files contain information from signals that were transmitted by DUCK and received at LISL. The range is calculated based on the total great circle path traveled by the signal from the DUCK transmit antenna, to the target, and then to the LISL receive antenna. The two radars form the foci of an ellipse. Bearing and radial velocity are calculated similarly to the radials (MUSIC and using the Doppler shift of the signal). An example of an elliptical coverage can be seen in Figure 15. Radials from various sites and ellipticals are used in conjunction to make total vectors.

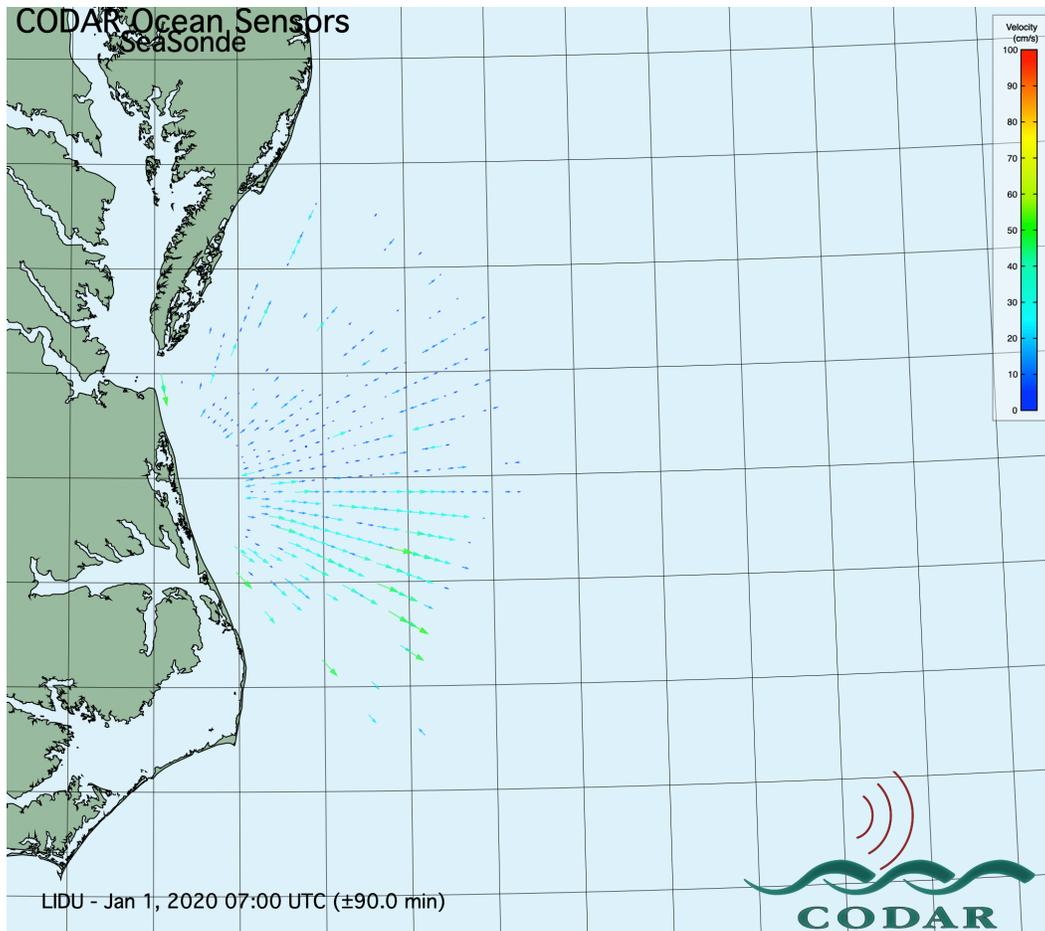


Figure 15: Elliptical Coverage Map of Signals Transmitted at DUCK and Received at LISL

C.3.2.3 Totals

When there are several sites with overlapping coverage, the data from these sites can be used to create a total vector map. The ellipticals and radials represent radial components of the velocity and therefore only contain partial information about the total surface current velocity. However, total velocity vectors (with direction and magnitude) can be calculated by combining multiple radials and/or ellipticals within a spatial area. The combination method used in this study is the least squares method outlined in Lipa and Barrick (1983).

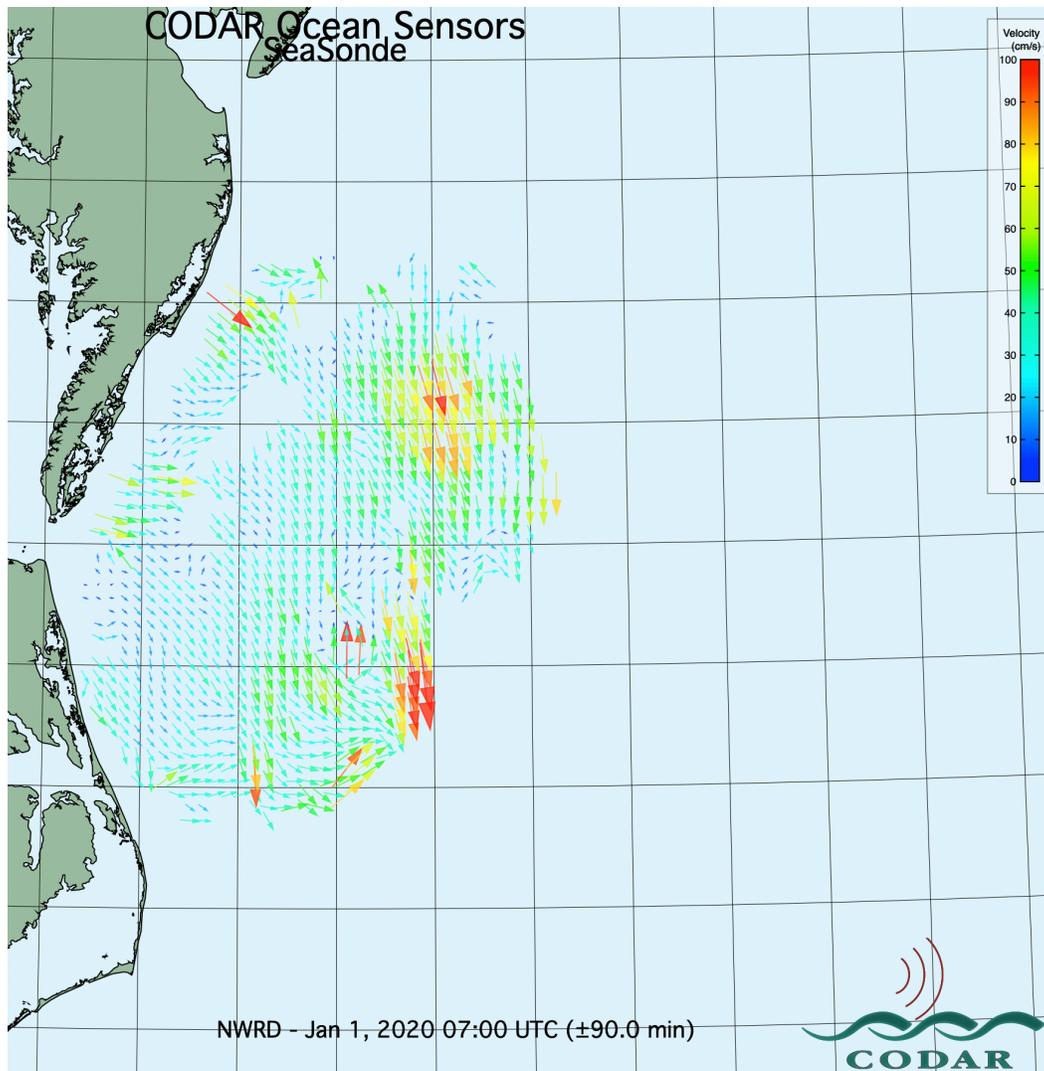


Figure 16: Total Vector Map Comprising Data from LISL, CEDR, and DUCK

C.3.3 Data Flagging

The locations of two turbines (installed after the collection of the data) off the shore of Virginia Beach were used to calculate the range bin where the WTI would occur in the monostatic as well as a bistatic signal of each radar. The monostatic location is the distance from the radar divided by the range bin size. The bistatic range R_{bin} of the turbines is given by:

$$R_{bin} = R^* + \frac{r_1 + r_2}{2}$$

where R^* is the range offset where the bistatic signal starts (in the case of monostatic signals the offset is 0), r_1 is the distance from the transmitting radar at site one to the turbine, and r_2 is the distance from the turbine to the receiver radar at site two.

C.3.4 Processing Range-Doppler Spectra

Three range bins that would be affected by the wind turbines were calculated. Three Doppler bins within the left Bragg and three symmetrical Doppler bins within the right Bragg were also selected for flagging. These 18 range-Doppler bins were flagged in the range-Doppler spectra.

The effect of the flagging can be seen in Figure 17. The top panel of this figure shows an unflagged input spectra from LISL while the bottom panel shows the flagged spectra. As can be seen in the flagged spectra, there are two 3 by 3 blocks missing within the right Bragg, one in monostatic, the other in the bistatic portion of the spectra.

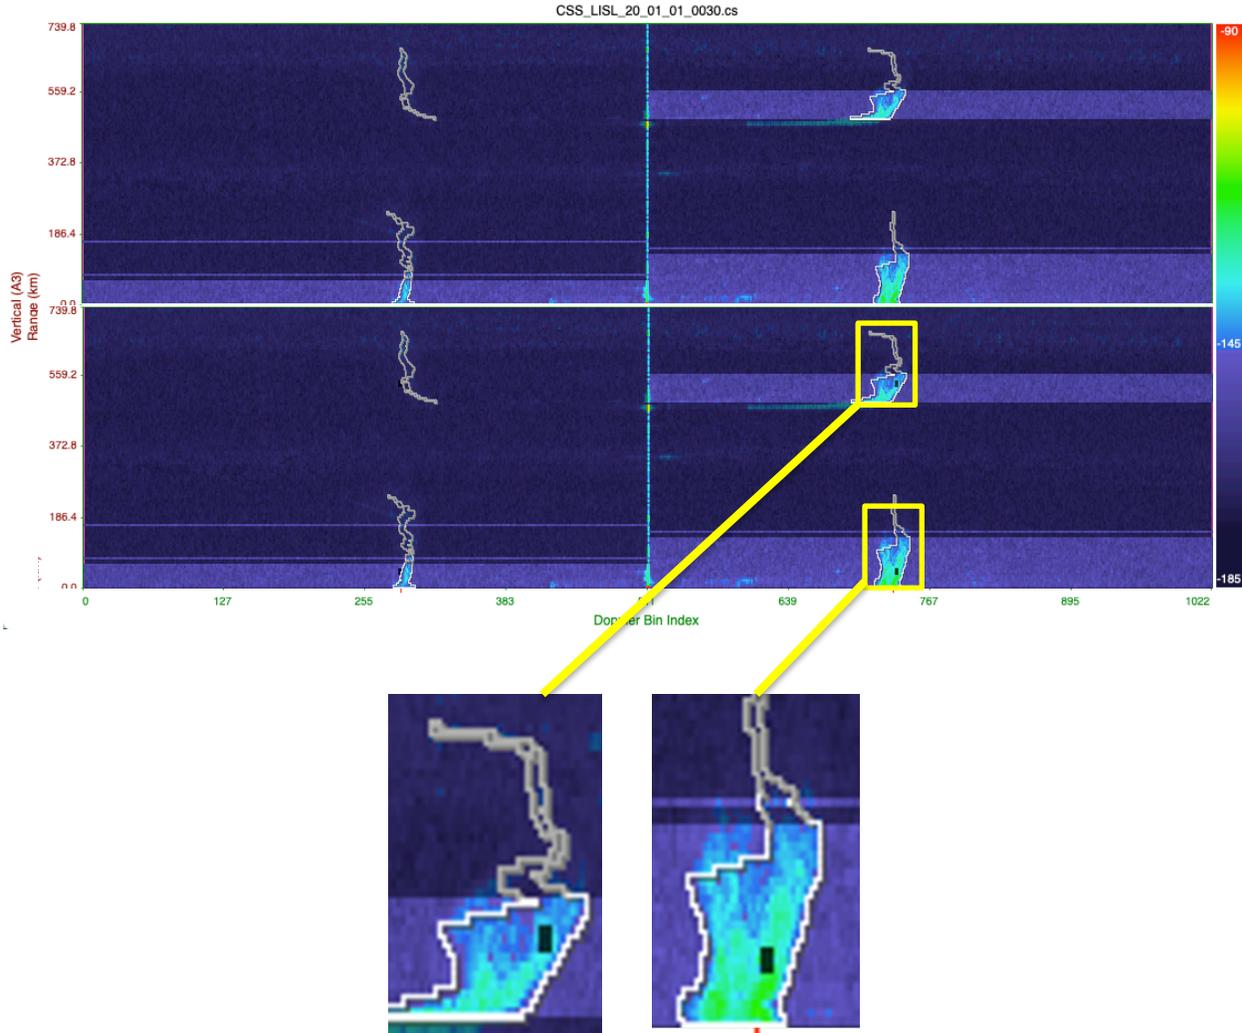


Figure 17: Range-Doppler Spectra. Unflagged spectra are shown in the top panel, and flagged spectra are shown in the bottom panel of the top figure. The bistatic (bottom left) and monostatic (bottom right) flagged portions are cropped and shown below.

Each spectra file with and without flagging was processed to produce radial vectors from each of the three sites and elliptical vectors were output from the DUCK to LISL signal. To investigate the improvement of using multiple radars and multi-static signals, total vectors were created from the resulting radial and elliptical current vectors in six ways:

- Totals from DUCK and CEDR radials all unflagged,
- Totals from DUCK, CEDR, and LISL radials all unflagged,
- Totals from DUCK, CEDR, and LISL radials and LISL-DUCK ellipticals all unflagged,
- Totals from DUCK and CEDR radials all flagged,
- Totals from DUCK, CEDR, and LISL radials all flagged, and
- Totals from DUCK, CEDR, and LISL radials and LISL-DUCK ellipticals all flagged,

The radials were compared; these were expected to have differences as radials are comprised of monostatic data. Missing data, therefore, cannot be replaced other than by averaging. Since the same data were missing from all spectra, averaging wouldn't fill in any meaningful data. Totals were also compared. Here, there is the possibility that the data from other sites would fill in gaps left by the flagged range-Doppler bins.

C.4 Results

C.4.1 Radials

Three metrics were analyzed at the radial level: the percent of vectors missing from the flagged dataset, the percent of radials that had a change in velocity, and the root mean squared difference in radial velocity among the changed vectors.

Radial files are organized into range-bearing bins. Each range-bearing bin in the radials from the data set without flagged range-Doppler bins was compared to the corresponding range-bearing bin from the data set with range-Doppler bins within the Bragg flagged. The percentage of vectors that were present in the unflagged data set but not present in the flagged data set is represented in the “% Lost” column in Table 6. If a vector was present in both the radial data sets, the radial velocities of the two vectors were compared. The percent of radial vectors that were different is represented in the “% Changed” column in Table 6. Using the radial vectors that were different, the root mean squared difference was calculated. This is represented in the “RMSD” column of Table 6.

Table 6: Statistical Results of Radial Vector Comparison Between Data With no Flags and Data which has been Flagged.

% Lost	% Changed	RMSD (cm/s)
0.98%	2.90%	5.82

C.4.2 Totals

Four metrics were used to analyze the total vector output from three combinations of sites listed above: the percent of vectors missing from the flagged data set, the percent of vectors that had a change in velocity, the root mean squared difference in magnitude of the changed vector velocities, and the standard deviation of the changes in velocity.

Total files are organized into a Latitude-Longitude grid. Each grid point in the totals from the unflagged data set was compared to the corresponding grid point from the flagged data set. The percentage of vectors

that were present in the unflagged data set but not present in the flagged data set is represented in the “% Lost” column in Table 7. If a vector was present in both datasets, the total velocities of the two vectors were compared. The percent of total vectors that were different is represented in the “% Changed” column in Table 7. Using the total vectors that were different, the root mean squared difference was calculated. The difference was taken by comparing the x and y components of the two total velocity vectors. The difference in x and difference in y were then squared and added together. The mean of these values was taken, and the squared root of that mean is represented in the “RMSD” column of Table 7.

Table 7: Statistical Results of Total Vector Comparison Between Unflagged and Flagged Data Sets

Sites	% Lost	% Changed	RMSD (cm/s)	Sigma (cm/s)
CEDR, DUCK	0.61%	11.50%	5.28	4.19
CEDR, LISL, DUCK	0.07%	16.63%	3.77	3.00
CEDR, LISL, DUCK, LISL-DUCK	0.06%	19.58%	3.45	2.76

The results in Table 7 show the expected result, the more redundant radial vector measurements the lower the error. Including extra radial or elliptical vectors by either adding a site or operating the sites multistatically greatly reduces the error and the amount of data lost at the total vector level. It should be noted that the percentage of total vectors that change increases as the number of sites used increases. This is expected as each site has flagged vectors thus there is more data omitted from processing as the number of sites increases. There is still an overall lower number of lost total vectors and a lower error. Recall also that these results are showing the impact of a single turbine placing WTI in the Bragg region of the spectra. If more turbines are in the radars’ surveillance area and place WTI in the Bragg region the errors could be much larger, but it is expected that they would still be reduced by using multiple radars.

C.5 Final Remarks

The results from this section show that the impact of flagging and removing data from HF radars with WTI can be reduced by using multiple radars. Furthermore, if it is not feasible to install new radar systems, errors introduced by data flagging can be reduced by operating the existing radar network in a multi-static mode. The mechanism for the improvement is clear and straightforward: the more redundant information the less severe the impact of removing some of the data.

Appendix D. Machine Learning Wind Turbine Interference Estimation

D.1 Background

High Frequency (HF) radars provide oceanographic data that is used in oil spill response, tsunami detection, search and rescue, and more. Recent years have seen an increase in plans for offshore wind turbines. Offshore turbines cause interference in the HF radar observations (Colburn et al., 2020; Teague & Barrick, 2012; Trockel, Rodriguez-Alegre, Barrick, & Whelan, 2018; Wyatt et al., 2011). Current Wind Turbine Interference (WTI) mitigation software finds the interference in the data and flags it so that those range-Doppler cells are not processed. The current approach (Trockel et al., 2021) utilizes the symmetrical nature of the interference to flag range-Doppler cells containing WTI and omit these from further processing. This approach works well when the radar's sweep rate is sufficiently high to prevent the aliasing of the WTI. At lower sweep rates, however, the interference aliases and the symmetrical nature of the problem disappears, rendering the current methods of WTI mitigation ineffective. Another limitation of current WTI mitigation schemes is that they are only able to predict the location of WTI, not the relative amplitude of the interference. Thus, the cells with interference must be excluded from further processing, and all the oceanographic data from those cells is lost.

The location and amplitude of the interference peaks are impacted by the turbine's rotation rate, yaw angle, and the variability in rotation rate over the radars' Doppler FFT integration period (Colburn et al., 2020; Teague & Barrick, 2012; Trockel, Rodriguez-Alegre, Barrick, & Whelan, 2018). The rotation rate of the turbine blades determines Doppler frequencies containing WTI. Figure 18 shows the Doppler frequencies of the first six positive and negative harmonic wind turbine peaks as the rotation rate varies. As the turbine rotation rate increases, the Doppler frequencies of the positive harmonics increases, and the negative harmonics decrease. When a turbine's rotation rate places a harmonic WTI peak in a frequency that exceeds the Nyquist frequency of the radar, the peak is aliased in the Doppler frequency and pushed out in range (for positive harmonics) or brought closer in range (for negative harmonics). The top plot of Figure 18 shows the aliasing occurring at lower rotation rates, around 2.5 rpm, for radars using a sweep rate of 1 Hz. While the bottom plot shows aliasing does not occur until much higher rotation rates, around 10 rpm, for radars using a 4Hz sweep rate. The gray rectangular regions indicate the typical Doppler frequencies containing the sea echo.

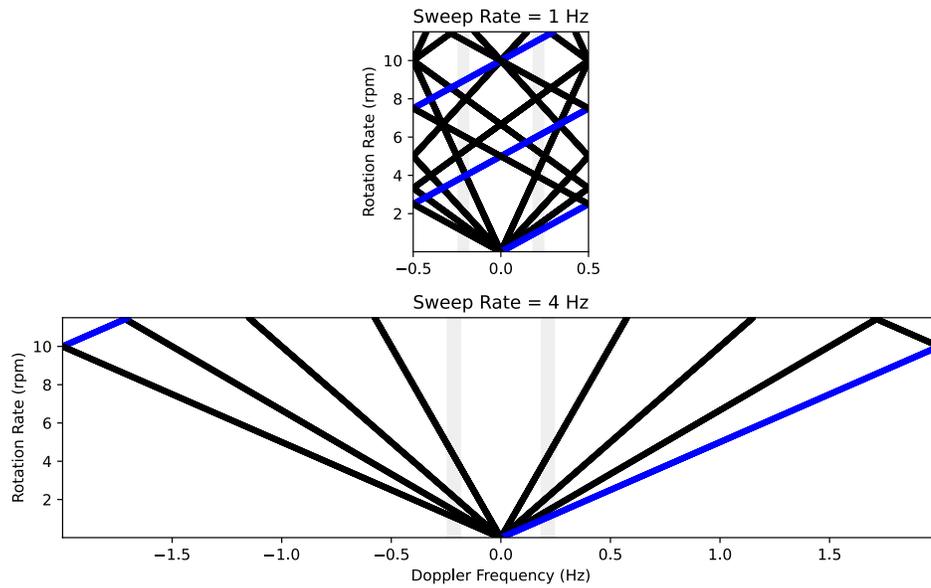


Figure 18: The Doppler frequencies of the first four positive and negative harmonic wind turbine interference peaks as the rotation rate varies. The top shows the impact of rotation rates for radars using a 1 Hz sweep rate, and the bottom plot shows the impacts on radars with a 4 Hz sweep rate. The gray rectangular regions indicate the typical Doppler frequencies containing the sea echo. The blue line shows the fourth positive harmonic and how aliasing loops around in Doppler.

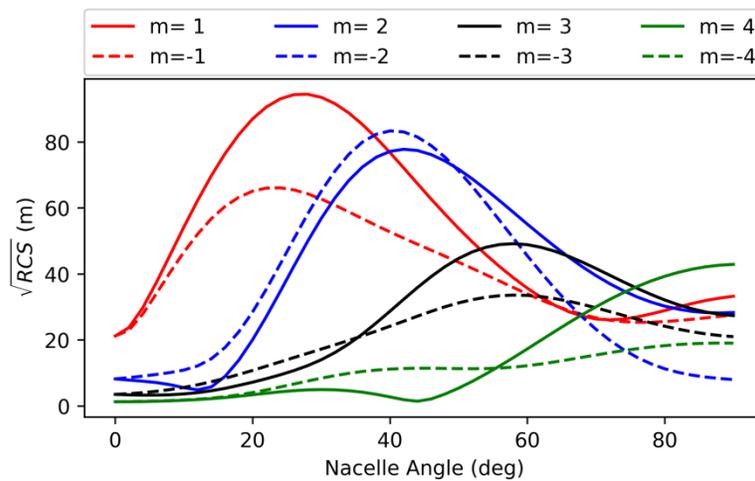


Figure 19: The influence of yaw angle on the amplitude of the first four positive and negative WTI harmonic peaks. The Nacelle (yaw) angle represents the angle of the nacelle relative to the radar. Each harmonic, indicated by m, is plotted in a different color. A Nacelle angle of 0 is when the turbine blades are parallel to the radar.

In addition to rotation rates, mitigation methods seeking to separate the WTI from the sea echo must have accurate estimations of the amplitude of the WTI peaks. The yaw angle of the turbine blades affects the

amplitudes of the WTI harmonic peaks (Teague & Barrick, 2012). Figure 19 shows the change in the amplitude of the harmonic components of the WTI as the yaw angle changes. The mitigation method developed by Trockel et al. (2021) did not estimate yaw angles and was unable to separate the WTI from the sea echo but instead relied on a flagging method to exclude data with WTI.

When a wind turbine's rotation rate is variable during the Doppler FFT integration period of a SeaSonde, the WTI peaks are spread out in Doppler. This spread causes both a wider affected area and a decrease in amplitude. Figure 20 shows an example of the WTI peaks spread out in Doppler frequency because of the wind turbine's rotation rate variability.

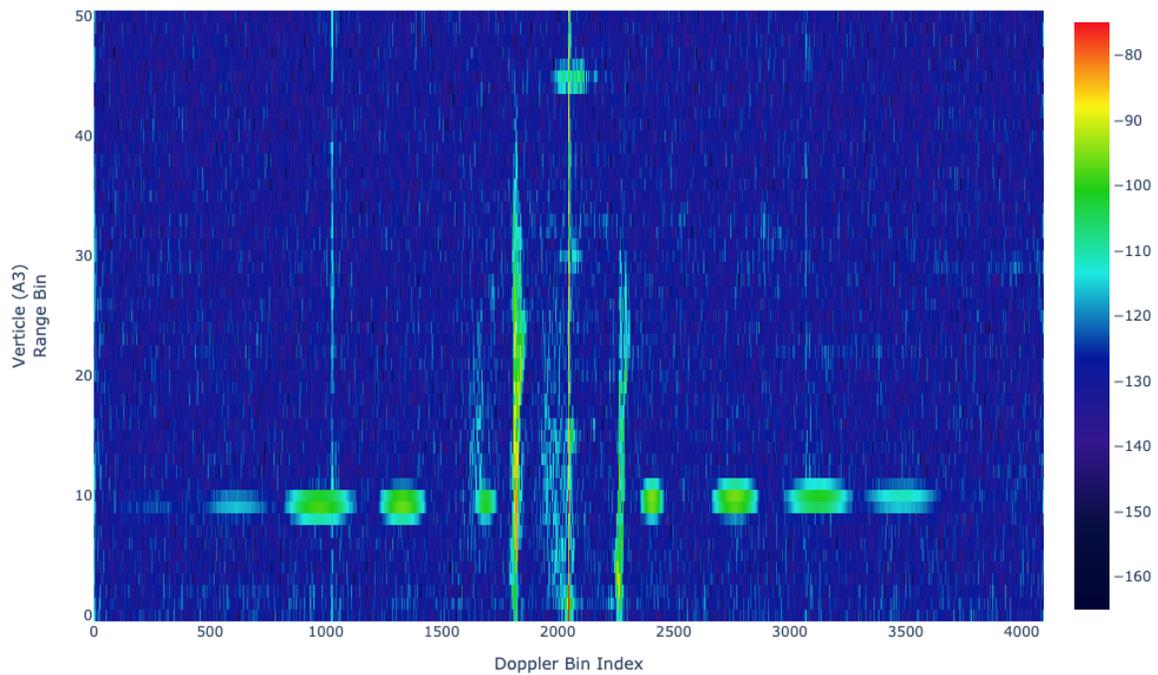


Figure 20: The spreading of simulated wind turbine interference peaks resulting from short time variability in a turbine rotation rate. The amplitude of the WTI has been increased to show the peaks clearly.

In this report we present our efforts to use empirical and machine learning (ML) methods to improve the identification of WTI in SeaSonde Doppler Spectra as well as estimate its amplitude so that it can be removed. Following Trockel et al. (2021), the method is split into two parts. First, the rotation rates, yaw

angles, and the rotation rate variability are estimated from the cross spectra. Second, the yaw angle and rotation rate variability are used to predict the ratios of the harmonic peaks.

An additional advantage of the proposed method is that it does not require the sweep rate of the radar to be increased to prevent aliasing. This is particularly advantageous when the radars are part of a network with many radars operating at the same frequencies. Increasing the sweep rate makes it difficult to set the timing alignments of the different radars in the network so that they do not interfere with each other. This appendix is organized as follows. Section 2 explains the new mitigation method and the evaluations used to assess the method. In section 3 the results of the evaluation are presented, followed by a discussion of the results in section 4. The report is concluded with final remarks in section 5.

D.2 WTI Estimation Method

We treated the problem of WTI characteristics in the cross spectra as a computer vision problem, using the cross spectra segment with WTI interference as the input image and deep learning models to identify the key characteristics of the wind turbines. This problem was split into three steps: classification, identification of key characteristics of the WTI, and prediction of the WTI amplitude.

D.2.1 Classification and RPM, RPM Variability, and Yaw Angle Estimation

We designed and implemented various machine learning models to classify the number of turbines spinning and predict turbine rotation rate, yaw angle, and variation in rotation rate. Each of these parameters affects the location or amplitude of the WTI in SeaSonde cross spectra. The models were created using Tensorflow's (Abadi et al., 2015) open-source deep learning libraries. The estimators were trained on four cases: SeaSonde with 1 Hz sweep rate and one turbine, SeaSonde with 1 Hz sweep rate and two turbines, SeaSonde with 4 Hz sweep rate and one turbine, and SeaSonde with 4 Hz sweep rate and two turbines.

Each of the cases had four associated models: one to classify the number of turbines spinning in the spectra, one to estimate the rotation rates, one to estimate the variation of the rotation rates over the Doppler FFT integration period, and one to estimate the yaw angles. A range-reduced portion of the SeaSonde cross spectra containing the WTI was used as input for each of the models. The range-reduced region included one range bin above and below the range cell(s) containing the wind turbine(s). For spectra collected with a 4 Hz sweep rate with no aliasing, this equates to a three-range bin slice. In the case of spectra with a 1 Hz sweep rate, this resulted in a five-range bin slice due to aliasing.

D.2.3 Training and Testing Datasets

For each model, we simulated approximately 60,000 spectra. The classifiers had equal numbers within each class. In the case of two turbines, 20,000 spectra were created with no turbine interference, 20,000 with only one turbine spinning, and 20,000 with both turbines spinning. The rotation rate variability, yaw, and rotation rate were randomly selected for the estimators. Yaw was selected from the range of 0 to 90 degrees. The Variation in the rotation rate was selected from a normal distribution with a mean of zero and a standard deviation of 0.33 rpm. The rotation rate variability was chosen based on our analysis of previous SCADA data from BLCK. The rotation rate was randomly sampled from 4 -1 1.5 rpm. The rotation rate bounds were chosen based on the bounds of existing turbines at BLCK. The amplitude was also randomly varied across all datasets to be from 5 dB below Bragg to 10 dB above. The input spectra were taken from the BLCK and CEDR radar sites, and simulated WTI was added.

D.2.4 WTI Peak Amplitude Estimation

The location in Doppler of the harmonic peaks can be found using the rotation rate. This allowed us to save the cross spectra value at the range-Doppler location of each harmonic peak. Using the harmonic peak amplitude values, we saved the ratios of the harmonic values relative to the first positive and first negative harmonic peaks. We then used regression models to predict the harmonic peak ratios given the turbines' yaw angle and the variation of the rotation rate during the Doppler FFT. It was anticipated that by accurately estimating the harmonic ratio, we could identify the amplitude of the interference peaks mixed with the sea echo and separate it from the sea echo.

D.2.4.1 Regression and K-Nearest Neighbors

We used two different regression models to predict the harmonic ratio values: K-nearest neighbors (KNN) and polynomial regression. A model was used to predict each of the following harmonic ratios:

- 1) harmonic peak negative three to harmonic peak negative one,
- 2) harmonic peak negative two to harmonic peak negative one,
- 3) harmonic peak positive two to harmonic peak negative one,
- 4) harmonic peak positive three to harmonic peak negative one,
- 5) harmonic peak negative three to harmonic peak positive one,
- 6) harmonic peak negative two to harmonic peak positive one,
- 7) harmonic peak positive two to harmonic peak positive one, and
- 8) harmonic peak positive three to harmonic peak positive one.

The harmonic ratio predictions were multiplied by the amplitude at the given location to find the amplitude of the first negative and positive harmonics. For example, if the first harmonic peak, H_{p_1} , is mixed with the

sea echo it can be estimated using the amplitude ratio of peaks one and two, $\frac{H_{p1}}{H_{p2}}$, and the amplitude of harmonic peak two by

$$H_{p1} = \frac{H_{p1}}{H_{p2}} H_{p2}.$$

D.2.5 Testing

The data was split into training and testing datasets. The classification models were tested for accuracy using testing data that was not included in the training. Accuracy was measured as the percent of correct estimations, in other words, the number of correct predictions of the number of turbines spinning divided by the number of measurements in the test dataset. The rotation rate, variation in rotation rate, and the yaw angle estimators' accuracies were measured by the root mean square error (RMSE) of the predicted value relative to actual values of the rotation rate, variation in rotation rate, and yaw angle in the test data. To measure the accuracy of amplitude predictions, we report the RMSE in volts of the predictions as well as the Pearson correlation between predicted and estimated values.

D.3 Results

D.3.1 Classification:

The models accurately predicted the number of turbines spinning in all cases. The accuracy of the predictions went down as more turbines were introduced (see Table 8). The 1 Hz system accuracy was 99.6% in identifying whether one turbine or no turbines were spinning and 94.7% accurate in identifying whether zero, one, or two turbines were spinning. The 4 Hz system accuracy was 98.7% in identifying whether one turbine or no turbines were spinning and 94.2% accurate in identifying whether zero, one, or two turbines were spinning.

Table 8: Accuracy of classification of how many turbines were spinning.

Simulation	Accuracy
1 Hz Single Turbine (0 turbines spinning, 1 turbine spinning)	99.6%
1 Hz Two Turbines (0 turbines spinning, 1 turbine spinning, 2 turbines spinning)	94.7%
4 Hz Single Turbine (0 turbines spinning, 1 turbine spinning)	98.7%
4 Hz Two Turbines (0 turbines spinning, 1 turbine spinning, 2 turbines spinning)	94.2%

D.3.2 Rotation rate, Yaw, Rotation Rate Variation Estimation:

The models were able to reliably predict rotation rate variation across all cases tested (see Table 9) with a high degree of accuracy. The errors in rotation rate estimation ranged from 0.039 rpm to 0.098 rpm. The error in rotation rate variation estimation ranged from 0.016 drpm to 0.042 rpm. The Yaw angle estimations were less accurate, with errors ranging from 4.61 degrees to 9.64 degrees.

Table 9: Accuracy of estimation of rotation rate, yaw angle, and variation. Mean error in prediction and standard deviation of the error shown.

Simulation	Rotation Rate Mean Error (std)	Yaw Angle Mean Error (std)	Variation Mean Error (std)
1 Hz Single Turbine	0.039 rpm (0.12)	4.61 degrees (6.98)	0.016 drpm (0.032)
1 Hz Two Turbines	0.073 rpm (0.15)	9.64 degrees (5.77)	0.036 drpm (0.032)
4 Hz Single Turbine	0.096 rpm (0.25)	5.85 degrees (8.97)	0.03 drpm (0.063)
4 Hz Two Turbines	0.098 rpm (0.21)	7.53 degrees (5.15)	0.42 pm (0.046)

D.3.3 Amplitude Estimation:

The amplitude is commonly measured in volts which yields an error whose magnitude is hard to interpret because of the scale. In Table 10 and Table 11, we report the RMSE in Volts and for interpretability, we also report correlation coefficients between predicted and true amplitude values. The amplitude estimation was correlated highly at 4 Hz using KNN and polynomial regression with correlation coefficients between estimated amplitude and true amplitude ranging from 0.87 to 0.95. The 1 Hz estimations were less correlated with correlations ranging from 0.68 to 0.80.

Table 10: The RMSE and Pearson Correlation of prediction of the amplitude of the first positive and first negative harmonic peaks using polynomial regression.

Simulation	Harmonic Peak negative 1 Predicted Amplitude vs True Amplitude RMSE in Volts (Pearson Correlation Coefficient)	Harmonic Peak Positive 1 Predicted Amplitude vs True Amplitude RMSE in Volts (Pearson Correlation Coefficient)
1 Hz Single Turbine	4.89e-9 (0.80)	4.84e-9 (0.72)
1 Hz Two Turbines	1.80e-9 (0.78)	1.81e-9 (0.78)
4 Hz Single Turbine	4.57e-8 (0.93)	4.55e-8 (0.90)
4 Hz Two Turbines	4.24e-8 (0.90)	4.21e-8 (0.87)

Table 11: The RMSE and Pearson Correlation of prediction of the amplitude of the first positive and first negative harmonic peaks using KNN regression.

Simulation	Harmonic Peak negative 1 Predicted Amplitude vs True Amplitude RMSE in Volts (Pearson Correlation Coefficient)	Harmonic Peak Positive 1 Predicted Amplitude vs True Amplitude RMSE in Volts (Pearson Correlation Coefficient)
1 Hz Single Turbine	2.78e-8 (0.76)	2.78e-8 (0.68)
1 Hz Two Turbines	3.82e-8 (0.74)	3.82e-8 (0.74)
4 Hz Single Turbine	2.88e-9 (0.95)	7.08e-9 (0.92)
4 Hz Two Turbines	1.40e-8 (0.91)	1.53e-8 (0.88)

D.4 Discussion

D.4.1 Limitations

The present methodology only works to find the harmonic interference's amplitude at its center (at one Doppler value). As can be seen in Figure 20, the turbine interference can often be spread into several Doppler bins. The shape and spikes in the amplitude of the peaks depend on how the rotation rate changed during the Doppler FFT integration period. Since the turbine rotation rate varies according to the wind speed, and wind speed varies unpredictably within a given integration period, this method cannot reliably be used to find the amplitude of the WTI harmonic peak in each Doppler cell during times the turbine's rotation rate is variable. The center of a WTI harmonic peak also depends on how the rotation rate varies; thus, estimating the structure of the signal within each WTI harmonic peak cannot be done with enough resolution in the frequency domain to separate it from the sea echo during times of high rotation rate variability.

The models presented in this report were also trained and tested using simulation, as our attempts to procure high-resolution SCADA data throughout this project were unsuccessful. A future report will explore the effects of using these models, trained primarily on simulated data, for real-time software tested in the field.

D.4.2 Improved Flagging

The use of ML for identifying rotation rate is promising, as it provides an accurate way to flag interference, especially for systems operating at 1Hz where existing methods cannot predict the turbine rotation rate. The high degree of accuracy of the rotation rates can be used in conjunction with current methods to effectively flag WTI from cross spectra at both 1 Hz and 4 Hz.

D.4.3 Effect of Rotation Rate Variability

The spectra used have a Doppler FFT integration period of 10 minutes. During these ten minutes, the wind turbine blade rotation rate can be subject to change. Changes in rotation rate over an integration cause the WTI peaks to spread. This spread causes both a wider affected area and a decrease in amplitude.

While investigating the structure of WTI in real-world data, we discovered that the way that the turbine blades vary over an integration period affects the characteristics of the interference peaks. For example, a rotation rate varying linearly from 2 rpm to 4 rpm will have a significantly different interference peak than one varying from 2 rpm to 4 rpm in a non-linear fashion. We found that simulating the variation in rotation rate with a nonlinear oscillating change produced interference peaks structurally similar to observed wind turbine interference peaks. For example, using a starting rotation rate of 2 rpm and an ending rotation rate of 4 rpm and adding small oscillations up and down as it increases to the final rotation rate, Figure 21 demonstrates that the way a turbine's rotation rate varies during the Doppler FFT integration period affects the shape and amplitude of the WTI. The oscillating increase and decrease of the rotation rate were used for the simulations, but the changes in rotation rate that occur in the field are more diverse and unpredictable. The resulting interference peaks are highly sensitive to this variability. Unfortunately, this means there is no clear way to predict the amplitude and structure of the peaks in the frequency domain with enough accuracy to subtract them away.

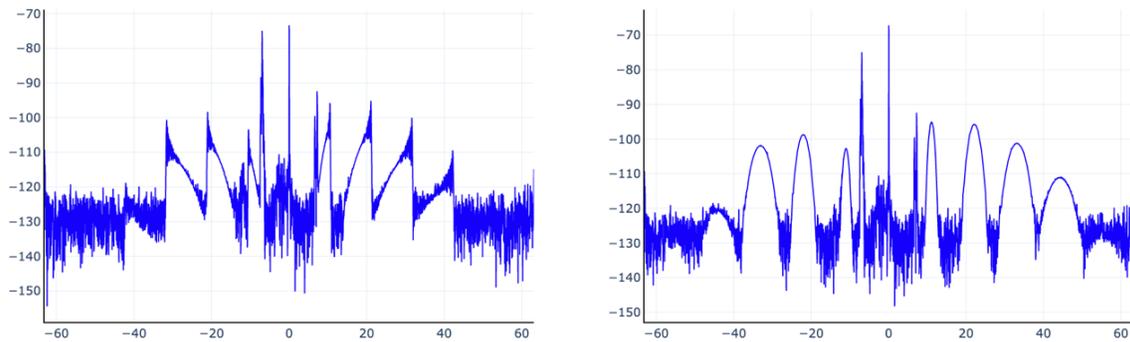


Figure 21: BLCK simulated spectra showing the effects of different kinds of variability in rotation rate on the characteristics of the WTI. The leftmost figure shows an oscillating increase in rotation rate from 6rpm to 7rpm over an integration period. The rightmost figure shows a linear increase in rotation rate from 6rpm to 7rpm over an integration period. The peaks are wider and lower in amplitude as the variance in rotation rate increases. The way the rotation changes determines the characteristics of the WTI harmonic peaks.

D.5 Final Remarks

The WTI interference amplitude estimation is limited to times when the rotation rate is steady. The amplitude of the WTI is also highly dependent on the nature of the rotation rate variability over the integration period. This limitation makes it impractical to subtract the entire interference peaks from the spectra when the rotation rate is changing. Furthermore, the only times the turbines at BLCK and CEDR are in the Bragg are during times when rotation rates are likely to be changing. This method is, therefore, impractical for the current size of turbines. However, as turbines increase in size, there could be a turbine system whose preferred rotation rate speed places the WTI peaks in the Bragg. This method could be used then as WTI is consigned to 1-3 Doppler bins when the variance in rotation rate is low, as is the case when turbines are spinning at their optimal rotation rate.

The use of machine learning for identifying wind turbine interference and estimating the rotation rate, yaw angle, and rotation rate variability shows promise. Models were able to estimate the attributes of the WTI with high accuracy and precision for radar operation with sweep rates of 1 Hz and 4 Hz. The models were also able to distinguish between times when no turbines were spinning, a single turbine was spinning, and two turbines were spinning. Previous attempts to estimate rpm were confined to radars operating with a 4 Hz sweep rate as they relied on the symmetrical nature of the harmonic peaks. This symmetry is lost with aliasing at lower sweep rates. Positive harmonic WTI peaks that alias are pushed out in range, while negative aliased harmonic peaks are pulled back in range. This phenomenon leads to unsymmetrical range bins, which was the foundation for previous analytical methods. The models explored in this report were able to predict the rotation rate regardless of the sweep rate. They were also able to predict the change in rotation rate and the turbines' yaw angle, two attributes that were not previously estimated by models.

Appendix E. Wind Turbine Interference Flagging Real-Time Software

E.1 Software Overview

This appendix explains the elements of the real-time wind turbine interference (WTI) mitigation software toolchain that was developed for this project. The real-time software uses machine learning to classify the number of turbines present in the cross spectra and estimates the rotation rates of the turbines present in a SeaSonde coverage area. These rotation rates are then corrected based on the location of the harmonic peaks found in the SeaSonde Doppler spectra, and the range-Doppler bins containing WTI are flagged to be excluded at the spectra averaging stage.

E.2 WTI Mitigation Software Toolchain

The software creates Doppler cross-spectra with flagged WTI. The software is configurable with separate configuration files for the radar configurations and WTI flagging. Using the turbine locations and the radar configuration, the tool selects the affected range bins and writes that information in a format compatible with the machine learning algorithms. In addition to the range bin containing the WTI, the software toolchain identifies WTI in the two range bins above and below any range bin containing the wind turbines. If there are bistatic signals, this is repeated for the bistatic portion of the spectra as well. Figure 23 shows a diagram of the portion of the SeaSonde data processing toolchain that was modified to enable WTI flagging.

Once SeaSondeAcquisition acquires enough radar sweeps to create an unaveraged Doppler cross spectra file (CSQ) the first step of the WTISpectraFlagger tool is its classification step. The tool's machine learning models have been trained to identify the number of turbines present in an range-reduced slice of the Doppler cross spectra. The model identifies the number of turbines present in the cross-spectra. The model identifies whether there are zero turbines spinning, one turbine spinning, or two turbines spinning.

If the classification step identifies no turbines, the tool proceeds to the next cross spectra. If there are turbines identified, the range region of the cross-spectra containing WTI is passed to another machine learning model which has been trained to estimate rotation rates. This model predicts the rotation rates for each of the turbines that were identified to be present in the cross-spectra. If the site is running bistatically, as is the case at LISL, this step is done for both the monostatic and the bistatic portions of the cross-spectra.

Once the rotation rates of the turbines have been estimated, the range region of the cross spectra with WTI, the location of the identifiable WTI is used to further adjust the estimated rotation rates. To account for the error in rotation rate assignment, because of rotation rate variability or other factors, the WTISpectraFlagger tool looks at a 60 Doppler bin window around the predicted harmonic WTI peaks associated with the rotation rates assigned by the machine learning algorithm, which are found outside of the Bragg region. The tallest peaks associated with each harmonic are saved and then compared to each other. If the spacing of the harmonic peaks is consistent with a single rotation rate, that rotation rate is calculated and returned. The widths of the peaks are also averaged together for flagging purposes. However, if there is no single rotation rate that explains the harmonic peaks, then the rotation rate is flagged as an incorrect assignment, and no range-Doppler bins associated with that rotation rate will be flagged. The WTI from a single harmonic interference peak often extends into multiple range bins, so this process is repeated with the range bins above and below the primary range bin. Each range bin that has the same harmonic structure is saved for flagging.

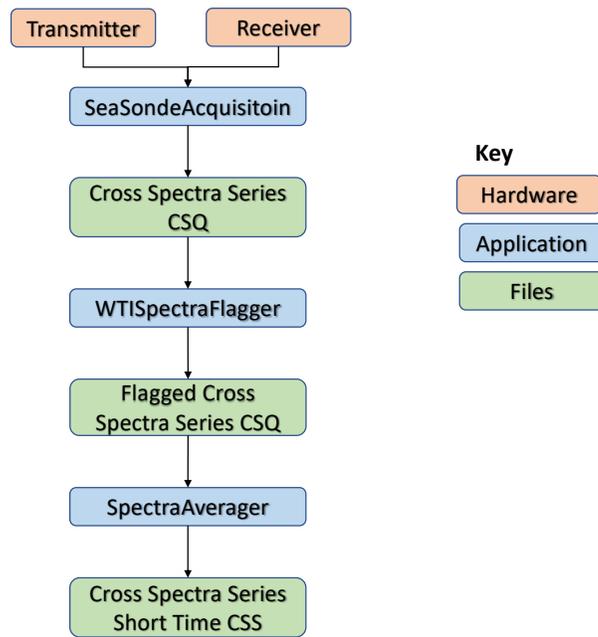


Figure 22: Software implementation Diagram.

Once the rotation rates have been verified and corrected, and the range Doppler bins with WTI as well as mean WTI harmonic peak width identified, the cross spectra file is flagged in the infected range bins. This tool uses rotation rate estimates from a turbines rotation rate developed by Trockel et al. (2018) to identify

the location in Doppler associated with the corrected rotation rate. In each range bin that was found to be infected, the Doppler bins associated with the rotation rate are flagged. The number of Doppler bins flagged depends on the mean width of the peaks found in the previous step.

The flagged cross spectra is then passed to the SpectraAverager tool and the remaining CODAR SeaSonde data processing software where the flagged range-Doppler bins are suppressed from averaging and omitted from further processing.

Appendix F. Effects of Machine Learning Wind Turbine Estimation and Multisite Operation on Sea Surface Measurements

F.1 Methods

To test the methods developed over the course of this project, we simulated WTI into the Bragg region of the range Doppler cross spectra at three radar sites with overlapping coverage. The analysis in this report was focused on cross spectra from the SeaSonde HFR manufactured by CODAR Ocean Sensors. These cross spectra were flagged using machine learning algorithms to predict the rotation rate, then processed into radials and ellipticals, where radials are the surface current measurements from a single radar and ellipticals are the measurements from a pair of radars operating in a multistatic mode. The radials and ellipticals from all the sites were combined into totals. The radials and totals were compared to each other to measure the effect of flagging as well as unmitigated WTI on HF measurements. The totals were combined using two sites, three sites, and three sites with ellipticals to measure the effect of using multiple sites with overlapping coverage to mitigate data loss and corruption.

F.1.1 Simulation

The data came from the radar stations at Cedar Island, VA (CEDR), Little Island, VA (LISL) and Duck, NC (DUCK) and the elliptical data between LISL and DUCK (LIDU). The study area is shown in Figure 23. The data covered from January 1 to January 17, 2020. This was before turbines were operating in the coverage area. There was a gap in the record from January 9-10 at CEDR, but otherwise, the data record was complete. For analysis purposes, only days when all sites had data were used. The variability in the number of radial vectors per hour for each of the stations is shown in Figure 24. These radial stations were chosen due to their proximity to two offshore turbines that are part of the Coastal Virginia Offshore Wind (CVOW) project. CVOW is located 43 km off the coast of Virginia Beach, VA. The initial phase of the project has two 6 megawatt Siemens Gamesa turbines. The locations of the two turbines are provided in Table 12.

The locations of the turbines relative to each radar station is provided in

Table 13. Because of Doppler aliasing in the spectra, the turbine interference can be found in ± 3 range bins from where they are centered. Analysis for this report was focused on these range cells.

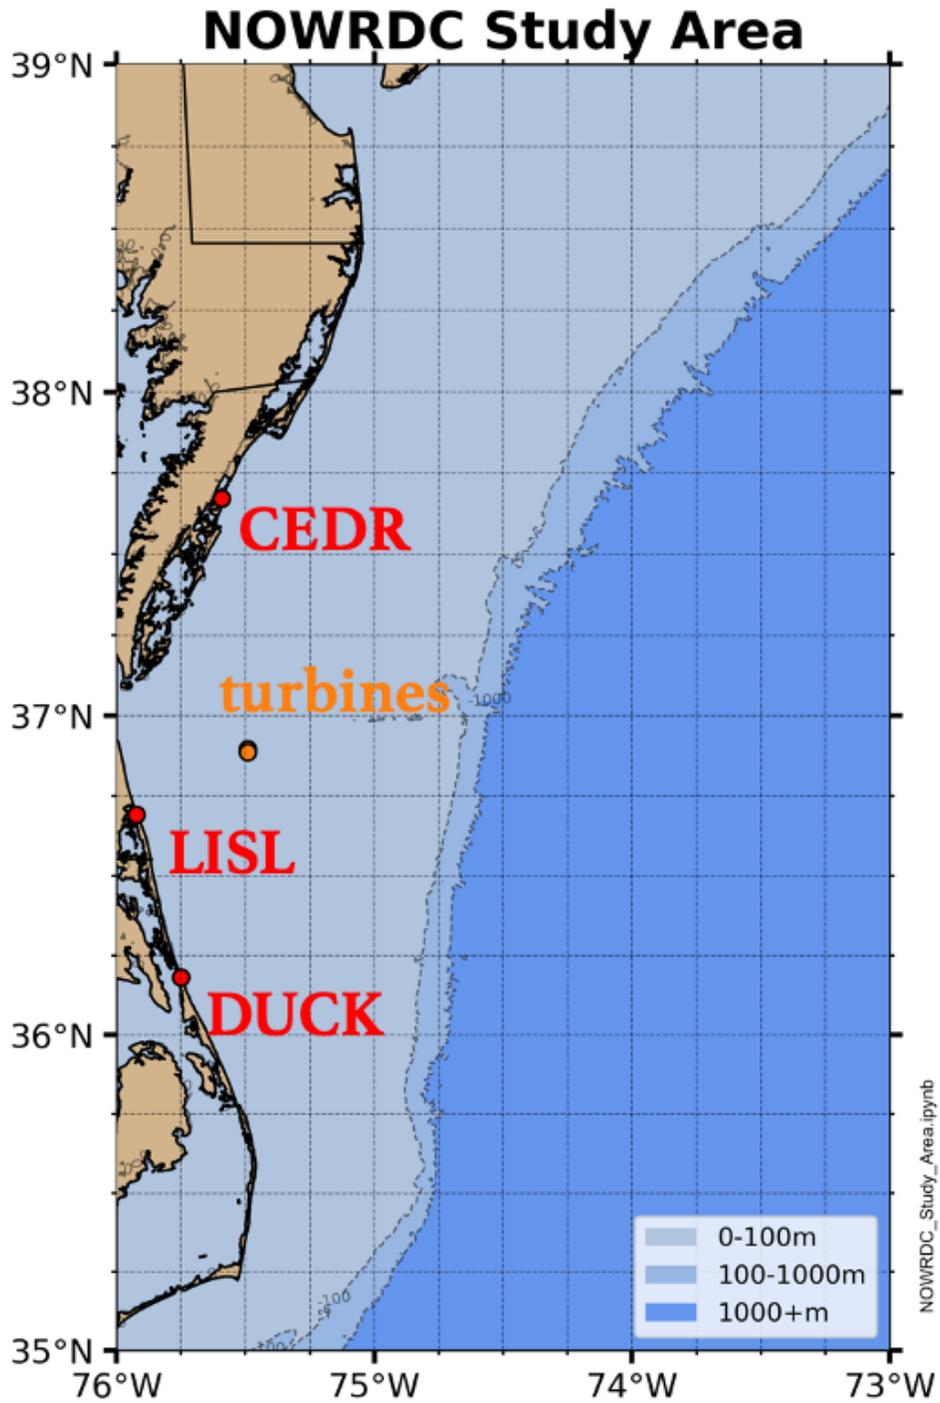


Figure 23: Map of the study area showing the location of the Cedar Island, VA (CEDR), Little Island, VA (LISL) and Duck, NC (DUCK) SeaSonde stations.

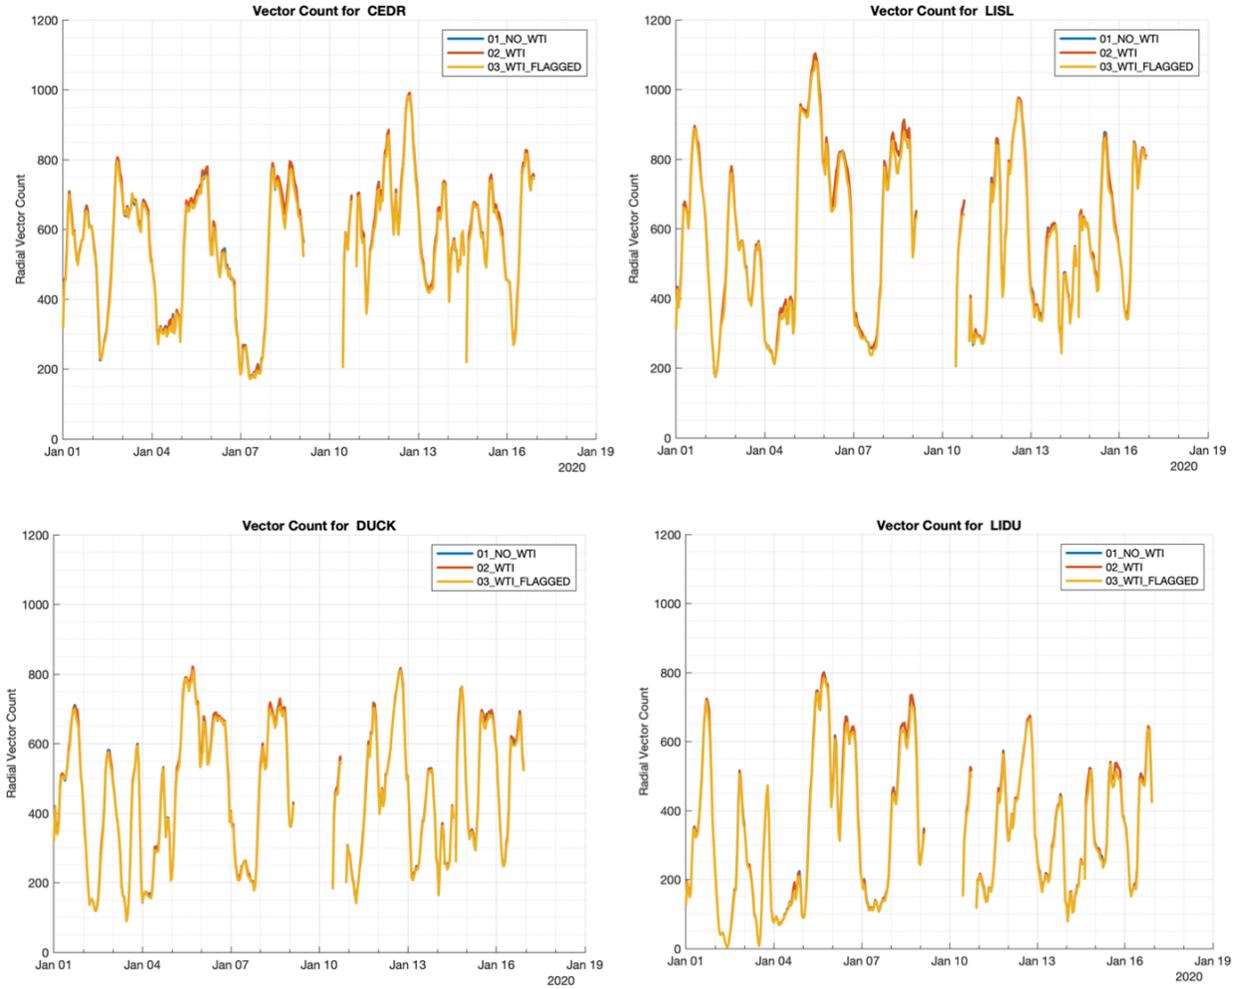


Figure 24: Radial vector count for each SeaSonde station CEDR (upper left), LISL (upper right), DUCK (lower left) and LIDU (lower right). Each panel contains the data from each test case no WTI (blue), WTI inserted into spectra (red) and WTI flagged (yellow).

Table 12: Location of the two turbines used in the study.

Turbine	Latitude	Longitude
1	36° 53' 12.58732" N	75° 29' 29.68120" W
2	36° 53' 46.66124" N	75° 29' 29.89972" W

Table 13: Range cells where the wind turbines are located relative to each radar station and range cells where WTI is expected.

Station	Range Cell of Turbines	Range Cells of Wind Turbine Interference
CEDR	15 and 16	12-19
LISL	8	5-11
DUCK	15	12-18

The simulation tool developed by Trockel et al. (2021) was modified for this project to include a variable rotation rate of the turbines. Two turbines were simulated at corresponding locations to the CVOW turbines in the overlapping coverage area of CEDR, LISL, and DUCK. To achieve a realistic range of amplitudes of the simulated WTI, we looked at the SCADA data and associated WTI dataset the radar at Block Island (BLCK). The WTI at BLCK tended to have WTI with a maximum amplitude bound between 5dB below the Bragg or sea echo, and an amplitude similar to the Bragg. We used this range of amplitudes to simulate the turbine interference at each of the sites. The rotation rate of the turbines determines the location of the interference in Doppler. To accurately test the mitigation techniques, we chose at least one of the turbine's rotation rates so that at least one of its first four positive and negative harmonic peaks would be placed in the Bragg region. Since these rotation rates are not the turbine's optimal rotation rate, and they would likely occur only during times of changing wind conditions and would not be consistent between cross spectra, we chose the rotation rates randomly within the ranges that would place some WTI in the Bragg. We also selected a variance of rotation from a normal distribution with 0 as the mean and 0.2 as the standard deviation, as this is what we observed in the BLCK SCADA data.

F.1.2 Flagging

To determine the rotation rate of the turbines, we used the ML models developed for this project. These models were trained on a data set that had similar amplitudes to those within the simulated test. As input, the models take the range-reduced portion of the cross-spectra that contains WTI. This is a range slice of three range bins above and below the locations of the turbines.

Using the output estimated rotation rates from the models, we flagged the dataset. Since one of the harmonic peaks was always placed in the Bragg, it was necessary to estimate the width in Doppler of that peak based on the other peaks found in the cross spectra. The widths in Doppler of each WTI harmonic peak not found in the Bragg were averaged together, and this was the number of Doppler bins that were flagged. Three range bins were flagged for each peak: one containing the turbine, one below, and one above. This process was repeated for each cross spectra in the simulation.

Once the flagging was completed, it became apparent that although the error in rotation rate estimate was small, this small error could shift the flags onto incorrect Doppler bins. The way the variable rotation rate was implemented also placed the strongest portion of interference away from the starting rotation rate. Because of these shortcomings, we developed a new method for flagging using the rotation rates and the analytical methods described in Trockel, Rodriguez-Alegre, Barrick and Whelan (2018).

The analytical flagging methods are done in two parts. First, symmetrical peaks in the cross spectra are identified and associated with harmonic peaks from estimated rotation rates. Second, the range-Doppler bins associated with the estimated rotation rates are flagged. The symmetrical nature of these peaks is only present for radars operating at 4Hz as lower sweep rates lead to aliasing of the WTI, which removes symmetry and, therefore, the ability to estimate rotation rates.

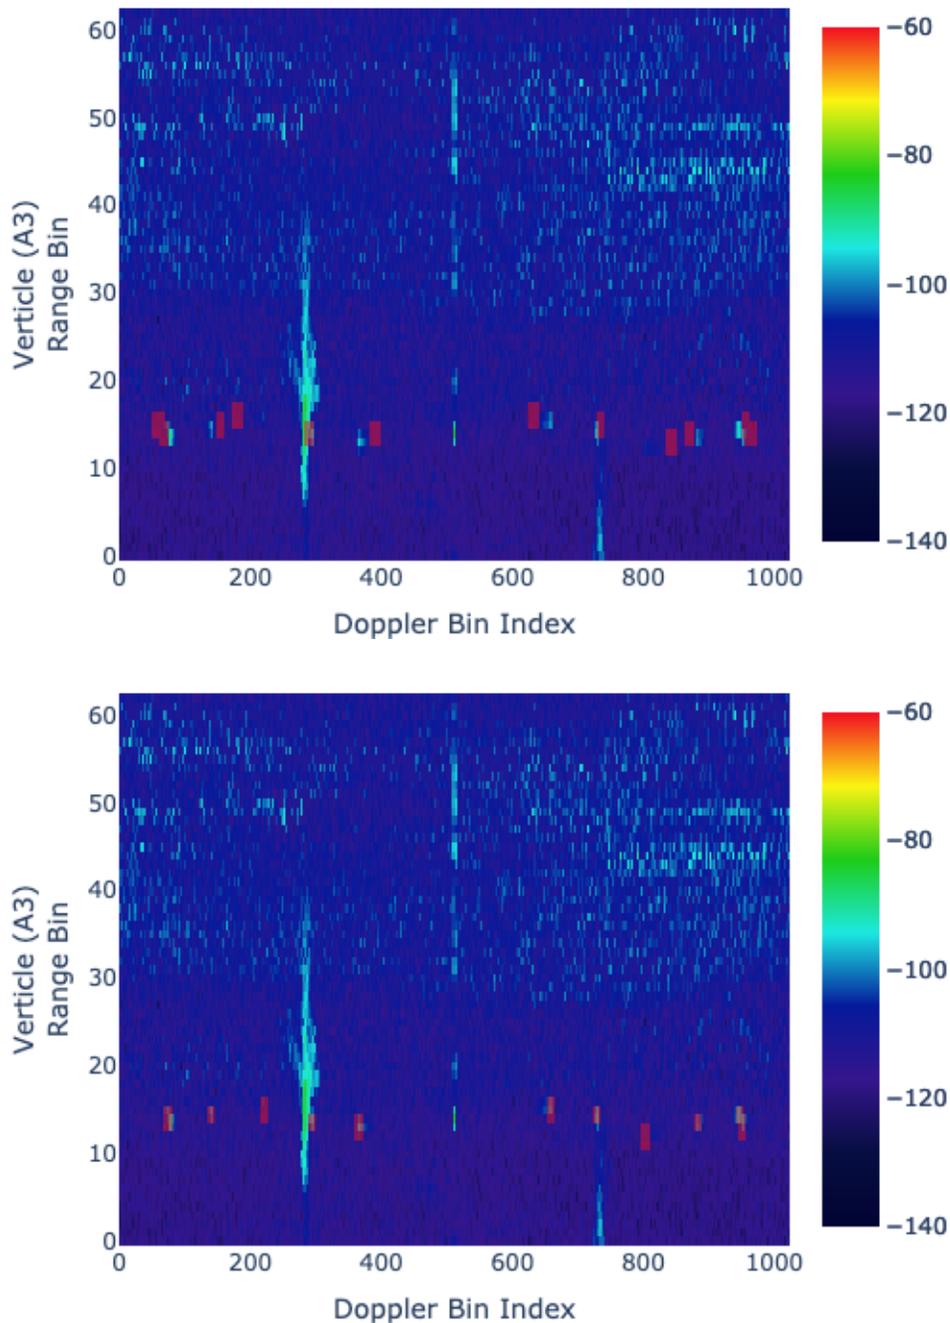


Figure 25: Flagging using different methods. Top shows flagging using only ML to estimate the rotation rate. The red squares indicate the location of the flagged range-Doppler bins. Bottom shows flagging using ML to estimate rotation rate, then corrected using analytical methods. The red bounding boxes match the peaks more accurately in Right. This results in different sections of the Bragg being flagged.

In this hybrid method between the ML and the analytical methods, the rotation rates estimated by the ML models are used to form a search-window around the areas associated with each harmonic peak. This removes the need for the initial symmetry used in the analytical method. The dominant peak within each

window is then found and the location of these peaks in range and Doppler are used to correct the initial rotation rate. This allows for a correction in the ML predictions based on the location of WTI peaks in range and Doppler. The corrected rotation rate is then used to flag infected range-Doppler bins. Using the hybrid method over the analytical method removes the need for higher sweep rates. This hybrid model was used to flag a new dataset for comparison. An example of the difference in flagging between the two methods is shown in Figure 25.

F.1.3 Surface Current Processing

The Doppler spectra were processed in four ways. The first method did not include WTI. Subsequent methods added WTI and tested different approaches to mitigation. No flagging was applied in method 2. In method 3, spectra data were flagged and removed using ML alone. In method 4, spectra data were flagged and removed using ML and the analytical methods described in section 2.2. Table 14 shows the four cases and their abbreviations in the figures. These four spectra data sets were processed into four sets of radials at each station and four sets of LIDU ellipticals. Radial vector counts are shown in Figure 24. Those radials and ellipticals were then processed into totals. The totals were processed using (1) only two sets of radials (CEDR and DUCK), (2) using all three sets of radials (CEDR, DUCK, and LISL), and (3) using all three sites' radials as well as the ellipticals from LISL and DUCK (CEDR, DUCK, LISL, and LIDU). The totals were processed in those three ways for each of the four radial/elliptical datasets.

Table 14: Summary of four methods of processing radials.

Method	Type	Abbreviation
1	No Wind Turbine Interference	01_NO_WTI
2	With Wind Turbine Interference	02_WTI
3	Wind Turbine Interference Flagged and Removed	03_WTI_FLAGGED
4	Wind Turbine Interference Flagged and Removed using the Machine Learning and Analytical Method	04_FLAGGED_ML_ANALYTICAL

F.2 Results

F.2.1 ML Assignment:

Table 4 reports the performance of the ML models at estimating rotation rates across the test data. The absolute mean error in rotation rate assignment, absolute median error in assignment, and the standard deviation of that error are reported. The absolute mean error in rotation rate estimation by the model was 0.14 rpm.

Table 15: Performance metrics of ML model in predicting rotation rates.

	Absolute Mean Error	Absolute Median Error	Standard Deviation
Model Performance	0.14 rpm	0.072 rpm	0.22pm

F.2.2 Radial Results

At the radial level, the radials produced by using spectra from methods 2, 3, and 4 were compared to the radials processed using method 1 to measure the amount of change introduced by each method. The metrics used to compare the datasets were: the number of vectors missing, the number of vectors added, the number of vectors changed, and the root mean squared difference (RMSE) in velocity among the changed vectors. These metrics were computed for each site and averaged together in Table 16.

Table 16: Radial statistics. Radials from CEDR, DUCK, and LISL were processed with four methods: 1. with no WTI added, 2. with WTI added and no mitigation performed, 3. with WTI added and WTI mitigation performed using only ML, and 4. with WTI added and mitigation performed using ML and analytical methods. Methods 2-4 were compared to method 1 for each site. The results were averaged together and average # vectors lost, # vectors added, # vectors changed, and RMSE in velocity among the changed vectors are reported here. The average total number of vectors compared was 45,070.

	WTI No Mitigation (method 2)	WTI ML Mitigation (method 3)	WTI ML and Analytical Mitigation (method 4)
# Vectors Lost	1,807	4,359	2,373
# Vectors Added	2,199	1,333	1,858
# Vectors Changed	13,769	15,252	14,210
RMSE in Velocity within Changed Vectors (cm/s)	7.97	8.28	7.85

Across the radials, methods 2, 3, and 4 had similar results. The mitigation using both ML and the analytical methods (4) combined outperformed the WTI mitigation using only ML (3). The RMSE in velocity between radials processed using method 2 and 1 was 7.97cm/s The RMSE in velocity between radials processed using method 3 and 1 was 8.28 cm/s The RMSE in velocity between radials processed using method 4 and 1 was 7.85 cm/s. Method 4 consistently performed at least as well as method 2, consistently outperformed method 3, and performed better than both during times of extreme change as can be seen in Figure 26.

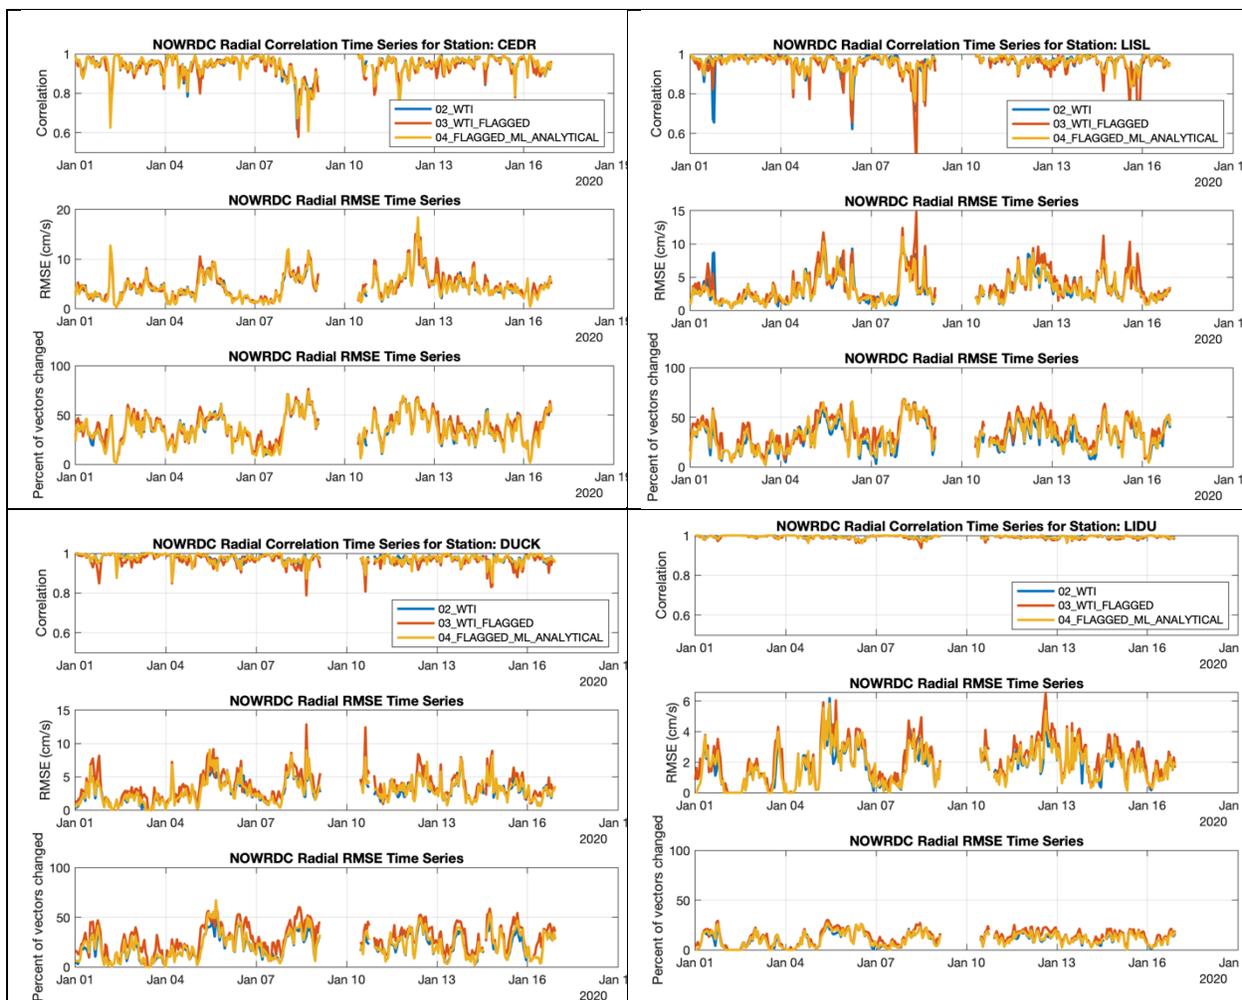


Figure 26: Statistics of the radial data for each site. Top left: CEDR, Top Right: LISL, Bottom Left: DUCK, Bottom Right: LIDU. Within each image, top panel: Correlation between the no WTI data and the three test cases, WTI inserted into spectra (blue), WTI flagged (red), WTI flagged with machine learning and analytical techniques (yellow). Middle panel: root mean square error (cm/s) between no WTI and the three cases. Bottom panel: Percentage of vectors in the affected range cells that were altered.

The radial data from each station was concatenated and the mean radial vector map was generated for each of the first three cases (01_NO_WTI, 02_WTI and 03_WTI_FLAGGED). The maps for the CEDR, LISL, DUCK, and LIDU data are shown in Figure 27. The maps show the average radial velocity as the color where red indicates currents away from the radar while blue shows currents towards the radar. The size of the dot in the maps represents the standard deviation of the velocity over the three-week period. There is no visually discernable difference between cases 1, 2 and 3 in any of these figures.

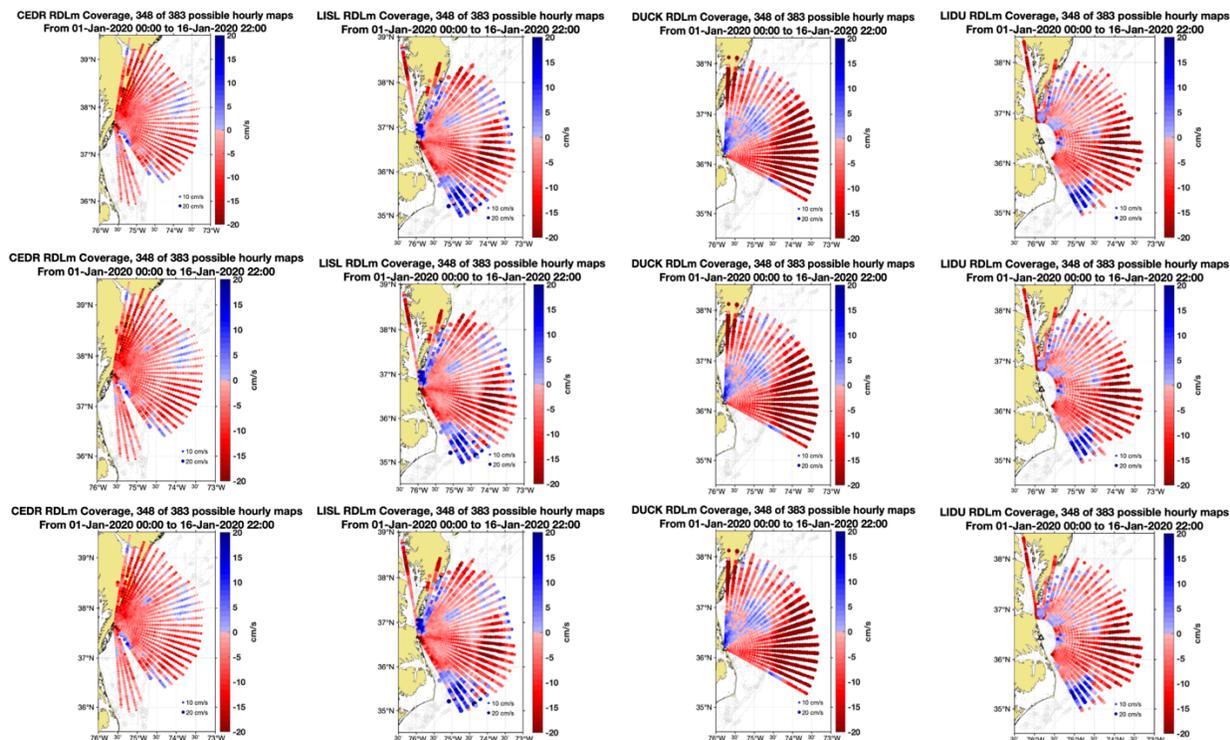


Figure 27: Average radial vector map for the CEDR station. Each map contains the data from each test case no WTI (top), WTI inserted into spectra (middle) and WTI flagged (bottom). The mean velocity is shown as the colorbar and the standard deviation is the size of the dot.

Figures 29-32 focus on the change in radial velocity for the range cells specified in

Table 13. The radial velocity from Method 1 (No WTI) was compared against the data from methods 2-4. The data was matched against range and bearing. An example of this analysis is shown in Figure 28. This shows the radial velocity from the 01_NO_WTI data set along the x axis compared against the 03_WTI_FLAGGED data set along the y axis. The correlation between the two data sets can be calculated along with the root mean square error (RMSE) and the percentage of vectors that were altered by the addition of the WTI or the flagging of the WTI.

Time series plots of this analysis are shown in Figure 26. Maps of the RMSE for the station data sets are provided CEDR (Figure 29), LISL (Figure 30), DUCK (Figure 31) and LIDU (Figure 32). The WTI is not restricted to the bearing of the turbines, as the mixing of WTI with ocean data in the cross spectra can distort direction finding algorithms. This means the entire range ring containing WTI has the possibility of being affected.

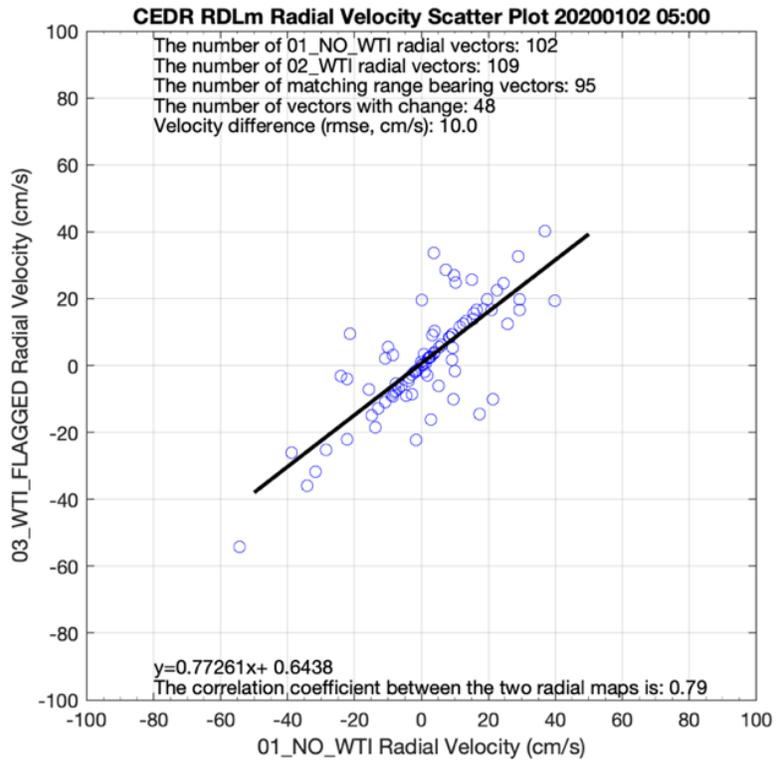


Figure 28: Scatter plot of radial velocity data matching in range and bearing between the no WTI case (x axis) and the WTI flagged case (y axis) for the CEDR station at January 02, 2020 05:00 GMT.

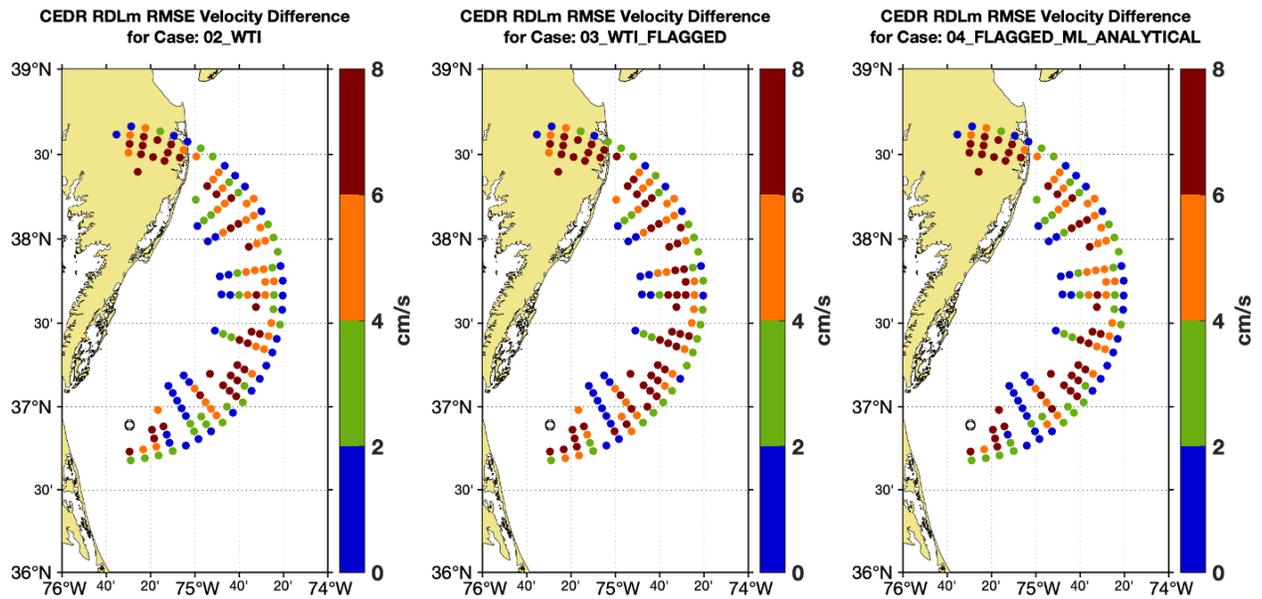


Figure 29: Root mean square error for the CEDR station for each test case WTI inserted into spectra (left), WTI flagged (middle), WTI flagged with machine learning and analytical techniques (right). The locations of the two wind turbines are shown as black circles.

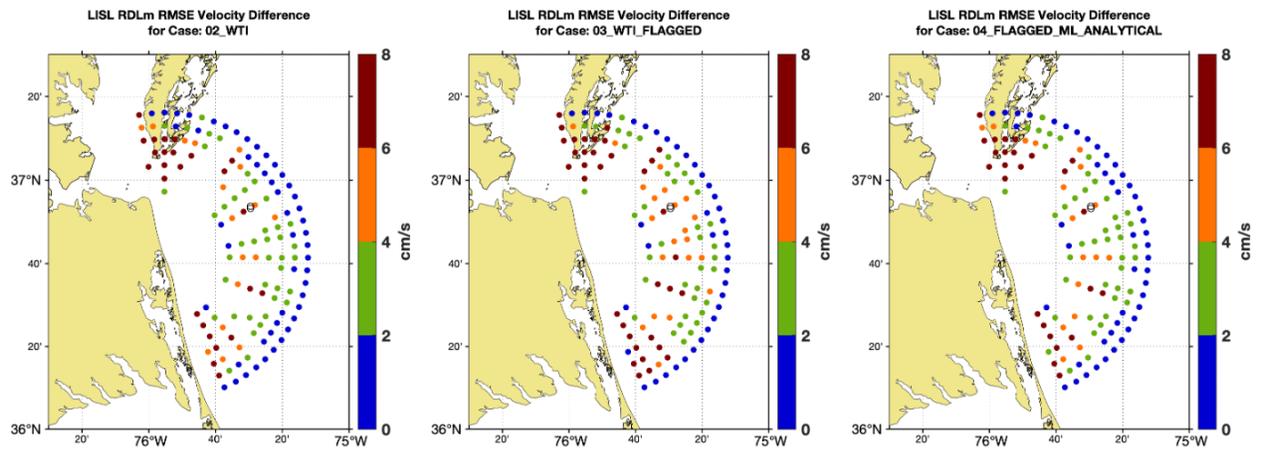


Figure 30: Same caption as Figure 29 for the LISL station.

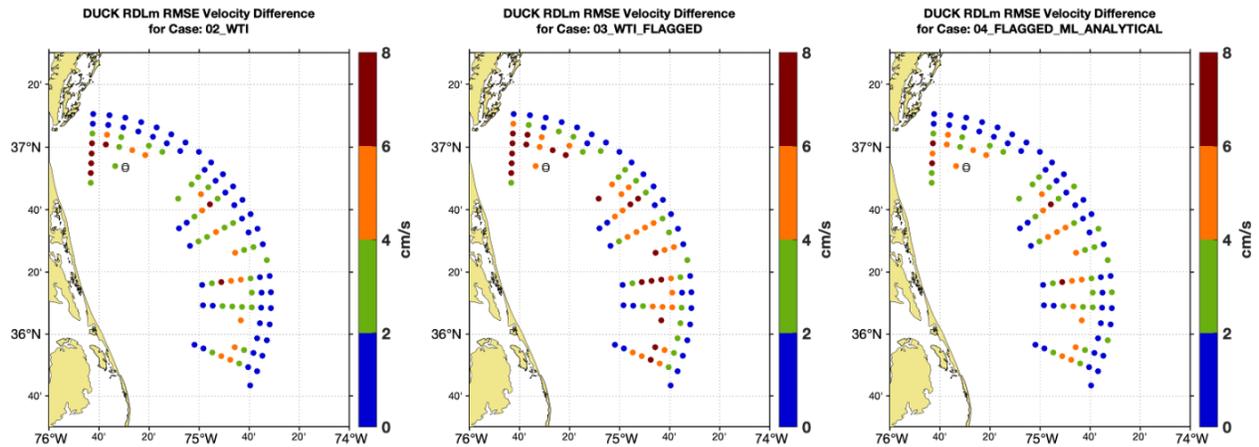


Figure 31: Same caption as Figure 29 for the DUCK station.

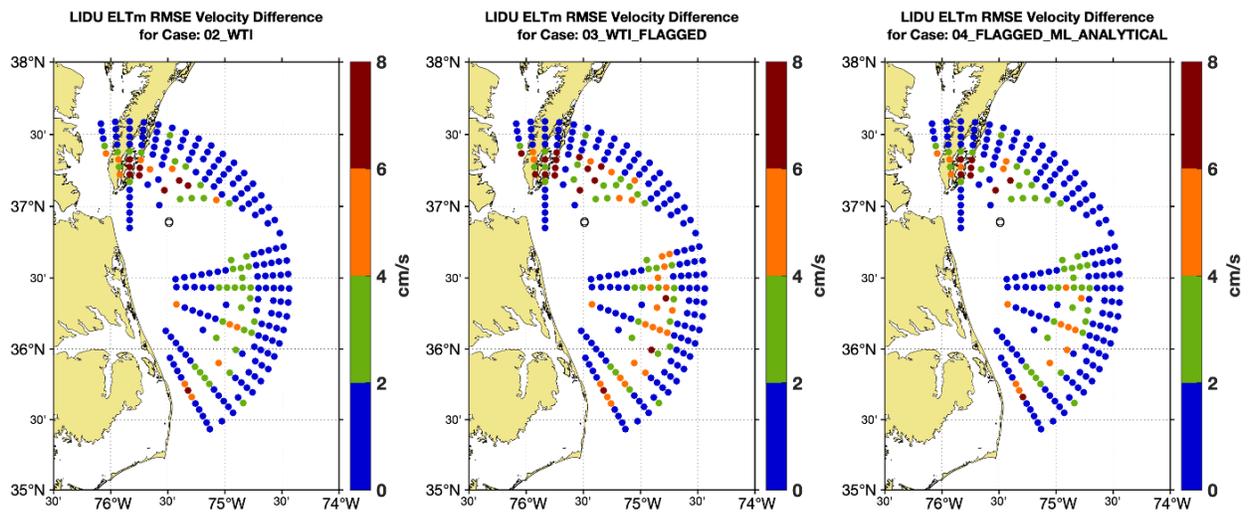


Figure 32: Same caption as Figure 29 for the LIDU station.

F.2.3 Totals Results

At the totals level, the number of vectors added, missing, and changed as well as the RMSE of the changed vectors were calculated. These metrics were computed for each level of processing: two sites, three sites, and three sites with ellipticals. The results are found in Table 17, Table 18, and Table 19.

Table 17: Totals stats using two sites. The same comparisons shown in Table 16 were repeated at the totals level for totals processed using radials from CEDR and DUCK. The total number of vectors in method 1 is 117,829. This same grid is used for comparisons in the case of three sites as well as three sites and ellipticals.

	WTI No Mitigation (method 2)	WTI ML Mitigation (method 3)	WTI ML and Analytical Mitigation (method 4)
# Vectors Lost	1,002	2,149	1,168
# Vectors Added	1,465	923	1,312
# Vectors Changed	29,630	31,316	29,806
RMSE in Velocity within Changed Vectors (cm/s)	10.85	10.8	10.94

Table 18: Totals stats using three sites. The same comparisons shown in Table 16 were repeated at the totals level for totals processed using radials from CEDR, DUCK, and LISL.

	WTI No Mitigation (method 2)	WTI ML Mitigation (method 3)	WTI ML and Analytical Mitigation (method 4)
# Vectors Lost	105	223	106
# Vectors Added	116	79	113
# Vectors Changed	37,259	40,802	37,821
RMSE in Velocity within Changed Vectors (cm/s)	5.48	6.45	5.58

Table 19: Totals stats using three sites and ellipticals. The same comparisons shown in Table 16 were repeated at the totals level for totals processed using radials from CEDR, DUCK, LISL, as well as ellipticals received at LISL from DUCK.

	WTI No Mitigation (method 2)	WTI ML Mitigation (method 3)	WTI ML and Analytical Mitigation (method 4)
# Vectors Lost	98	183	94
# Vectors Added	101	68	97
# Vectors Changed	40,569	45,532	41,351
RMSE in Velocity within Changed Vectors (cm/s)	5.33	6.02	5.38

Across the totals, the WTI dataset without mitigation and WTI datasets with flags had similar results. The RMSE in velocity relative to the dataset from method 1 for totals combined using two sites was 10.85 cm/s for the method 2, 10.80 cm/s for the method 3, and 10.94 cm/s for method 4. The RMSE in velocity relative to the dataset from method 1 for totals combined using three sites was 5.48 cm/s for method 2, 6.45 cm/s for method 3, and 5.58 cm/s for method 4. The RMSE in velocity relative to the dataset from method 1 for

totals combined using three sites and ellipticals was 5.33 cm/s for method 2, 6.02 cm/s for method 3, and 5.38 cm/s for method 4.

The overall change in totals across the entire coverage area caused by the two simulated turbines and the different mitigation methods was small. Visually, only three total vector maps had identifiably changed vectors. Figure 33 shows one such map.

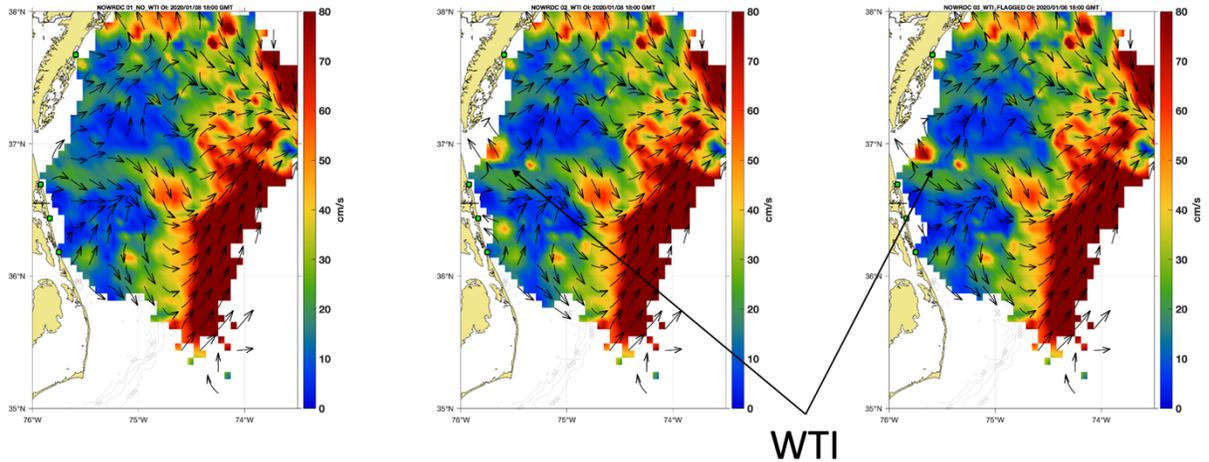


Figure 33: Total vector map with change detected by processing using method 2 and 3. From left to right the figure shows totals processed using method 1, method 2, and method 3.

F.3 Discussion

The results of this analysis show that the best thing that can be done for mitigation of WTI is to operate multiple sites with overlapping coverage. The RMSE in velocity between datasets with WTI and without WTI among changed vectors improved from 10.85cm/s to 5.33 cm/s with the inclusion of another site and ellipticals to the totals.

It is also clear that flagging of transient peaks with a low SNR does not significantly change the reliability of the data, especially at the totals level. One possible reason for this is that the cross spectra all had different rotation rates. This means that each peak is transient between the cross spectra. The comparisons were done using radials and ellipticals, which averaged together several cross spectra. Since each of the cross spectra had WTI at different Doppler values, the impact of the WTI was averaged out. This is true in the case of flagged WTI as well. The randomness in the rotation rate of the simulations was chosen by design. The turbines located in the field of view of the radars have an optimal rotation rate which places the WTI outside of the Bragg. The only times WTI is in the Bragg are during times when the turbines are ramping up or

down due to wind conditions. This means that the peaks are unlikely to remain in the same location in the Bragg across contiguous Doppler cross spectra collect in the field and were consistent with observed WTI at BLCK. The WTI peaks also had a small SNR relative to the Bragg. This makes the effects of the WTI more likely to be averaged out across several spectra. This was again by design, as the WTI we have observed in the field has similar levels of SNR to those simulated in this analysis. It is also hard to quantify the overall error introduced by missing vectors as well as added vectors. Flagging reduces the number of vectors added but tends to increase the number of vectors missing. To measure the effect of all errors together, a drifter study would need to be performed.

These findings are good news for HFR operators as they indicate the effects of WTI created by turbines currently in operation can be mitigated by using multiple radars with overlapping coverage. As turbines increase in size, the transient nature of the interference within the Bragg could change. If the optimal rotation rate of larger turbines places the WTI in the Bragg, there would be stable interference across multiple contiguous cross spectra which would not be averaged out. With larger turbines and more turbines spinning at the same rate, the SNR of the WTI would also increase. In these cases, the methods developed across this project to accurately predict rotation rate and flag the cross spectra may prove necessary. To fully explore the effects of WTI mitigation on conditions such as those, it would be necessary to perform a drifter analysis to measure the true difference in velocity calculations of larger turbines and larger turbine farms. Future work funded through the NOAA Ocean Technology Transition (OTT) plans to explore these effects using drifters.

The combination of ML and analytical methods previously developed to flag the WTI is superior to either method individually. The improved accuracy and flexibility of this method to be used at any sweep rate is encouraging moving forward.

Appendix G. Field Test

G.1 Field Test Preparation

This section explains the software field test conducted on February 8, 2023, and its outcomes. The field test was conducted using the wind turbine interference (WTI) flagger developed by CODAR Ocean Sensors Ltd. The software included two different machine learning models: a model used to classify the number of turbines spinning, and another model used to estimate the rotation rate of the turbines.

The machine learning (ML) models developed for the field test at the Little Island Park (LISL) radar station testing were trained on historical data from LISL from January 2020, before turbines were rotating in the field of view of the radar. Simulated WTI was added to the range-Doppler spectra from LISL and used to train and test the ML models. A comparison between the training data and the recent LISL data indicated an increase in noise at the site since the collection of the historical data set. To aid the ML models in accounting for the increase in noise, a new simulated WTI data set was created using the LISL data from 2020 with added noise. In addition to adding background noise, simulated WTI of two wind turbines with random rotation rates, variance, yaw, and amplitude was added to the Doppler spectra files. The magnitude of the WTI was constrained relative to the maximal Bragg peak. At the Block Island radar site (BLCK), the maximum amplitude of the WTI from the wind turbines was found to have a signal to noise ratio (SNR) in the range from 5 dB below the maximal Bragg peak to the same amplitude of the maximal Bragg. Since the turbines at LISL are further from the radar, the range of simulated WTI for the training data set was scaled to be between 15 dB below the Bragg to at the Bragg.

Once the WTI data set was simulated, each range-Doppler spectra was mixed with random noise. This noise was added to the range slices, which contained the simulated WTI. The noise was added randomly across range and Doppler to simulate the noise present at LISL. See Figure 34 for an example of the noise added.

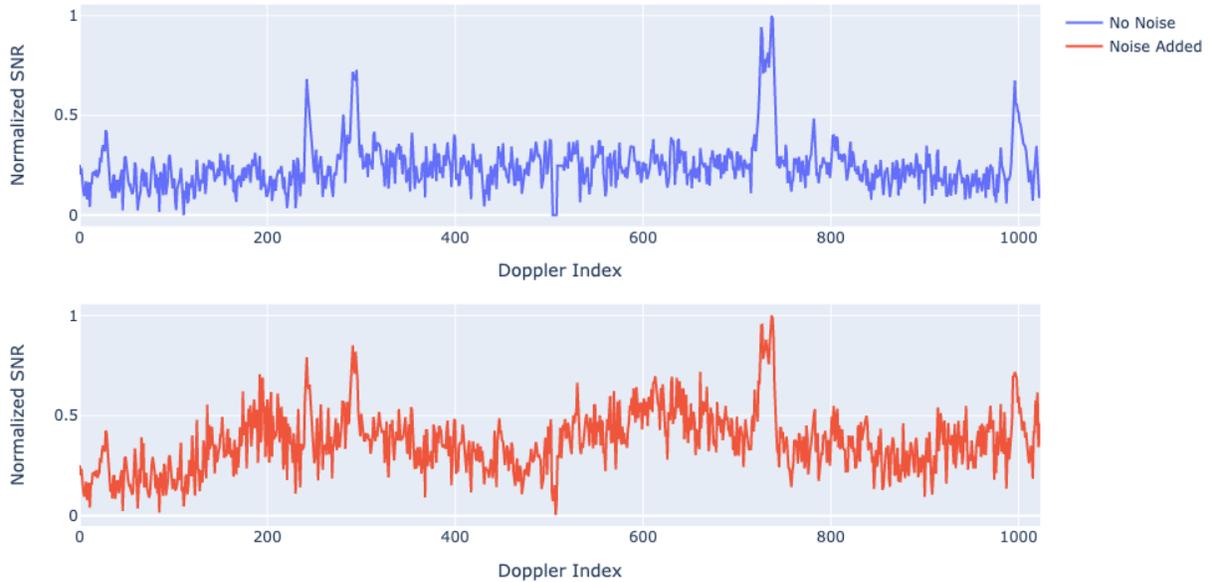


Figure 34: Range slice showing normalized signal to noise ratio of loop three with no noise added (blue) and noise added (red).

Several different classification and prediction models were trained and tested. We tested to see if separate models were necessary for the prediction of rotation rates when the two turbines were spinning at the same rate, as opposed to when they were spinning at different rates. Turbine power curves are designed such that the turbines of the same model in a wind farm will often have very similar rotation rates. To use these separate models, one for predicting rotation rates when the two turbines were spinning at similar rates, and one where the two turbines were spinning at different rates, a separate model was used to classify the two scenarios. In other words, the classifier identified if no turbines were present, two turbines were spinning at the same rate, or two turbines were spinning at different rates. If two turbines were identified to be spinning at a similar rate, the spectra would be passed to a trained model for similar rotation rates. If the model identified the two turbines to be spinning at different rates, the spectra would be passed to a model that was trained on turbines with different rotation rates. This method was tested against a simpler method where an initial model was a binary classifier that identified if turbines were present in the spectra at all, and a second model would predict two rotation rates. If the rotation rates were similar, they would still be fed in separately. After testing, although the rotation rate prediction step had better results with separate models, the larger classification error of the separate model approach led us to choose the simpler single-model design for the field test. The results of the test can be seen in Table 20.

Table 20: results of different machine learning models. “Separate Models” represents two different prediction models: one for when the two turbines are rotating at similar rates, another for when they are different. “Single Prediction Model” represents a single model used for both scenarios. Classification accuracy is the percentage of correct classifications.

	Classification Accuracy	RPM Estimation RMSE at Times when Turbines had Similar Rotation Rates.	RPM Estimation RMSE at Times when Turbines had Different Rotation Rates.
Separate Models	85%	0.09	0.13
Single Prediction Model	98%	0.13	0.18

These models were then tested against several different simulated data sets containing varying levels of added noise and different WTI parameters, to see if they would perform well under conditions that they were not trained for. Table 20 shows mean error across all simulations. Since we were unable to obtain SCADA data, we trained our models to be as versatile as possible, in the hope that they would perform well under real-world conditions. The models consistently performed well both at classifying and estimating rotation rates. The errors introduced into the assignments by larger amounts of noise were corrected by the added step of correcting the rotation rate estimates as outlined in Trockel, Rodriguez-Alegre, Barrick and Whelan (2018) and described in the Software Design brief (Deliverable 2.2).

To further test the flexibility of the models, simulated WTI was added to cross spectra from the Cedar Island (CEDR) radar and models were trained on CEDR then tested against the LISL data set. These models that were trained on a different site also performed well at LISL. This gives us confidence that the models can classify and predict rotation rates reliably across a wide range of noise and radar conditions.

The machine learning models were then tested against recent LISL data. The data did not have visible WTI and the models were consistently able to classify the spectra as not having turbines. Once the models were fully trained and tested, they were put into the software and deployed on LISL for real-time testing.

G.2 Field Test

The software described in the software design brief was deployed at LISL on February 8, 2023 and ran for 13 hours. In total, 49 range-Doppler spectra were processed by the tool. The cross spectra were then passed to the SeaSonde processing toolkit to produce radials. Figure 35 shows a SeaSonde range-Doppler spectra from LISL during that test. There is no visible WTI. The software successfully classified all the spectra as having no WTI and did not flag them. A radial map displaying average velocities during the testing period is shown in Figure 36.

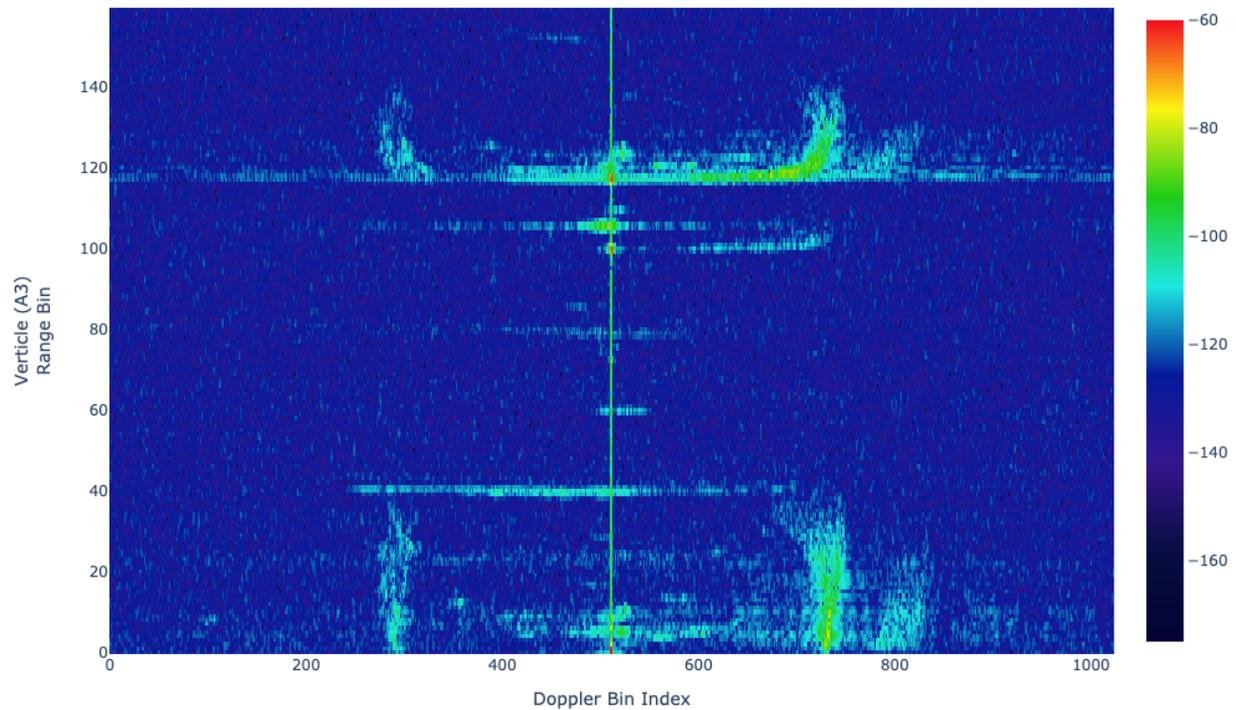


Figure 35: An example cross spectra file at LISL during the field test. There are no observable WTI peaks.

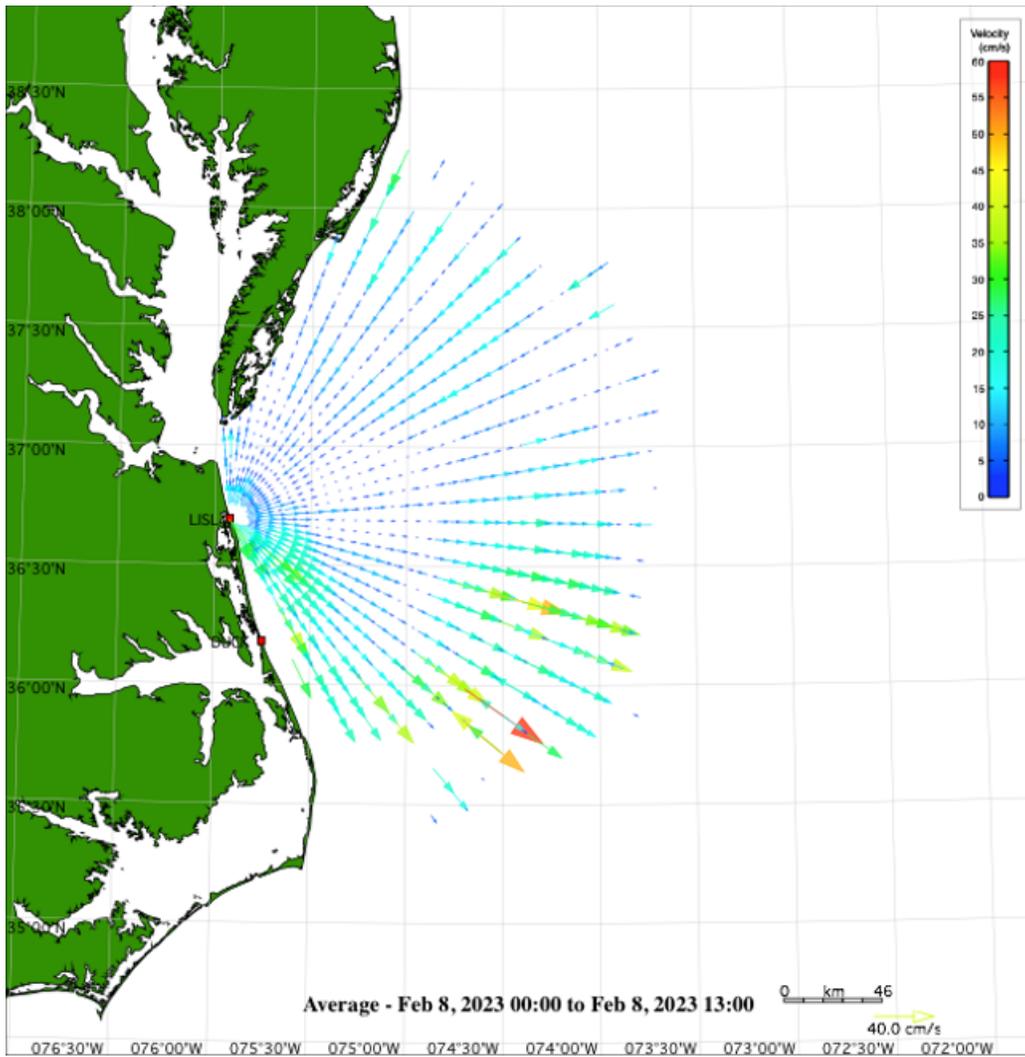


Figure 36: Average radial velocities at LISL during the field test.

The WTI flagging software field test was successful. It integrated with the existing CODAR SeaSonde processing software and produced results consistent with our expectations.

G.3 Discussion

The flexibility of the trained machine learning algorithms to different levels of noise is encouraging. The machine learning algorithms were tested on a site that it was not trained on, as well as datasets with simulated noise added to them. The models were still able to classify WTI reliably when present and estimate the rotation rates of the turbines in the field of view of the radar. This finding is encouraging because it indicates that properly trained machine learning models will be insensitive to the background noise conditions at a radar site. This also provides added confidence to the prospect of using WTI simulations to augment datasets. Furthermore, even in extreme cases where SCADA data is not available at all, this field test implies that simulated WTI can be used exclusively for the training of ML models as well. For the test performed in this study, training was done using only simulated WTI, and the flagging on the field radar worked as anticipated.

The field test also confirms the analysis and conclusions from report 3.1, namely that the wind turbines installed in the U.S. are unlikely to cause large issues with the SeaSonde surface current measurements. However, it should be remembered that each wind turbine model will have a potentially different optimal rotation rate at which it will rotate most of the time. The optimal rotation rates of the offshore wind turbines in the U.S. place WTI outside of the Bragg region. Additionally, since the radar cross-section of a wind turbine does not scale with the range from the radar, like the sea surface, the amplitude of the WTI drops off rapidly with distance from the radar, limiting the impact of wind turbines further from the radar. This field test confirms that turbines placed at further range bins do not produce large levels of WTI. These findings are also encouraging, as they mean WTI will not often interfere with ocean data from HFR systems.

As turbines are installed closer to the radar, or turbines with power curves which place WTI into sea echo portion of the Doppler spectra more often are deployed, we are confident that the mitigation schemes developed throughout this project and implemented in this real-time software package will continue to help mitigate the adverse effects of WTI that is mixed with the SeaSonde oceanographic observations. Turbines placed closer to the radar would have WTI of larger amplitude. Turbines with different power curves which place WTI in the Bragg during their optimal rotation rate would mix WTI with the Bragg region at fixed Doppler frequencies the majority of the time. The mitigation software has been developed to handle both cases and as larger turbines and offshore wind farms with more turbines are built, we anticipate an increase in the utility of the software described here.

References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Zhifeng Chen, C. C., Greg S. Corrado, A. D., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jozefowicz, R., Jia, Y., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., . . . Zheng, X. (2015). *TensorFlow*. In <https://www.tensorflow.org/>
- Barrick, D. E., Lipa, B. J., Lilleboe, P. M., & Isaacson, J. (1994). Gated FMCW DF radar and signal processing for range/Doppler/angle determination. In: Google Patents.
- Burke, G. J., Poggio, A., Logan, J., & Rockway, J. (1979). Numerical electromagnetic code (NEC). 1979 IEEE International Symposium on Electromagnetic Compatibility,
- Colburn, R. J., Randolph, C. A., Drummond, C., Miles, M. W., Brody, F. C., McGillen, C. D., Krieger, A. S., & Jankowski, R. E. (2020). Radar Interference Analysis for Renewable Energy Facilities on the Atlantic Outer Continental Shelf. OCS Study, McLean. *Sterling, VA: US Department of the Interior, Bureau of Ocean Energy Management. OCS Study BOEM.*
- Harlan, J. (2009). A plan to meet the nation's needs for surface current mapping. *NOAA IOOS*, 64.
- Lipa, B., & Barrick, D. (1983). Least-squares methods for the extraction of surface currents from CODAR crossed-loop data: Application at ARSLOE. *IEEE Journal of Oceanic Engineering*, 8(4), 226-253. <https://doi.org/10.1109/JOE.1983.1145578>
- Lipa, B., Whelan, C., Rector, B., & Nyden, B. (2009). HF Radar Bistatic Measurement of Surface Current Velocities: Drifter Comparisons and Radar Consistency Checks. *Remote Sensing*, 1(4), 1190-1211. <http://www.mdpi.com/2072-4292/1/4/1190>
- Teague, C. C., & Barrick, D. E. (2012). Estimation of wind turbine radar signature at 13.5 MHz. 2012 Oceans,
- Trockel, D., Rodriguez-Alegre, I., Barrick, D., & Whelan, C. (2018). Impact Assessment and Mitigation of Offshore Wind Turbines on High Frequency Coastal Oceanographic Radar. *Sterling, VA: US Department of the Interior, Bureau of Ocean Energy Management. OCS Study BOEM*, 53, 49.
- Trockel, D., Rodriguez-Alegre, I., Barrick, D., Whelan, C., Vesesky, J., & Roarty, H. (2018). Mitigation of Offshore Wind Turbines on High-Frequency Coastal Oceanographic Radar. OCEANS 2018 MTS/IEEE Charleston,
- Trockel, D., Trockel, J., & Whelan, C. (2021). Coastal High Frequency Radar Wind Turbine Interference Mitigation. Sterling, VA: US Department of the Interior, Bureau of Ocean Energy Management. 37 p. Report No.: OCS Study BOEM 2021-081. *Contract(140M0120C0002)*.
- Trockel, D., & Whelan, C. (2021). *Coastal High Frequency Radar Wind Turbine Interference Mitigation* (OCS Study BOEM 2021-081). https://espis.boem.gov/final%20reports/BOEM_2021-081.pdf
- Wilson, S. G., & Fischetti, T. R. (2010). *Coastline population trends in the United States 1960 to 2008*. US Department of Commerce, Economics and Statistics Administration, US . . .
- Wyatt, L., Robinson, A., & Howarth, M. (2011). Wind farm impacts on HF radar current and wave measurements in Liverpool Bay. OCEANS 2011 IEEE-Spain,