BeFORECAST

Deliverable: Virtual met masts from deep learning and data fusion

Deliverable No.: D2.1

This project has received funding from the Energy Transition Fund of the Belgian Federal Public Service for Economy, SMEs, and Energy (FOD Economie, KMO, Middenstand en Energie)

DOCUMENT CONTROL SHEET

DOCUMENT INFORMATION	
Document title	(Virtual) met masts from deep learning and data fusion
Author(s), (organization)	Pieter-Jan Daems, Timothy Verstraeten, Rebeca Marini, Stijn Ally, Jan Helsen (Vrije Universiteit Brussel), Wim Munters (VKI), Daniele D'Ambrosio, Hasan Yazicioglu (3E)
Deliverable No.	D2.1
Work Package No.	WP2
Dissemination level	Public
Date of issue	November 2025

DOCUMENT SUMMARY

This document aims to present the results performed in T2.1, T2.3 and T2.6, which have as overarching goal to calibrate data sources such as reanalysis data using field-data (e.g. LiDARs, scada data). To this end, the following two steps are followed:

1. Assessment of the accuracy of the reanalysis data w.r.t. the inflow conditions observed by the scada data. As reanalysis datasets are available on a long-term basis (25+ yrs), they are commonly used for wind resource assessment. Therefore, it is essential to validate their behaviour at different locations with real measurements. As different non-linear effects play a role here (wake, blockage), a deep learning model is used to perform the calibration.
2. Assessment of the accuracy of the scada data compared to the nacelle LiDAR measurements. As scada data is commonly one of the most important tools for assessing turbine power losses, the accuracy of this data source is of significant importance. As the relationship proved less complex compared to the one in step (1), linear models and gradient boosting models are used.

Table of Contents

1 Introduction

This document aims to present the results performed in T2.1, T2.3 and T2.6, which have as overarching goal to calibrate data sources such as reanalysis data using field-data (e.g. LiDAR, scada data). In this way, virtual met masts are created that can assess the expected wind field at a given time using continuously available data (e.g. scada). In particular, the following two steps are followed:

1. Assessment of the accuracy of the reanalysis data w.r.t. the inflow conditions observed by the scada data. As reanalysis datasets are available on a long-term basis (25+ yrs), they are commonly used for wind resource assessment. Therefore, it is essential to validate their behaviour at different locations with real measurements. As different non-linear effects play a role here (wake, blockage), a deep learning model is used to perform the calibration.
2. Assessment of the accuracy of the scada data compared to the nacelle LiDAR measurements. As scada data is commonly one of the most important tools for assessing turbine power losses, the accuracy of this data source is of significant importance. As the relationship proved less complex compared to the one in step (1), linear models and gradient boosting models are used.

1.1 Structure

To meet the above mentioned goals, Section 2 discusses the need of calibration of reanalysis data, and achieves this by means of a deep learning model. Afterwards, Section 3 discusses the calibration of scada data by means of the LiDAR measurement campaigns performed in T1.2.

2 Calibration and downscaling of reanalysis data [VUB, 3E, VKI]

The goal of this section is to give insights into the calibration of reanalysis data by means of scada data. Only a brief summary of the outcomes is given in this section, for a full description of the applied methods, references are given to published articles (Section 5).

2.1 Need for calibration [VUB]

Within the Beforecast project, several datasets have been gathered which capture the wind conditions within the Belgian offshore zone:

- Two wind speed model databases, ERA5 and NORA3
- A combination of measurement campaign data (nacelle LiDAR), wind farm data (scada), and open-source data (Meetnet Vlaamse Boeien). These are shown in Figure 1.

These different datasets offer an opportunity to assess differences in their behaviour, and to see if these can readily be explained by the modelled physics. For example, reanalysis models do not incorporate the presence of wind turbines, meaning that effects such as wake velocity deficits and blockage are not considered and should therefore be compensated for. Reanalysis data is at present day commonly used for wind resource assessment, and it is therefore important to have a framework to transform this reanalysis data into local measurements by compensating for these phenomena. The other approach is to apply models that directly model these, but this comes at a greater computational cost than the deep learning models that are established within T2.1, and often requires significant expertise. This is not the case with the established deep learning model, which is computationally cheap to execute after its initial training step.

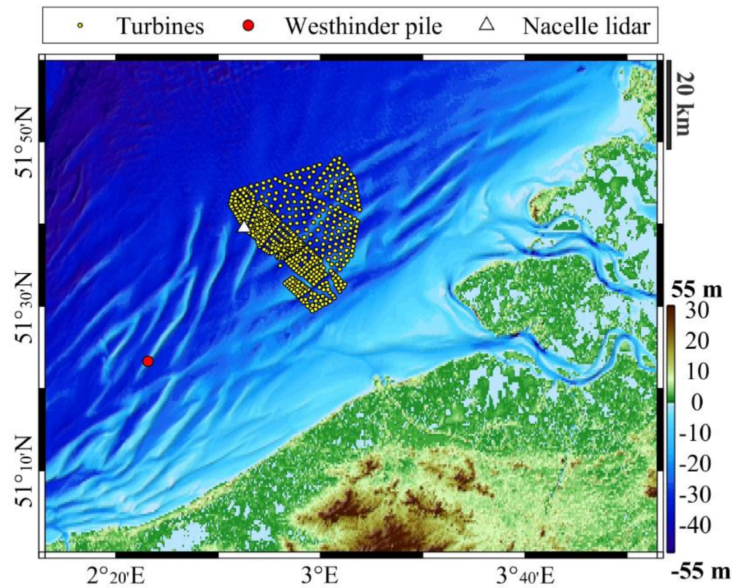

Figure 1: Overview of the available measurement locations within the Beforecast project, consisting of open-source data (e.g. Westhinder pile), dedicated measurement campaigns (e.g. nacelle LiDAR campaigns within T1.2), and wind farm data (e.g. scada data). In addition, long-term reanalysis data (ERA5, NORA3) is available.

To evaluate the influence of a neighbouring wind farm on the accuracy of reanalysis data, the bias between SCADA measurements, NORA3 and ERA5 modelled wind data is analysed over time. The results, visualised through quantile plots in Figure 2, reveal a temporal change in the correlation after the construction of neighbouring wind farms. The quantile plots demonstrate that before the introduction of neighbouring wind farms, high correlations are obtained between the reanalysis data and the scada observations (R^2>0.9) for both NORA3 and ERA5. After construction of a neighbouring wind farm in Year 9 (Figure 2d), the R^2 drops to 0.8, indicating a worsening alignment between modelled and measured wind speeds. Furthermore, it can be noted that the observations start changing significantly from the 1:1 relationship that is ideally observed.

These findings highlight the influence of wake interactions introduced by neighbouring wind farms on the accuracy of reanalysis datasets. The drop in correlation underscores the need for refined models or correction mechanisms to account for these complex interactions, establishing virtual met masts. In the remainder of the document, different methodologies are discussed to establish this.

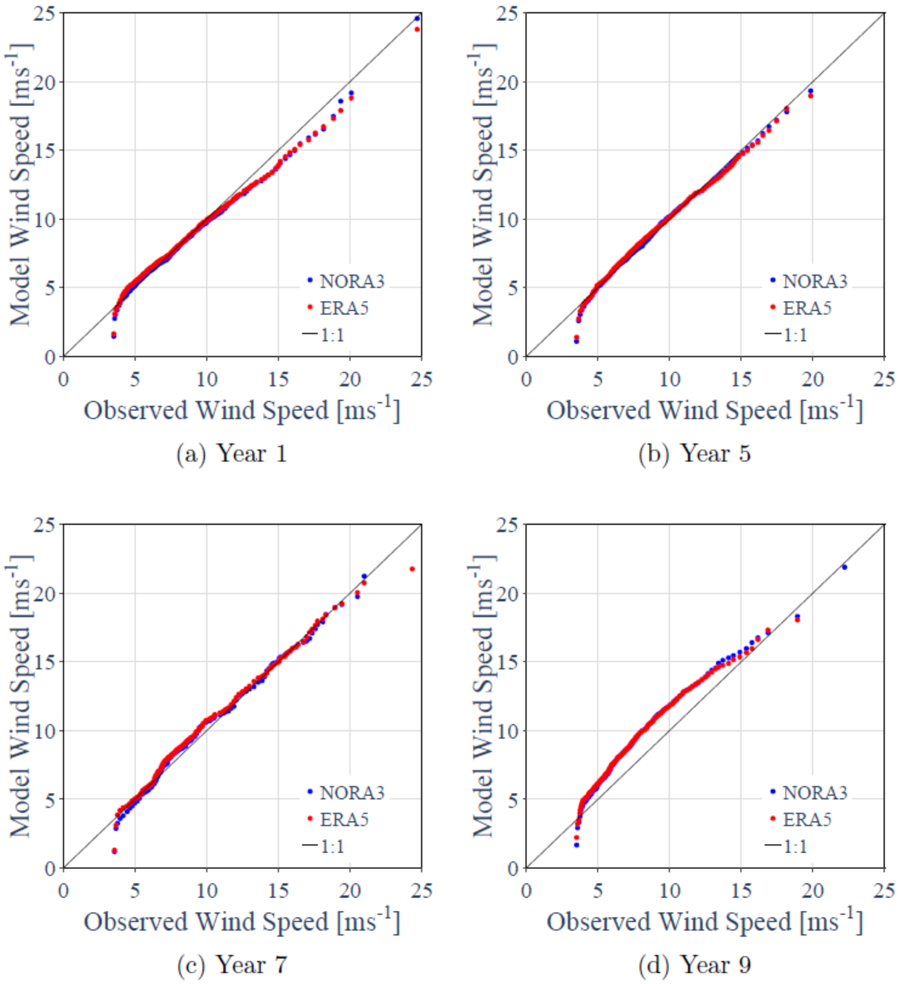

Figure 2: Quantile-quantile plots along time of observed data (SCADA) in WF1 compared with the model data (NORA3 in blue and ERA5 in red) for different years before, (a) and (b), and after the introduction of a single row of wind turbines (c) and the construction of a wind farm (d).

2.2 Calibration by means of deep learning [VUB]

The calibration of the reanalysis data is performed by means of a convolutional neural network, of which details can be found within [2]. In short, this neural network is used to map the inflow conditions, measured by the scada data to the reanalysis data (i.e. ERA5) to consider the presence of neighbouring wind farms. An example of the estimated correction factors can be seen in Figure 3. It can be noted that the ERA5 data overpredicts the observed wind speed in almost all cases, as wake velocity deficits are not considered. Smaller correction factors are obtained in case directions are analysed where a significant amount of neigboring wind turbines are present. This calibration was not only performed on reanalysis data, but also on the wind forecasts provided by two commercial parties. Different error metrics have then been compared before and after calibration, considering the scada measurements of the free-flow turbines as reference. This is shown in Figure 4 for different time horizons. It can be noted that the different errors (ME, MAE, RMSE) are lower for the three different time horizons, and that the correlation (R^2 value) is higher after correction with the developed model. This highlights its proper working and importance for power forecasting (meteo forecasts) and/or power potential assessment (reanalysis data) based on these data sources.

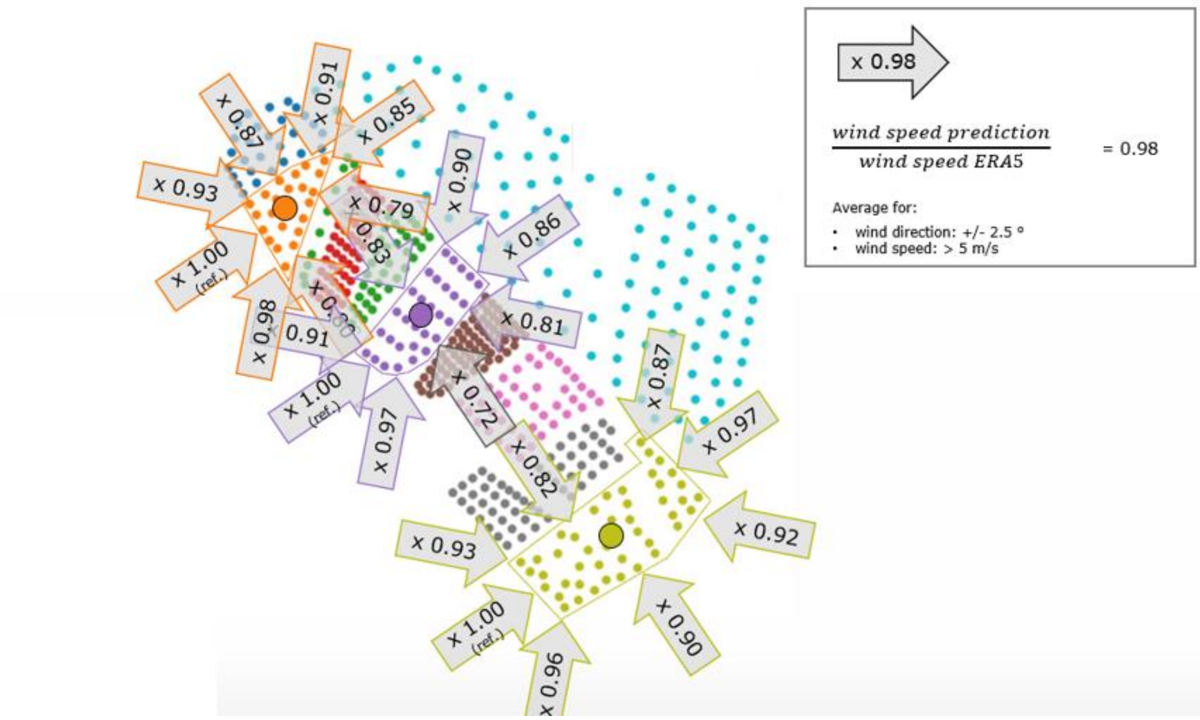

Figure 3: Correction factors of the model for different wind farms and wind directions.

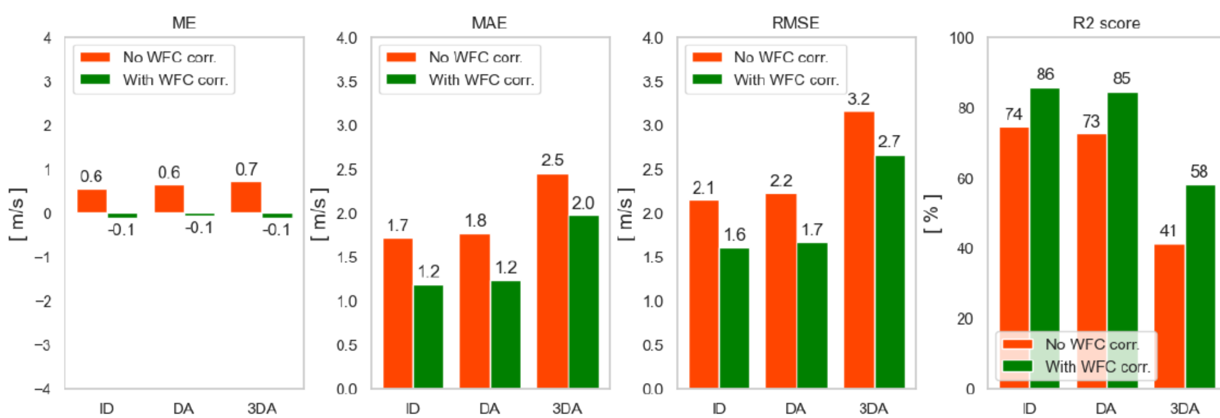

Figure 4: Error metrics for intra-day (ID), day-ahead (DA) and three-days-ahead (3DA) wind speed forecasts for Wind Farm 1. The orange bars are the metrics for the wind speed forecasts of a commercial weather forecast provider. The green bars are the metrics after application of the WFC model.

2.3 Downscaling of reanalysis data by means of SRCNN and WRF [3E, VKI]

Deep leaning models based on super-resolution convolutional neural networks (SRCNN) can be leveraged to downscale reanalysis data to the high-resolution mesoscale range. Reanalysis data, such as ERA5, provide atmospheric variables at a horizontal resolution of tens of kilometers. These data are employed as boundary forcing to WRF simulations to further downscale the atmospheric variables in the mesoscale range. Due to their high computational requirements and consequent cost, these physics-based numerical weather prediction simulations typically run to downscale only one year of data. SRCNN-based models can help overcome this limitation in that they can be trained with only one year of WRF data and used to downscale multiple years of ERA5 data at a much lower computational cost.

In T2.6, the 2km resolution WRF mesoscale simulations of the entire Belgian offshore zone for the year 2023 provided by VKI were employed to train the SRCNN model developed by 3E, named deep learning-WRF (DLWRF). To do that, first a re-gridding process of the input WRF data was performed to comply with the DLWRF input grid requirements and its standard resolution, which is 3km. This re-gridding process is illustrated in Figure 5. Two DLWRF models were then trained with the available WRF simulations, one including the wind farm parametrization (WFP) of the entire Belgian offshore zone and a second one without any wind farms in the domain.

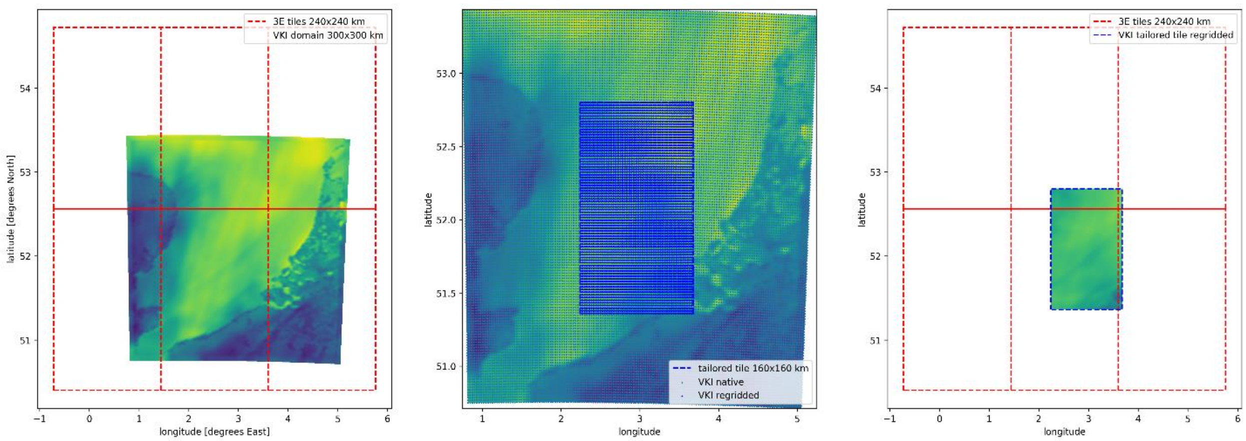

Figure 5: Re-gridding process of the VKI WRF dataset. Comparison between VKI simulation domain and DLWRF tiling system (left). Output of the re-gridding process and resulting domain selection in blue w.r.t. initial VKI domain (centre). Resulting tile used for DLWRF training compared to the native DLWRF tiling system (right).

The trained models were then used to downscale or predict WRF-like data from ERA5 reanalysis data at the centroid of Norther wind park. Downscaled variables included wind speed, wind direction and air density at several height levels, as well as additional diagnostic atmospheric variables typical of the WRF output. A model-to-model comparison was performed by calculating selected error metrics, namely bias, MAE, and R^2, for the predicted DLWRF time series with respect to the WFR wind speed data available for validation at the closest location to the Norther centroid for the year 2023. Table 1 shows the metrics resulting from the wind speed time series comparison carried out for models with and without WFP.

Model	Bias (m/s)	MAE (m/s)	R^2
DLWRF	0.3	1.48	0.83
DLWRF-WFP	1.5	1.97	0.7

Table 1: Wind speed error metrics comparison calculated with respect to WRF data for models with and without WFP for 2023.

This comparison demonstrates that the SRCNN-based downscaling model (DLWRF) accurately reproduces physics-based WRF time series at the mesoscale, with deviations within acceptable tolerances. To support this result, a time series comparison of the wind speed from the different models is also shown in Figure 6.

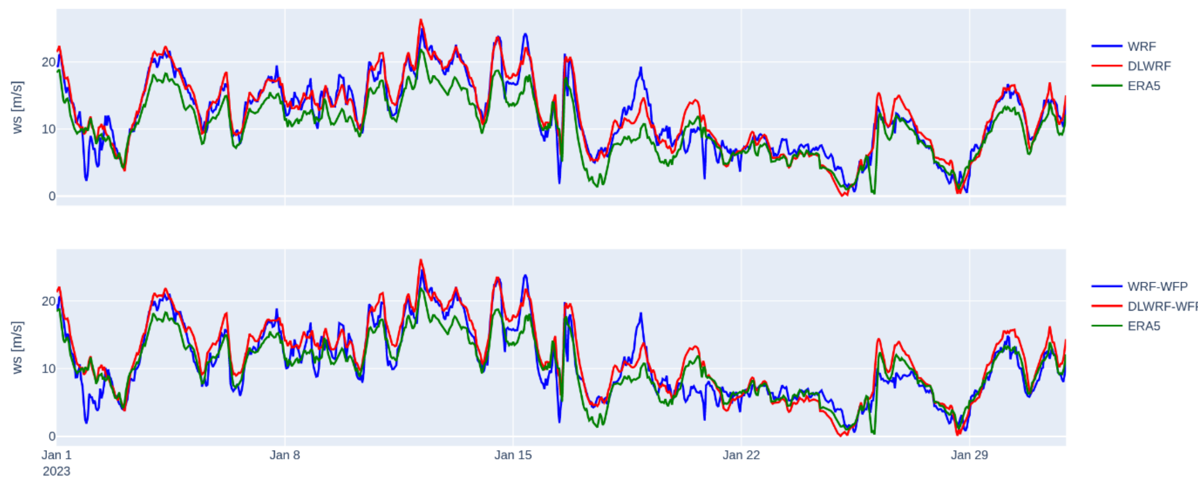

Figure 6: Wind speed time series comparison with ERA5 data for models with WFP (bottom) and without WFP (top). Shown period is January 2023.

Finally, sector-wise mean wind speed was computed and compared for all the models to assess the direction-binned performance of the DLWRF model and its accuracy in reproducing the wind speed reduction due to the waked flows from the neighboring wind farms. Table 2 shows this comparison between the WFP and non-WFP models.

Model	Sector-wise Mean Wind Speed (m/s)												
	15-45	45-75	75-105	105-135	135-165	165-195	195-225	225-255	255-285	285-315	315-345	345-15	all
ERA5	8.2	6.9	6.3	5.6	6.5	8.7	9.8	10.1	8.9	7.6	7.2	7.2	8.5
WRF	9.7	7.8	7.8	7.4	7.5	10.9	11.7	11.7	10.8	9.4	8.9	8.1	10.1
DLWRF	10.5	8.7	7.9	7.4	7.4	9.6	11.1	12.0	11.5	10.3	10.3	10.0	10.4
WRF-WFP	7.2	6.5	6.8	6.2	6.7	9.9	10.6	10.4	9.3	7.4	6.5	6.4	8.5
DLWRF-WFP	9.8	7.9	7.6	7.0	7.3	9.5	11.0	12.0	11.1	9.7	9.7	9.0	10.1

Table 2: Sector-wise mean wind speed comparison between ERA5 and models with and without WFP for 2023. Wind speed time series were binned in 30-degree wind direction sectors.

3 Calibration of turbine anemometer data [VUB]

3.1 Use of LiDAR measurement campaigns for anemometer calibration [VUB]

The anemometer of the wind turbine is located behind the rotor plane, and is therefore affected by the blockage and wake of the wind turbine. The wind turbine manufacturer therefore does not log the raw anemometer measurements, but compensates them with a transfer function to represent behavior as if the anemometer is measuring within the free-flow. It nevertheless needs to be validated that this indeed matches the correct behavior. This can be done with nacelle LiDAR measurements, which measure the wind speed at different distances from the rotor plane. An example of the wind speed evolution at different distances can be seen in Figure 5, where the fitted curve represents the one obtained using a blockage model. According to IEC61400-12-1, wind turbine performance (e.g. power curve) is to be assessed using measurements at a distance of 2.5 rotor diameters from the rotor plane. The measurements at this location are compared with scada measurements in the following sections in order to assess their accuracy.

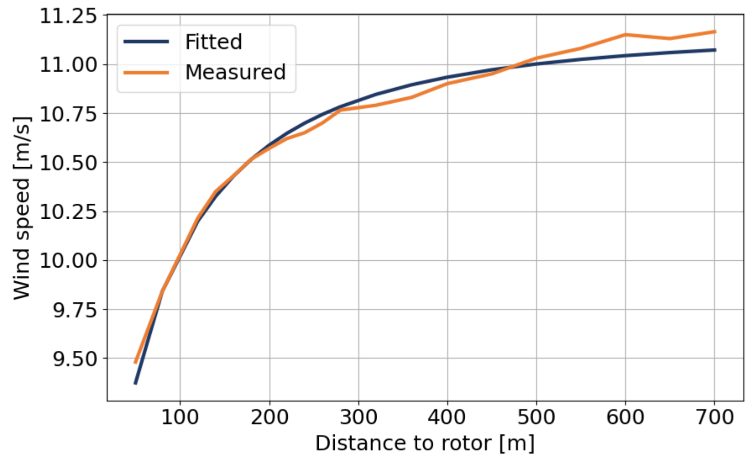

Figure 5: Wind speed at different distances with respect to the rotor plane, as measured by the LiDAR in T1.2

3.2 Calibration approaches [VUB]

In total, three different calibration approaches are explored within the project:

1. Linear model that maps the scada wind speed measurements to LiDAR measurements.
2. Linear model that estimates the wind speed based on the active power production using a training set of LiDAR measurements. This approach only works for wind speeds below rated power and in case no control takes place
3. Gradient boosting model that estimates the wind speed based on the rotor speed and pitch angle of the wind turbine based on a training set of LiDAR measurements. By incorporating these two measurements, the estimation is able to be performed across the entire range of wind speeds

The first approach directly tries to compensate the anemometer measurement, whereas the other two approaches try to establish an independent wind speed measurement based on variables which are directly measured by the wind turbine (rotor speed, pitch angle, active power). Within Beforecast, measurement campaigns have been organized at two different locations. Within cSBO SeaFD, another nacelle LiDAR has been used to perform a measurement campaign at two turbines within the same farm. In the following sections, the results of the first and third method are further explored, as these both are valid in the entire range of wind speeds.

3.2.1 Linear model (scada wind speed – LiDAR reconstructed wind speed)

Figure 6 shows a direct comparison between the scada and LiDAR measurements. Ideally, these two are highly correlated and have, on average, identical behavior. This is the case for turbine 4 (Red), where the fitted curve ($y = 1.01x - 0.09$) closely resembles the idealized expected curve ($y=x$) with high correlation ($R^2 = 0.95$). For the other three turbines, larger deviations are present. As the first dataset (blue) is of limited size, it is opted to mainly focus on the second and third dataset (grey, orange). These are measured at two identical turbines, which can be noted by the similar fit. Both these datasets have an overestimation of the scada wind speed at low wind speeds (below cut-in, <5m/s) and an underestimation at high wind speeds (above rated speed, >12.5m/s). Especially the former behavior cannot easily be compensated for easily using only the wind speed measurements, as there is very little correlation between the two data sources between 2 and 5m/s. Although this might not have the biggest impact on operational strategies, as the turbines are not operational at these low wind speeds, it still shows the importance of properly validating the behavior of the transfer function.

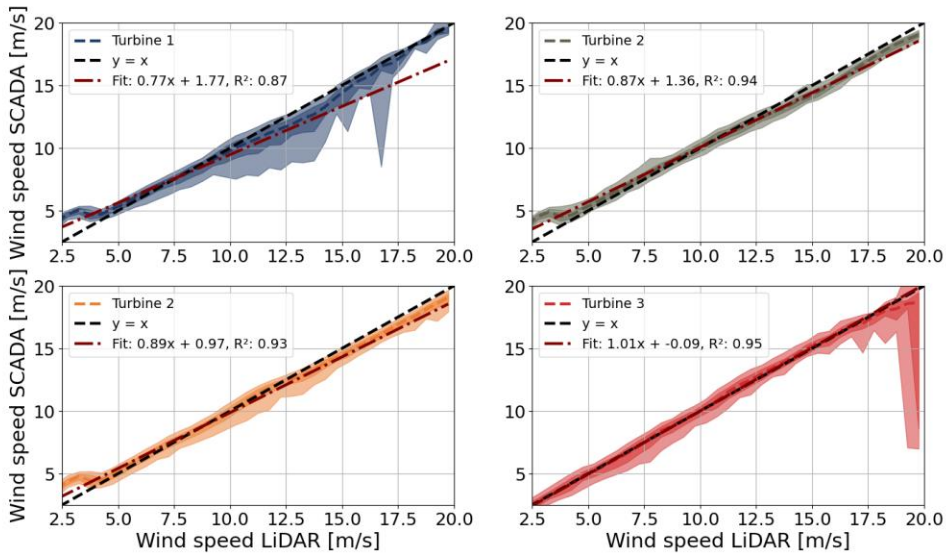

Figure 6: Comparison between the scada and LiDAR wind speed measurements for the four different datasets available in the project. The dark and light shaded areas respectively correspond with the P25-P75 and P10-90 uncertainty bounds.

3.2.2 Gradient boosting model (scada rotor speed and pitch angle – LiDAR reconstructed wind speed)

In order to establish the gradient boosting model, an initial dataset is first established to train the model. This dataset contains the rotor speed and pitch angle of the turbine, and the corresponding wind speed measured by the LiDAR. This model is then applied to a second dataset in order to validate the capability of the model on new data. This is shown in Figure 7. It can be noted that high correlations ($R^2 > 0.9$) are obtained for all turbines, and that they closely resembles the expected curve (y=x).

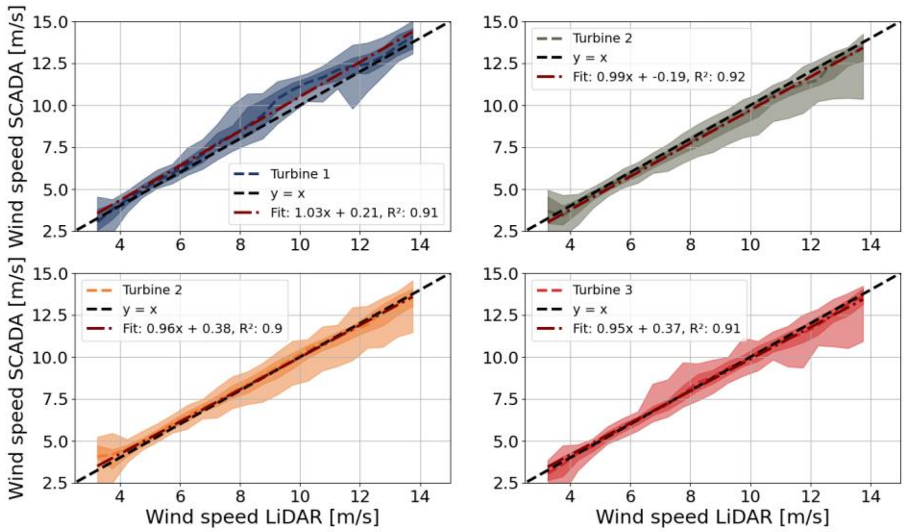

Figure 7: Comparison between the scada wind speed measurements and the ones predicted by the gradient boosting model based on the rotor speed and pitch angle measurements

Up until now, each model has been trained using data from its own measurement campaign. Ideally, the insights from one measurement campaign can nevertheless be transposed to other wind turbines within the same wind farm to obtain a (corrected) wind speed estimate using readily available parameters for all turbines. A proof-of-concept is done for this in the following manner:

1. Measurement data is gathered from turbine 2 (grey) and 3 (orange).
2. The data of turbine 2 is used to train the model (as previously explained and shown in Figure 6)
3. The model of turbine 2 is applied to the scada data of turbine 3, yielding a wind speed estimate.
4. The wind speed estimate of step (3) is compared with the LiDAR observation from the measurement campaign at turbine 3.

Using this approach, high correlations are still obtained ($R^2 > 0.85$), albeit lower than when using the measurement at the level of the turbine itself. This nevertheless highlights the potential of a single measurement campaign to gain insights into the regions where the anemometer is less reliable, and to calibrate an independent model based on scada parameters which are known with more certainty.

4 Conclusions

This document has given a summary overview on different methods to calibrate (a) reanalysis data to scada data, for different wind farms in the Belgian offshore zone (b) scada data to LiDAR data based on data collection during short-term measurement campaigns. This document especially focused on showcasing the need for calibration, and some outcomes highlighting the capability for calibrating one dataset based on another using different algorithms. In the attached references, more technical details can be found regarding the developed algorithms.

5 References and dissemination

Publications

[1] Marini, Rebeca, et al. “An assessment of the metocean conditions in the Belgian offshore zone” *Journal of Physics: Conference Series (2025)*

[2] Ally, Stijn, et al. "Modular deep learning approach for wind farm power forecasting and wake loss prediction." *Wind Energy Science Discussions* 2024 (2024): 1-33.

Conferences

Marini, Rebeca " An assessment of the metocean conditions in the Belgian offshore zone." *EERA Deepwind conference (January 15th – 17th 2025), Trondheim, Norway*