Spatio-Temporal Wind Speed Forecasting using Graph Networks and Novel Transformer Architectures *

Lars Ødegaard Bentsen
Department of Technology Systems
University of Oslo
Kjeller, Norway
lars.bentsen@uio.its.no

Narada Dilp Warakagoda
Department of Technology Systems
University of Oslo
Kjeller, Norway

Roy Stenbro
Institute for Energy Systems (IFE)
Kjeller, Norway

Paal Engelstad
Department of Technology Systems
University of Oslo
Kjeller, Norway

Abstract

To improve the security and reliability of wind energy production, short-term forecasting has become of utmost importance. This study focuses on multi-step spatio-temporal wind speed forecasting for the Norwegian continental shelf. A graph neural network (GNN) architecture was used to extract spatial dependencies, with different update functions to learn temporal correlations. These update functions were implemented using different neural network architectures. One such architecture, the Transformer, has become increasingly popular for sequence modelling in recent years. Various alterations of the original architecture has been proposed to better facilitate time-series forecasting, of which this study focused on the Informer, LogSparse Transformer and Autoformer. This is the first time the LogSparse Transformer and Autoformer have been applied to wind forecasting and the first time any of these or the Informer have been formulated in a spatio-temporal setting for wind forecasting. By comparing against spatio-temporal Long Short-Term Memory (LSTM) and Multi-Layer Perceptron (MLP) models, the study showed that the models using the altered Transformer architectures as update functions in GNNs were able to outperform these. Furthermore, we propose the Fast Fourier Transformer (FFTransformer), which is a novel Transformer architecture based on signal decomposition and consists of two separate streams that analyse trend and periodic components separately. The FFTransformer and Autoformer were found to achieve superior results for the 10-minute and 1-hour ahead forecasts, with the FFTransformer significantly outperforming all other models for the 4-hour ahead forecasts. Finally, by varying the degree of connectivity for the graph representations, the study explicitly demonstrate how all models were able to leverage spatial dependencies to improve local short-term wind speed forecasting.

Keywords Spatio-Temporal Wind Forecasting · Multi-Step · Transformers · Graph Neural Networks

1 Introduction

In the context of the global climate debate, wind has emerged as a prominent renewable energy resource to accelerate the depletion of fossil fuel based energy production [1]. Nevertheless, in contrast to conventional power plants, wind resources are inherently intermittent in the short-term, which pose significant challenges for operators and grid planning [2]. To alleviate some of these and help facilitate large-scale adoption of wind power, accurate wind forecasting has become of critical importance.

Wind forecasting methods can be categorised as either physical or statistical. Physical models are based on detailed physical laws that model the atmosphere and typically aim to increase the resolution of coarse numerical
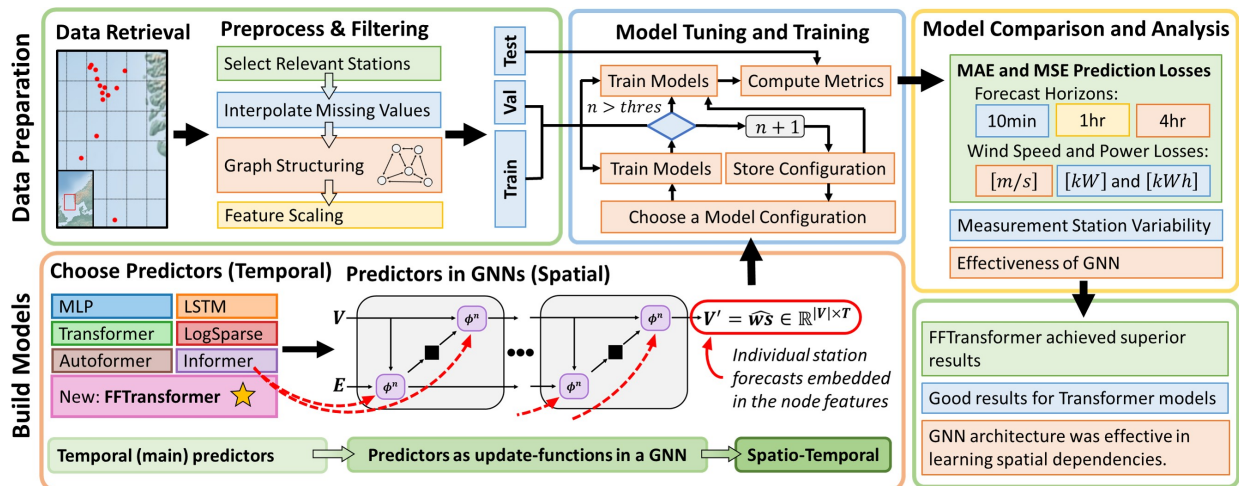

Figure 1: Graphical Abstract - Schematic illustration of the methodology and research outline.

weather prediction (NWP) models. A challenge with physical models is that they come at a very high computational cost, making them less viable for local short-term forecasting [3]. Statistical and machine learning (ML) methods on the other hand, leverage historical data to optimise model parameters. Even though the training of ML models might be a time exhaustive process, such models are very quick during inference, meaning that forecasts can be obtained in near real-time.

Autoregressive moving average methods, such as the autoregressive integrated moving average (ARIMA) model, are robust and easy to implement, making them popular for wind forecasting. Kavasseri *et al.* [4] studied *fractional*-ARIMA models to perform one and two day-ahead wind speed forecasts, managing to outperform both a persistence and an ARIMA model for four potential wind generation sites in North Dakota. The persistence model is a commonly used benchmark in wind-speed forecasting, where the forecasted values, $\hat{ws}_{t+1}$ are simply taken as the last recorded value ws_t, i.e. $\hat{ws}_{t+1} = ws_t$. Since wind speed time-series are characterised by both long-term trends and high frequency variation, Singh *et al.* [5] proposed the RWT-ARIMA model, combining the ARIMA model with wavelet transform (WT) to decompose the signal into multiple sub-series with different frequency characteristics. Various decomposition techniques have been studied for wind time-series, showing the potential benefits of introducing some signal decomposition into the forecasting models, such as complete ensemble empirical mode decomposition [6], variational mode decomposition [7] and wavelet packet decomposition [8, 9].

Support vector regressors (SVR) and K-nearest neighbour (KNN) algorithms have also been popular within wind forecasting [10, 11, 12]. The KNN-algorithm is based on finding similar points in the available data and can be fast in both training and testing. SVR have been shown to yield very good forecasting performance [11], but do not scale well for larger datasets, resulting in longer computation times [10].

Neural networks lie at the heart of modern ML and have become increasingly popular for wind speed forecasting, due to their ability of modelling non-linear relationships. Multi-Layer Perceptrons (MLP) have been successfully used both in isolation [13, 14] and in combination with other methods [15, 16]. Recurrent (RNN) and convolutional neural networks (CNN) represent the quintessential deep learning (DL) architectures for sequence modelling, and are typically favoured over MLPs for wind forecasting [17]. The long short-term memory (LSTM) unit is an alteration of the vanilla RNN architecture, where gating mechanisms and skip connections are introduced to mitigate the problem of vanishing or exploding gradients [18]. Li *et al.* [19] proposed a hybrid architecture, using empirical wavelet transform and the regularised extreme learning machine, together with an LSTM as the main predictor. Due to the recurrent architecture, the LSTM network relies on encoding all the relevant input information into a fixed-length memory cell, which can cause information loss and limit the network's ability to retain information across longer sequences. Within the context of natural language processing (NLP), the attention mechanism was introduced by Bahdanau *et al.* [20], to allow the networks to directly attend to previous hidden states according to their importance. Many studies have focused on integrating attention mechanisms with LSTMs to further help the models learn long-term dependencies that can improve forecasting performance [21, 22].

Oord *et al.* [23] proposed the WaveNet architecture for generating raw audio waveforms. The main ingredient of WaveNet is dilated causal convolution, which is a 1D convolutional operation where the causality ensures that the model cannot violate the sequence ordering, while the dilation increases the receptive field by skipping input values with a certain step. Dilated causal convolution is well suited for time-series modelling and has been successfully

deployed for some wind forecasting studies [24, 25]. Other popular DL-based architectures used for wind forecasting also include deep belief networks [26], RNNs with Gated Recurrent Units [27] and the N-BEATS model [28].

Building on the attention mechanism, the Transformer was proposed by Vaswani *et al.* [29], as a new sequence transduction model, particularly focused on NLP. The Transformer is fundamentally different from previous models in that it does not rely on recurrence or convolution, making it better at learning long-term dependencies. However, since the complexity scales quadratically with sequence length, various alterations of the original architecture have been proposed to alleviate the computational limitations and make the models better suited for time-series forecasting, such as the Longformer [30], FEDformer [31], Temporal Fusion Transformer [32], LogSparse Transformer [33], Informer [34], Reformer [35] and Autoformer [36]. Both [37] and [38] managed to outperform LSTM models for wind forecasting by using a Transformer, comprised of an encoder and decoder, as in [29]. A bidirectional LSTM-Transformer model achieved superior results compared to a gated recurrent unit (GRU) and an LSTM model in [39]. Wang *et al.* [40] proposed a model based on the Informer together with convolutional layers that extract temporal features at different frequencies, to forecast the average wind power over the next three hours. The Spatial-Temporal Graph Transformer Network (STGTN) extends the previous research by leveraging both spatial and temporal correlations within a wind farm, to more accurately forecast wind speeds at a turbine level, 10 minutes - 1 hour ahead [41].

Spatial dependencies can be important for improving meteorological forecasts, such as for wind. Some studies organise spatial data into an ordered grid, where the features for a particular physical location are assigned to a particular cell [42, 43, 44]. A CNN is then used to extract spatial features, while another network, such as a CNN [42] or LSTM [43], is used to learn temporal correlations. However, the strict ordering of the input data required for CNNs might not be able to effectively represent the underlying spatial relationships. Graph neural networks (GNN) can better facilitate arbitrary spatial ordering and have therefore been popularly used for spatio-temporal forecasting. Khodayar *et al.* [45] made forecasts at different wind sites simultaneously, where each site was represented by a node in an undirected graph, using a Graph Convolutional Network (GCN) and an LSTM to extract spatial and temporal features, respectively. Similarly, Stańczyk *et al.* [46] also used a GCN, but with a CNN to learn temporal dependencies. Instead of using static edge features, Wang *et al.* [47] constructed edge features based on the time-varying spatial correlation between wind sites and used a GCN for wind farm cluster power forecasting. Finally, the M2STAN model was also proposed for spatio-temporal multi-step wind power forecasting [48], which employs a Graph Attention Network (GAT) and a Bidirectional GRU for spatial and temporal correlation modelling, respectively,

along with a Transformer network for multi-modal feature fusion.

This paper focuses on spatio-temporal multi-step forecasting based on recent developments in DL. With multi-step forecasting, the study aims to output more informative time-series. Considering a scenario of forecasting one hour ahead, some studies give a single prediction for the average wind speed over the entire period. However, sometimes one might also be interested in the development of the wind within that hour, i.e. is the wind speed increasing or decreasing, when do the highest or lowest wind speeds occur and what are the expected peaks. Furthermore, since wind power is proportional to the wind speed cubed, $P \propto ws^3$, higher resolution wind speed forecasts are necessary to more accurately obtain the expected wind power for a particular time period. Bearing this in mind, this study decided to focus on multi-step forecasting with 10-minute time resolution for all forecasts of 10-minutes, 1-hour and 4-hours ahead. The main contributions of this paper can be summarised as:

- We show the effectiveness of a generic framework for multi-step spatio-temporal forecasting, with GNNs to capture spatial correlations and optional update functions to learn temporal dependencies, such as a Transformer or LSTM network.

- Test and compare the performance of different Transformer architectures for use in wind forecasting, namely the vanilla Transformer, LogSparse Transformer, Informer and Autoformer. This is the first paper to formulate many of these in a GNN setting and the first to apply such models to wind forecasting.

- We propose a new alteration of the Transformer architecture, namely the Fast Fourier Transformer (FFTransformer), and show its competitive performance in wind forecasting. The novel architecture is based on wavelet decomposition and an adapted Transformer architecture consisting of two streams. One stream analyses periodic components in the frequency domain with an adapted attention mechanism based on fast Fourier transform (FFT), and another stream similar to the vanilla Transformer, which learns trend components.

2 Theory

2.1 Multi-Layer Perceptron

MLPs are feed forward networks that learn weights, θ, which map the input to output, $y \approx f(x|\theta)$. A chain structure, where multiple layers are stacked, give depth to the model, as $\hat{y} = f^{(3)}(f^{(2)}(f^{(1)}(x|\theta_1)|\theta_2)|\theta_3)$, for a model with two hidden layers. To improve the learning ability of the model, non-linearities such as the ReLU or sigmoid functions, are applied to the neuron outputs. Optimal weights are determined by minimising a differentiable

Lars Ø. Bentsen *et al.*

loss function, using backpropagation, which will update network weights by propagating gradients of the weights with respect to the loss function, back through the network [49].

2.2 LSTM

Even though sequence analysis using DL have largely been dominated by RNNs, the original architecture is prone to exploding or vanishing gradients, resulting in significant information loss for longer sequences. The long short-term memory (LSTM) unit introduces gating mechanisms and skip connections to alleviate some of these shortcomings [18]. An illustration showing the internal workings of an LSTM unit is given in Fig. 2, where an input gate controls the contribution of the new input, x_t, to the memory, while the forget and output gates control what information to be kept in memory and encoded in the output, h_t, respectively.

2.3 Transformer Architectures

The Transformer architecture was proposed as a model for sequence analysis, without any recurrence or convolution, but instead based on the attention mechanism [29]. If we consider a time-series forecasting task with d input features to predict a single output feature, the Transformer model should take an input, $x^{(e)} \in \mathbb{R}^{S \times d}$, and produce an output, $y \in \mathbb{R}^{P \times 1}$, where, S, is the look-back window and, P, the forecasting horison. The original architecture consists of an encoder and a decoder, as shown in Fig. 4, where the encoder should encode a longer input into a hidden state representation for the decoder to decode. Inputs to the decoder, $x^{(d)} \in \mathbb{R}^{(L+P) \times d}$, are typically shorter than those to the encoder, containing the last L elements of the encoder inputs, where $L < S$, and some placeholders are used in place for the last P-indices, which correspond to the forecasting locations.

Inputs are first linearly transformed to E-dimensional space and then added some positional encoding, so that the model can make use of the sequence order, without any recurrence. Unless otherwise is stated, it will be assumed that the positional encoding added to the encoder and decoder inputs will be the sine-cosine positional encoding proposed in the original architecture [29]. The multi-head attention (MHA) block in the encoder employs full self-attention, where each attention operation can attend to the full input sequence. Scaled dot-product attention takes Q, K and V as inputs, which represent the queries, keys, and values, respectively. Dot products between keys and queries are computed, before being passed to a softmax function to obtain the attention weights, which are then multiplied by the values to produce the final outputs. This process can be summarised as

$$\text{Attn}(Q,K,V) = \text{softmax}(\frac{(QW^Q)(KW^K)^T}{\sqrt{d_k}})(VW^V), \quad (1)$$

where $W^{(\cdot)} \in \mathbb{R}^{E \times d_{(\cdot)}}$ are weights of the different linear transformations (LT), E is the hidden dimensionality and, d_k, the dimension of keys and queries. Performing multi-head attention, with h separate projections, allow the module to simultaneously attend to different information in the input series. Outputs from the MHA are then concatenated and linearly projected to produce a single output.

A residual connection and layer normalisation are applied to the outputs, before the signal is passed to an MLP, typically with a single hidden layer, which is applied to each position in the series separately. Multiple encoder lay-
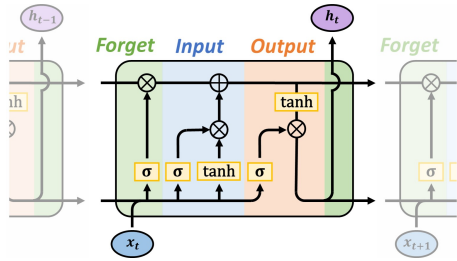

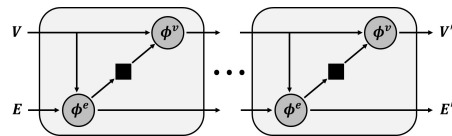Figure 2: Visualisation of an RNN with LSTM units having forget, input and output gating mechanisms.

Figure 3: Visualisation of the GNN architecture for graphs with edge and node features. $\phi^{(\cdot)}$ and ■ represent the update functions and edge-to-node aggregation, respectively.
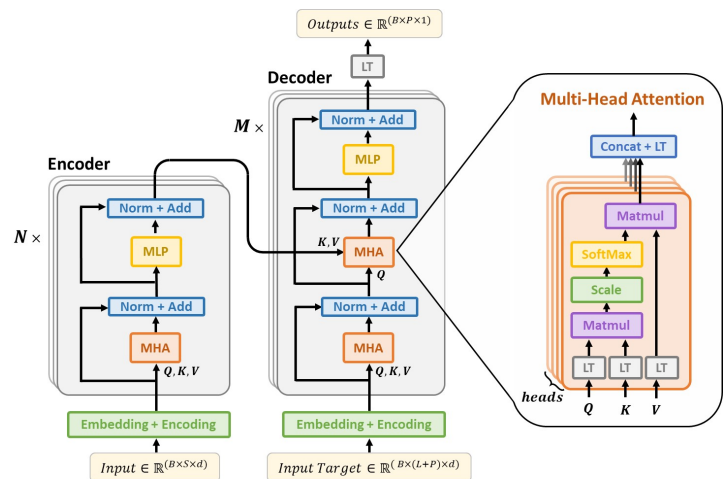

Figure 4: The Transformer architecture [29] in an encoder-decoder setting, adapted to facilitate forecasting rather than categorical outputs, with multivariate inputs and univariate outputs. B in the inputs and outputs refer to the batch size, for a single prediction $B = 1$.

4

ers, with different weights, are stacked to add depth to the model and hence a stronger function approximation ability.

The decoder follows almost the same architecture as the encoder, but with an additional MHA block referred to as the cross-attention, which precedes the MLP. The cross-attention module is the same as the other MHA blocks, but takes the encoder outputs as inputs to the keys and values. Additionally, the first MHA block in the decoder use masking to prevent information flow from subsequent positions.

Even though the attention mechanism alleviates the problem of information loss for longer sequences, by allowing every position to directly attend to all other positions, its complexity scales quadratically with sequence length. Furthermore, since the Transformer was originally developed for machine translation and other NLP tasks, various modifications have emerged, which aim to both combat the complexity limitations and further adapt the architecture to facilitate time-series forecasting problems. A short summary of some of these will now be given in the subsequent sections.

2.3.1 LogSparse Transformer [33]

The LogSparse Transformer introduce two novel alterations. Sparse attention, where each position is only allowed to attend to other positions with an exponential step size and itself, significantly reduces the space complexity. Since the point-wise attention operation described in Sec. 2.3 is insensitive to local context, causal 1D-convolution was used to compute keys and queries, instead of point-wise linear transformations. The modified transformation of keys and queries, which will here be referred to as convolutional attention, might be particularly advantageous for time-series forecasting, as local context could be very important for signals characterised by high-frequency fluctuations or noise.

2.3.2 Informer [34]

Instead of introducing sparsity through a fixed pattern decided by heuristic methods, the *ProbSparse* self-attention mechanism proposed for the Informer, introduces sparsity by locating the most dominant queries and only allows keys to attend to these. Dominant queries are taken as those that maximise a surrogate for the KL-divergence between a uniform distribution and the query's attention probability distribution. Furthermore, self-attention distilling in the encoder highlights dominant attention by halving cascading layer inputs through 1D-convolution and MaxPooling, which makes the model much more efficient for very long sequences.

2.3.3 Autoformer [36]

Unlike the Informer and LogSparse Transformer, Wu *et al.* [36] proposed significant alterations to both the overall Transformer architecture and the attention module, to better facilitate time-series forecasting. Instead of the scaled dot-product attention, the Autoformer introduces the Auto-Correlation mechanism, which use keys and queries to decide on the most important time-delay similarities through autocorrelation and a time-delay aggregation, which roll the series according to the selected time delays, before adding them together to produce the outputs. Different to point-wise attention, the Auto-Correlation mechanism finds dependencies based on periodicity and is specifically designed for time-series forecasting. Series decomposition is also applied after every Auto-Correlation and MLP module, using average pooling to decompose the signal into trend and seasonal components.

2.4 Graph Neural Networks

A graph can be defined as a tuple with node- and edge-specific features, $G = (V, E)$ [50]. For the spatial forecasting problem, node features, $v_i \in V$, can represent attributes associated with a particular measurement station, while edge features, $e_{ij} \in E$, contain information on the relationship between two nodes, i and j, such as the Euclidean distance between two stations. A GNN is comprised of stacked graph blocks, which perform per-edge and per-node updates in the following order:

$$e'_{ij} = \phi^e(e_{ij}, v_i, v_j) \tag{2}$$

$$v'_j = \phi^v(v_j, \bar{e}'_j), \tag{3}$$

where $\phi^{(\cdot)}$ are the update functions, which could be represented by a neural network such as an LSTM, MLP or Transformer. $(\cdot)'$ and $(\bar{\cdot})$ represent updated and aggregated features, respectively. Aggregated edge features, $\bar{e}'_j$, are computed using an aggregation function, $\rho^{e \to v}$, as

$$\bar{e}'_j = \rho^{e \to v}(E'_j), \text{ where } E'_j = \{e'_{ij} | \forall i \in R_j\}, \tag{4}$$
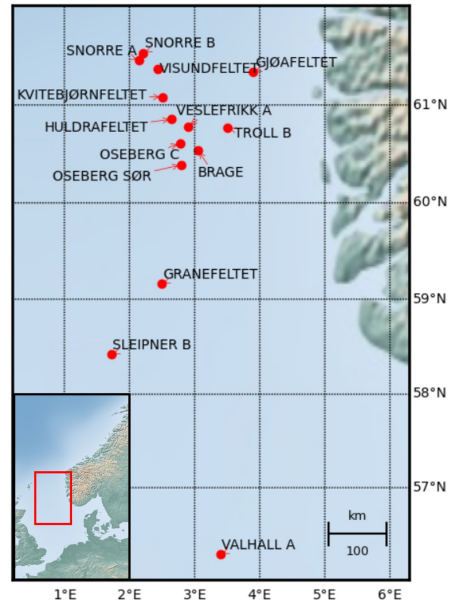

Figure 5: The 14 measurement stations in the North-Sea used to construct the dataset.

which could for example be a sum or mean operation. R_j is an index set containing all nodes sending to a node j, and define the connectivity of the graph. A visualisation of the GNN architecture is given in Fig. 3.

3   Methodology

3.1   Dataset

Due to the future potential for off-shore wind energy [51], this study decided to focus on off-shore wind speed forecasting, using meteorological measurements recorded on the Norwegian continental shelf. The data was made available by the Norwegian Meteorological Institute and is openly available through the Frost API[1]. Ten-minute averaged measurements on air temperature, air pressure, dew point, relative humidity, wind direction and speed and maximum wind speed in the 10-minute interval were used as input features to forecast the wind speed time-series. Wind direction was decomposed into its sine and cosine components in order to fully capture the circular characteristics. For every time-step, we would therefore have a feature matrix, $f_t \in \mathbb{R}^{N \times 8}$, corresponding to the eight recorded measurements for N available stations. The forecasting problem then becomes, $\hat{V}_t, \hat{V}_{t+1}, ..., \hat{V}_{t+P} = F(f_{(t-1)}, f_{(t-2)}, ..., f_{t-T})$, where F is the model, P the number of future time-steps to forecast, $\hat{V}_t \in \mathbb{R}^{N \times 1}$ the predicted wind speeds at time, t, and T is the look-back window, i.e. the number of previous time-steps used to make the forecasts. For the period between June 23, 2015, and February 28, 2022, 14 out of 25 available stations had some periods with measurements on all the relevant meteorological variables, and are shown in Fig. 5.

The first 60% of the data was used for training, whereas the following 20% and remainder were used for validation and testing, respectively. If measurements for a single time-step were missing for a particular station, linear interpolation was used to fill the missing entries. However, if there were consecutive time-steps missing, the station would not be considered for these periods. For different

periods, there would be a variable number of stations that had available data, meaning that the models should be able to take any subset of the 14 stations as input.

3.2   Fast Fourier Transformer

Since many time-series, such as wind measurements, are characterised by both trend and periodic components, we propose the Fast Fourier Transformer (FFTransformer), based on multilevel discrete wavelet decomposition (MDWD) and an adapted attention mechanism, named FFT-Attention. Each DWD layer decomposes a signal into high- and low-frequency components, referred to as the detail and approximate signals, respectively. In a multi-layer setting, the approximate component from one layer is fed as input to the next, resulting in multiple detail and a single approximate signal for the final outputs, as shown in Fig 6d, where 'D1', 'D2' and 'D3' are the detail signals and 'A4' the approximate signal. The decomposition of a wind speed time-series is shown in Fig. 6b, where inverse MDWD is applied independently to the outputs from the MDWD. Each input feature is decomposed separately, using the Daubechies wavelet Db4, which has been shown to be suitable for wind forecasting [52]. By applying fast Fourier transform (FFT) to each time-series in Fig 6b, it can be seen in Fig. 6c that detail (D1, D2, D3) and approximate (A4) components yield very different frequency characteristics. The four detail components exhibited clear periodic information, with peaks at different frequencies, while the approximation, 'A4', contained most of the low frequency trend information, clearly visualised in Fig. 6a.

Because of the different frequency characteristics, the FFTransformer is comprised of two streams, one which analyse signals with clear periodicity and another that should learn trend components, i.e. the detail and approximate signals in Fig. 6b, respectively. As shown in Fig. 7, the right-hand stream in both the encoder and decoder were exactly the same as for the original Transformer in Fig. 4. However, to better facilitate the analysis of periodic signals in the left-hand stream, we introduce the FFT-Attention mechanism, which performs attention in the frequency domain. In particular, as the first stage of FFT-Attention,

---

[1]`https://frost.met.no/index.html`

(a) Approximation and Input   (b) Inverse MDWD   (c) ABS(FFT(Inverse MDWD))   (d) MDWD Architecture
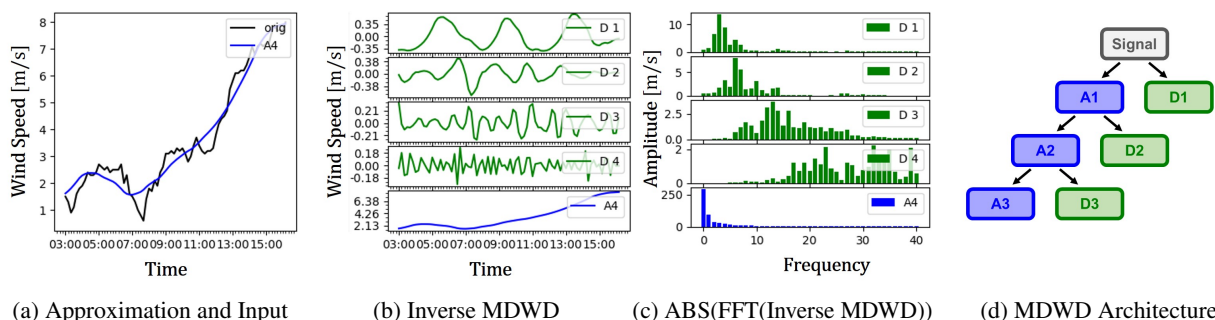
Figure 6: MDWD applied to a wind speed time-series. The FFT outputs clearly exhibit the different frequency characteristics of the sub-components and plotting the approximation together with the input signal visualise how this component retains the trend information.
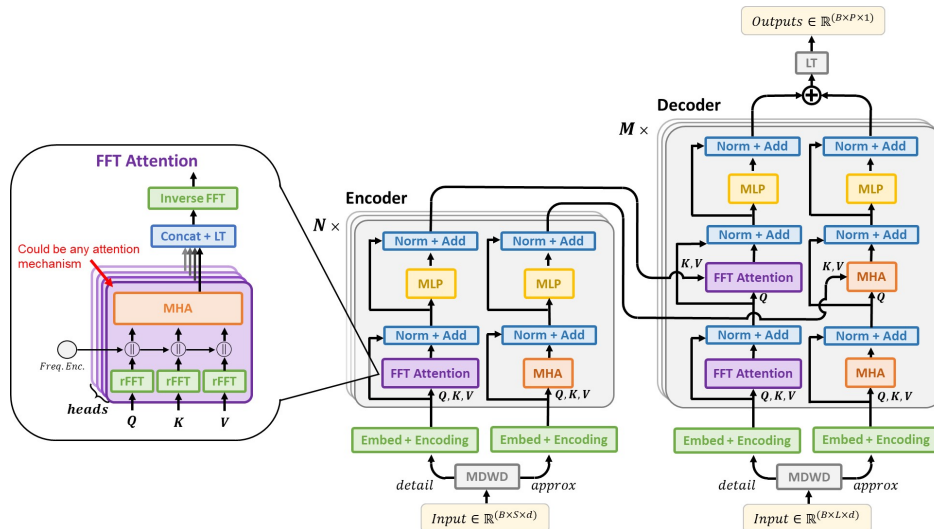
Figure 7: FFTransformer in an encoder-decoder setting. Two streams focus on trend and periodic signal components separately. The FFT-Attention is employed for the periodic analysis in the frequency-domain, using an MHA mechanism together with FFT.

FFT is applied to the key, query and value inputs. Real and imaginary components of the FFT outputs are concatenated with the frequency values, to provide information on the corresponding frequency for the values in a particular position, similar to the motivation behind positional encoding, before being fed to an MHA block. Outputs from the FFT-Attention module are then concatenated and projected, as for the other attention mechanisms, and finally, inverse FFT return values to the time domain. Any attention mechanism could be used in place of all the MHA blocks in Fig. 7, such as the *ProbSparse* or convolutional attention. The final network outputs from both streams (FFT Attention and Trend) are at last added and linearly transformed to produce the predictions.

After some experimentation, it was found that not adding temporal or positional encoding to the detail components (i.e. inputs to the FFT-Attention) yielded the best results. This seemed sensible, as adding positional encoding in the time-domain would not translate to the frequency domain. Instead, the concatenation of frequency values to the FFT outputs, served somewhat the same purpose as the positional encoding did for the trend components. Nevertheless, since this method was fairly trivial, further research studying better encoding strategies for the frequency domain could be desirable.

3.3 Spatio-Temporal Framework

All models were constructed in an encoder-only setting, i.e. without decoders, and used as update functions, $\phi^{(\cdot)}$, in two-layer GNNs. To facilitate the GNN framework, the dataset was constructed as graphs. Measurement stations were represented by nodes, with the historical time-series of the eight meteorological variables assigned to the input node features. Since there would be a variable number of

available measurement stations at different times, the input graphs would not be the same at different times, but with a variable number of nodes. This was desirable, as it meant that we did not have to discard the data for all stations or interpolate missing values for a particular time-interval where a few stations had missing data, as would be the typical case for a CNN. Considering a node, i, its input features were $v_i \in \mathbb{R}^{1\times(S+P)\times 8}$, where S is the historical look-back window, i.e. the number of previous time-steps used to predict the next P time-steps. Placeholders were used for the last P indices in the inputs, corresponding to the forecast locations in the outputs, set to the mean or last recorded value. Zero-values were used as placeholders for the seasonal inputs to the Autoformer, as well as for the detail components in the FFTransformer model. It should be noted that placeholders were added subsequent to the 'MDWD' and prior to the 'Embed + Encoding block' in Fig. 7, meaning that inputs to the MDWD were $v_i \in \mathbb{R}^{1\times S\times 8}$. Difference in latitude and longitude between stations were assigned to the input edge features as $e_{ij} = [(\text{lon}_j - \text{lon}_i), (\text{lat}_j - \text{lat}_i)]$, for two stations, i and j.

3.4 Experimental Set-up

All features were scaled separately using a standard scaler, with zero mean and unit variance. Three forecasting horizons were considered for the multi-step forecasting problem, namely 1-, 6- and 24-step forecasts, corresponding to 10-minute, 1-hour and 4-hour periods. Since the study considered multi-step forecasting, a prediction was made for every 10-minute interval also in the 1- and 4-hour ahead settings, instead of simply average wind speed forecasts over the respective periods. Models were trained to minimise the mean squared error (MSE) and hyperparameter tuning was conducted in Optuna [53]. Every spatio-temporal

(ST) GNN model consisted of two stacked graph blocks and were trained for 30 epochs using an Adam optimiser, batch size of 32, learning rate of 0.001, decayed by 20% every epoch, and a 5% dropout rate. A look-back window of 32 time-steps was used for the 10-minute and 1-hour ahead forecasts, increased to 64 for the 4-hour forecasts. Similarly, the dimensionality of hidden layers were set to 64 for the 10-minute and 1-hour ahead settings, but reduced to 32 for the 4-hour setting. All transformer-based models employed 8 heads for the MHA blocks and used edge update functions, $\phi^{(e)}$, represented by a single vanilla Transformer encoder layer [29]. The additional, model-specific, parameters obtained from the tuning procedure will now be summarised.

ST-MLP: Both edge and node update functions, $\phi^{(\cdot)}$ in eq. (2) and (3), represented by two-layer MLPs, with ReLU activation functions. Inputs were embedded using a linear transform, but without positional encoding.

ST-LSTM: Update functions represented by LSTMs and without positional encoding due to the recurrent architecture. The ST-LSTM model was slightly more sensitive to the architectural configuration than the other models and additional tuning was conducted to decide on the final parameters. The dimensionality of hidden layers were set to 16, 32 and 64, for the 1-, 6- and 24-step forecasts, respectively. Single-layer LSTMs were used for the 1- and 6-step forecasts, while two layers were used for the 24-step setting together with a higher dropout rate of 0.15.

ST-Transformer: Single encoder layer from the vanilla Transformer was used as update functions, with ReLU activation function.

ST-LogSparse: Node update functions, $\phi^{(v)}$, represented by a LogSparse Transformer encoder layer, with a kernel size of 6 for the 1D convolutional attention mechanisms and GELU activation functions. Sparsity was introduced into the attention mechanism as in the original formulation [33], with a restart length of 16.

ST-Informer: Informer encoder, without self-attention distilling to ensure constant sequence lengths, was used as node-update functions, with a sampling factor of $c = 3$ for the *ProbSparse Attention* and GELU activation functions.

ST-Autoformer: Autoformer decoder, without cross Auto-Correlation, was used for the node update functions. A GELU activation was used and the series decomposition blocks had a moving average kernel of 25.

ST-FFTransformer: Node update functions represented by FFTransformer encoder with GELU activation functions and four layer MDWD, meaning that the series for each input feature was decomposed into four detail and one approximate signal, as visualised in Fig. 6b. The MHA modules in the trend (right-hand) stream in Fig.7, employed convolutional *ProbSparse* attention, with both a kernel size and sampling factor of 3. Different to the original *ProbSparse Attention*, which focuses on dominant queries, it was found marginally beneficial to instead lo-cate dominant keys. Convolutional *ProbSparse Attention* was also used for the MHA module in the FFT-Attention, but finding top queries, as in the original *ProbSparse* formulation. For non-dominant query locations, the original *ProbSparse Attention* formulation assigns mean values to the outputs [34]. However, when using the *ProbSparse Attention* for the FFT-Attention, we instead set the outputs for non-dominant locations to zero to introduce sparsity and only select a subset of the possible frequencies. This was thought desirable, as to avoid outputs with a very large number of frequencies of small or similar amplitudes, which could be considered noisy.

Persistence: The persistence model was used as a benchmark for which to compare all the other models against. For the spatio-temporal setting, the persistence model will use the last recorded wind speed for a station as the forecast over the entire horizon, as $\hat{V}_t = \hat{V}_{t+1} = ... = \hat{V}_{t+P} = V_{t-1}$, using the notation in Sec. 3.1. Even though this is quite a trivial method for making forecasts, the model can achieve fairly accurate results in the short-term and is therefore used as an important baseline to outperform.

4 Results and Discussion

4.1 Forecasting Error

To evaluate the predictive performance of the different models, we start by comparing the mean absolute (MAE) and squared (MSE) errors, given by the following equations:

$$\text{MAE} = \frac{1}{n} \sum_{i=0}^{n} |y_i - \hat{y}_i| \tag{5}$$

$$\text{MSE} = \frac{1}{n} \sum_{i=0}^{n} (y_i - \hat{y}_i)^2, \tag{6}$$

where n is the total number of samples and, $\hat{y}$, the model predictions, which should be close to the targets y. For each forecasting horizon, every model was trained five times, with the average errors on the test set given in Table 1. Before computing the metrics, the predictions and labels were transformed back to a meters per second scale for better interpretability of the results. The percentage improvement values in Table 1 are relative to the persistence model, and are provided in order to highlight the relative forecasting performances.

Looking at the ST-Transformer model in Table 1, it was evident that the vanilla Transformer-based model did not report any remarkable advancements. Compared to the ST-MLP model, the ST-Transformer only showed marginal improvements for the 6- and 24-step forecasts, while being on-par with the persistence and ST-LSTM models for the single-step setting. For the 1- and 6-step forecasts, the ST-MLP model outperformed the ST-LSTM model, while the ST-LSTM only yielded slightly better results than ST-MLP model for the longer 24-step forecasts. This might indicate that for the immediate short-term forecasts,

Table 1: Test losses in m/s for different models and forecast horizons.

Model	MSE (% Improvement)						MAE (% Improvement)					
	1-step (10 min)		6-steps (1 hrs)		24-steps (4 hrs)		1-step (10 min)		6-steps (1 hrs)		24-steps (4 hrs)	
Persistence	0.4138	0.0 %	1.0282	0.0 %	2.9712	0.0 %	0.4277	0.0 %	0.6852	0.0 %	1.1922	0.0 %
ST-MLP	0.3938	4.8 %	0.9108	11.4 %	2.3951	19.4 %	0.4265	0.3 %	0.6588	3.9 %	1.0987	7.8 %
ST-LSTM	0.4040	2.3 %	0.9430	8.3 %	2.3707	20.2 %	0.4336	-1.4 %	0.6739	1.7 %	1.0964	8.0 %
ST-Transformer	0.4022	2.8 %	0.8923	13.2 %	2.3302	21.6 %	0.4346	-1.6 %	0.6561	4.3 %	1.0866	8.9 %
ST-LogSparse	0.3937	4.9 %	0.8492	17.4 %	2.3290	21.6 %	0.4309	-0.7 %	0.6369	7.0 %	1.0838	9.1 %
ST-Informer	0.3948	4.6 %	0.8552	16.8 %	2.2321	24.9 %	0.4321	-1.0 %	0.6383	6.8 %	1.0559	11.4 %
ST-Autoformer	**0.3454**	**16.5 %**	0.8203	20.2 %	2.3990	19.3 %	**0.4046**	**5.4 %**	**0.6231**	**9.1 %**	1.0907	8.5 %
ST-FFTransformer	0.3492	15.6 %	**0.8147**	**20.8 %**	**2.1336**	**28.2 %**	0.4098	4.2 %	0.6245	8.9 %	**1.0329**	**13.4 %**

all the previous time-steps might be relevant, resulting in the fully connected architecture of the ST-MLP achieving better accuracies than the ST-LSTM. However, for longer input sequences and forecast horizons, the LSTM architecture might be better at encoding the useful information, resulting in the improving performance of the ST-LSTM model, compared to the ST-MLP, for the longer forecasting settings.

The ST-LogSparse and ST-Informer performed consistently better than the ST-Transformer model across all forecasting horizons in terms of both MSE and MAE, which showed the potential improvements brought by the *ProbSparse* and convolutional attention mechanisms for wind forecasting. Even though the ST-LogSparse and ST-Informer models reported slightly inferior forecasting performance in terms of MAE for the single-step forecasts, compared to the persistence model, both showed approximately a five percent improvement in MSE. Since MSE penalise large errors more heavily than the MAE metric, it meant that for the single-step forecasts, the persistence model had on average fewer slightly smaller errors, but a larger number of drastically wrong predictions than the ST-LogSparse and ST-Informer models.

In general, all Transformer-based models attained better results than the ST-MLP and ST-LSTM models for the multi-step forecasts. The ST-Autoformer and ST-FFTransformer performed remarkably well for the 1- and 6-step forecasts, achieving more than three times the improvement reported

for the third best model in terms of MSE for the 1-step forecasts, and nearly a doubling compared to the ST-MLP for the 6-step forecasts. Even though the ST-Autoformer model performed very well for the 1- and 6-step forecasts, its performance was seen to degrade for the 24-step setting, where it was inferior to all other Transformer-based models. The ST-FFTransformer on the other hand, continued to improve, achieving superior results compared to all other models also for the 4-hour forecasts. Both the ST-FFTransformer and ST-Autoformer achieved the best accuracy for three settings each. For the 1- and 6-step ahead forecasts, the difference between the two models was very small, while the ST-FFTransformer substantially outperformed the ST-Autoformer for the longer forecasting horizon of four hours. As a result, we therefore argue for the competitive forecasting performance of our proposed FFTransformer architecture.

Fig. 8 shows some predictions under the different forecast horizons, for some randomly chosen time-periods for the Kvitebjørnfeltet measurement station in Fig. 5. Overall, it was difficult to conclude on major differences between the models based on the plotted time-series. Nevertheless, it was observed that for the 6-step setting, the ST-MLP and ST-Transformer models often gave near-constant predictions for all six time-steps, while the ST-LSTM, ST-LogSparse, ST-Informer, ST-FFTransformer and ST-Autoformer were able to produce slightly more diverse time-series. For the 24-step (4 hrs) setting, all models

(a) One-Step Forecasts (10-minute) (b) Six-Step Forecasts (1-hour) (c) 24-Step Forecasts (4-hour)
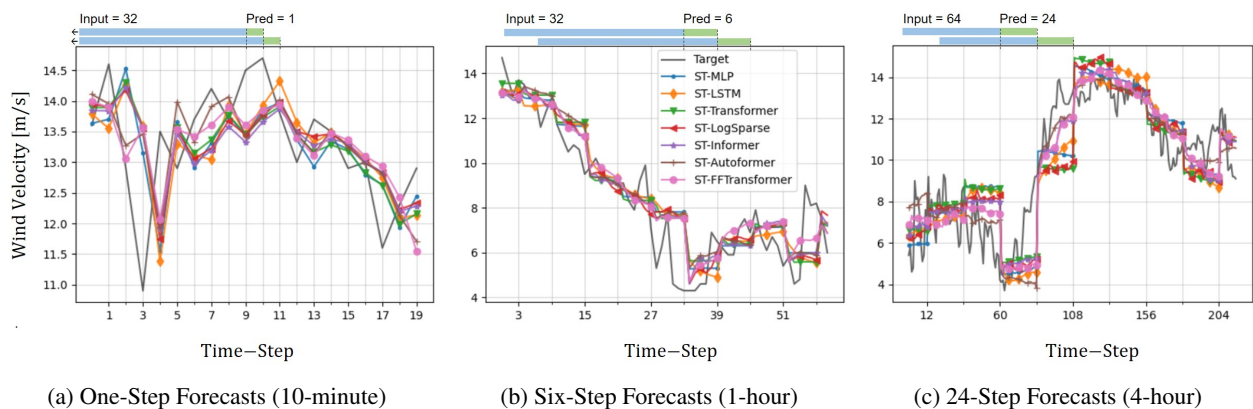
Figure 8: Prediction examples of the different models under the 1-, 6- and 24-step forecasts. The vertical grid line spacings correspond to the prediction horizons, i.e. in (b), 6-step predictions are shown with vertical grid spacing 6.

Table 2: Estimated power errors for different models and forecast horizons.

Model	MAE [kW] (% Improvement)						Interval MAE [kWh] (% Improvement)					
	1-step (10 min)		6-steps (1 hrs)		24-steps (4 hrs)		1-step (10 min)		6-steps (1 hrs)		24-steps (4 hrs)	
Persistence	166.1	0.0 %	265.4	0.0 %	459.2	0.0 %	27.684	0.0 %	228.8	0.0 %	1604.2	0.0 %
ST-MLP	165.1	0.6 %	253.2	4.6 %	422.0	8.1 %	27.509	0.6 %	209.4	8.5 %	1387.0	13.5 %
ST-LSTM	166.0	0.1 %	258.5	2.6 %	418.9	8.8 %	27.669	0.1 %	212.0	7.3 %	1386.0	13.6 %
ST-Transformer	166.4	-0.2 %	250.7	5.5 %	417.7	9.1 %	27.729	-0.2 %	206.9	9.5 %	1373.0	14.4 %
ST-LogSparse	164.7	0.8 %	245.5	7.5 %	417.4	9.1 %	27.453	0.8 %	203.5	11.0 %	1371.1	14.5 %
ST-Informer	165.1	0.6 %	245.7	7.4 %	405.7	11.7 %	27.519	0.6 %	204.4	10.7 %	1347.3	16.0 %
ST-Autoformer	**156.3**	**5.9 %**	239.8	**9.7 %**	416.9	9.2 %	**26.046**	**5.9 %**	196.3	**14.2 %**	1383.0	13.8 %
ST-FFTransformer	158.5	4.6 %	239.9	9.6 %	**396.5**	**13.7 %**	26.409	4.6 %	197.4	13.7 %	**1300.6**	**18.9 %**

showed more near-constant predictions, as seen in Fig. 8c, where some forecasts exhibit step-like changes at every prediction interval (i.e. 24 time-steps). Forecasting accurate time-series for longer horizons is challenging and the near-constant predictions indicate that the models sometimes struggled to confidently predict sharp increases or decreases in wind speeds. Having additional features and measurement stations or much deeper networks together with larger datasets, could potentially enable the models to more accurately learn these long-term changes. However, some of the near-constant predictions might also reflect the inherent stochasticity associated with wind time-series, meaning that in some scenarios it might be inconceivable to attain accurate multi-step forecasts for longer horizons, based solely on previous time-series and no physical information. Nevertheless, for other time-steps, the ST-LSTM, ST-FFTransformer, ST-Autoformer and ST-Informer models seemed much better than the ST-MLP, ST-Transformer and ST-LogSparse at capturing the overall trends within the intervals in Fig. 8c.

To investigate the physical interpretation of the forecasting results in relation to wind energy production, two additional MAE metrics were computed and provided in Table 2, which correspond to the estimated errors in kW and kWh. Powers were estimated based on the NREL 5 MW reference wind turbine for offshore system development [54]. By transforming the true and predicted wind speeds to powers using the reference turbine's power curve and then calculating the MAEs, the results show crude estimates for the average power errors using the different models. For the first metric in Table 2, results were fairly similar to those discussed in Table 1, but arguably more interpretable, in terms of understanding the consequence of differing predictive performances and potential risks associated with the proposed models.

Instead of looking at the point-wise power difference between the true and predicted values, an operator might be primarily concerned with the total overall energy for a future time interval. For the Interval MAE, predictions were again transformed to power values, before the values associated with a particular forecast interval were summed. Finally, since each step was associated with a 10-minute interval, summed values were divided by six, to convert them to the total energy estimates in kWh. The Interval MAE therefore provides an estimate of the difference between the predicted and true total energies produced for
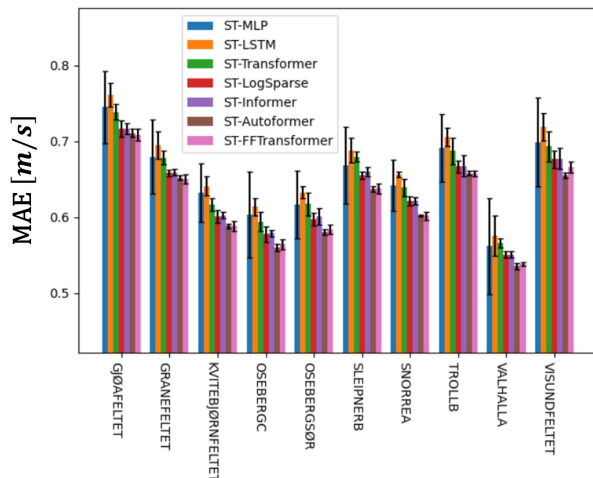

Figure 9: MAEs for the different available measurement stations in the test set under the 6-step forecasting setting. The error bars are the ± 2 standard deviations from the five training iterations of each model.

the relevant forecast horizon. The performance of different models was similar to previous results, but with all models showing greater improvements over the persistence model. Considering the ST-FFTransformer, it managed to reduce the error by 300 kWh for the 4-hour forecasts, corresponding to a 19% improvement over the persistence model. This was fairly notable, as it also meant an additional ≈ 5% improvement compared to most other models, and illustrated the potential cost savings and benefits of more accurate forecasting models.

4.2 Station Predictability

Since the models made forecasts for all 14 stations in Fig. 5, simultaneously, it was thought interesting to inspect the average errors for each station independently. Fig. 9, shows the average MAEs for every station under the 6-step forecasting setting, along with error bars of ±2σ, where σ is the standard deviation computed based on five separately trained models. It was seen that there was significant variability in how well the models were able to make forecasts for the different stations, with MAEs ranging from 0.53 m/s for some stations, to 0.74 m/s for others. Even though there were distinct differences in the performance

of different models, as discussed in the previous section, stations associated with higher or lower MAEs seemed consistent across all models. The data only contained historic information on the specific measurement stations in Fig. 5, with a fixed physical layout, meaning that spatial features would not change and that the meteorological data was likely to follow a particular distribution, specific for the area considered. As a result, a station might be inherently easier to forecast than others, due to its location relative to surrounding stations and due to the dominant wind fields for this specific geographical area potentially being preferable for forecasting at a particular location.

Inspecting the error bars in Fig. 9, it was evident that the ST-MLP model showed considerably more variability than the ST-LSTM and Transformer-based models. Since this variability was uncontrollable, it meant that one could not reliably conclude on the ST-MLP model's exact performance, which was undesirable. In contrast, all Transformer-based models showed much smaller variability across training iterations, meaning that a user could potentially be more confident in the expected errors associated with forecasts made by these models.

4.3 Graph Connectivity

For the results discussed so far, all input data was constructed as complete graphs, meaning that all nodes had edge features sending to all other nodes in the network and itself. This trivial method for formulating the graphs could result in very large inputs, significantly increasing the computational and memory costs. For instance, if a graph had 14 nodes, it would result in $14^2 = 196$ edges. Some of the edges might be redundant, meaning that the relevant spatial information could be provided by a subset of the edges, in addition to excess information potentially

making training more challenging. Even though this study did not focus on discussing better connectivity strategies for wind forecasting or using learnable adjacency matrices, we conduct a brief investigation into whether some of the connections could potentially be removed. Fig. 10 plot the different MAEs on the test data, as we increase the number of connections in the input graphs. 'n_closest' refers to the number of closest neighbours for which we allow a node to connect to. Looking at Fig. 5, this meant that if 'n_closest' was set to three, Valhall A would only have edges sending to it from Sleipner B, Granefeltet and Oseberg Sør. The first row in Fig. 10, show the MAEs for different 'n_closest' values, while the second row contains the same information, but normalised by MAE_0, which was the MAE when 'n_closest' $= 0$, i.e. the MAE for predictions made without any spatial information. While the first row provides information on the relative differences between models, the second row visualises the percentage improvements gained from including additional edges for a particular model.

First, MAEs were seen to rapidly decrease as we increased the number of edges, before converging to constant values when more than around five neighbours were considered for the edges. The sharp decrease indicated that the models were able to successfully leverage spatial correlations to improve forecasts, proving the effectiveness of the proposed GNN architectures. Nevertheless, since MAEs converged to constant values for non-complete graphs, it indicated that a number of connections could potentially be removed without impairing predictive performance. Further work would therefore be desirable to investigate better methods for which to construct the graphs or learn optimal connectivity for spatio-temporal wind forecasting.
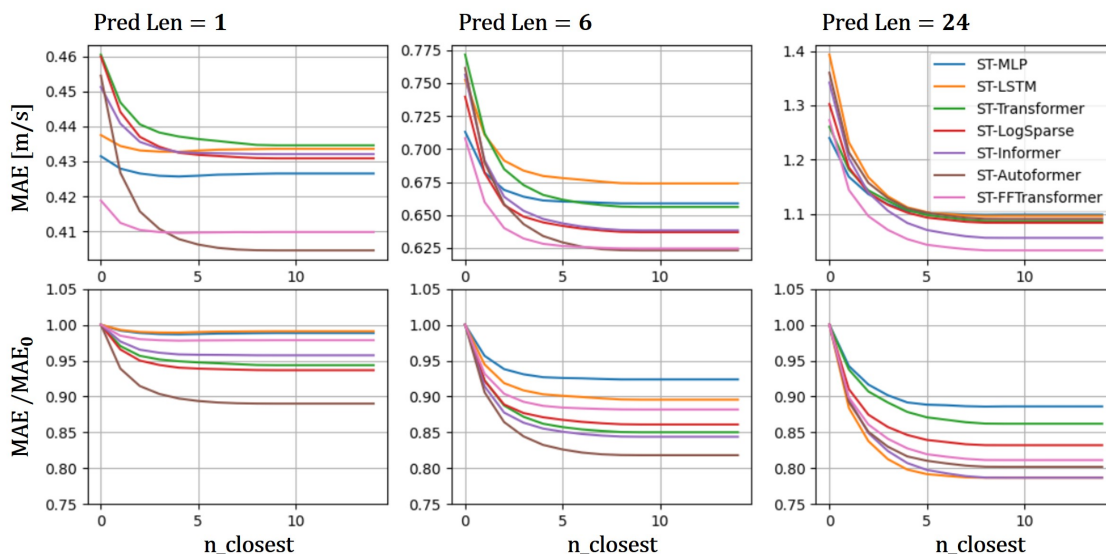

Figure 10: Test MAEs for different degree of graph connectivity, with 'n_closest' referring to the maximum number of edges sending to a particular node. The bottom row are the MAEs normalised by the MAEs for 'n_closest' = 0.

11

Looking at the second row of Fig. 10, it was seen that the percentage reduction in MAEs was greater for the longer forecasting horizons. For the immediate short-term (i.e. prediction length of 1), spatial correlations were thought less important than for the 1- and 4-hour forecasts, due to the large physical distances between nodes resulting in wind fields not having time to propagate in the immediate short-term. The added benefit of having long-range connections between nodes far apart, was therefore greater for the 6- and 24-step settings, than for the 1-step ahead forecasts. Comparing the 6- and 24-step forecasts, the latter also converged slightly later, which was likely due to the longer term forecasts taking advantage of information from nodes further away from the target. The percentage change in MAEs was also greater for the 24-step setting than for the 6-step, even though the difference was not as big as between the 1- and 6-step settings, which indicated that the 24-step forecasts might have benefited from even larger geographical information than was available. Overall, it was concluded that all models were able to take advantage of spatial information in making forecasts, with the Transformer-based architectures generally showing slightly larger improvements than the ST-MLP model.

5   Conclusion

In recent years, Transformer-based models have presided over sequence-based deep learning, often superseding recurrent or convolutional models. Nevertheless, research employing these architectures for wind forecasting has been scarce. This study considered different Transformer architectures as the main predictor for spatio-temporal multi-step forecasting, focusing on the LogSparse Transformer, Informer and Autoformer. This is the first time many of these have been applied to wind forecasting and placed in a spatio-temporal setting using GNNs. Additionally, the novel FFTransformer was proposed, which is based on signal decomposition using wavelet transform and an adapted attention mechanism in the frequency domain. Results show that the FFTransformer architecture was very competitive, achieving results on par with the Autoformer-based model for the 1- and 6-step forecasts, while significantly outperforming all other models for the longer 24-step forecasts. Even though the vanilla Transformer architecture generally did not yield significant improvements over an MLP model, it was seen that the convolutional attention in the LogSparse Transformer and the *ProbSparse Attention* of the Informer, were able to slightly improve prediction performance. By estimating the associated prediction errors in kW and kWh, we showed the potential physical effects of different forecasting performances with regards to the power grid, with the FFTransformer model showing an additional 5 % improvement over all other models for the 4-hour forecasts. Nevertheless, obtaining the powers based on the NREL 5 MW reference turbine, the method was fairly trivial and it would be desirable to further test the different models on real wind power datasets. By removing graph connections in the input data,

we showed that the proposed GNN architectures were successful in leveraging spatial correlations to improve local forecasts. However, it was also seen that some connections in the input data might be redundant, calling for additional research into more efficient approaches for graph connectivity in the context of wind forecasting. Since the FFTransformer model is not restricted to a particular signal decomposition technique or attention mechanism, slight alterations from the particular set-up used in this study might also be relevant and could be easily implemented and tested to facilitate different applications. We therefore hope that this study will spark further research into modifications and other applications of the FFTransformer, as well as investigation into the applicability of different Transformer-based architectures for use in wind forecasting.

CRediT authorship contribution statement

Lars Ø. Bentsen: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Project administration, Resources, Software, Validation, Visualization, Writing - original draft. **Narada Dilp Warakagoda**: Project administration, Supervision, Writing - review. **Roy Stenbro**: Project administration, Supervision. Writing - review. **Paal Engelstad**: Project administration, Supervision. Writing - review.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

References

[1] Inci Okumus and Ali Dinler. Current status of wind energy forecasting and a hybrid method for hourly predictions. *Energy Conversion and Management*, 123:362–371, 2016.

[2] GAM Van Kuik, Joachim Peinke, Rogier Nijssen, Denja Lekou, Jakob Mann, Jens Nørkær Sørensen, Carlos Ferreira, Jan-Willem van Wingerden, David Schlipf, Pieter Gebraad, et al. Long-term research

challenges in wind energy–a research agenda by the european academy of wind energy. *Wind energy science*, 1(1):1–39, 2016.

[3] Wen-Yeau Chang et al. A literature review of wind forecasting methods. *Journal of Power and Energy Engineering*, 2(04):161, 2014.

[4] Rajesh G Kavasseri and Krithika Seetharaman. Day-ahead wind speed forecasting using f-arima models. *Renewable Energy*, 34(5):1388–1393, 2009.

[5] SN Singh, Abheejeet Mohapatra, et al. Repeated wavelet transform based arima model for very short-term wind speed forecasting. *Renewable energy*, 136:758–768, 2019.

[6] Ramon Gomes da Silva, Matheus Henrique Dal Molin Ribeiro, Sinvaldo Rodrigues Moreno, Viviana Cocco Mariani, and Leandro dos Santos Coelho. A novel decomposition-ensemble learning framework for multi-step ahead wind energy forecasting. *Energy*, 216:119174, 2021.

[7] Binrong Wu, Lin Wang, and Yu-Rong Zeng. Interpretable wind speed prediction with multivariate time series and temporal fusion transformers. *Energy*, 252:123990, 2022.

[8] Anbo Meng, Jiafei Ge, Hao Yin, and Sizhe Chen. Wind speed forecasting based on wavelet packet decomposition and artificial neural networks trained by crisscross optimization algorithm. *Energy Conversion and Management*, 114:75–88, 2016.

[9] Hui Liu, Hong-qi Tian, Di-fu Pan, and Yan-fei Li. Forecasting models for wind speed using wavelet, wavelet packet, time series and artificial neural networks. *Applied Energy*, 107:191–208, 2013.

[10] Kathrine Lau Jørgensen and Hamid Reza Shaker. Wind power forecasting using machine learning: State of the art, trends and challenges. In *2020 IEEE 8th International Conference on Smart Energy Grid Engineering (SEGE)*, pages 44–50. IEEE, 2020.

[11] Alireza Zendehboudi, M Abdul Baseer, and R Saidur. Application of support vector machine models for forecasting solar and wind energy resources: A review. *Journal of cleaner production*, 199:272–285, 2018.

[12] Ilhami Colak, Seref Sagiroglu, and Mehmet Yesilbudak. Data mining and wind power prediction: A literature review. *Renewable Energy*, 46:241–247, 2012.

[13] A Sfetsos. A novel approach for the forecasting of mean hourly wind speed time series. *Renewable energy*, 27(2):163–174, 2002.

[14] Anurag More and MC Deo. Forecasting wind with neural networks. *Marine structures*, 16(1):35–49, 2003.

[15] Jing Shi, Jinmei Guo, and Songtao Zheng. Evaluation of hybrid forecasting approaches for wind speed and power generation time series. *Renewable and Sustainable Energy Reviews*, 16(5):3471–3480, 2012.

[16] Zhen-hai Guo, Jie Wu, Hai-yan Lu, and Jian-zhou Wang. A case study on a hybrid wind speed forecasting method using bp neural network. *Knowledge-based systems*, 24(7):1048–1056, 2011.

[17] Ghadah Alkhayat and Rashid Mehmood. A review and taxonomy of wind and solar energy forecasting methods based on deep learning. *Energy and AI*, 4:100060, 2021.

[18] Jürgen Schmidhuber, Sepp Hochreiter, et al. Long short-term memory. *Neural Comput*, 9(8):1735–1780, 1997.

[19] Yanfei Li, Haiping Wu, and Hui Liu. Multi-step wind speed forecasting using ewt decomposition, lstm principal computing, relm subordinate computing and iewt reconstruction. *Energy Conversion and Management*, 167:203–219, 2018.

[20] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.

[21] Pengbo Li, Xiangwen Wang, and Junjie Yang. Short-term wind power forecasting based on two-stage attention mechanism. *IET Renewable Power Generation*, 14(2):297–304, 2020.

[22] Bin Huang, Yuying Liang, and Xiaolin Qiu. Wind power forecasting using attention-based recurrent neural networks: a comparative study. *IEEE Access*, 9:40432–40444, 2021.

[23] Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499*, 2016.

[24] Xiaochong Dong, Yingyun Sun, Ye Li, Xinying Wang, and Tianjiao Pu. Spatio-temporal convolutional network based power forecasting of multiple wind farms. *Journal of Modern Power Systems and Clean Energy*, 10(2):388–398, 2021.

[25] Kumar Shivam, Jong-Chyuan Tzou, and Shang-Chen Wu. Multi-step short-term wind speed prediction using a residual dilated causal convolutional network with nonlinear attention. *Energies*, 13(7):1772, 2020.

[26] HZ Wang, GB Wang, GQ Li, JC Peng, and YT Liu. Deep belief network based deterministic and probabilistic wind speed forecasting approach. *Applied Energy*, 182:80–93, 2016.

[27] Zhewen Niu, Zeyuan Yu, Wenhu Tang, Qinghua Wu, and Marek Reformat. Wind power forecasting using attention-based gated recurrent unit network. *Energy*, 196:117081, 2020.

[28] Dominik Putz, Michael Gumhalter, and Hans Auer. A novel approach to multi-horizon wind power fore-

casting based on deep neural architecture. *Renewable Energy*, 178:494–505, 2021.

[29] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

[30] Iz Beltagy, Matthew E Peters, and Arman Cohan. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*, 2020.

[31] Tian Zhou, Ziqing Ma, Qingsong Wen, Xue Wang, Liang Sun, and Rong Jin. Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting. *arXiv preprint arXiv:2201.12740*, 2022.

[32] Bryan Lim, Sercan Ö Arık, Nicolas Loeff, and Tomas Pfister. Temporal fusion transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting*, 37(4):1748–1764, 2021.

[33] Shiyang Li, Xiaoyong Jin, Yao Xuan, Xiyou Zhou, Wenhu Chen, Yu-Xiang Wang, and Xifeng Yan. Enhancing the locality and breaking the memory bottleneck of transformer on time series forecasting. *Advances in neural information processing systems*, 32, 2019.

[34] Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 11106–11115, 2021.

[35] Nikita Kitaev, Łukasz Kaiser, and Anselm Levskaya. Reformer: The efficient transformer. *arXiv preprint arXiv:2001.04451*, 2020.

[36] Haixu Wu, Jiehui Xu, Jianmin Wang, and Mingsheng Long. Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. *Advances in Neural Information Processing Systems*, 34:22419–22430, 2021.

[37] Lei Wang, Yigang He, Lie Li, Xiaoyan Liu, and Yingying Zhao. A novel approach to ultra-short-term multi-step wind power predictions based on encoder–decoder architecture in natural language processing. *Journal of Cleaner Production*, 354:131723, 2022.

[38] Kai Qu, Gangquan Si, Zihan Shan, XiangGuang Kong, and Xin Yang. Short-term forecasting for multiple wind farms based on transformer model. *Energy Reports*, 8:483–490, 2022.

[39] Zhifeng Liu, Feng Ding, and Fujing Wan. Wind speed forecasting method based on deep learning strategy using long short term memory neural network and transformer model. In *2021 IEEE 23rd Int Conf on High Performance Computing & Communications; 7th Int Conf on Data Science & Systems; 19th Int Conf on Smart City; 7th Int Conf on Dependability in Sensor, Cloud & Big Data Systems & Application (HPCC/DSS/SmartCity/DependSys)*, pages 2288–2293. IEEE, 2021.

[40] Hai-Kun Wang, Ke Song, and Yi Cheng. A hybrid forecasting model based on cnn and informer for short-term wind power. *Frontiers in Energy Research*, page 1041, 2022.

[41] Xiaoxin Pan, Long Wang, Zhongju Wang, and Chao Huang. Short-term wind speed forecasting based on spatial-temporal graph transformer networks. *Energy*, 253:124095, 2022.

[42] Tianyu Hu, Wenchuan Wu, Qinglai Guo, Hongbin Sun, Libao Shi, and Xinwei Shen. Very short-term spatial and temporal wind power forecasting: A deep learning approach. *CSEE Journal of Power and Energy Systems*, 6(2):434–443, 2019.

[43] Qiaomu Zhu, Jinfu Chen, Dongyuan Shi, Lin Zhu, Xiang Bai, Xianzhong Duan, and Yilu Liu. Learning temporal and spatial correlations jointly: A unified framework for wind speed prediction. *IEEE Transactions on Sustainable Energy*, 11(1):509–523, 2019.

[44] Yun Wang, Runmin Zou, Fang Liu, Lingjun Zhang, and Qianyi Liu. A review of wind speed and wind power forecasting with deep neural networks. *Applied Energy*, 304:117766, 2021.

[45] Mahdi Khodayar and Jianhui Wang. Spatio-temporal graph deep neural network for short-term wind speed forecasting. *IEEE Transactions on Sustainable Energy*, 10(2):670–681, 2018.

[46] Tomasz Stańczyk and Siamak Mehrkanoon. Deep graph convolutional networks for wind speed prediction. *arXiv preprint arXiv:2101.10041*, 2021.

[47] Fei Wang, Peng Chen, Zhao Zhen, Rui Yin, Chunmei Cao, Yagang Zhang, and Neven Duić. Dynamic spatio-temporal correlation and hierarchical directed graph structure based ultra-short-term wind farm cluster power forecasting method. *Applied Energy*, 323:119579, 2022.

[48] Lei Wang and Yigang He. M2stan: Multi-modal multi-task spatiotemporal attention network for multi-location ultra-short-term wind power multi-step predictions. *Applied Energy*, 324:119672, 2022.

[49] Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. *Deep learning*, volume 1. MIT press Cambridge, 2016.

[50] Peter W Battaglia, Jessica B Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018.

[51] H Díaz and C Guedes Soares. Review of the current status, technology and future trends of offshore wind farms. *Ocean Engineering*, 209:107381, 2020.

[52] JP da S Catalão, Hugo Miguel Inácio Pousinho, and Victor Manuel Fernandes Mendes. Short-term wind

power forecasting in portugal by neural networks and wavelet transform. *Renewable energy*, 36(4):1245–1251, 2011.

[53] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 2623–2631, 2019.

[54] Jason Jonkman, Sandy Butterfield, Walter Musial, and George Scott. Definition of a 5-mw reference wind turbine for offshore system development. Technical report, National Renewable Energy Lab.(NREL), Golden, CO (United States), 2009.