

PAPER • OPEN ACCESS

## Wind Park Power Prediction: Attention-Based Graph Networks and Deep Learning to Capture Wake Losses

To cite this article: Lars Ødegaard Bentsen *et al* 2022 *J. Phys.: Conf. Ser.* **2265** 022035

View the [article online](#) for updates and enhancements.

You may also like

- [Graphene nanonet for biological sensing applications](#)  
Taekyeong Kim, Jaesung Park, Hye Jun Jin et al.
- [Graph networks for molecular design](#)  
Rocío Mercado, Tobias Rastemo, Edvard Lindelöf et al.
- [Learning physical properties of anomalous random walks using graph neural networks](#)  
Hippolyte Verdier, Maxime Duval, François Laurent et al.

### ECS Toyota Young Investigator Fellowship



For young professionals and scholars pursuing research in batteries, fuel cells and hydrogen, and future sustainable technologies.

At least one \$50,000 fellowship is available annually.  
More than \$1.4 million awarded since 2015!



Application deadline: January 31, 2023

**Learn more. Apply today!**

# Wind Park Power Prediction: Attention-Based Graph Networks and Deep Learning to Capture Wake Losses

Lars Ødegaard Bentsen<sup>1</sup>, Narada Dilp Warakagoda<sup>1</sup>, Roy Stenbro<sup>2</sup> and Paal Engelstad<sup>1</sup>

<sup>1</sup>Department of Technology Systems, University of Oslo, PO Box 70, 2027 Kjeller, Norway.

<sup>2</sup>Institute for Energy Systems (IFE), PO Box 40, 2027 Kjeller, Norway.

E-mail: [lars.bentsen@its.uio.no](mailto:lars.bentsen@its.uio.no)

**Abstract.** With the increased penetration of wind energy into the power grid, it has become increasingly important to be able to predict the expected power production for larger wind farms. Deep learning (DL) models can learn complex patterns in the data and have found wide success in predicting wake losses and expected power production. This paper proposes a modular framework for attention-based graph neural networks (GNN), where attention can be applied to any desired component of a graph block. The results show that the model significantly outperforms a multilayer perceptron (MLP) and a bidirectional LSTM (BLSTM) model, while delivering performance on-par with a vanilla GNN model. Moreover, we argue that the proposed graph attention architecture can easily adapt to different applications by offering flexibility into the desired attention operations to be used, which might depend on the specific application. Through analysis of the attention weights, it was showed that employing attention-based GNNs can provide insights into what the models learn. In particular, the attention networks seemed to realise turbine dependencies that aligned with some physical intuition about wake losses.

## 1. Introduction

Wind energy has been subject to rapid expansion in recent years, with a record 82 GW of new wind capacity installed globally, and a year-on-year growth of 53% in 2020 [1]. Due to the intermittent nature of wind energy, research into wind power modelling and forecasting will remain important to establish wind as a more reliable energy resource.

Turbines have traditionally been characterised by their power curves, a simple relationship that indicates the expected power production at different wind speeds. However, with the construction of large wind farms, interaction losses from turbines in near vicinity to each other become prominent, resulting in actual production values deviating from power curve predictions. The wind velocity deficit and increased turbulence generated downstream of a turbine characterise the turbine wake, which result in turbines situated downstream seeing a reduction in the power generated. Significant efforts have been spent trying to account for these wake losses, modelling the complicated flow patterns that arise downstream within a farm. A simple, but widely used method, has been the Jensen wake model, which calculates the velocity deficit downstream of a single wind turbine [2]. High resolution computational fluid dynamics



can yield high fidelity results [3], but a problem with such numerical methods have been their high computational cost.

DL-based models have the potential to yield accurate results, while also significantly reduce run times for trained models. Various studies have reported success in training artificial neural networks (ANN) to predict wind farm power from the free stream wind velocity and direction [4, 5]. A two-dimensional power curve was introduced by Yan et al. [5], where the expected power could be determined from the wind speed and direction, thereby taking wake losses into account. Ti et al. [6] trained an MLP to predict spatial velocity deficit and added turbulence downstream, aiming to output high resolution flow fields similar to those of numerical methods.

Tastu et al. [7] considered spatial dynamics by using measurements from a number of farms scattered across Denmark, to make their predictions for a target wind farm. Similarly, both [8] and [9] proposed spatio-temporal models to make predictions at a number of wind sites jointly. Ghaderi et al. [9] modelled the spatial characteristics by utilising graphical representations, where the weather stations were represented by nodes and the physical distance between them as edges in a graph, using a Long Short-Term Memory (LSTM) network to make the forecasts.

Although Recurrent (RNN) and Convolutional Neural Networks (CNN) have predominantly been used for time-series analysis [10, 11, 12], research has also emerged in applying such techniques to capture spatial features. CNNs have shown superior performance in extracting features from image data. In [13] and [14], the authors aimed to leverage some of these attributes, by organising the input data in a grid structure, where each pixel contained information for a particular physical location. Despite yielding good results, a drawback of such an approach is that the strict ordering of the input data might pose some challenges for complex farm layouts.

Graphs have been used to represent data characterised by complex relations, where the data is not naturally ordered in a strict grid-like structure. Similar to [9], the authors of [15], [16] and [17] also aimed to improve predictions by considering measurements from different wind farms, but instead employed graph convolutional models. Bleeg [18] used a GNN as a surrogate model for a steady-state Reynolds-averaged Navier–Stokes (RANS) model to learn turbine interaction losses within a single wind farm. Yu et al. [19] constructed a spatio-temporal model to predict wind power for individual turbines using a combination of GNNs and RNNs, trained on data for four offshore wind farms. To improve on the vanilla GNN model, Park et al. [20] proposed the physics-induced GNN (PGNN), where the intensity of different turbine interactions was regulated by computing weights using a physics-induced wake function. The authors showed the potential benefits of introducing relative weighting of upstream turbines to learn interaction losses, using data simulated in FLORIS [21]. Albeit not fully learned from the data, the weighting introduced in the PGNN draws some resemblance to an attention mechanism, as it regulates the relative impact of neighbouring turbines when predicting the power for a target turbine.

Attention was first introduced by Bahdanau et al. [22], and has become an integral part of sequence modelling using DL. With regards to wind, attention mechanisms have mainly been used to extract temporal tendencies in time-series data, and thereby improve RNN-based forecasting models [23, 24, 25]. Due to the success of attention-based sequence models, there has been a surge in the application of such techniques to other DL domains, such as the development of Graph Attention Networks [26, 27, 28, 29]. Nevertheless, these networks have yet to find their way into the wind modelling community, as the authors of this paper cannot report on finding any research discussing the use of attention-based GNNs applied to wind energy.

The main contributions of this paper can be summarised as follows:

- A BLSTM- and an MLP-model, where inputs are organised around individual turbines, are used to predict the power for a target turbine given the location of upstream neighbours.
- We propose a modular graph network architecture with attention mechanisms applied to any desired components of a GNN model. To the knowledge of the authors, this is the first




Figure 1: General structure of an LSTM network. The *Forget*, *Input* and *Output* notation aims to illustrate how the different gating mechanisms operate.




Figure 2: Single-layer bidirectional architecture, with LSTM units parametrised by  $\theta_1$  and  $\theta_2$ .

paper that proposes a framework where attention can be applied to any major operation in a GNN.

- Attention weights are investigated to assess the potency of the proposed GNN framework.

This study shows the advantages of a modular GNN-attention framework. By analysing various DL-based models, the study also aims to compare the performance of different techniques and show new ways in which DL can be applied to wind power and wake modelling.

## 2. Theory

### 2.1. Bidirectional LSTM

Recurrent Neural Networks (RNN) represent the quintessential DL model for sequence analysis. RNNs recurrently apply a shared function over the inputs to learn temporal dependencies, but the basic RNN architecture is prone to vanishing or exploding gradients, which makes it difficult for the network to learn dependencies for long input sequences. Long short-term memory (LSTM) cells employ a gated recurrent unit with skip-connections, which allow gradients to flow for many time-steps, combating this problem [30]. A general illustration of the LSTM block can be seen in Fig. 1, comprised of input, forget and output gates. The forget gate allows the unit to restrain what information to be kept in the memory,  $s$ . The input gate encodes the current input,  $x_t$ , and controls its contribution to the memory, while the output allows the network to focus on relevant parts of the current state to compute the output,  $h$ .

Since inputs are processed sequentially, the RNN architecture does not take future context into consideration when making current predictions. A bidirectional LSTM (BLSTM) could be used, where a second network simultaneously processes the input sequence in reverse, before outputs from both networks are used to compute the final outputs, as visualised in Fig. 2 [31].

### 2.2. Attention

Learning to focus on relevant parts of the input, attention aims to allocate computational resources wisely and in some applications improve the interpretability of the models. Suppose an input sequence  $x_1, x_2, \dots, x_L \in \mathbb{R}^n$ . Each input,  $x_i$ , is matched with all inputs by computing attention scores, using some shared attention mechanism,  $a$ , often a single- or multi-layered perceptron [29, 32]. A softmax function normalises the attention scores to generate probabilities:

$$\alpha_{ij} = \frac{\exp(a(x_i, x_j))}{\sum_{k=1}^L \exp(a(x_i, x_k))}. \quad (1)$$

Finally, the updated values,  $z_i \in \{z_1, z_2, \dots, z_L\}$ , are obtained by taking a weighted sum of the inputs, where a value function,  $F$ , is used to transform the inputs, along with a non-linearity,  $\sigma$ :

$$z_i = \sigma\left(\sum_{j=1}^L \alpha_{ij} F(x_j)\right). \quad (2)$$

Multi-head attention has also been found beneficial for many applications, where the process in eq. (1) and (2) is repeated, and their outputs concatenated [33].

### 2.3. Graph Neural Networks

Even though CNNs are excellent at analysing image data, a fundamental limitation of these models, is that these networks cannot naturally adapt to irregular domains. Graph Neural Networks (GNN) aim to combat this by creating a more general framework to analyse problems thought to have important relational structures, but lack an inherent grid-like ordering. The general concepts of GNNs will now be presented, based on the framework outlined in [34].

A graph can be defined as a tuple containing global, edge and node features,  $G = (u, E, V)$ .  $E$  and  $V$  are the sets containing the edge and node features  $e_{ij} \in E$  and  $v_i \in V$ .  $v_i$  are attributes specific to a node  $i$ , while  $e_{ij}$  are attributes associated with ‘sending’ from node  $i$  to  $j$ , such as the distance between two nodes. There is only a single global feature for each graph, stored in  $u$ . Similar to the stacking of layers in CNNs, GNNs employ a series of graph blocks which extract relevant features sequentially. Global, edge and node features are updated using three functions within each block, to perform per-edge, per-node and global updates, as:

$$e'_{ij} = \phi^e(e_{ij}, v_i, v_j, u) \quad (3)$$

$$v'_j = \phi^v(v_j, \bar{e}'_j, u) \quad (4)$$

$$u' = \phi^u(\bar{e}', \bar{v}', u), \quad (5)$$

where  $\phi^{(\cdot)}$  represent the update functions, here assumed to be MLPs,  $(\cdot)'$  indicate updated and  $(\bar{\cdot})$  aggregated features. Updated features are fed as inputs to the next graph block. Three aggregation functions map edges to nodes, edges to globals and nodes to globals as follows:

$$\bar{e}'_j = \rho^{e \rightarrow v}(E'_j), \text{ where } E'_j = \{e'_{ij} | \forall i \in R_j\} \quad (6)$$

$$\bar{e}' = \rho^{e \rightarrow u}(E'), \text{ where } E' = \{e'_{ij} | \forall (i, j) \in U\} \quad (7)$$

$$\bar{v}' = \rho^{v \rightarrow u}(V'), \text{ where } V' = \{v'_i | \forall i \in I\}. \quad (8)$$

$\rho^{(\cdot) \rightarrow (\cdot)}$  are the different aggregation operations, e.g. mean or sum.  $R_j$ ,  $U$  and  $I$  are index sets containing indices for all nodes sending to node  $j$ , all edge indices  $(i, j)$  and all node indices, respectively. These sets define the connectivity of the graph. A visualisation of an edge-to-node sum aggregation,  $\rho^{e \rightarrow v}$ , can be seen in Fig. 3.

## 3. Methodology

### 3.1. Data Description

The wind farm data was simulated in FLORIS [21], using a Gaussian-based wake model [35]. A range of random farm configurations were simulated for different wind speeds and directions, varying the layout and number of turbines from 28 to 150. 18,280 different scenarios were simulated by varying the farm layout, number of turbines and wind conditions, resulting in 1,200,862 power values for individual turbines. All data was transformed by scaling features within  $(0, 1)$ , before being split into test, train and validation sets. The test data was obtained by randomly sampling 20% of the original data. The remaining 80% of the data was partitioned into training and validation sets, using the same 80-20% split.




Figure 3: Visualisation of the  $\rho^{e \rightarrow v}$  operation for an arbitrary graph using summation aggregate function.




Figure 4: Visualisation of the two upstream neighbours for a turbine  $i$  for a particular wind direction, with the relevant wake distances and angles shown.

### 3.2. MLP and BLSTM Models

The first two DL-based wake models were trained to predict the power production for individual turbines, with information on the locations of upstream neighbours and the wind speed. Turbines that were within  $\pm 30^\circ$  (i.e.  $|\alpha| < 30^\circ$  with reference to Fig. 4) were considered upstream neighbours. Depending on a turbine's location within a farm and the wind direction, its number of upstream neighbours would vary for different time-steps, resulting in variable length inputs. It should be noted that even though we refer to 'time' here, temporal correlations were not considered, but used to indicate the different inputs with distinct wind conditions and farm layouts. Considering for example a turbine,  $i$ , situated downstream of two turbines,  $j$  and  $k$ , for a particular wind direction (depicted in Fig. 4), the BLSTM model would look as:

$$\hat{P}_t^{(i)} = \mathbf{h}(ws_t, \mathbf{g}(d^{ij}, \sin(\alpha_t^{ij}), \cos(\alpha_t^{ij}), d^{ik}, \sin(\alpha_t^{ik}), \cos(\alpha_t^{ik}))), \quad (9)$$

where  $\hat{P}_t^{(i)}$  is the predicted power for turbine  $i$  at time  $t$ ,  $ws_t$  the free stream wind velocity,  $d^{ij}$  the distance between turbines  $i$  and  $j$ , and  $\alpha_t^{ij}$  the angle of  $j$  relative to the wind direction.  $\mathbf{h}$  and  $\mathbf{g}$  are functions represented by MLP and BLSTM networks, respectively. The MLP model is very similar to eq. (9), but with padded inputs to ensure constant length inputs and without the transform  $\mathbf{g}$  represented by the BLSTM network.

Additionally, two baseline models were used to compare against, where one would be trained to predict total farm power and another to predict power for individual turbines. The inputs to these models only contained information on the total number of turbines in the farm and the free stream wind velocity. Since the baseline MLP models did not have any information on farm layout or upstream neighbours, these models should not be able to accurately predict wake losses. If the other models were able to achieve significantly lower prediction errors than the baselines, it would be concluded that they captured, at least part of, the wake losses.

### 3.3. GNN Model with Attention

For the GNN models, a graph was created for each wind condition and farm layout, where turbines are represented by nodes in the graph, with free-stream wind velocity,  $ws_t$ , as input node and global features. Interactions between nodes were encoded in the edges, with input features as;  $d_{ij}$ ,  $\sin(\alpha_t^{ij})$  and  $\cos(\alpha_t^{ij})$ . An edge exists, sending from turbine  $i$  to  $j$ , if turbine  $i$  is upstream of turbine  $j$ . The number of edges and connectivity of a graph would therefore change for different wind directions. The GNN produce updated graphs, trained to predict power for individual turbines embedded in the updated node features, and total farm power in the global.

Graph Attention Networks (GAT) were first introduced by Veličković et al. [29], where attention was used to update nodes in a graph. Edge features might also be important for certain applications, and some studies have therefore tried to extend the concept of GATs to facilitate such features [26]. Incorporating attention into the model, our first aim was to be able to elegantly handle a variable number of upstream neighbours. Secondly, the attention weights can be seen as the relative importance placed on different parts of the input by the network. For this particular application, the hope was that the network would learn weights to focus on the most influential upstream turbines to determine wake losses, which should hopefully align with some physical intuition about wind farm dynamics and enhance interpretability.

The proposed method is made up of a modular architecture, where attention can be applied to the different update or aggregate operations outlined in Sec. 2.3. A visualisation of the proposed graph block can be seen in Fig. 5, where the original architecture without attention can be retrieved if MLPs and desired aggregate functions are used in place of the highlighted update and aggregation boxes. Depending on the application, the framework allows for various graph attention architectures to be tested and attention to be applied to the relevant parts of the network for a particular task. For instance, if attention was only to be applied in the node update function in Fig. 5, one would retrieve a similar architecture as the GAT [29].

Considering the **NodeAttn** block in Fig. 5, the network will attend to all nodes sending to a target node. For a turbine  $j$ , this means computing weights for all  $i \in R_j$ . Such an operation can be considered as masked attention, since the model does not compute attention weights over the entire input, but only for the nodes that are in the neighbourhood of  $j$ . For this particular study, it was decided to use scaled dot product attention, as in [33]:

$$\alpha_{ij} = \frac{\exp((K^T[u \parallel v_i \parallel e'_{ij}])^T Q^T[\bar{e}'_j \parallel v_j]/\sqrt{d})}{\sum_{\forall k \in R_j} \exp((K^T[u \parallel v_k \parallel e'_{kj}])^T Q^T[\bar{e}'_j \parallel v_j]/\sqrt{d})}, \quad (10)$$

where  $\parallel$  denotes the concatenation operator.  $Q$  and  $K$  are weights for linear transformations to compute  $d$ -dimensional keys and queries. However, it should be noted that instead of using the keys and queries mapping in scaled dot product attention, an MLP or other function could still be used, as in eq.(1). The outputs from the **NodeAttn** block,  $z_j$ , were then computed as

$$z_j = \sigma\left(\sum_{\forall i \in R_j} \alpha_{ij} F^T v_i\right). \quad (11)$$

For the **EdgeAttn**, the same operations were performed, but now attending over the edges that send to a target edge. To obtain a representation where edges send to other edges via

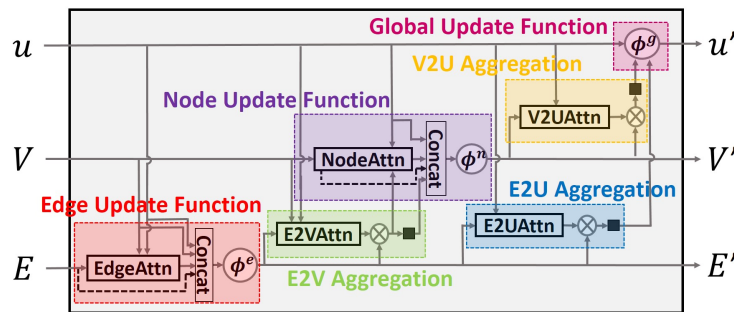


Figure 5: The proposed attention-based graph block. To retrieve the original graph block, without attention, all attention blocks would be omitted, denoted by **-Attn**.  $\blacksquare$  and  $\phi$  denote the aggregation operation and transforms represented by MLPs, respectively.




Figure 6: Edge-transformed graph representation. The same operations as in the **NodeAttn**, could then be applied in the **EdgeAttn** in Fig.5.

Table 1: MAE Test Results.

Model	MAE Turbine	MAE Farm
BS_Farm	-	0.015820
BS_Turb	0.052266	-
MLP	0.005652	0.000790
BLSTM	0.005331	0.000720
O-Graph	<b>0.001426</b>	0.000511
N-Graph	0.001452	0.000526
F-Graph	0.001441	<b>0.000495</b>

nodes, the input graph was transformed to a temporary edge-graph, where an edge  $e_{ij}$  sends to  $e_{jk}$  via the node  $v_j$ . The edge-graph mapping was inspired by the works in [26], producing a new graph with edge features mapped to nodes. Different to [26], the proposed mapping works for directed graphs, as visualised for an arbitrary graph in Fig. 6. The same attention mechanism used for the **NodeAttn** could then be used for the **EdgeAttn** operation.

Considering the attention operations for the aggregations, similar operations were used, but with different neighbourhoods. For the **E2VAttn**, masked attention was performed over all edges sending to the same node. Since global features are shared across a single graph, **V2UAttn** and **E2UAttn** perform attention over the full input, i.e. not masked, using all nodes or edges.

## 4. Results and Discussion

### 4.1. Wind Power Prediction Accuracy

To assess the performance of the proposed models in predicting wind power production, the Mean Absolute Error (MAE) metric was used:

$$\text{MAE} = \frac{1}{T} \sum_{t=1}^T |\hat{y}_t - y_t|, \quad (12)$$

where  $y$  are true and  $\hat{y}$  predicted values, for  $T$  samples. Two losses were computed for predicting individual turbine and total farm powers, denoted "MAE Turbine" and "MAE Farm" in Table 1. The MLP and BLSTM models improved the accuracy ten- and twenty-fold for individual turbine and farm predictions, respectively, compared to the baseline models, denoted "BS\_" in Table 1. The BLSTM outperformed the MLP model, with an approximate 5 and 10% improvement for the turbine and farm predictions. Besides a slightly inferior performance, a drawback of the MLP model was that it would not be able to make predictions on wind farms with a greater number of turbines than what was present in the training data. Nevertheless, comparing the models to the baselines, it was concluded that both architectures were able to capture wake losses, though the BLSTM seemed potentially better suited for the variable length inputs.

Considering the graph models in Table 1, all achieved significantly better prediction performance than the other models. Here, O-Graph is the vanilla GNN model without attention. F-Graph employ all attention modules in Fig. 5 for the first graph block layer, while the N-Graph only use attention for the edge-to-node aggregation and the node-update function. The reason for including both the F- and N-Graph models, was to demonstrate the modularity of the proposed graph attention framework. The O-Graph model reported the smallest MAE of 0.00143 for individual turbine predictions, while the F-Graph yielded the best results for total farm predictions, with an MAE of 0.000495. The prediction performance of the N-Graph model was close to the other graph models, with MAEs of 0.00145 and 0.000526. Despite slightly different prediction performances between the three graph models, we argue that the comparable results demonstrate the functioning of the proposed framework in Fig. 5.



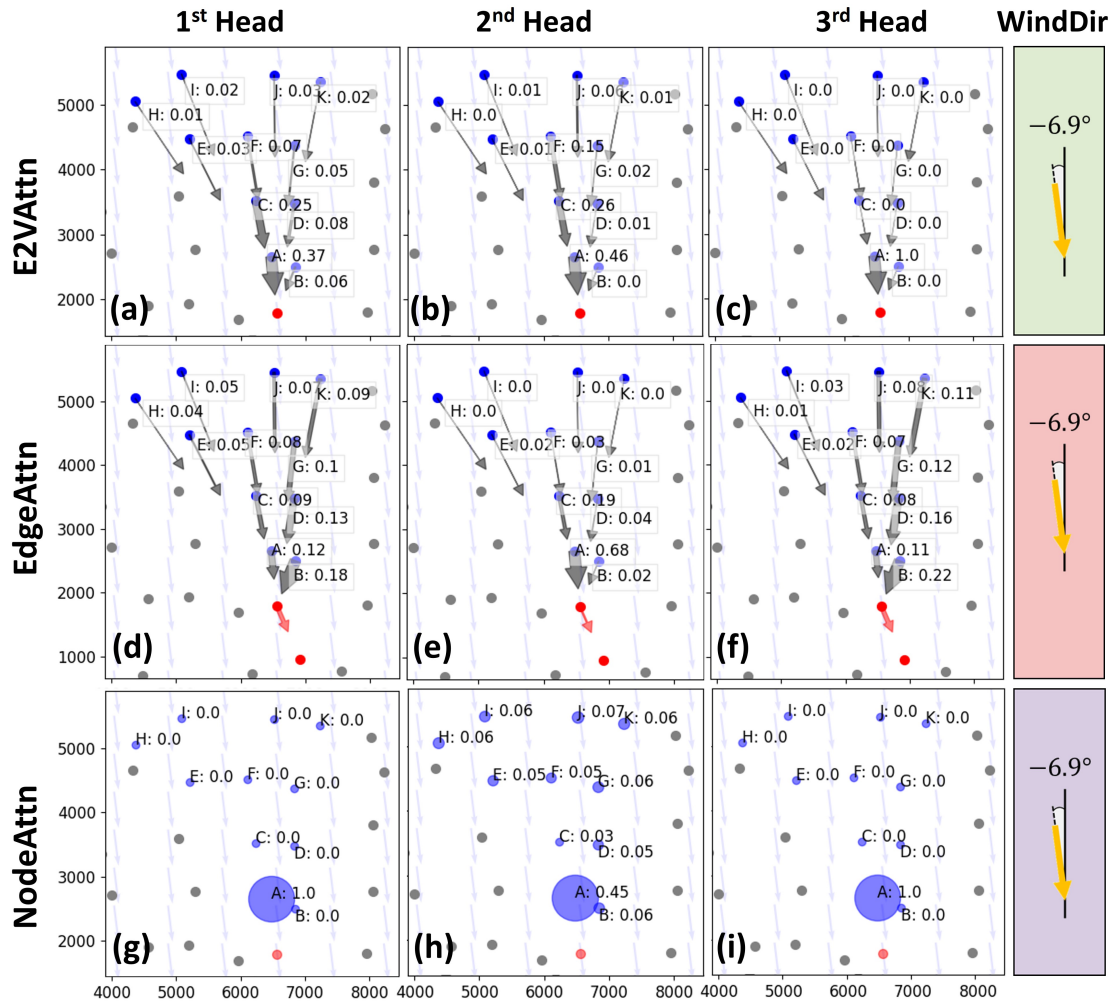


Figure 7: Attention weights from the F-Graph model for a randomly chosen farm configuration and wind condition. Red colouring represents a target wind turbine or edge, while blue represents upstream neighbours. **E2VAttn**, **EdgeAttn** and **NodeAttn** refers to the blocks in Fig. 5.

#### 4.2. Analysis of Attention Weights

Each of the five attention operations in the F-Graph model used three heads, computing three sets of attention weights for each operation. Fig. 7 shows the attention weights from the edge-to-node aggregation, edge- and node-update functions, for a randomly selected farm layout and wind condition during testing, while global attention weights were excluded from the plots due to the size of the farm. Furthermore, another scenario was also investigated, shown in Fig. 8, in which the turbines were situated in a two-dimensional grid, with the wind direction set to zero degrees. The dots in the figures represent turbine locations, with red being the target and blue its upstream neighbours for the particular wind conditions. Arrows represent edges, sending from the upstream turbines to the target. The width of the arrows and size of the circles visualise the relative weighting computed by the **E2VAttn**, **EdgeAttn** and **NodeAttn**, respectively, given by the corresponding numbers in Fig. 7 and 8.

Considering the edge-to-node attention weights, it was seen that the heads seemed to focus on slightly different aspects of the input. The main focus was placed on the edge sending from the turbine directly upstream and in close vicinity to the target, i.e. the edge sending from turbine

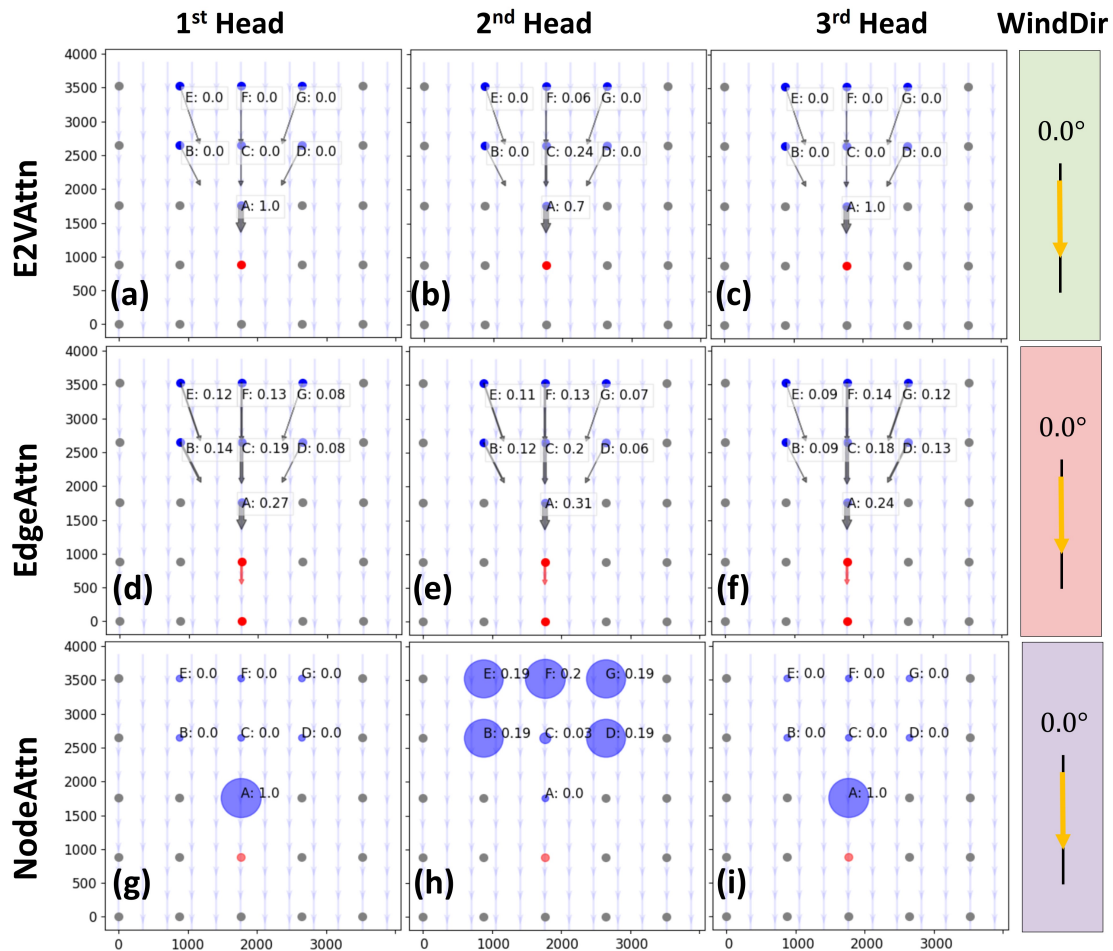


Figure 8: Attention weights from the F-Graph model for an ordered wind farm layout with 100% wake overlap. Red colouring represents a target wind turbine or edge, while blue represents upstream neighbours. **E2VAttn**, **EdgeAttn** and **NodeAttn** refers to the blocks in Fig. 5.

A in Fig. 7 and 8 a-c. From Fig. 7, it was evident that the most important edges were not necessarily those sending from the turbines closest to the target, as hardly any weighting was placed on the edge from turbine B. This was interesting, as it showed that the networks were able to learn more complex behaviours, than simply how close turbines were to the target. Looking at the second head **E2VAttn**, incrementally lower attention was placed on edges moving directly upstream from the target, as 0.46, 0.26 and 0.15, for edges A, C and F in Fig. 7. Hardly any weighting was placed on other edges, confirmed by the weights in Fig. 8. Intuitively, the second head might be thought to help determine wake losses by capturing changes when moving directly upstream from the target wind turbine.

Moving to the **EdgeAttn**, the second head weights showed similar patterns, with weights tapering off when moving upstream from the sending target turbine. Since this attention mechanism was used to update edges, there was effectively two target turbines, one sending and one receiving. Looking at the first and third head weights in plots d and f of Fig. 7, weights were more evenly distributed, but still with the most focus placed on turbines upstream of the target edge. Different to the **E2VAttn**, more attention was placed on edges sending from turbines moving clock-wise away from the upstream line, i.e. B, D, G and K. In particular, while the edge from turbine B was negligible for the **E2VAttn**, it was highly significant for

the **EdgeAttn** in Fig. 7. This behaviour was likely due to the location of the receiving target turbine, with the turbines receiving significantly more weighting since they were influential to the power of the receiving target turbine. Observing that the location of the receiving turbine was important and the learned behaviour for the first two attention networks were different was interesting, as it showed that the network could learn complicated structures in the data.

Again, attention was mainly placed on the turbine directly upstream of the target for the **NodeAttn** weights. Two of the heads placed the entire weighting on this turbine, while the second head also attended to the other turbines. For the ordered farm layout in Fig. 8, the second head did not place any weighting on turbine A. No difference was evident from the first and third head weights, which could imply two things; either, that one head focuses on dependencies that were not present for the particular layouts and wind conditions in Fig. 7 and 8, or that one of the three heads for the **NodeAttn** were redundant. The distributed weighting for the second head might indicate that an averaging of the features for upstream turbines provides important information to determine the power production for a target turbine.

Even though the attention-based models were not able to significantly improve on MAEs, investigating the attention weights showed that the proposed modules were able to learn meaningful dependencies in the data. An important finding, was that the attention mechanisms did not learn a simple average of all the input features, as this would mean that the mechanisms would replicate a similar behaviour to the O-Graph model. Although we were able to deduce some physical intuition about what the F-Graph model learned, the conclusions drawn were quite simple. Nevertheless, one might expect that attention mechanisms could be used more extensively to aid in applications where we want to learn about and understand more complicated physical properties or patterns, that are normally inaccessible to the users of DL models.

Generally, it was observed that the model seemed exceedingly confident, often placing almost all the weighing on a small subset of the inputs, entirely disregarding others. For the **E2VAttn** weights, this was evident, as turbine B received almost no weighting in Fig. 7. One could therefore argue that the model might be too sensitive, as slight offsets in the inputs could result in the model focusing on the wrong features. Even though the high confidence allowed for perspicuous interpretations of the results, it was thought that this was an inherent reflection of the data used. Wind farm data simulated in FLORIS lack uncertainty, which would be a prominent feature of any real dataset, due to sensor drifts, non-stationary atmosphere conditions and non-uniform wind field across a farm. If the models were instead trained on real site measurements, it is likely that attention weights might be more evenly distributed, with less values close to 0 or 1, as information from only a small number of upstream turbines might be a poor reflection of global characteristics. For future research, it would be particularly interesting to investigate how the models operate when trained on real farm measurements or data simulated using more accurate numerical methods that considers non-stationary atmosphere. Utilising real datasets, it is also thought that the superior performance of the GNN models, with regards to MAEs, might slightly curtail, as other features than interaction losses might become more pronounced. Nevertheless, for this preliminary study, the rapid simulations and clean data obtained using FLORIS was important as it allowed for better analysis of the attention weights, enabling investigation into whether the proposed framework behaved as desired.

## 5. Conclusion

In this study, it was seen that graph models can outperform other DL architectures in predicting turbine powers using data simulated in FLORIS. A novel framework was proposed to apply attention to all the relevant update- or aggregate-functions in a GNN model and the framework achieved results on-par with vanilla GNN models. Through analysis of the attention weights, it was seen how the networks learned interactions between turbines and were able to produce results that aligned well with some physical intuition about wake losses. For future research,

it will be interesting to investigate how attention mechanisms can help to learn more complex physical behaviours or improve prediction performance of graph models applied to wind-based technologies. Furthermore, the modular framework allows full flexibility for the user and it is a hope that it can aid in the development of graph-attention models for a range of different research applications in the future.

### Acknowledgements

This work was in part financed by the research project ELOGOW (Electrification of Oil and Gas Installation by Offshore Wind). ELOGOW (project nr: 308838) is funded by the Research Council of Norway under the PETROMAKS2 program with financial support of Equinor Energy AS, ConocoPhillips Skandinavia AS and Aibel AS. The authors would also like to thank the developers of FLORIS [21] for enabling wind farm simulations used for this study.

### References

- [1] GWEC 2021 Gwec — global wind report 2021 Available Online: <https://gwec.net/global-wind-report-2021/> (accessed 13-Dec-2021)
- [2] Jensen N O 1983
- [3] Sanderse B, Van der Pijl S and Koren B 2011 *Wind Energy* **14** 799–819
- [4] Manobel B, Sehnke F, Lazzús J A, Salfate I, Felder M and Montecinos S 2018 *Renew. Energy* **125** 1015–20
- [5] Yan C, Pan Y and Archer C L 2019 *Wind Energy* **22** 1421–32
- [6] Ti Z, Deng X W and Yang H 2020 *Appl. Energy* **257** 114025
- [7] Tastu J, Pinson P, Trombe P J and Madsen H 2013 *IEEE Trans. Smart Grid* **5** 480–89
- [8] Grover A, Kapoor A and Horvitz E 2015 *Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Min.* 379–86
- [9] Ghaderi A, Sanandaji B M and Ghaderi F 2017 *arXiv preprint arXiv:1707.08110*
- [10] Hanifi S, Liu X, Lin Z and Lotfian S 2020 *Energies* **13** 3764
- [11] Hong Y Y and Rioflorida C L P P 2019 *Appl. Energy* **250** 530–9
- [12] Memarzadeh G and Keynia F 2020 *Energy Convers. Manag.* **213** 112824
- [13] Kou P, Wang C, Liang D, Cheng S and Gao L 2020 *IET Renew. Power Gener.* **14** 2416–28
- [14] Zhu Q, Chen J, Shi D, Zhu L, Bai X, Duan X and Liu Y 2019 *IEEE Trans. Sustain. Energy* **11** 509–23
- [15] Chen R, Liu J, Wang F, Ren H and Zhen Z 2020 982–87
- [16] Khodayar M and Wang J 2018 *IEEE Trans. Sustain. Energy* **10** 670–81
- [17] Stańczyk T and Mehrkanoon S 2021 *arXiv preprint arXiv:2101.10041*
- [18] Bleeg J 2020 *JPCS* **1618** 062054
- [19] Yu M, Zhang Z, Li X, Yu J, Gao J, Liu Z, You B, Zheng X and Yu R 2020 *FGCS* **113** 145–57
- [20] Park J and Park J 2019 *Energy* **187** 115883
- [21] NREL 2021 Floris. version 2.4 URL <https://github.com/NREL/floris>
- [22] Bahdanau D, Cho K H and Bengio Y 2015 *3rd ICLR*
- [23] Fu X, Gao F, Wu J, Wei X and Duan F 2019 *ICDMW* 149–54
- [24] Li P, Wang X and Yang J 2020 *IET Renew. Power Gener.* **14** 297–304
- [25] Niu Z, Yu Z, Tang W, Wu Q and Reformat M 2020 *Energy* **196** 117081
- [26] Wang Z, Chen J and Chen H 2021 Egat: Edge-featured graph attention network *Int. Conf. Artif. Neural Netw.* (Springer) pp 253–264
- [27] Georgousis S, Kenning M P and Xie X 2021 *IEEE Access*
- [28] Lee J B, Rossi R A, Kim S, Ahmed N K and Koh E 2019 *TKDD* **13** 1–25
- [29] Veličković P, Cucurull G, Casanova A, Romero A, Liò P and Bengio Y 2018 *ICLR*
- [30] Hochreiter S and Schmidhuber J 1997 *Neural Comput.* **9** 1735–1780
- [31] Schuster M and Paliwal K K 1997 *IEEE Trans. Signal Process.* **45** 2673–2681
- [32] Xu K, Ba J, Kiros R, Cho K, Courville A, Salakhudinov R, Zemel R and Bengio Y 2015 *ICML* 2048–57
- [33] Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez A N, Kaiser L u and Polosukhin I 2017 *NIPS* **30**
- [34] Battaglia P W, Hamrick J B, Bapst V, Sanchez-Gonzalez A, Zambaldi V, Malinowski M, Tacchetti A, Raposo D, Santoro A, Faulkner R *et al.* 2018 *arXiv preprint arXiv:1806.01261*
- [35] King J, Fleming P, King R, Martínez-Tossas L A, Bay C J, Mudafort R and Simley E 2020 *Wind Energy. Sci. Discuss.* **2020** 1–22