A Graph Auto-encoder Framework for Spatio-temporal Anomaly Detection of Corrosion across a Fleet of Offshore Wind Turbines Using ICCP Data

Orkun Temel[1], Saeid Hedayatrasa[1], Joachim Verhelst[1], Bram De Baere[2] and Stefan Hendricx[1]

[1] *Flanders Make, Leuven, Belgium*
orkun.temel@flandersmake.be
saeid.hedayatrasa@flandersmake.be joachim.verhelst@flandersmake.be stefan.hendricx@flandersmake.be

[2] *Norther Offshore Wind Farm, Ostend, Belgium*
bram.debaere@norther.be

ABSTRACT

Offshore wind farms are exposed to severe marine conditions, which can lead to long-term structural integrity concerns due to corrosion-induced degradation processes. Here, we propose a spatio-temporal anomaly detection methodology using the Impressed Current Cathodic Protection (ICCP) data from an offshore wind farm. First, we employ a graph autoencoder (GAE) to infer the spatial variations in the measurements. We construct a graph based on the spatial proximity between wind turbines, where nodes and edges correspond to wind turbines and distance between turbines. Then, the latent representation of the measurements obtained by the GAE, are passed to a long-short term memory (LSTM) model, which infers the temporal evolution of measured signal and predict the next state. Finally, we perform anomaly detection using a combined scoring that includes graph reconstruction errors, latent prediction errors and observation-space prediction errors. Our results highlight the potential of integrating graph-based and sequence-based approaches for industry-relevant anomaly detection and demonstrate that the proposed methodology can identify turbines and corresponding time periods exhibiting deviations from fleet-level behavior.

1. INTRODUCTION

Offshore areas exhibit stronger and steadier wind resources compared to the onshore wind energy sites as a result of meteorological forcing mechanisms (Porchetta et al., 2021). Recent developments in fixed bottom offshore wind platforms and in the design of floating wind turbine platforms has extended the deployment range of wind energy systems in terms of water depth and geographical range (Santhakumar, Smart, Noonan, Meerman, & Faaij, 2022). Hence, offshore wind energy has become one of the key components of low-carbon, systems (Ren, Verma, Li, Teuwen, & Jiang, 2021). Moreover, offshore wind farms can also reduce some land related constraints such as land-use conflicts and noise pollution in comparison to onshore wind sites. Nevertheless, the geographical advantage of offshore wind farms can also lead to operational reliability and lifecycle performance related issues. Considering that offshore wind farms are subjected to limited accesibility and located in regions where weather conditions can affect logistics, it is essential to develop methodologies that can take their increased downtime sensitivity and maintenance costs (Ren et al., 2021; Martin, Lazakis, Barbouchi, & Johanning, 2016; Tremps, Yeter, & Kolios, 2024). Therefore, robust monitoring methodologies for the early detection of anomalous system behaviour is vital for offshore systems, where even minor degradation processes can lead to major implications for asset health management.

Corrosion remains as one of the most destructive and structurally consequential degradation mechanisms that can affect offshore wind turbine systems (Brijder et al., 2022; Price & Figueira, 2017). Monopile foundations, the most widely used foundation type, are particularly vulnerable when exposed to harsh marine environments characterized by variable salinity, oxygen levels, and cyclic wetting conditions. Moreover, the interaction between near-surface winds and waves can impose persistent cyclic stresses that exacerbate corrosive exposure and promote damage initiation via corrosion–fatigue mechanisms (Adedipe, Brennan, & Kolios, 2016; Okenyi, Bodaghi, Mansfield, Afazov, & Siegkas, 2022). Furthermore, localized surface phenomena such as pitting can also trigger fatigue crack initiation and substantially affect the remaining

---

useful life and reliability of offshore wind systems (Rokhlin, Kim, Nagy, & Zoofan, 1999; Shittu et al., 2020). Existing corrosion managemenent methodologies for offshore systems mainly rely on either corrosion allowance and periodic inspection and assessment processes (Price & Figueira, 2017; Erdogan & Swain, 2021; Kalovelonis, Gortsas, Tsinopoulos, & Polyzos, 2025). Despite that these approaches are well established, they are limited in temporal continuity and spatial representativeness. In contrast, cathodic protection systems (such as ICCP) can continuously regulate electrochemical protection conditions and provide operational time series (such as current, voltage and reference potential), which can be physically coupled to protection demand and environmental conditions. Hence, the exploitation of ICCP operational data with implications for asset health management can lay the foundations for state of the art methodologies.

Early detection of anomalous behaviour in corrosion related monitoring data is essential for state of the art asset health management tools for offshore systems. However, corrosion related anomaly detection is intrinsically both a spatial and temporal problem. The electrochemical response due to corrosion will be driven both by seasonal timescales affecting environmental conditions but also depends on the water depth (hence the spatial location) of the wind turbine. Moreover, understanding the population level behaviour can help understanding the turbine level anomalies. Whereas, purely turbine level analytics can mischaracterize behaviour that is only abnormal in a spatial context (Black, Yeter, Häckell, & Kolios, 2024). Similarly, purely spatial approaches may misinterpret temporal changes that indicate degradation or ICCP system faults (Erdogan & Swain, 2022). Hence, joint modeling of spatial (population/fleet level) and temporal (turbine level) dynamics is essential. For the spatial behaviour, graph based representations can provide a mathematical foundation to understand population level dependencies. On the other hand, recurrent neural networks such as LSTM can also provide a basis for modeling temporal dependence in the sequences of observed data (Hochreiter & Schmidhuber, 1997). There is also a recent interest in coupling graph-based approaches with sequential methods to capture complex spatio-temporal dependencies (Zamanzadeh Darban et al., 2025; Zheng, Ma, Chen, Zhang, et al., 2024). Hence, we also propose a coupled Graph Autoencoder and LSTM Autoencoder framework for the detection of corrosion-relevant anomalies using ICCP operational data.

Existing anomaly detection approaches for offshore wind monitoring systems have traditionally focused either on temporal sequence modelling at the turbine level using recurrent neural networks such as LSTMs (Qian, Tian, Kanfoud, Lee, & Gan, 2019; Cui, Bangalore, & Bertling Tjernberg, 2021), or on spatial population-based graph representations using graph neural networks and graph autoencoders (Pinciroli, Baraldi, & Zio, n.d.; Jiang et al., 2026). More recently, there has been increasing interest in coupling spatial and temporal learning strategies for wind turbine condition monitoring (Jin et al., 2024). Temporal approaches can capture evolving degradation behaviour at individual turbines but may fail to distinguish localized anomalies from fleet-wide operating regimes. Conversely, graph-based spatial methods can identify deviations from fleet-level behaviour but may not adequately capture temporal regime transitions and slowly evolving dynamics. The present study therefore proposes a coupled Graph Autoencoder (GAE) and Long Short-Term Memory (LSTM) framework that combines fleet-level spatial graph learning with turbine-level latent-space temporal forecasting for fully unsupervised anomaly detection using ICCP monitoring data.

2. METHODOLOGY

Our anomaly detection framework consists of a Graph Autoencoder (GAE) which learns spatial relationships among turbines and produces latent graph representations (Z), and a latent-space LSTM model which captures temporal evolution of these representations and predicts future turbine states. The overall architecture is illustrated conceptually in Fig.1. The GAE encoder and decoder are first trained independently using spatial reconstruction losses and subsequently frozen. The LSTM is then trained exclusively on the fixed latent representations generated by the pretrained encoder. This sequential training strategy allows the spatial fleet-level representations learned by the GAE to remain unchanged during temporal learning and enables a clearer separation between spatial and temporal anomaly contributions.

Considering an ICCP dataset containing T time steps, N turbines (or monitoring locations), and F derived features, the input tensor is defined as:

$$\mathbf{X} \in R^{T\times N\times F}, \tag{1}$$

The ICCP measurements within the wind farm are represented as a graph $G = (V, E)$, that turbines correspond to nodes V and edges E so encodes spatial proximity between turbines. The spatial proximity is quantified by using a k-nearest-neighbour scheme:

$$A_{ij} = \exp\left(-\frac{d_{ij}^2}{2\sigma^2}\right), \tag{2}$$

Here A_{ij} is the weight of the edge between turbines i and j, d_{ij} represents the distance between the coordinates, and σ is the kernel width, which controls the spatial decay of the edge weights.

For each time step, the turbine feature matrix $\mathbf{X}_t \in R^{N\times F}$ is processed by a two-layer Graph Convolutional Network (GCN) similar to (Kipf & Welling, 2016). The graph con-

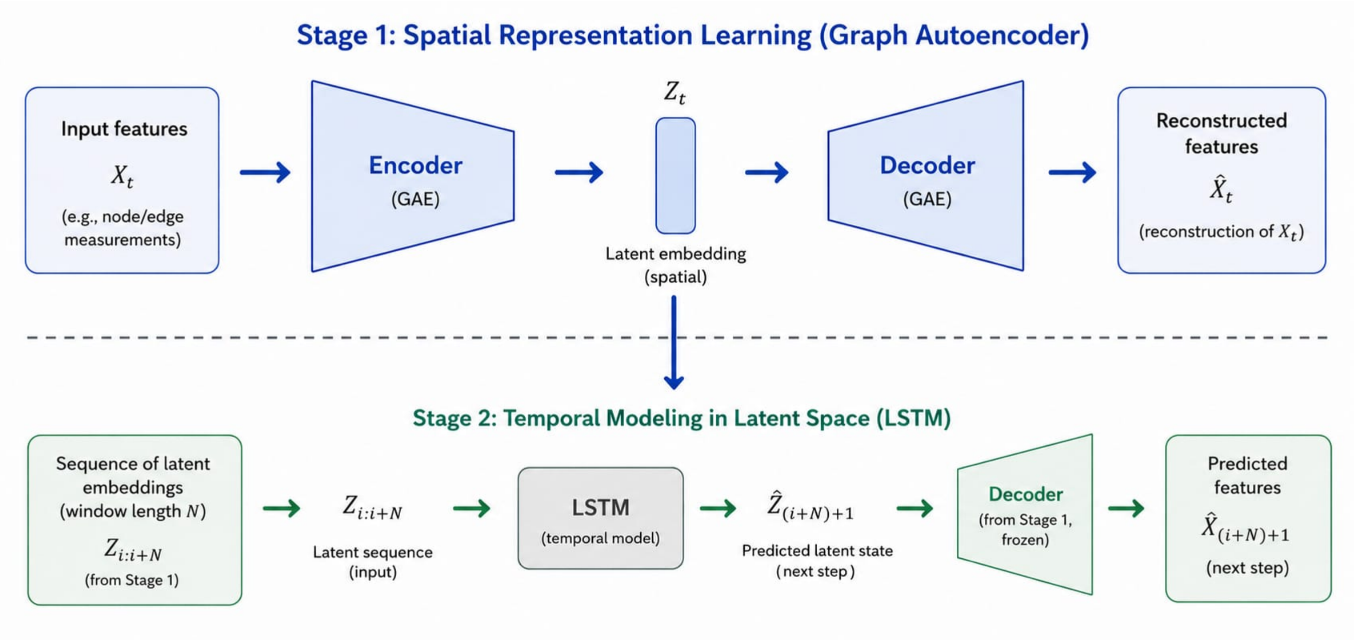

Figure 1. The proposed spatio-temporal anomaly detection framework with two stages: a Graph Autoencoder with an LSTM.

volution operation is defined as

$$\mathbf{H}^{(l+1)} = \sigma\left(\tilde{\mathbf{D}}^{-1/2}\tilde{\mathbf{A}}\tilde{\mathbf{D}}^{-1/2}\mathbf{H}^{(l)}\mathbf{W}^{(l)}\right), \tag{3}$$

Where $\tilde{\mathbf{D}}$ is the degree matrix, $\mathbf{W}^{(l)}$ are weights and $\sigma(\cdot)$ is an activation function. The encoder maps the node features into a latent space

$$\mathbf{Z}_t = f_\theta(\mathbf{X}_t, \mathbf{A}), \tag{4}$$

where $\mathbf{Z}_t \in R^{N\times d}$ corresponds to the latent embedding of all turbines at time t.

Then we make use of a node-wise multilayer perceptron (MLP) decoder to reconstruct the original features from these embeddings:

$$\hat{\mathbf{X}}_t = g_\phi(\mathbf{Z}_t). \tag{5}$$

And GAE is trained by minimizing the reconstruction loss:

$$\mathcal{L}_{GAE} = \|\mathbf{X}_t - \hat{\mathbf{X}}_t\|_2^2. \tag{6}$$

Once the GAE is trained, the encoder is applied to the entire dataset to produce a latent tensor:

$$\mathbf{Z} \in R^{T\times N\times d}. \tag{7}$$

A sliding window of length L is constructed from these latent graphs. For each window, the model receives a sequence,

$$\{\mathbf{Z}_{t-L+1}, \ldots, \mathbf{Z}_t\} \tag{8}$$

and predicts the next latent state $\hat{\mathbf{Z}}_{t+1}$. The temporal model is implemented based on a Long Short-Term Memory (LSTM) network (Hochreiter & Schmidhuber, 1997).

In the next step, the predicted latent representation is mapped back to the observation space through the frozen decoder:

$$\hat{\mathbf{X}}_{t+1} = g_\phi(\hat{\mathbf{Z}}_{t+1}). \tag{9}$$

And the LSTM is trained using=:

$$\mathcal{L}_{LSTM} = \|\mathbf{Z}_{t+1} - \hat{\mathbf{Z}}_{t+1}\|_2^2 + \beta\|\mathbf{X}_{t+1} - \hat{\mathbf{X}}_{t+1}\|_2^2, \tag{10}$$

After training, anomalies are identified by comparing predicted and observed turbine states. We consider three error components: 1) Graph reconstruction error from the GAE, 2) Latent prediction error from the LSTM and 3) Observation-space prediction error after decoding. A combined anomaly score for turbine i at time t is defined as

$$S_{i,t} = w_1 E_{i,t}^{GAE} + w_2 E_{i,t}^{LAT} + w_3 E_{i,t}^{OBS}, \tag{11}$$

where $E_{i,t}^{GAE}$, $E_{i,t}^{LAT}$ and $E_{i,t}^{OBS}$ correspond to the normalized graph reconstruction, latent prediction, and observation-space prediction errors, respectively, and w_k are weighting coefficients. For this study, we set the weighting coefficients to $w_1 = w_2 = w_3 = 0.33$ so that each error component contributes equally to the anomaly score.

3. SYNTHETIC DATASET

The available ICCP dataset does not contain any confirmed failure labels or verified anomaly events. Hence, it is suitable for unsupervised anomaly detection but it cannot be used to quantify detection performance and parameter sensitivity analysis. Therefore, we first construct a synthetic ICCP-like dataset that mimics the time variation of reference potential over a year and provide explicit ground-truth labels for different anomaly classes.

The synthetic dataset represents a fleet of N turbines observed over T time steps and assigned a two-dimensional spatial coordinate $\mathbf{p}_i = (x_i, y_i)$. The spatial distance between turbines i and j is denoted by d_{ij}. We assume that the fleet-level correlations can be defined as a function of spatial proximity. To introduce spatially correlated variations across the fleet, we define a covariance matrix using a squared-exponential kernel,

$$\Sigma_{ij} = \exp\left(-\frac{d_{ij}^2}{2\ell^2}\right), \tag{12}$$

where ℓ is the spatial correlation length scale. Spatially correlated noise is then generated from this covariance structure, so that nearby turbines exhibit more similar background fluctuations than distant turbines.

The healthy reference potential signal for turbine i at time t is constructed as

$$s_{i,t}^{0} = s_{\text{base}} + s_{\text{fleet}}(t) + b_i + \epsilon_{i,t}^{\text{sp}} + \epsilon_{i,t}^{\text{ind}} + d_i(t), \tag{13}$$

where s_{base} is the nominal reference potential, $s_{\text{fleet}}(t)$ is a slowly varying fleet-level background component, b_i is a turbine-specific offset, $\epsilon_{i,t}^{\text{sp}}$ is the spatially correlated noise component, $\epsilon_{i,t}^{\text{ind}}$ is independent measurement noise, and $d_i(t)$ is a slowly varying turbine-specific drift term. The fleet-level component is generated as a smoothed stochastic process,

$$s_{\text{fleet}}(t) = A_{\text{fleet}} \frac{\mathcal{S}_{W_f}\left[r(t)\right]}{\text{std}\left(\mathcal{S}_{W_f}\left[r(t)\right]\right)}, \tag{14}$$

where $r(t)$ is a random process, $\mathcal{S}_{W_f}[\cdot]$ denotes moving-average smoothing over a window W_f, and A_{fleet} controls the amplitude of the fleet-wide variability. Similarly, the turbine-specific drift term is generated as

$$d_i(t) = A_{\text{drift}} \frac{\mathcal{S}_{W_d}\left[r_i(t)\right]}{\text{std}\left(\mathcal{S}_{W_d}\left[r_i(t)\right]\right)}, \tag{15}$$

where $r_i(t)$ is an independent random process for each turbine and W_d is the drift smoothing window. In addition to the non-anomalous background variability, we also add short-lived downward transients to mimic the system behaviour of the ICCP system that leads to short time-scaled changes of temporary protection level in response to environmental perturbations or controller responses. For a transient starting at time t_0, the imposed perturbation is

$$\Delta s_{i,t}^{\text{tr}} = -A_{\text{tr}} \exp\left[-\gamma(t - t_0)\right], \qquad t_0 \le t \le t_0 + \tau_{\text{tr}}, \tag{16}$$

where A_{tr} is the transient amplitude, γ controls the decay rate, and τ_{tr} is the duration of the event.

We consider four main major anomaly types in our synthetic dataset. First, a temporal anomaly is introduced at a single turbine by adding a gradually increasing drift, an oscillatory component, and additional noise:

$$\Delta s_{i,t}^{\text{temp}} = \alpha(t) + A_{\text{osc}} \sin\left(\frac{2\pi t}{P_{\text{osc}}}\right) + \epsilon_{i,t}^{\text{temp}}, \qquad t \in [t_0^{\text{temp}}, t_1^{\text{temp}}], \tag{17}$$

where $\alpha(t)$ is a linear drift term, A_{osc} is the oscillation amplitude, P_{osc} is the oscillation period, and $\epsilon_{i,t}^{\text{temp}}$ is additional noise.

Second, a spatial anomaly is introduced as a persistent offset at one turbine over a fixed time interval,

$$\Delta s_{i,t}^{\text{sp}} = A_{\text{sp}}, \qquad t \in [t_0^{\text{sp}}, t_1^{\text{sp}}], \tag{18}$$

where A_{sp} is the imposed spatial offset. This anomaly represents a turbine that deviates from the fleet-level spatial pattern without necessarily exhibiting a complex temporal evolution.

Third, a spatio-temporal anomaly is introduced over a local cluster of neighbouring turbines. For turbine i within the affected cluster, the perturbation is defined as

$$\Delta s_{i,t}^{\text{st}} = \lambda_i \beta(t) + \epsilon_{i,t}^{\text{st}}, \qquad t \in [t_0^{\text{st}}, t_1^{\text{st}}], \tag{19}$$

where $\beta(t)$ is a time-dependent drift affecting the cluster, $\epsilon_{i,t}^{\text{st}}$ is additional noise, and λ_i is a spatial decay factor that decreases with distance from the central affected turbine. This case is designed to test whether the framework can detect anomalies that are both spatially localized and temporally evolving.

Finally, a fleet-wide global shift is introduced as an anomalous regime change,

$$\Delta s_{i,t}^{\mathrm{glob}} = A_{\mathrm{glob}} \sin\left(\pi \frac{t - t_0^{\mathrm{glob}}}{t_1^{\mathrm{glob}} - t_0^{\mathrm{glob}}}\right), \qquad t \in [t_0^{\mathrm{glob}}, t_1^{\mathrm{glob}}], \tag{20}$$

which is applied to all turbines simultaneously.

The final synthetic reference potential is therefore given by

$$s_{i,t} = s_{i,t}^{0} + \Delta s_{i,t}^{\mathrm{tr}} + \Delta s_{i,t}^{\mathrm{temp}} + \Delta s_{i,t}^{\mathrm{sp}} + \Delta s_{i,t}^{\mathrm{st}} + \Delta s_{i,t}^{\mathrm{glob}}. \tag{21}$$

The ground-truth anomaly label is then defined as

$$y_{i,t} = I\left[\Delta s_{i,t}^{\mathrm{temp}} \neq 0 \ \vee \ \Delta s_{i,t}^{\mathrm{sp}} \neq 0 \ \vee \ \Delta s_{i,t}^{\mathrm{st}} \neq 0\right], \tag{22}$$

where $I[\cdot]$ is the indicator function. The transient spikes are therefore not included in the anomaly label.

The generated synthetic ICCP dataset, including the injected temporal, spatial, spatio-temporal, and global fleet-wide regimes, is shown in Fig. 2. The top panel of Fig. 2 illustrates the synthetic turbine layout together with the spatial location of the injected anomaly classes. Temporal anomalies are introduced at a single turbine and primarily affect the temporal evolution of the signal without strong fleet-level spatial deviation. Spatial anomalies are imposed as persistent offsets at one turbine relative to the surrounding fleet behaviour. In contrast, spatio-temporal anomalies are injected over a local cluster of neighbouring turbines and therefore combine both spatially localized and temporally evolving behaviour. The middle panel presents the corresponding ground-truth anomaly classification map used for model validation. Different colours represent the temporal, spatial, and spatio-temporal anomaly regimes, together with a fleet-wide global shift. The bottom panel shows representative synthetic ICCP reference potential signals corresponding to the different anomaly classes together with the fleet mean behaviour. The generated signals reproduce several characteristics observed in real ICCP monitoring data, including fleet-level correlated variability, localized deviations, transient disturbances, and slowly evolving temporal changes.

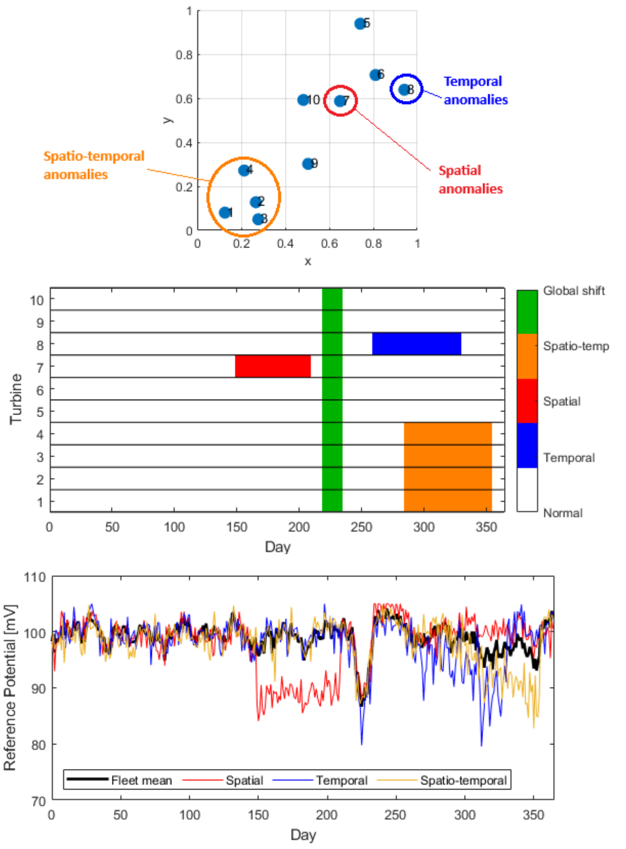

Figure 2. Top panel: Synthetic turbine layout and location of the injected anomaly classes. Middle panel: Ground-truth anomaly classification map used for validation of the anomaly detection framework. Bottom panel: Example synthetic ICCP reference potential signals for different anomaly regimes.

4. MODEL VALIDATION AND SENSITIVY ANALYSIS

Using the generated synthetic dataset, we first evaluate the added value of using a coupled GAE/LSTM framework vs a purely spatial Graph Autoencoder (GAE) approach. We perform an ablation analysis comparing a purely spatial Graph Autoencoder (GAE) approach against the coupled GAE/LSTM framework. Figure 3 presents the corresponding anomaly scores and detection maps for the synthetic ICCP dataset. In the GAE-only configuration, the anomaly score is computed exclusively from the graph reconstruction error, whereas the coupled GAE/LSTM framework combines latent-space temporal prediction errors and observation-space reconstruction errors. For this initial analysis, we used a graph with $k = 3$ nearest neighbours, a latent embedding dimension of $d = 16$, a hidden graph-convolution dimension of 32, and an LSTM window length of 30 days. The anomaly score threshold was initially fixed to the 85^{th} percentile of the training distribution without any additional anomaly score smoothing or filtering. The initial parameter choices were selected empirically to provide a stable baseline configuration before performing the full parameter sensitivity analysis presented later in this section.

The top panel of Fig. 3 compares the evolution of the mean anomaly score for the GAE-only and GAE/LSTM configurations. The middle and lower panels present the corresponding detection maps relative to the known ground-truth anomalies.

Despite a relatively high level of short-time scale false detections, both approaches successfully identify the localized spatial and spatio-temporal anomalies. However, a major difference emerges for the fleet-wide global shift occurring around day 220. In the GAE-only configuration, the graph autoencoder largely interprets this regime as normal behaviour because the shift occurs coherently across the entire fleet and therefore preserves the relative spatial structure between turbines. Consequently, the graph reconstruction error remains relatively small. In contrast, the coupled GAE/LSTM framework successfully detects the global shift because the latent-space temporal dynamics deviate significantly from the previously learned temporal evolution of the system. This result demonstrates that the temporal modeling component provides complementary information to the spatial graph reconstruction and highlights the importance of combining both spatial and temporal learning mechanisms for detecting regime changes in ICCP monitoring data.

To evaluate the robustness of the proposed framework with respect to model and anomaly-scoring parameters, we performed an additional sensitivity analysis consisting of 30 independent experiments with latent-space, temporal, and post-processing configurations. The evaluated parameters included the latent embedding dimension, graph-convolution hidden dimension, LSTM hidden dimension, temporal window length, latent-versus-observation anomaly weighting, smoothing settings, and minimum-duration filtering. The detection threshold was fixed to the 85^{th} percentile of the training distribution in all experiments in order to isolate the effect of the remaining parameters.

The sensitivity analysis was organized into three groups of experiments. The first group consisted of baseline configurations without temporal smoothing or duration filtering. The second group introduced temporal smoothing of the anomaly score using moving-average windows of different lengths. Finally, the third group combined temporal smoothing with a minimum-duration constraint requiring anomalies to persist for several consecutive timesteps before being classified as anomalous. The objective of this progressive analysis was to evaluate how temporal post-processing influences the trade-off between false positives and detection sensitivity. The complete set of tested configurations and corresponding performance metrics is provided in Appendix A.

The obtained results indicate that temporal post-processing substantially improves the robustness of the anomaly detection framework. Baseline configurations without smoothing or duration filtering produced moderate detection performance, with F1-scores typically between approximately 0.56 and 0.59. Introducing temporal smoothing increased recall substantially and improved the overall F1-score to approximately 0.61–0.64, indicating that coherent anomaly structures are more reliably detected when isolated fluctuations in the anomaly score are suppressed. The best overall performance was obtained when smoothing and minimum-duration filtering were combined, with several configurations achieving F1-scores above 0.75 and accuracies exceeding 0.92.

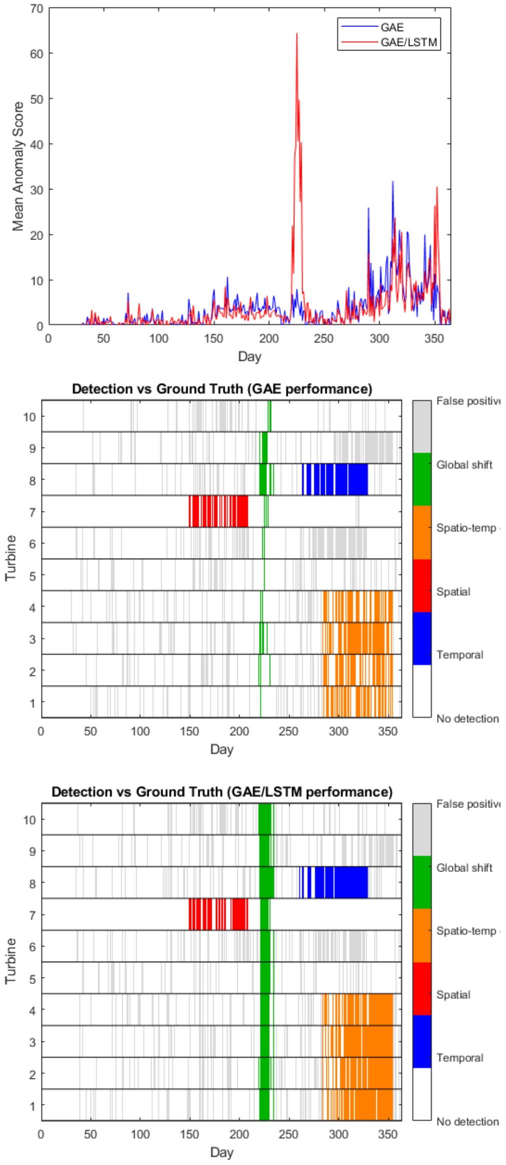

Figure 3. Top panel: Synthetic turbine layout and location of the injected anomaly classes. Middle panel: Ground-truth anomaly classification map used for validation of the anomaly detection framework. Bottom panel: Example synthetic ICCP reference potential signals for different anomaly regimes.

In contrast, variations in the latent embedding dimension, graph-convolution hidden dimension, LSTM hidden dimension, and temporal window length produced comparatively smaller changes in the detection metrics. Although moderate differences in precision and recall were observed across configurations, the overall behaviour of the framework remained relatively stable over a broad range of architectural parameter choices. These results suggest that the proposed framework is not strongly dependent on precise hyperparameter tuning and that the learned spatial-temporal representations are reasonably robust to moderate variations in model complexity.

Overall, the sensitivity analysis indicates that the proposed framework benefits primarily from enforcing temporal consistency in the anomaly decision process. This behaviour is physically consistent with ICCP monitoring systems, where isolated transient fluctuations may occur due to short-lived environmental or controller-related disturbances, whereas meaningful anomalous regimes are expected to persist over longer time periods. Consequently, temporal smoothing and minimum-duration filtering substantially reduce false positives while preserving the detection of sustained anomalous behaviour.

Among the tested configurations, run 25 yielded the best overall trade-off between precision, recall, F1-score, and accuracy, achieving an F1-score of approximately 0.76 and an accuracy exceeding 0.92. This configuration employed a latent embedding dimension of 16, graph hidden dimension of 16, LSTM hidden dimension of 16, temporal window length of 10 days, temporal smoothing with a window length of 3 samples, and a minimum anomaly duration of 7 samples. The corresponding anomaly detection results for this optimal configuration are presented in Fig. 4.

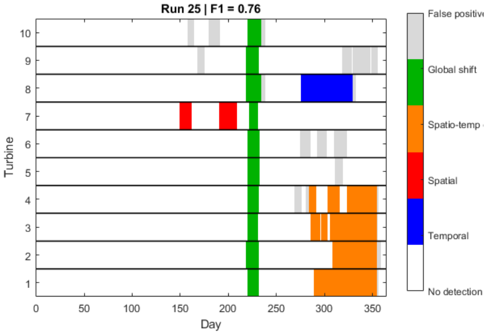

Figure 4. Anomaly results with the optimal configuration found in total 30 experiment runs.

5. APPLICATION TO ICCP DATA

Here, we demonstrate the proposed anomaly detection framework using ICCP monitoring data from an offshore wind farm. ICCP systems operate by imposing an external electrical current between inert anodes mounted on the steel structure to shift the electrochemical potential of the steel to more negative values. This shift effectively suppresses the corrosion reaction as the steel surface then behaves as a cathode instead of an anode in an electro-chemical cell. This potential is measured using reference electrodes installed near the structure. As a result, ICCP monitoring systems typically provide time series of voltage, injected current and reference potential measurements.

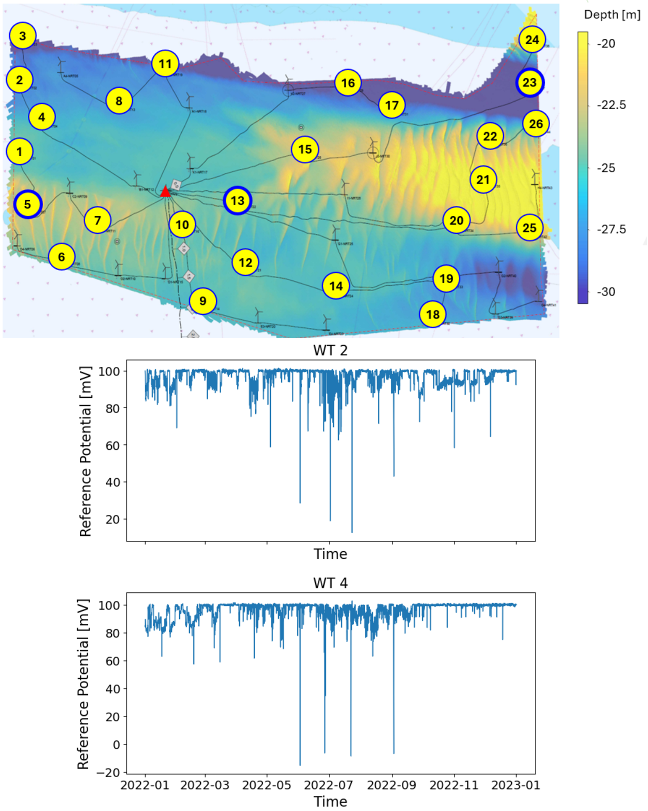

Figure 5. ICCP dataset used in this study. Top panel: spatial distribution of wind turbines in the offshore wind farm and subset of turbines used in the analysis. Lower panels: examples of reference potential time series measured at two turbines over the one year observation period.

Here we consider ICCP data collected from an offshore wind farm consisting of 44 wind turbines. However, not all turbines provide continuous data coverage over the full observation period. Hence, we restrict the analysis to a subset of 26 turbines which provide uninterrupted measurements over approximately one year. The spatial distribution of turbines and the subset used in the analysis are illustrated in Fig. 5. The lower panels of the figure also show examples of the reference potential time series recorded at several turbines.

The reference potential provides the most direct indicator of the electrochemical protection state of the steel structure among the available ICCP measurements. It represents the electrical potential difference between protected steel and a reference electrode placed in the surrounding seawater, and therefore directly reflects the electrochemical conditions governing corrosion processes. In contrast, injected current or system voltage are indirect control variables influenced by environmental conditions, coating state, and controller adjustments. In the dataset analysed here the ICCP controller maintains a nominal reference potential setpoint of approximately 100 mV relative to the free corrosion potential. Because this signal directly reflects the electrochemical protection state, deviations from its nominal behaviour may indicate abnormal conditions such as coating degradation, sensor malfunction, electrical faults in the ICCP system, or changes in environmental conditions affecting the electrochemical interface.

To extract physically meaningful characteristics from the signal, we construct a set of derived features computed over rolling time windows. The first feature corresponds directly to the raw input (daily averaged reference potential):

$$x_{i,t}^{(1)} = s_{i,t}. \tag{23}$$

Slow drifts in the reference potential may indicate progressive degradation mechanisms such as coating damage or changes in the electrochemical environment. To capture such behaviour we compute the slope of the signal over a rolling window of 10 days (hereafter denoted as L) over one year:

$$x_{i,t}^{(2)} = \frac{\sum_{k=0}^{L-1}(k-\bar{k})(s_{i,t-k}-\bar{s}_{i,t})}{\sum_{k=0}^{L-1}(k-\bar{k})^2}, \tag{24}$$

where $\bar{s}_{i,t}$ denotes the mean value of the signal within the window. This feature measures the local trend of the reference potential and allows the detection of gradual deviations from the nominal protection level.

Short-term fluctuations of the potential may arise from electrical disturbances, environmental variability or unstable control behaviour. To quantify the variability of the signal in a robust manner we compute the rolling median absolute deviation within a daily window

$$x_{i,t}^{(3)} = \text{median}\left(|s_{i,t-k}-\tilde{s}_{i,t}|\right), \quad k = 0, \ldots, L-1, \tag{25}$$

where $\tilde{s}_{i,t}$ denotes the median potential within the rolling window. Variations in this feature may indicate intermittent disturbances, oscillatory controller behaviour or electrode instability.

Finally, we compute the lag-1 autocorrelation coefficient within a 21-day rolling window

$$x_{i,t}^{(4)} = \frac{\sum_{k=1}^{L-1}(s_{i,t-k}-\bar{s}_{i,t})(s_{i,t-k+1}-\bar{s}_{i,t})}{\sum_{k=0}^{L-1}(s_{i,t-k}-\bar{s}_{i,t})^2}. \tag{26}$$

This feature captures the temporal persistence of the signal. Changes in autocorrelation structure may reflect regime transitions in the ICCP system dynamics, sensor sticking behaviour or oscillatory control responses.

The resulting feature vector for turbine i at time t is therefore

$$\mathbf{x}_{i,t} = \left[s_{i,t}, \text{slope}_{i,t}, \text{MAD}_{i,t}, \rho_{i,t}\right], \tag{27}$$

where $\rho_{i,t}$ denotes the lag-1 autocorrelation.

Figure 6 presents the distribution of detected anomalies across the wind farm together with a detailed view of the most anomalous turbine (middle and bottom panels). The top panel shows the anomaly fraction for each turbine over the one-year observation period. Each node represents a turbine and the color indicates the percentage of timestamps flagged as anomalous by the model. Most turbines exhibit near-zero anomaly fractions, indicating highly consistent behaviour across the wind farm during the analysed period. However, a limited number of turbines display elevated anomaly fractions, with the highest anomaly fraction reaching approximately 23.9%. Interestingly, the most anomalous turbines are located within the same local region of the wind farm, suggesting that the detected deviations may be associated with localized operational, environmental, or corrosion-related effects rather than purely random fluctuations.

The middle panel presents the reference potential time series of the most anomalous turbine together with the fleet mean and the associated variability envelope. During a substantial fraction of the observation period, the turbine exhibits a moderate offset relative to the fleet-average behaviour while still remaining relatively smooth and temporally stable. Although this spatial deviation alone could potentially be interpreted as anomalous in a purely graph-based reconstruction framework, the coupled GAE/LSTM architecture does not flag these periods because the temporal evolution of the signal remains largely consistent with previously learned behaviour. Toward the final months of the dataset, however, the signal begins to exhibit both larger deviations from the fleet-average magnitude and increased short-timescale variability.

Correspondingly, the bottom panel shows a substantial increase in the anomaly score during this later period, indicating that the turbine progressively deviates from the learned spatio-temporal behaviour of the fleet. This behaviour suggests that the temporal modelling stage contributes to suppressing stable but persistent spatial offsets while emphasizing evolving temporal inconsistencies and dynamically changing regimes. Nevertheless, it must be noted that the lack of labelled degradation events or direct corrosion measurements means that the detected anomalies should presently be interpreted as statistically unusual spatio-temporal patterns rather than confirmed physical faults.

To better understand the behavior of the detected anomalies, Fig. 7 shows the reference potential signal together with the derived features for the most anomalous turbine (WT16). The gray shaded bands indicate time windows classified as anoma-

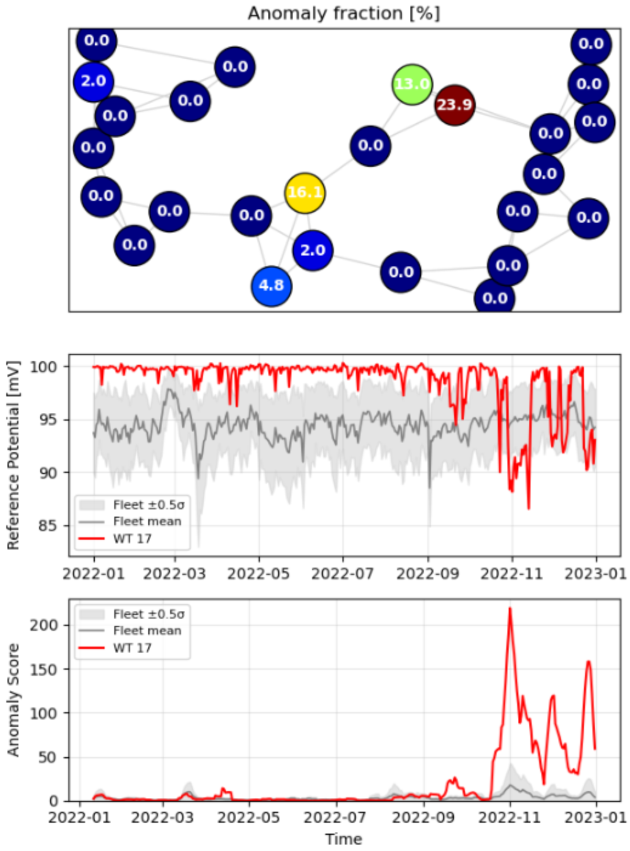

Figure 6. Anomaly detection results for the ICCP dataset. Top panel: spatial distribution of anomaly fractions across turbines in the wind farm. Middle panel: reference potential time series for the most anomalous turbine in comparison to the fleet mean. Bottom panel: corresponding anomaly score in comparison to the fleet mean.

lous by the proposed framework.

During the first part of the year (approximately January to October), the reference potential remains largely stable around the nominal protection level of approximately 100 mV, with occasional short-lived drops in potential. The derived features during this period remain relatively small: the rolling slope fluctuates around zero, the rolling Medium Absolute Deviation (MAD) remains close to its baseline level, and the lag-1 autocorrelation exhibits moderate variability. The anomalies detected in this regime correspond primarily to isolated potential excursions in the raw signal. Such events are consistent with transient disturbances commonly observed in offshore cathodic protection systems, including temporary electrode perturbations, environmental fluctuations, or short control responses of the ICCP system.

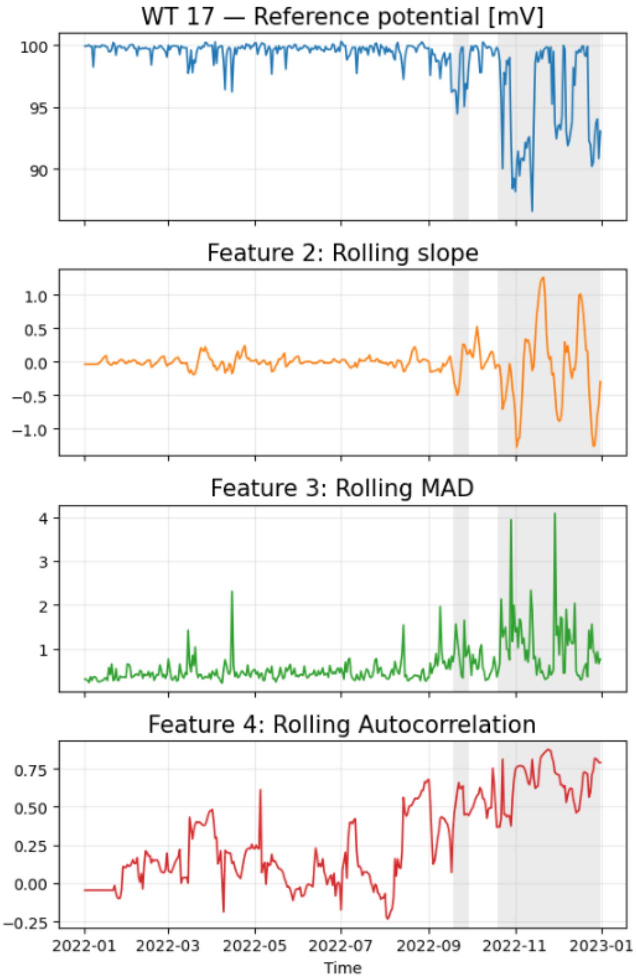

Figure 7. Diagnostic representation of the detected anomalies for the most anomalous turbine (WT16). From top to bottom: reference potential signal, rolling slope, rolling MAD, and rolling lag-1 autocorrelation. Gray shaded regions indicate time windows classified as anomalous by the proposed framework.

However, we observe a different behaviour during the second part of the observation period, starting around early autumn. In this regime the signal exhibits a noticeable increase in variability and temporal persistence. The rolling MAD increases significantly, indicating stronger short-term fluctuations in the reference potential. For the same time period, rolling slope shows larger oscillations, which might suggest the presence of local trends in the signal. The lag-1 autocorrelation also increases toward values close to unity, reflecting a higher degree of temporal persistence in the signal dynamics. These simultaneous changes in multiple features indicate a shift in the dynamical regime of the ICCP monitoring signal rather than isolated disturbances. As a result, the anomaly detection framework identifies a larger number of anomalous windows during this period. The feature-level representation therefore provides insight into the nature of the detected anomalies by distinguishing between short-lived signal perturbations and longer-term structural changes in the behaviour of the ICCP system. These feature variations further suggest that the anomalous periods identified by the framework are unlikely to correspond solely to intrinsic ICCP control behaviour or short-timescale operational switching, since the observed temporal scales are substantially longer and involve persistent changes across multiple derived features. Such behaviour may potentially reflect seasonal environmental forcing effects; however, the absence of additional environmental measurements such as seawater temperature, salinity, wave activity, or hydrodynamic conditions prevents a more robust interpretation of the role of seasonality in the present study.

6. CONCLUSION AND PERSPECTIVES

We presented an initial implementation of a spatio-temporal anomaly detection framework for ICCP monitoring data in offshore wind farms based on a coupled Graph Autoencoder and latent-space LSTM architecture. The framework was applied to long-term ICCP monitoring data from an offshore wind farm, focusing on reference potential measurements from a subset of turbines with continuous observations. Feature engineering was designed to capture multiple aspects of signal behaviour, including local drift, short-term variability, and temporal persistence. The results show that the majority of turbines exhibit highly consistent behaviour with low anomaly fractions, while a small number of turbines display recurrent deviations from the learned patterns. Detailed feature-level diagnostics further indicate that the detected anomalies include both transient disturbances in the ICCP signal and shifts in the dynamical regime of the system characterized by increased volatility, stronger temporal persistence, and local trends. These findings demonstrate that combining graph-based spatial learning with temporal sequence modelling provides an effective approach for identifying abnormal behaviour in cathodic protection monitoring systems without requiring extensive sensor deployment or predefined fault labels. Moreover, the proposed framework illustrates the potential of population-based monitoring strategies for offshore wind farms, where information from multiple turbines can be used to detect minor deviations that would be difficult to identify through single-turbine analysis alone.

The graph adjacency is defined based on spatial proximity between turbines. However, future extensions could incorporate more advanced graph construction strategies based on correlated turbine responses, operational similarity, environmental forcing, or statistical dependency measures such as correlation coefficients. Such extensions could allow the framework to capture more complex interactions between environmental forcing, operational conditions, and corrosion protection behaviour while potentially improving the physical interpretability of detected anomalies. Moreover, potential seasonal environmental effects should also be considered when interpreting long-term ICCP monitoring behaviour.

The sensitivity analysis further demonstrated that anomaly scoring and post-processing parameters can significantly influence the final detection performance. In particular, percentile-based thresholds in combination with temporal smoothing and minimum-duration filtering provided substantially more stable detection behaviour and reduced the occurrence of isolated false-positive detections. Based on the present synthetic benchmark experiments with moderate smoothing windows and minimum-duration constraints yielded robust performance across a broad range of model configurations. Nevertheless, the optimum parameter selection remains application-dependent and may vary depending on the temporal variability, environmental forcing, and operational characteristics of a given offshore wind farm. Consequently, future studies should ideally rely on ICCP datasets containing labelled degradation events or maintenance records to enable systematic calibration and optimization of anomaly scoring weights, thresholds, and post-processing strategies. Such labelled datasets would also facilitate transition from fully unsupervised anomaly detection toward semi-supervised performance optimization and more physically interpretable fault classification. Moreover, additional comparisons against fully temporal non-graph architectures constitute an interesting direction for future work.

ACKNOWLEDGMENT

This research is funded by the EU Project "Wholistic and integrated digital tools for extended lifetime and profitability of offshore wind farms (WILLOW)" (https://willow-project.eu/), under grant agreement number 101122184 in call HORIZON-CL5-2022-D3-03-04 — Integrated wind farm control. An AI-based language model was used as an editorial tool for typesetting and formatting.

REFERENCES

Adedipe, O., Brennan, F., & Kolios, A. (2016). Review of corrosion fatigue in offshore structures: Present status and challenges in the offshore wind sector. *Renewable and Sustainable Energy Reviews*, *61*, 141–154. doi: 10.1016/j.rser.2016.02.017

Black, I. M., Yeter, B., Häckell, M. W., & Kolios, A. (2024). Assessing structural homogeneity and heterogeneity in offshore wind farms: A population-based structural health monitoring approach. *Ocean Engineering*, *311*, 118842. doi: 10.1016/j.oceaneng.2024.118842

Brijder, R., Hagen, C. H. M., Cortés, A., Irizar, A., Thibbotuwa, U. C., Helsen, S., et al. (2022). Review of corrosion monitoring and prognostics in offshore wind turbine structures: Current status and feasible approaches. *Frontiers in Energy Research*, *10*, 991343. doi: 10.3389/fenrg.2022.991343

Cui, Y., Bangalore, P., & Bertling Tjernberg, L. (2021). A fault detection framework using recurrent neural networks for condition monitoring of wind turbines. *Wind Energy*, *24*(11), 1249–1262.

Erdogan, C., & Swain, G. (2021). Conceptual sacrificial anode cathodic protection design for offshore wind monopiles. *Ocean Engineering*, *235*, 109339. doi: 10.1016/j.oceaneng.2021.109339

Erdogan, C., & Swain, G. (2022). The effect of macro-galvanic cells on corrosion and impressed current cathodic protection for offshore monopile steel structures. *Ocean Engineering*, *265*, 112575. doi: 10.1016/j.oceaneng.2022.112575

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, *9*(8), 1735–1780. doi: 10.1162/neco.1997.9.8.1735

Jiang, G., Tang, H., Yue, J., Ding, X., He, Q., & Xie, P. (2026). Conditional graph autoencoder with embedded prior knowledge for wind turbine fault diagnosis. *IEEE Transactions on Instrumentation and Measurement*.

Jin, X., Lv, S., Kong, Z., Yang, H., Zhang, Y., Guo, Y., & Xu, Z. (2024). Graph spatio-temporal networks for condition monitoring of wind turbine. *IEEE Transactions on Sustainable Energy*, *15*(4), 2276–2286.

Kalovelonis, D. T., Gortsas, T. V., Tsinopoulos, S. V., & Polyzos, D. (2025). A novel design methodology for sacrificial anode cathodic protection systems using numerical modeling: A case study of offshore wind turbine monopile foundations. *Ocean Engineering*, *318*, 120169. doi: 10.1016/j.oceaneng.2024.120169

Kipf, T. N., & Welling, M. (2016). Variational graph auto-encoders. *arXiv preprint arXiv:1611.07308*.

Martin, R., Lazakis, I., Barbouchi, S., & Johanning, L. (2016). Sensitivity analysis of offshore wind farm operation and maintenance cost and availability. *Renewable Energy*, *85*, 1226–1236. doi: 10.1016/j.renene.2015.07.078

Okenyi, V., Bodaghi, M., Mansfield, N., Afazov, S., & Siegkas, P. (2022). A review of challenges and framework development for corrosion fatigue life assessment of monopile-supported horizontal-axis offshore wind turbines. *Ships and Offshore Structures*, 1–15. doi: 10.1080/17445302.2022.2140531

Pinciroli, L., Baraldi, P., & Zio, E. (n.d.). Early anomaly detection in wind turbines by causality-based graph attention networks. *Available at SSRN 6466344*.

Porchetta, S., Temel, O., Warner, J. C., Muñoz-Esparza, D., Monbaliu, J., van Beeck, J., & van Lipzig, N. (2021). Evaluation of a roughness length parametrization accounting for wind–wave alignment in a coupled atmosphere–wave model. *Quarterly Journal of the Royal Meteorological Society*, *147*(735), 825–846.

Price, S. J., & Figueira, R. B. (2017). Corrosion protection systems and fatigue corrosion in offshore wind structures: Current status and future perspectives. *Coatings*, *7*(2), 25. doi: 10.3390/coatings7020025

Qian, P., Tian, X., Kanfoud, J., Lee, J. L. Y., & Gan, T.-H. (2019). A novel condition monitoring method of wind turbines based on long short-term memory neural network. *Energies*, *12*(18), 3411.

Ren, Z., Verma, A. S., Li, Y., Teuwen, J. J. E., & Jiang, Z. (2021). Offshore wind turbine operations and maintenance: A state-of-the-art review. *Renewable and Sustainable Energy Reviews*, *144*, 110886. doi: 10.1016/j.rser.2021.110886

Rokhlin, S. I., Kim, J.-Y., Nagy, H., & Zoofan, B. (1999). Effect of pitting corrosion on fatigue crack initiation and fatigue life. *Engineering Fracture Mechanics*, *62*(4–5), 425–444. doi: 10.1016/S0013-7944(98)00101-5

Santhakumar, S., Smart, G., Noonan, M., Meerman, H., & Faaij, A. (2022). Technological progress observed for fixed-bottom offshore wind in the EU and UK. *Technological Forecasting and Social Change*, *182*, 121856. doi: 10.1016/j.techfore.2022.121856

Shittu, A. A., Mehmanparast, A., Shafiee, M., Kolios, A., Hart, P., & Pilario, K. (2020). Structural reliability assessment of offshore wind turbine support structures subjected to pitting corrosion-fatigue: A damage tolerance modelling approach. *Wind Energy*, *23*, 2004–2026. doi: 10.1002/we.2542

Tremps, L., Yeter, B., & Kolios, A. (2024). Review and analysis of the failure risk mitigation via monitoring for monopile offshore wind structures. *Energy Reports*, *11*, 5407–5420. doi: 10.1016/j.egyr.2024.05.026

Zamanzadeh Darban, M., et al. (2025). A novel anomaly detection method for multivariate time series based on spatial-temporal graph learning. *Journal of King Saud University – Computer and Information Sciences*. doi: 10.1007/s44443-025-00024-3

Zheng, Y., Ma, S., Chen, J., Zhang, J., et al. (2024). Spatio-temporal anomaly detection for multivariate time series based on graph neural network. *Information Fusion*, 102255. doi: 10.1016/j.inffus.2024.102255

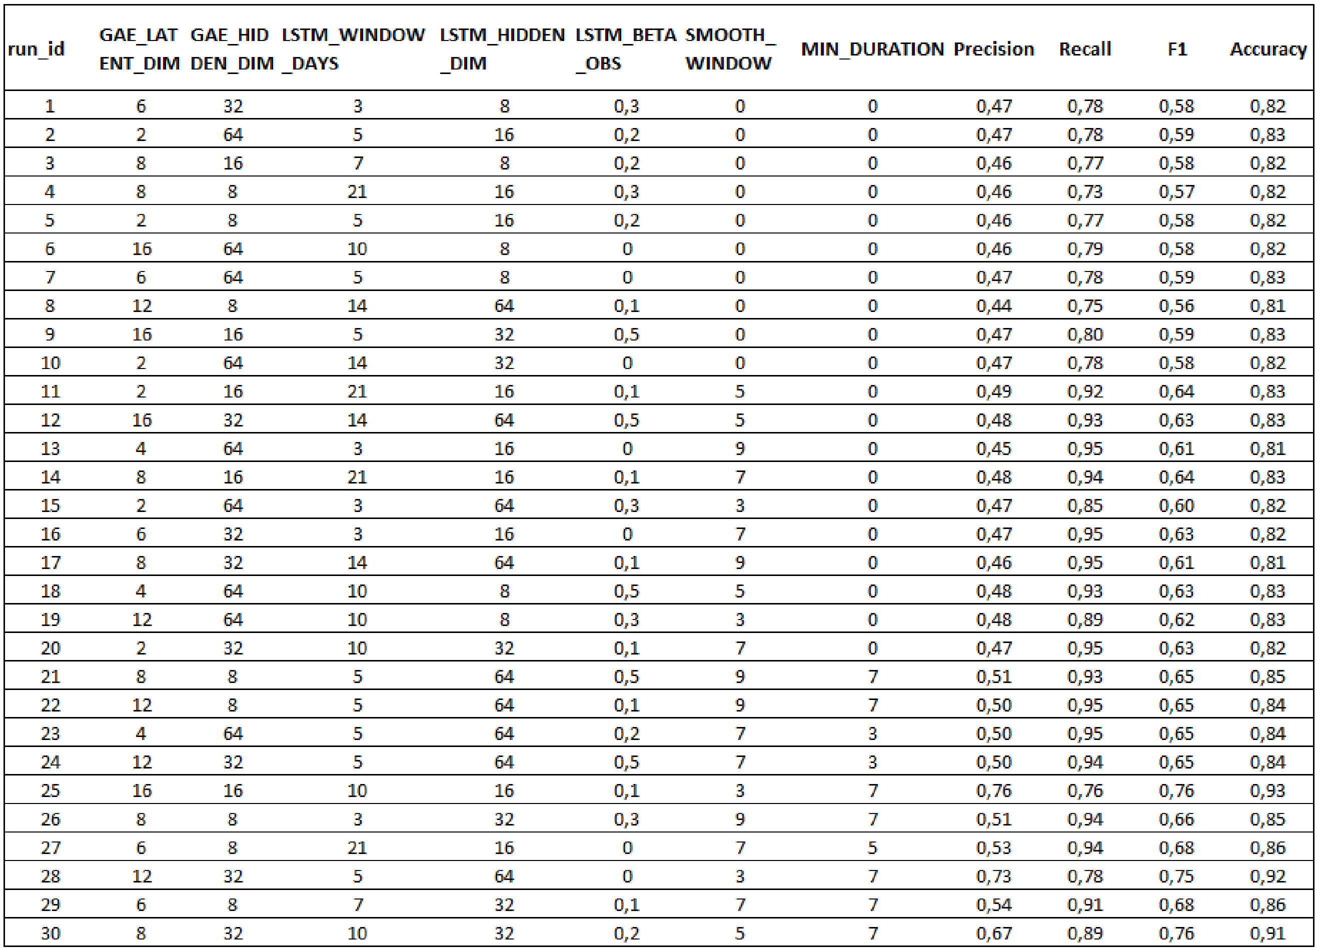

run_id	GAE_LATENT_DIM	GAE_HIDDEN_DIM	LSTM_WINDOW_DAYS	LSTM_HIDDEN_DIM	LSTM_BETA_OBS	SMOOTH_WINDOW	MIN_DURATION	Precision	Recall	F1	Accuracy
1	6	32	3	8	0,3	0	0	0,47	0,78	0,58	0,82
2	2	64	5	16	0,2	0	0	0,47	0,78	0,59	0,83
3	8	16	7	8	0,2	0	0	0,46	0,77	0,58	0,82
4	8	8	21	16	0,3	0	0	0,46	0,73	0,57	0,82
5	2	8	5	16	0,2	0	0	0,46	0,77	0,58	0,82
6	16	64	10	8	0	0	0	0,46	0,79	0,58	0,82
7	6	64	5	8	0	0	0	0,47	0,78	0,59	0,83
8	12	8	14	64	0,1	0	0	0,44	0,75	0,56	0,81
9	16	16	5	32	0,5	0	0	0,47	0,80	0,59	0,83
10	2	64	14	32	0	0	0	0,47	0,78	0,58	0,82
11	2	16	21	16	0,1	5	0	0,49	0,92	0,64	0,83
12	16	32	14	64	0,5	5	0	0,48	0,93	0,63	0,83
13	4	64	3	16	0	9	0	0,45	0,95	0,61	0,81
14	8	16	21	16	0,1	7	0	0,48	0,94	0,64	0,83
15	2	64	3	64	0,3	3	0	0,47	0,85	0,60	0,82
16	6	32	3	16	0	7	0	0,47	0,95	0,63	0,82
17	8	32	14	64	0,1	9	0	0,46	0,95	0,61	0,81
18	4	64	10	8	0,5	5	0	0,48	0,93	0,63	0,83
19	12	64	10	8	0,3	3	0	0,48	0,89	0,62	0,83
20	2	32	10	32	0,1	7	0	0,47	0,95	0,63	0,82
21	8	8	5	64	0,5	9	7	0,51	0,93	0,65	0,85
22	12	8	5	64	0,1	9	7	0,50	0,95	0,65	0,84
23	4	64	5	64	0,2	7	3	0,50	0,95	0,65	0,84
24	12	32	5	64	0,5	7	3	0,50	0,94	0,65	0,84
25	16	16	10	16	0,1	3	7	0,76	0,76	0,76	0,93
26	8	8	3	32	0,3	9	7	0,51	0,94	0,66	0,85
27	6	8	21	16	0	7	5	0,53	0,94	0,68	0,86
28	12	32	5	64	0	3	7	0,73	0,78	0,75	0,92
29	6	8	7	32	0,1	7	7	0,54	0,91	0,68	0,86
30	8	32	10	32	0,2	5	7	0,67	0,89	0,76	0,91

Figure Appendix A. Summary of 30 test runs