Ocean Systems Engineering, *Vol. 15, No. 3 (2025) 345-360*
https://doi.org/10.12989/ose.2025.15.3.345

ML-based sensor reduction for condition monitoring of wind turbine pitch bearings

Brian Rudie Jr. and Moohyun Kim*

Department of Ocean Engineering, Texas A&M University, College Station, TX, 77840, USA

(Received July 7, 2025, Revised September 8, 2025, Accepted September 12, 2025)

Abstract. Wind turbines often exhibit component failures before completion of their typical 20-year design life. Unexpected part failures increase the associated costs of their Operations & Maintenance (O&M). In turn, this raises the associated Levelized Cost of Electricity (LCOE), making this method of power generation less competitive compared to traditional methods. One of the components of particular concern is the slew or "pitch" bearing connecting the root of the blades to the rotor hub. The Drivetrain Reliability Collaborative (DRC) of the National Renewable Energy Lab (NREL) has begun investigations into pitch bearing reliability. One outcome of this is a collection campaign on the 1.5MW Wind Turbine at the NREL Flatirons Campus to observe variations in pitch bearing strains during real operation. This entailed outfitting the turbine with additional instrumentation such as strain gauges in the rotor hub. The present study intends to extend the applicability of the DRC1.5 field tests by relating the strain signals to standard operational output. Machine Learning (ML) techniques include supervised learning by Artificial Neural Networks (ANN) and Long-Short-Term Memory (LSTM), as well as Principal Component Analysis (PCA). The same DRC test data sets were applied to ANN and LSTM and their results are compared. Discussions of results describe which generalize best for the purpose of sensor reduction, and which of the operational signals are most indicative of bearing strain. The results showed that both ANN and LSTM predicted future (or nonfunctional) sensor signals well with slightly higher accuracy by LSTM. Post processing of time series predictions can then track the progression of fatigue damage without additional sensors. Examining the prediction results details the model performance and highlights the relevance for wind turbine condition monitoring. Incorporation of learning techniques is presented as a systematic approach that can be replicated to simplify and optimize real monitoring strategies.

Keywords: applied ML; condition monitoring; pitch bearings; sensor reduction; wind turbines

1. Introduction

In the midst of a global energy transition, wind power offers a viable alternative to responsibly meet a substantial portion of the electrical needs in the modern world. Constructing wind turbines offshore offers the added benefit of stronger, steadier, unobstructed wind flows at speeds that are not as prevalent on land. Investment into offshore wind also creates new jobs in installation, maintenance, and specialized shipbuilding.

However, the opportunities are not without trade-offs, and these structures cannot be scaled up

*Corresponding author, Professor, E-mail: m-kim3@tamu.edu

http://www.techno-press.org/?journal=ose&subpage=7 ISSN: 2093-6702 (Print), 2093-677X (Online)

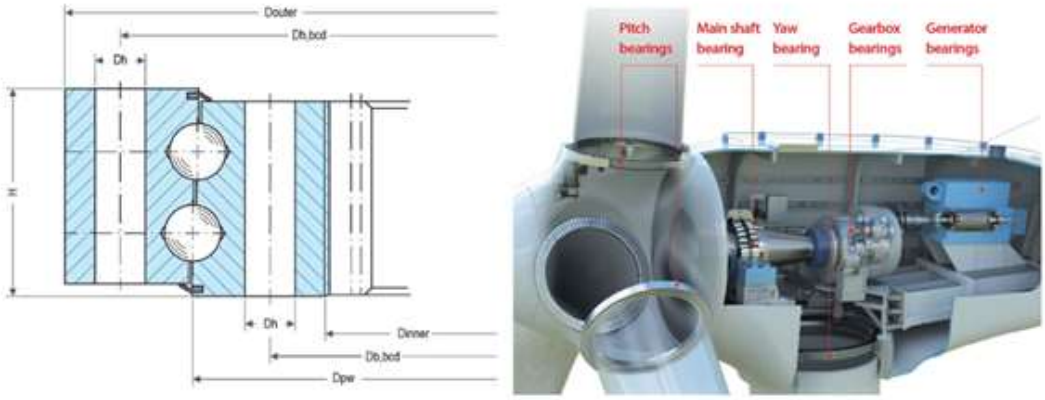

Fig. 1 Pitch Bearing Cross Section and Location Within Turbine, from (Rezaei *et al*. 2013)

without limit. The size and flexibility of modern wind turbines have allowed them to exceed the limits of some assumptions in established modeling tools (Veers *et al.* 2022). Wind turbines have a typical design life of 20 years but incur cumbersome maintenance costs as drivetrain components have exhibited failure earlier than designed for. One component of particular concern is the pitch bearing that connects the root of the turbine blades to the rotor hub (Fig. 1).

Industry surveys report failure rates as high as 12% within the 20-year design life, which is expected to increase with larger blade diameters. When factoring in variables such as component price, labor, crane time, and operational losses, replacing a pitch bearing is estimated to cost between $200,000 and $350,000 (Hornemann *et al.* 2019). In most cases, drivetrain components are designed well within existing standards but exhibit premature failures nonetheless (Keller *et al*. 2022). There is an increasing desire to better understand the mechanisms behind premature bearing failure as they are now, before it potentially becomes more prevalent in larger turbines installed in less accessible locations.

Condition Monitoring (CM) is a highly relevant approach for evaluating the health of the drivetrain components and supporting structure. The use of CM enables greater deployment of predictive maintenance schedules. Alternative maintenance strategies include predictive maintenance and the "run-til-failure" approach. Of these alternatives, the former may be redundant if done too frequently, and the latter incurs costly downtime. Condition monitoring entails using the signals of installed sensors to surmise the remaining useful life (RUL) in various components of the structure or machine. Problems of this nature lend themselves nicely to the recent advances and applications of Machine Learning (ML) (e.g., Jebari *et al.* 2025). More specifically, sensor reduction can be carried out using regression techniques, a subset of supervised machine learning. Such models allow for rapid processing of a multitude of sensor readings, relating these signals even when an analytic solution is computationally expensive or infeasible.

Wind turbine pitch bearings experience significant fatigue loads throughout their operational lifetime, leading to wear, degradation, and potential failure. Accurate prediction of fatigue life is essential for optimizing maintenance schedules, reducing downtime, and improving reliability. However, traditional condition monitoring and predictive maintenance methods rely on high-dimensional sensor data, which can be costly to process and challenging to interpret. Additionally, extracting meaningful insights from sensor signals remains difficult due to noise,

variability in operating conditions, and the complex relationship between sensor readings and fatigue progression.

Relating strains in these bearings to standard wind turbine sensors can boost predictive capabilities and supplement current monitoring techniques. There is a need for a systematic approach to reduce the dimensionality of sensor data while preserving key predictive features, ensuring that fatigue life estimates remain accurate and interpretable. This study seeks to address this challenge by investigating effective methods for sensor reduction and fatigue prediction that balance computational efficiency with predictive performance.

The present study examines the behavior of the General Electric 1.5-MW SLE wind turbine at the NREL Flatirons campus near Boulder, Colorado. Consistent with other literature that examines this turbine, it is hereafter referred to as the DOE1.5. Wind turbines rated at 250 kW or more, such as the DOE1.5, utilize rolling-element bearings (REBs) for their yaw and pitch motions (Harris *et al.* 2009). The pitch bearings examined currently are eight-point contact bearings, supplied by Kaydon with part number of 16230001 (Kaydon 2019).

The additional instrumentation and collection campaign are the result of a joint effort between NREL and ONYX Insight. The raw sensor data was shared with the present author for analysis, but the present author did not contribute to the instrumentation or collection of the data. To the knowledge of the author, the NREL is the first to make pitch bearing data of this nature publicly available. An example image of the installed strain gauges are included in Fig. 2 below.

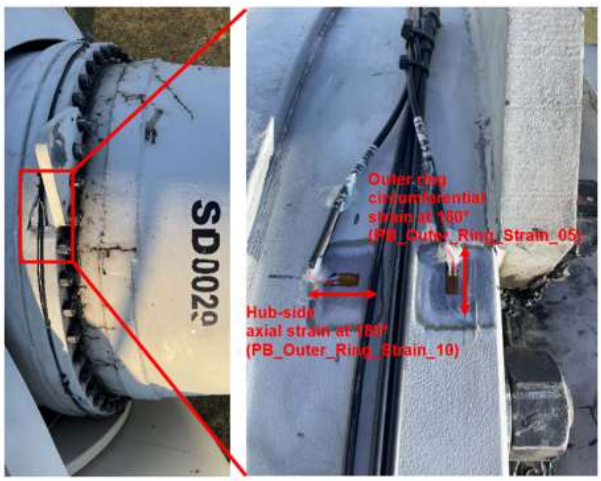

(a) Strain Gauges on Outer Bearing Ring

(b) Strain Gauges on Inner Bearing Ring

Fig. 2 Example of Installed Strain Gauges, Photos by Jesse Graeter ONYX Insight, (a) from NREL 85944, 85948 and (b) from NREL 85949

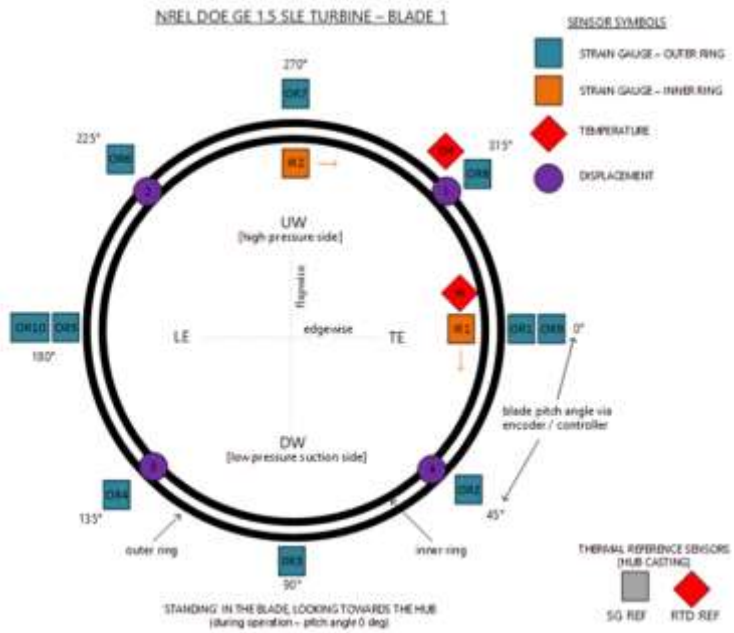

Fig. 3 Strain Gauge Placement Along Blade 1 Pitch Bearing, Reprinted from (Keller *et al*. 2024)

The installed sensors include 10 strain gauges on the stationary outer ring, 2 strain gauges on the rotating inner ring, and 1 for reference mounted directly on the hub casting. Almost all strain gauges were oriented to measure strain in the circumferential direction, except one that measures axial strain. An encoder kept track of the blade pitch. One resistance thermometer was placed on the inner and outer ring each, as well as another on the hub casting for reference. The arrangement of the strain gauges about the pitch bearing is shown in Fig. 3.

In recent years, applied machine learning and model-based processing have seen simultaneous growth in academic interest and popularity. Their combined application has shown promise for intelligent monitoring of wind turbine operation. Using SCADA data, Sheng et al. employ PCA and clustering techniques to identify gearbox fault indicators from SCADA data (Sheng *et al*. 2011). Similarly, Fu et al. performed feature extraction by CNN as input to an LSTM model to estimate gearbox bearing temperature using SCADA data (Fu *et al*. 2019). Hlaing *et al*. (2024) introduce the use of Bayesian Neural Network for structural health monitoring via surrogate modeling, observing strain gauge outputs of only the “fleet leader” turbine (Hlaing *et al*. 2023).

2. Methodology

Within the data, each signal is a time series whose magnitude, units, and measurement type are already associated with their respective sensors. Training a model on labeled data is a subset of machine learning known as supervised learning. More specifically, using input measurements to model the continuous outputs of the strain gauges is a multi-valued regression problem. This kind of learning problem is distinct from time-series forecasting, as predictions are only made at the current timestamp using inputs at or before the current timestamp. For the present study, this may

be applied to relate pitch bearing strains to standard operational measurements.

A regression problem such as this is a subset of supervised learning, for which several architectures may be used. The standard Artificial Neural Network (ANN) is a broadly applicable architecture that can make fast predictions once trained (e.g., Kwon *et al.* 2023, 2025). By default, feedforward architectures such as the ANN do not make use of timing as context in predictions. The Long-Short Term Memory (LSTM) is a modified Recurrent Neural Network (RNN) with improved capabilities for noticing time-dependent contexts. For the purpose of making predictions from time series data, with the potential for gaps or uneven sampling, these model types are expected to perform well. If the signals contain trends or nonstationary statistics, the LSTM will be able to more completely describe the behavior.

Model based processing is a technique to simulate system behavior by relating inputs and outputs of interest. A complex time dependent process, $G(t)$, can be emulated by a model, $\hat{G}$, that will give the same outputs, $\hat{y}(t)$, from the same combination of inputs, $u(t)$ (Garlick *et al.* 2009). In general, this may be applied to estimate all outputs using only inputs. A slight modification to this approach can be used to relate some process outputs to others in the case that all aren't known. The learning problem of the presented study is formulated under the assumption that several of the process outputs will be provided, those being standard sensors in normal turbine operation. In doing so, it is assumed that a reduced set of sensors still sufficiently capture the factors that influence the output of the dependent variables (Murrell *et al.* 2018). All of the sensor data examined is an output of the wind turbine operation but may be further subdivided into inputs and outputs of a process model, $\hat{G}$, based on which are less commonly installed and of current concern (Fig. 4).

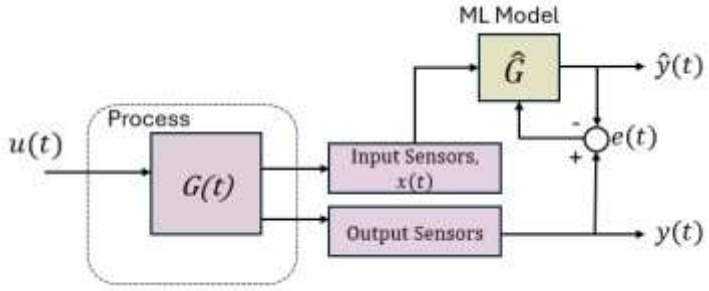

(a) When additional sensors are installed

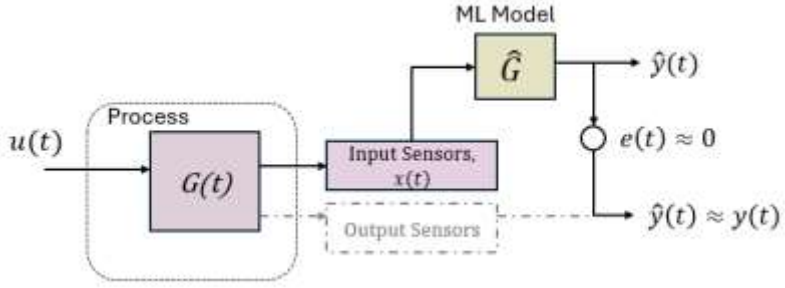

(b) During normal operation

Fig. 4 Visual Intuition for the (a) training and (b) predicting phases of ML-based sensor reduction

All computational analysis was carried out in Python. Libraries such as Scikit-learn (Chollet *et al*. 2018) and TensorFlow (Abadi *et al*. 2016) were used to build the ANN and LSTM. The NumPy, Pandas, and npTDMS libraries were most helpful when reading in and preprocessing the sensor data. The fatpack repository by Gunnstein was later used for the cycle counting and fatigue calculations (Norheim 2023).

3. Preprocessing

The preprocessing phase includes validating sensor output, calibrating to engineering units, and scaling the values prior to training the model. This step covers the cleaning and trimming of raw sensor data, as well as the conversion between file types. The .TDMS files from the DAQ store the waveforms in channels. Moreover, these channels have properties such as a scale and offset. In Python, these channels are extracted with the corresponding properties and stored as columns of a Pandas dataframe object. During preprocessing, it is important to split the data before scaling. Otherwise, the scaling of the test set will influence the training data, which may cause model performance to appear unrealistically high. After splitting, applying the same scaling to both training and testing data ensures consistency in prediction range, and prevents data leakage. At this splitting step, both the ANN and LSTM utilize 80% of these data points for training and the remaining 20% for testing. As the data is comprised of time series signals, they are not shuffled in training.

The initial dataset consists of 125 signals over twenty-four days between June 25 and December 13 of 2024. Naturally, not all signals will be linearly independent, and not all are relevant to the present study. The trimming of the signals was primarily done by visual inspection with the present learning objective in mind. The only instances of interest are when the target sensors produce meaningful output. After July 6th, the strain gauge signals were frozen, meaning the sensors were either broken or not installed for the remainder of the collection period. Therefore, analysis focuses on those between June 25th and July 6th. The remaining section was scanned to verify there were no gaps, NaNs or other missing values. At a sampling frequency of 1 Hz, there are 1,047,391 valid data points in this range.

The reduced set is refined for a specific learning problem – which in this case is to monitor accumulation of bearing damage. For this reason, sensors with little variance are removed, as they do not contribute to the learning. A summary of all signals used for predictions are listed in Table 1. These measurements comprise the operational status, control signals, drivetrain behavior, bending moments, wind, atmosphere, and the yaw and azimuth of the turbine.

Scaling and normalization is only done with a specific reason to preserve the signal information. A low pass filter of 10 Hz is applied to remove high frequency noise. Spike detection and removal is carried out with a simple padded Hampel filter. The data is not altered further for training purposes. A standard scaler was applied so that all signals have a zeroed mean and unity variance (Eq. (1)). This does not enforce an upper or lower limit on the range, but results in most readings falling between -3 and 3. Normalization ensures that no measurement has a greater relative influence due to the sheer magnitude of their signal.

$$X_{scaled} = \frac{X-\mu}{\sigma} \tag{1}$$

Table 1 Reduced Set of Input Signals Used to Predict Target Output Signals

Input Signals ("Features")	
Hum1	AzimuthAO
Temp1	HS Pinion Trq
Hum2	V_link_batt_voltage
Temp2	Blade_1_Flap
Windspeed_38m	Blade_1_Edge
Windspeed_55m	Blade_2_Flap
WD1_87m	Blade_2_Edge
Air_Press_2	Blade_3_Flap
Air_Press_1	Blade_3_Edge
Wind_Direction_38m	Pitch_Blade_2
WindSpeed_80m	Rotor_Azimuth_Pitch
Precipitation	Rotor_Azimuth_Roll
Windspeed_87m	PB_Displacement_Axial_1
Tower Base Bend 1	PB_Displacement_Axial_2
Tower Base Bend 2	PB_Displacement_Axial_3
Tower Base Torque	PB_Displacement_Axial_4
ActivePower	PB_Surface_Temperature_OR
ReactivePower	PB_Surface_Temperature_IR
PowerFactor	PB_Thermal_Reference_Temp
LaserDistMinimum	Yaw_Encoder
Windspeed_Acoustic_10m	Pitch_Blade1
TowerTopBending_0	EDAS_Health_Out
TowerTopBending_90	Power_Out
TowerTopTorque	Pitch_Out
TowerTopDCAccel_NS	RPM_Out
TowerTopDCAccel_EW	WS_Nac_Out
HSS RPM	WS_Prim_Out
NacelleWindSpeed	WS_10m_Out
Mainshaft_Downwind_Bend_0	Yaw_Posn_Out
Mainshaft_Downwind_Bend_90	Brake_Out
Mainshaft_Downwind_Torque	WD_Mod_Active

Output Signals ("Labels")
PB_Outer_Ring_Strain_01
PB_Outer_Ring_Strain_02
PB_Outer_Ring_Strain_03
PB_Outer_Ring_Strain_04
PB_Outer_Ring_Strain_05
PB_Outer_Ring_Strain_06
PB_Outer_Ring_Strain_07
PB_Outer_Ring_Strain_08
PB_Outer_Ring_Strain_09
PB_Outer_Ring_Strain_10
PB_Inner_Ring_Strain_01
PB_Inner_Ring_Strain_02
PB_Thermal_Reference_Strain

ML applications are not limited to high dimensional regression problems. Another capability is dimensionality reduction, as is done in Principal Component Analysis (PCA). PCA leverages the fact that there is often covariance between the columns of data. Covariance is the amount of information that any j^{th} column tells about the k^{th} column in matrix $A = |a_{m,n}|$ (Eq. (2)).

$$cov(col_j, col_k) = \frac{1}{N}\sum_{i-1}^{N} a_{ij}a_{ik} = \frac{1}{N} col_j^T col_k \tag{2}$$

PCA is an attempt to map d-dimensional data onto k-dimensional space, where $d > k$. This is carried out through Singular Value Decomposition (SVD) and a truncation of combined terms. For any matrix, even if not symmetrical, left and right singular matrices U and V can be determined to write the matrix as follows (Eq. (3))

$$A = U\Sigma V^T = \sigma_1 u_1 v_1^T + \sigma_2 u_2 v_2^T + \cdots + \sigma_N u_N v_N^T \tag{3}$$

The result is a linear combination of inputs that is able to adequately describe the variations of the dependent variables. These linear combinations comprise rotations that 'point' in the directions of highest data variation. The accuracy of the approximation is controlled by the number of terms retained. Examining a cumulative sum of explained variance, we see that over 95% of the variance in the dataset can be explained using 25 principal components (Fig. 5). Of course, these would each be linear combinations of the original 62 inputs, but the reduction indicates that many measurements observe similar driving mechanisms.

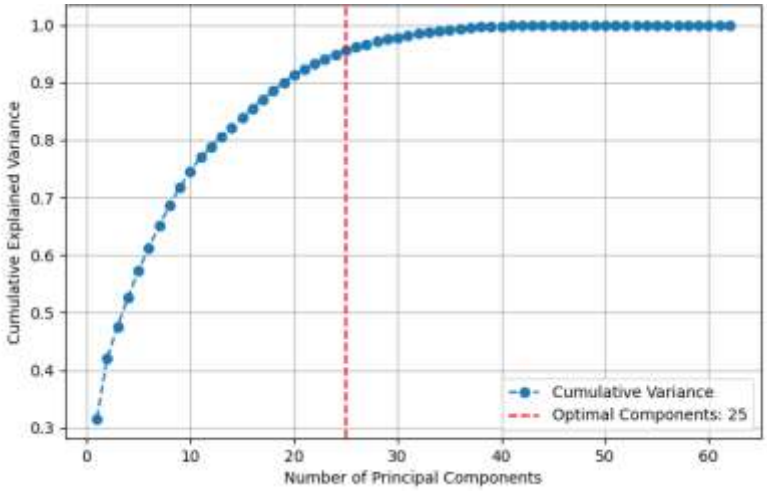

Fig. 5 Number of Principal Components Versus Explained Variance

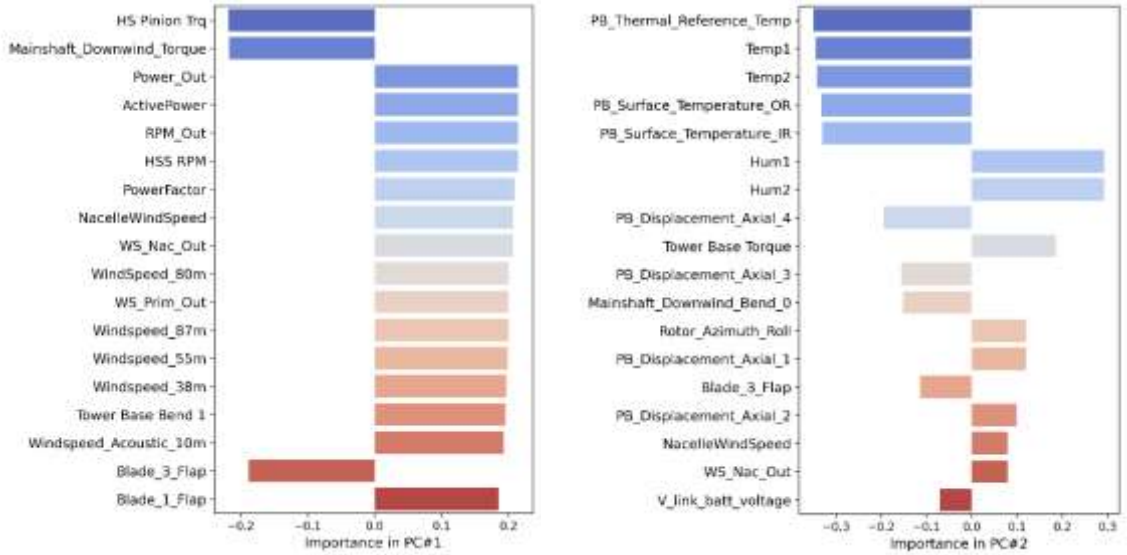

(a) (b)

Fig. 6 Bar Plots for the Highest Feature Importance Towards the (a) First and (b) Second Principal Components

Moreover, PCA provides a breakdown of how influential each signal is in the linear combinations. Such breakdowns for the first two principal components are provided in Fig. 6. These are ordered such that the first component points in the direction of most variation, the second components to the direction with the second-most variation, and so on.

Table 2 Working Learner Configurations

Model		**Dimensions**				**Hyperparameters**	
ANN	**Layer**	Input	Hidden	Hidden	Output	**Activation**	Tanh
	(Shape)	62	64	32	13	**Epochs**	20
LSTM	**Layer**	Input	LSTM	Dense		**Activation**	Linear
	(Shape)	(10, 62)	100	13		**Epochs**	10

4. Model validation

Several combinations of architecture and hyperparameters were applied for both ANN and LSTM. Those chosen for demonstration and further calculations were selected for acceptable accuracy metrics but also reasonable training times. The number of training epochs and model complexity were controlled by observing the training and validation loss. Higher accuracies are likely achievable with larger models and additional training epochs but begin to run the risk of overfitting to the dataset. For both models, the 'Adam' optimizer is used, and loss is defined by the Mean Squared Error (MSE). For this learning problem, some reliable parameter combinations are provided below for each (Table 2). Note that the extra dimension in shape at the input layer of the LSTM for lookback.

The following pages summarize the predictions for the twelve strain gauges over the entire testing set when made by the ANN and LSTM (Figs. 7 and 8, respectively). Without physical interpretation, this section is solely intended to demonstrate how well the models can estimate sensor signals as if they were not installed. To this effect, it is worth pointing out that the units of the vertical axis have been left as the voltage ratios of the quarter-bridge strain gauges, $[V/V]$. Similar predictions for reference strain gauge are also made and stored separately but excluded from the plots for conciseness.

Fig. 7 shows the last 20% of each strain gauge time series saved for testing. The ANN does not consider temporal context, and predicts each signal above based only on the inputs at that same timestamp. Nonetheless, the red appears to generally track well with the green. Similarly, calling the LSTM to predict over the same test period gives the following (Fig. 8).

The predictive accuracy of the models is quantified using Mean Absolute Error (MAE), Mean-Squared Error (MSE), Root-Mean-Squared Error ($RMSE$), and Coefficient of Determination (R^2). Each metric has its merits in describing predictive accuracy. For instance, the RMSE is in the original units of the strain gauge signals, and the R^2 offers a comparable score between 0.0 and 1.0. The test accuracy corresponding to the predictions shown in Figs. 7 and 8 are given in Table 3.

Table 3 Predictive Accuracy of the ANN versus LSTM

Model Type	MAE	MSE	$RMSE$	R^2
ANN	1.16×10^{-5}	2.75×10^{-10}	1.66×10^{-5}	0.9142
LSTM	1.18×10^{-5}	2.24×10^{-10}	1.59×10^{-5}	0.9241

Looking at everything at once, it can be observed that the ANN is able to predict the magnitudes at an appreciable accuracy. In light of the problem context, the strains on the outer ring near the trailing edge of the turbine blade pose the highest risk for crack initiation and propagation (Schleich and Brand 2025). A closer inspection of the waveforms at outer ring position #7 shows that the ANN sufficiently captured the oscillations as well as the magnitude (Fig. 9).

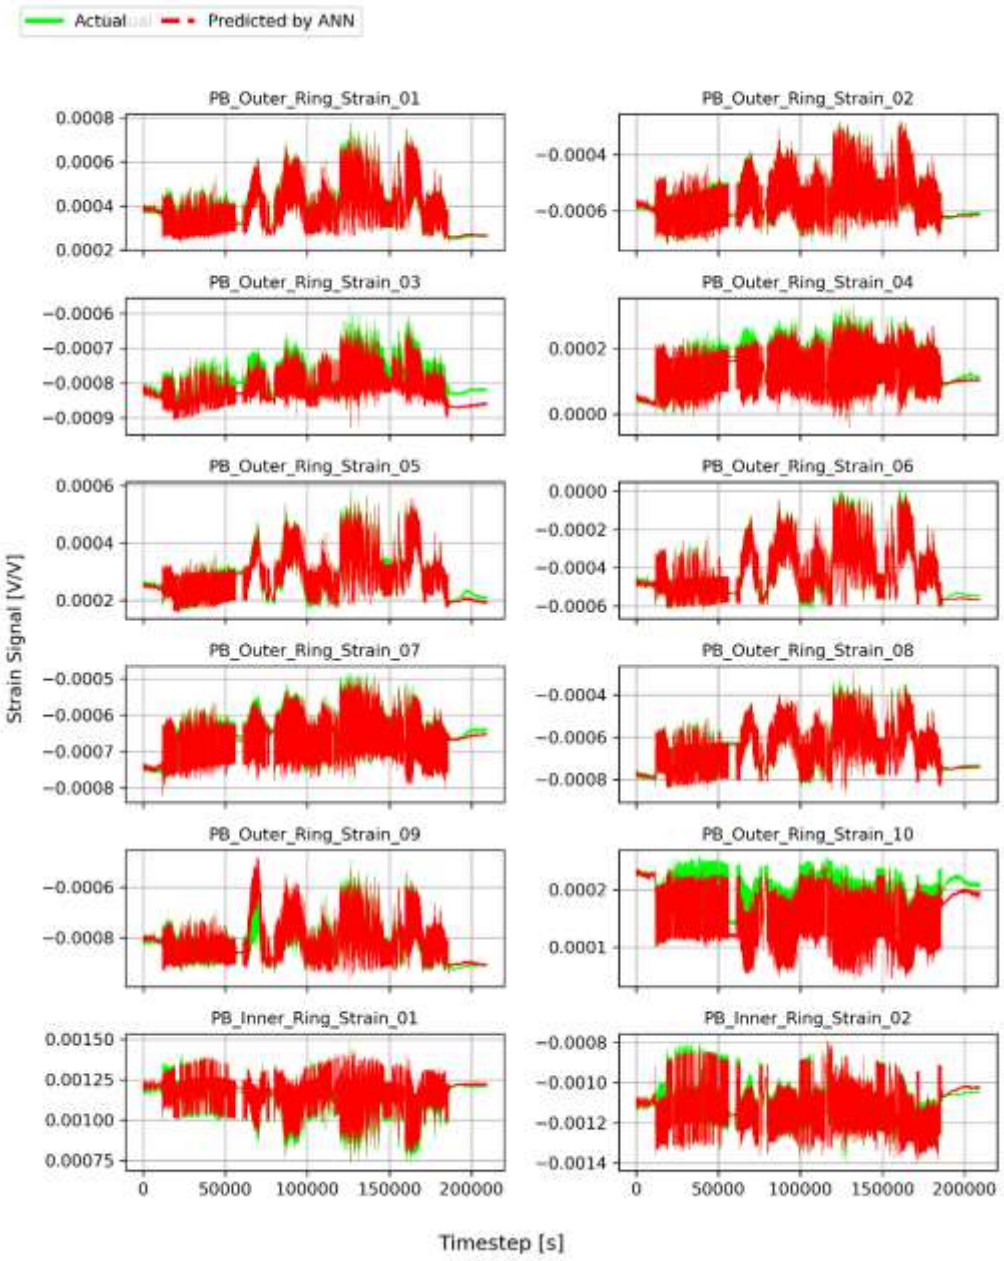

Fig. 7 Time Series Predictions of Each Strain Gauge by the ANN Over Entire Test Set

From Figs. 9(a) and 9(b), both model types appear to match the real strain behavior, even slightly less so by the LSTM at this instant. Regardless, the LSTM remains slightly more accurate over the entire testing range. The relatively small improvement between ANN and LSTM is a surprising result – it shows the general strength of the ANN in multi-variable regression problems. This may also indicate that there is either no observable trending in the data, or that the LSTM lookback was too short for this to take effect.

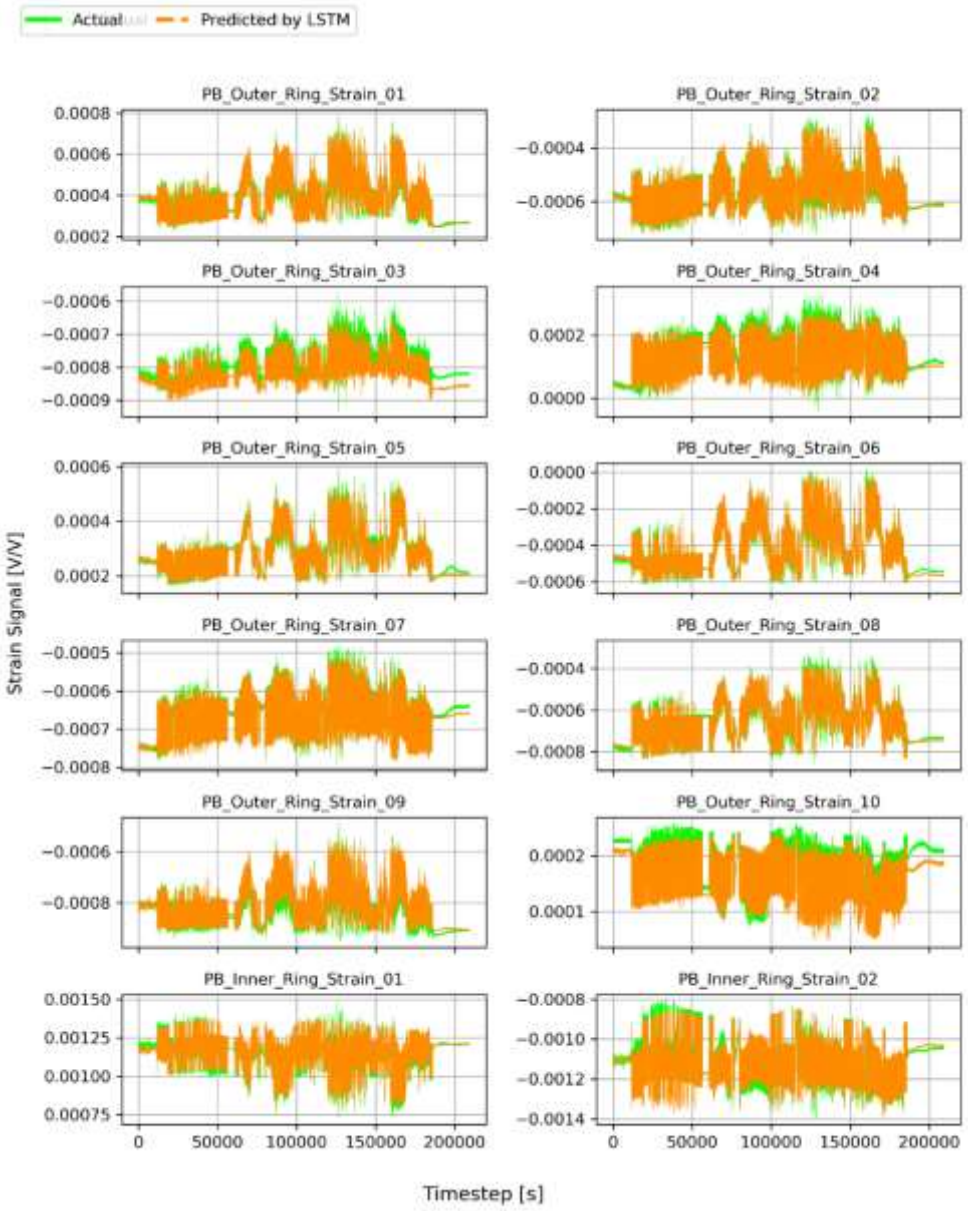

Fig. 8 Time Series Predictions of Each Strain Gauge by the LSTM Over Entire Test Set

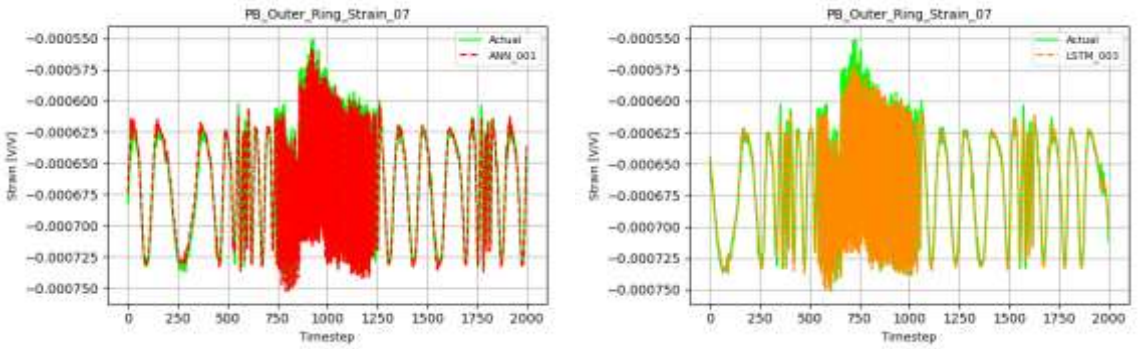

(a) ANN Predictions for Strain Gauge 07 (b) LSTM Predictions for Strain Gauge 07

Fig. 9 Close Up of Predictions Showing ML Pattern Matching Made by the (a) ANN Model and (b) LSTM Model

Table 4 Effect of Further Lookback on LSTM Predictive Accuracy

Lookback	MAE	$RMSE$	R^2
10	1.18×10^{-5}	1.59×10^{-5}	0.9241
100	1.12×10^{-5}	1.5×10^{-5}	0.9346

The utility of the LSTM is not as apparent both cases but may still exist in the second. It is then worth testing if a larger sliding window is helpful for surrogate modeling and condition monitoring. A larger lookback simply entails reshaping of the input data in training but quickly increases computational load. Strictly speaking, the lookback is the amount of prior timesteps the LSTM considers for its current prediction. This does not necessarily mean the accuracy will scale with memory. Nonetheless, training the same LSTM architecture with a lookback of 100 instead of 10 does slightly improve prediction accuracy (Table 4).

5. Results

Having established confidence in the predictive capabilities, the reason for doing so in the first place would be to then extend predictions to a time when there is no labeled data. The provided GE1.5 dataset contained a large portion in which the strain gauges were either not installed or did not provide meaningful output. This portion was left out in training but can now demonstrate a realistic use case.

With no real measurements for feedback, predictive accuracy cannot be calculated like before. A new LSTM is trained with the same data as earlier, but with inputs further reduced to those available in this later chunk of the data. Examples of predictions at a time when strain gauges were not available are shown in Fig. 10.

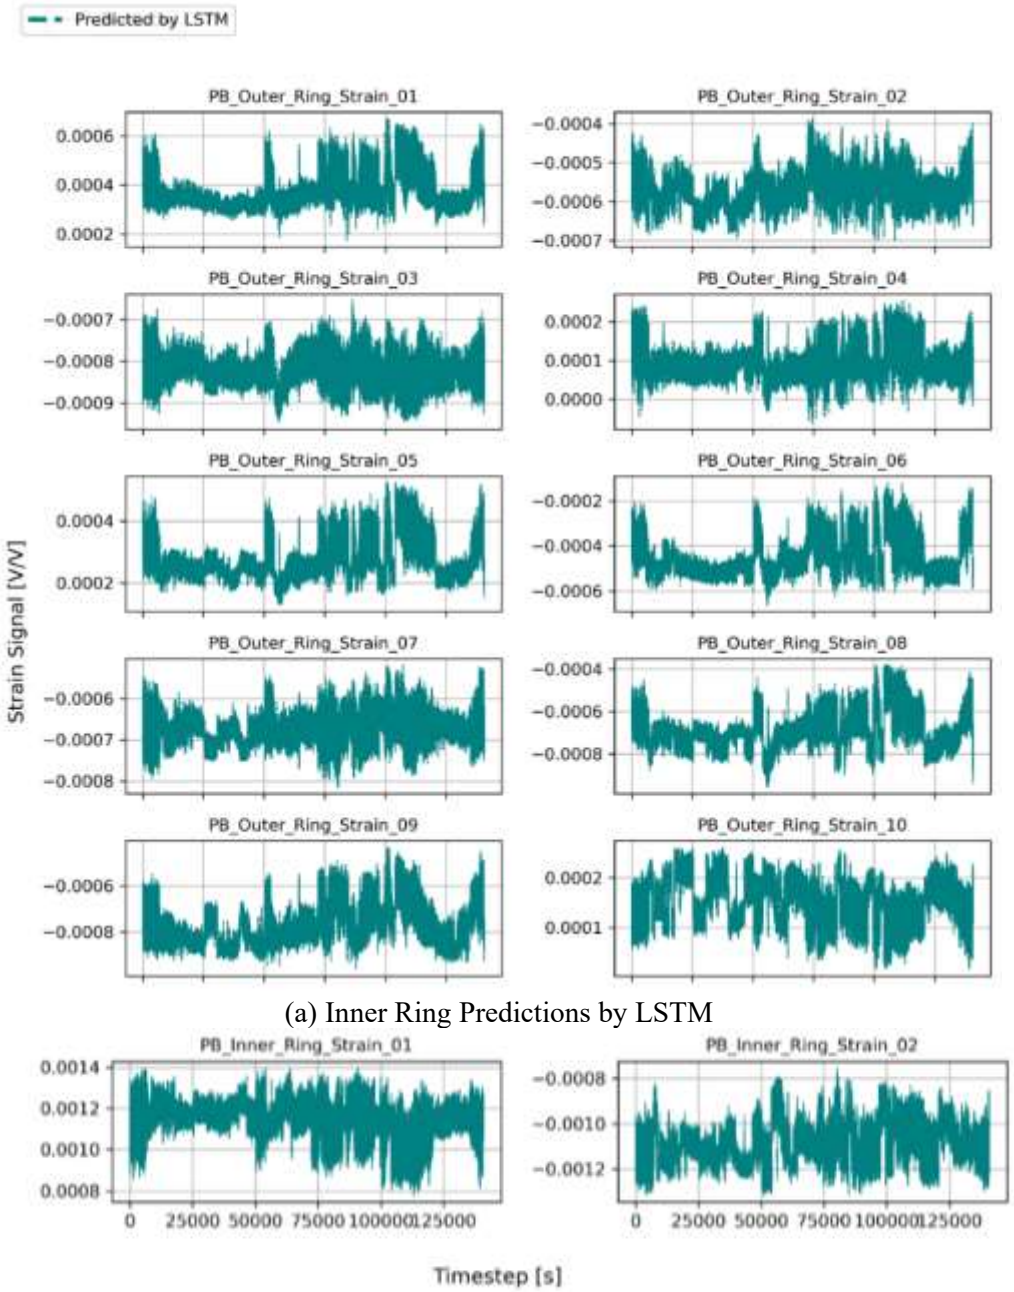

(a) Inner Ring Predictions by LSTM

(b) Outer Ring Predictions by LSTM

Fig. 10 Time Series Predictions by LSTM on the (a) Inner and (b) Outer Rings When Real Strain Gauges Were Unavailable

6. Conclusions

The purpose of this study was to demonstrate the use of supervised learning for sensor reduction and condition monitoring. This is a promising use case for machine learning in the engineering context. These results exemplify a specific scenario in which machine learning can be used to predict strain gauge signals using the output of other turbine signals. From here, the time series are to be post processed to track the progression of fatigue damage. This would most likely be done by cycle counting to identify hysteresis loops in the strain fluctuations. Its utility is demonstrated here in wind turbine pitch bearings, but the approach could offer similar insight in a variety of engineering disciplines. In this case, intelligent condition monitoring serves to reduce maintenance costs related to wind power.

A machine learning framework was introduced to use prior observation in order to predict the current strain level within wind turbine pitch bearings once previously installed sensors are either faulty or removed. The framework was applied using both ANNs and LSTMs for the learner. Principal Component Analysis was employed as a sensor selection methodology tailored to wind turbine operational loads and bearing fatigue. The existence of determinate or causal relationships between measurements were explored, and particular input sensitivities quantified. The learning mechanisms of ML techniques are critically examined, discussing the extent of their utility in the context of surrogate modeling. The same test-data sets were applied to ANN and LSTM and their results showed that both ANN and LSTM predicted future (or nonfunctional) sensor signals well with slightly higher accuracy by LSTM. The LSTM prediction accuracy can further be enhanced by choosing optimal lookback size.

Many incredible things are possible in the modern day due to powerful super-computers and techniques such as these, so long as the underlying mechanisms and the nature of the results are well understood. The LSTM is a powerful architecture, but the accuracy of the ANN speaks to its capabilities for short term prediction. Before all else, the questions that data driven analysis can answer depend on the nature of the data available.

Acknowledgements

The authors gratefully acknowledge Dr. Jon Keller and Dr. Shawn Sheng of the National Renewable Energy Laboratory (NREL) for their generous support in providing access to the sensor data used in this study and for their valuable feedback on intermediate results. Their insights contributed significantly to the development of this work.

References

Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., Kudlur, M., Levenberg, J., Monga, R., Moore, S., Murray, D.G., Steiner, B., Tucker, P., Vasudevan, V., Warden, P., Wicke, M., Yu, Y. and Zheng, X. (2016), "{TensorFlow}: a system for {Large-Scale} machine learning," in 12th USENIX symposium on operating systems design and implementation (OSDI 16), 265-283.

Chollet, F. (2018), "Keras: The python deep learning library", Astrophysics Source Code Library, record ascl:1806.022, June.

Florian Schleich, M.B. (2025), "Investigation of cracked blade bearing rings", Presented at the 2025

Drivetrain Reliability Collaborative.
Fu, J., Chu, J., Guo, P. and Chen, Z. (2019), "Condition monitoring of wind turbine gearbox bearing based on deep learning model", *IEEE Access*, **7**, 57078-57087.
Garlick, W.G., Dixon, R. and Watson, S. (2009), "A model-based approach to wind turbine condition monitoring using SCADA data".
Harris, T., Rumbarger, J. and Butterfield, C.P. (2009), "Wind turbine design guideline DG03: Yaw and pitch rolling bearing life", tech. rep., National Renewable Energy Lab. (NREL), Golden, CO (United States).
Hlaing, N., Morato, P.G., Santos, F.D.N., Weijtjens, W., Devriendt, C. and Rigo, P. (2024), "Farm-wide virtual load monitoring for offshore wind structures via bayesian neural networks", *Struct. Health Monit*., **23**(3), 1641-1663. https://doi.org/10.1177/14759217231186048.
Hornemann, M. (2019), "Top failure observations from the field", Presented at the Drivetrain Reliability Collaborative.
Jebari, O.1., Kwon, D.S., Kim, S.J., Jin, C. and Kim, M.H. (2025), "Machine learning-based mooring failure detection for FPSOs: A two-step ANN approach", *J. Mar. Sci. Eng*., **13**(4), 791. https://doi.org/10.3390/jmse13040791.
Kaydon Bearings, Kaydon Bearing Solutions for Slewing Ring Bearings, Catalog 390. Kaydon Corporation, Muskegon, MI, 2019.
Keller, J. and Guo, Y. (2022), "Rating of a pitch bearing for a 1.5-megawatt wind turbine", tech. rep., National Renewable Energy Laboratory (NREL), Golden, CO (United States).
Keller, J., Graeter, J., Roadman, J. and Skinner, M. (2024), "Instrumentation for the investigation of pitch bearing design and reliability", tech. rep., National Renewable Energy Laboratory (NREL), Golden, CO (United States).
Kim, K., Parthasarathy, G., Uluyol, O., Foslien, W., Sheng, S. and Fleming, P. (2011), "Use of SCADA data for failure detection in wind turbines", *Energy Sustainability*, **54686**, 2071-2079. https://doi.org/10.1115/ES2011-54243.
Kwon, D.S., Kim, S.J., Jin C.K., Kim, M.H., Guha, A., Esenkov, O.E. and Ryu, S. (2023), "Inverse estimation of vertical current velocity profile using motions of an FPSO and artificial neural network", *Ocean Eng*., **285**, 115343. https://doi.org/10.1016/j.oceaneng.2023.115343.
Kwon, D.S., Kim, S.J., Jin, C.K. and Kim, M.H. (2025), "Inverse estimation of directional wave spectrum from FPSO motion sensor data using artificial neural networks", *J. Mar. Sci. Eng*., **13**(69).
Murrell, N., Bradley, R., Bajaj, N., Whitney, J.G. and Chiu, G.T.C. (2018), "A method for sensor reduction in a supervised machine learning classification system", *IEEE/ASME Trans. Mechatron*., **24**(1), 197-206. https://doi.org/10.1109/TMECH.2018.2881889.
Norheim, G. (2023), "fatpack: Example usage of python fatigue package", https://github.com/Gunnstein/fatpack/blob/master/example.py, 2023. Accessed: 202504-03.
Rezaei, A., Guo, Y., Keller, J. and Nejad, A.R. (2023), "Effects of wind field characteristics on pitch bearing reliability: a case study of 5 mw reference wind turbine at onshore and offshore sites", *Forsch. Ingenieurwes*, **87**(1), 321-338. https://doi.org/10.1007/s10010-023-00654-x.
Veers, P., Dykes, K., Basu, S., Bianchini, A., Clifton, A., Green, P., Holttinen, H., Kitzing, L., Kosovic, B., Lundquist, J.K., Meyers, J., O'Malley, M., Shaw, W.J. and Straw, B. (2022), "Grand challenges: wind energy research needs for a global energy transition", *Wind Energy Science Discussions*, **7**(6), 2491-2496. https://doi.org/10.5194/wes-7-2491-2022.

PL

Nomenclature

Abbreviation	Description
DAQ	Data Acquisition System
DRC	Drivetrain Reliability Collaborative
GE	General Electric
LCOE	Levelized Cost of Energy
LSTM	Long Short-Term Memory
ML	Machine Learning
NREL	National Renewable Energy Laboratory
O&M	Operations & Maintenance
PCA	Principal Component Analysis
REB	Rolling Element Bearing
RPM	Revolutions per Minute
RUL	Remaining Useful Life
SCADA	Supervisory Control and Data Acquisition
SVD	Singular Value Decomposition
TAMU	Texas A&M University
WT	Wind Turbine