

Predictive Yaw System Control of Wind Turbines Through Long Short Term Memory Neural Networks

Adarsh Devakottai Raghavan¹

Received May 25, 2026

Accepted September 26, 2026

Electronic access October 31, 2026

Wind turbine yaw control seeks to maintain alignment between the rotor and incoming wind direction while limiting the energy and mechanical costs associated with nacelle movement. Predictive control using machine learning based wind direction forecasts may potentially improve this trade off by anticipating changes instead of responding to observed misalignment. This study evaluates whether an LSTM based wind direction forecasting strategy provides an advantage over simpler forecasting approaches for turbine yaw control. One year of five minute wind data from Amarillo Texas was used to develop a chronological training, validation, and locked test framework. The LSTM was additionally evaluated across five random seeds and its error difference from persistence was assessed using block bootstrap confidence intervals. Forecasts were then incorporated into an actuator constrained yaw control simulation accounting for yaw movement and actuator energy consumption. Persistence achieved a locked test circular MAE of 1.29° , outperforming the LSTM, which achieved 2.09° . Across five LSTM seeds, mean test MAE was $2.02 \pm 0.24^\circ$. The LSTM's higher error relative to persistence was supported by a 95% block bootstrap confidence interval of $0.63\text{--}1.00^\circ$ for the MAE difference. Nevertheless, both active controllers substantially improved energy capture relative to a fixed nacelle. Persistence achieved 5,447.825 MWh of net energy and the LSTM 5,445.065 MWh compared with the 4,372.157 MWh for the fixed nacelle case and 5,503.685 MWh for idealized perfect yaw. These results show that for the short prediction horizon and wind regime examined, increased forecasting complexity did not provide a measurable energy advantage over persistence. The findings highlight the importance of evaluating predictive models according to their downstream control performance rather than forecasting accuracy alone.

Introduction

Wind energy is a major source of renewable energy that can be harvested from the kinetic energy of moving air. This can be done through the utilization of a wind turbine. Wind turbines are mechanical devices which convert moving air into electrical energy. Wind energy has grown significantly in recent decades and is one of the fastest growing renewable energy sources¹.

A yaw system is a mechanism inside the wind turbine which allows the wind turbine to rotate horizontally in order to face the direction of air flow to ensure maximum energy capture²⁻⁴. Yaw misalignment occurs when the turbine heading differs from the incoming wind direction. Even moderate misalignment can reduce aerodynamic power capture, making yaw control an important component of turbine operation⁵⁻⁷. However, yaw movement is not free. Nacelle rotation requires electrical energy and subjects the yaw system to mechanical loading and accumulated travel⁸. Consequently, an effective yaw controller must balance alignment accuracy against the energy and mechanical costs associated with movement^{2,8,9}.

Conventional yaw control is primarily reactive. The controller observes the current wind direction or yaw error and

commands the nacelle to reduce that error. This approach is simple, but cannot anticipate a directional change before it is observed. If the wind direction changes rapidly, the finite speed of the yaw actuator can temporarily leave the rotor misaligned^{10,11}. Predictive control offers an alternative by estimating future wind direction and beginning the yaw movement before the change occurs^{3,9}.

Long Short-Term Memory (LSTM) neural networks are recurrent neural network architectures designed to model temporal dependencies in sequential data¹². Their memory mechanisms make them suitable for time series forecasting, including meteorological variables such as wind speed and direction^{12,13}. However, predictive control does not necessarily outperform reactive control. The value of a forecast depends on the magnitude and persistence of future changes, the forecast error, and the response characteristics of the physical actuator^{9,14}.

This study investigates whether LSTM wind direction prediction can provide a practical yaw control advantage over simpler strategies under real wind conditions. Rather than evaluating prediction accuracy alone, the study incorporates the forecast into an actuator constrained yaw simulation and evaluates generated energy and the energy consumed by yaw movement. A persistence strategy is included as an important

¹ Sherwood High School, India

null baseline as slowly varying wind direction time series may contain substantial information in the most recently observed direction⁹.

The five minute prediction horizon was selected as the dataset sampled at five minute intervals and therefore provides a consistent forecasting problem. However the yaw actuator was simulated as a substantially finer internal timestep so that the nacelle could not rotate instantaneously between intervals¹⁵. The forecast is updated every 5 minutes while the actuator responds continuously subject to its rotation speed and movement constraints. Consequently, potential benefits of prediction will arise from beginning new yaw movements towards predicted direction before next measured wind direction becomes available.

The contribution of this study is not the introduction of new LSTM architecture or yaw control algorithm. Instead the study provides a controlled case study examining whether the additional predictive information provided by an LSTM translates into a measurable downstream energy benefit when compared with persistence under identical actuator constraints.

The research question is: Can an LSTM based predictive yaw controller improve net simulated energy capture relative to persistence based yaw control and fixed nacelle baseline under real wind conditions?

The study focuses on one year of five minute meteorological data from Amarillo, Texas. The purpose is not to establish the universal superiority of LSTM control, but to determine whether predictive information provides a measurable benefit under the wind conditions and actuator assumptions represented in this case study.

Background

Wind turbine power capture depends partly on the alignment between the rotor and the incoming wind direction. When the nacelle heading differs from the wind direction, yaw misalignment reduces the effective aerodynamic power captured by the rotor^{5,7,16,17}. Active yaw control can reduce this misalignment, but nacelle rotation also consumes energy and contributes to mechanical wear^{2,8}. Yaw control thus requires a trade off between maintaining alignment and limiting unnecessary movement¹⁶.

Predictive yaw control may improve this trade off by estimating future wind direction rather than responding only after a directional change has occurred. If future changes can be predicted accurately, the nacelle could begin moving toward the expected direction before the full misalignment develops. However, the practical value of prediction depends on both forecast accuracy as well as the physical characteristics of the yaw actuator^{7,9,18}.

Long short term memory (LSTM) networks are recurrent neural network architectures designed to model temporal de-

pendencies in sequential data¹². Unlike conventional feed forward neural networks, LSTMs can incorporate information from previous observations when generating a prediction. Their memory mechanisms therefore make them suitable for time series forecasting tasks in which historical observations may contain certain information about future behavior¹². LSTM based approaches thus seem to be influential in wind forecasting applications^{7,19}. LSTM based approaches have been applied extensively to wind forecasting, including short term wind speed and wind power forecasting. However, these studies primarily focus on forecasting performance^{16,20}. Forecasting accuracy does not automatically imply improved turbine level control performance.

For wind direction prediction, however, increased model complexity does not guarantee improved performance. Wind direction is a circular variable, and short term atmospheric behavior may contain substantial persistence, meaning that the most recently observed direction can provide a strong forecast of the consecutive one^{9,21}. A more complex forecasting model is therefore useful if it can extract predictive information beyond those contained in recent observations^{9,22}. This makes persistence an important baseline when evaluating the practical advantages of LSTM based forecasting.

Methodology

Dataset

The study uses meteorological data from the Wind Integration National Dataset (WIND) Toolkit developed by the National Renewable Energy Laboratory (NREL)¹. The dataset provides wind-resource information at five-minute temporal resolution across the United States⁴. Data for Amarillo, Texas, covering the full 2019 calendar year, were used for this study.

The dataset contains approximately 105,120 five-minute observations. Wind speed and wind direction at a 100 m hub height were used as the primary meteorological inputs. The full-year dataset provides seasonal variation in both variables while maintaining a consistent temporal sampling interval.

Amarillo was selected as a case-study location because of its substantial wind resource and inland wind environment. The purpose of using one location was to investigate the interaction between forecast accuracy, wind-direction behavior, and yaw-control performance under a defined set of atmospheric conditions.

Data Preprocessing

The raw data required several steps before it was used to train the LSTM model. The data was inspected for missing values and impossible readings such as wind direction out of the 0°–360° range and wind speed values which were greater than 50

m/s. Any missing values were filled using a linear interpolation between nearby timesteps, and no physically impossible readings were found for this dataset.

Wind speed and direction were selected as the primary model inputs as they describe the local wind state most directly relevant to turbine operation. Wind direction provides the primary information required for yaw alignment, while wind speed influences the potential aerodynamic energy available from the flow. Other variables were not included in the model so that the experiment remained focused on whether recent wind state information could provide useful predictive yaw information.

An important step in preprocessing was encoding the wind direction. Because wind direction is a circular variable, a numerical representation has major discrepancies at the $0^\circ/360^\circ$ border. To get rid of this problem, wind direction was decomposed into two components using trigonometric transformations: a sine component, calculated as $\sin(\theta \times \pi/180)$, and a cosine component, calculated as $\cos(\theta \times \pi/180)$, where θ is the wind direction in degrees. This preserves the circular geometry of the directional data and has been widely adopted in wind direction forecasting literature. Wind speed was also normalized to the range $[0,1]$ through Min Max scaling so that a gradient behaviour was ensured in model training. This sine cosine encoding approach is widely used in wind direction forecasting literature⁹.

The 2019 dataset was divided chronologically into training, validation and locked test periods so that future observations could not influence model training or controller tuning. The model was trained using the chronological 2019 dataset. For each individual prediction, the LSTM received the preceding 24 observations, which corresponds to a 120 minute historical window and predicted the wind direction at the following five minute timestep. Thus, 120 minutes describes the model's historical input window rather than the total duration of the dataset.

The final locked test set used for the controller evaluation contained 21,024 samples. All controller inputs and physical wind observations used in the final actuator simulation were aligned to these same timestamps.

LSTM Model Architecture & Training

The model was implemented using TensorFlow and Keras in Python. The architecture consisted of two stacked LSTM layers followed by a connected output layer. The first layer contained 64 units and returned its hidden state sequence, enabling the second layer to process temporal context across all 24 input timesteps. The second layer contained 32 units and only returned the final hidden state which provided a single summary representation of the input sequence. Dropout regularization with a rate of 0.2 was applied after each LSTM by

randomly deactivating a fraction of units during each training. A fully connected layer with 16 units and ReLU activation was next which resulted in the final output layer consisting of 2 units corresponding to the predicted sine and cosine components of future wind direction.

A 120 minute historical window was selected to provide the model with sufficient temporal context while maintaining manageable sequence length. At five minute resolution, the window contains 24 observations and represents approximately 2 hours of preceding wind behaviour. The study does not establish that 120 minutes is the optimal history length. Alternative lengths may produce different results and are identified as a direction for future work.

The model was compiled using the Adam optimizer and mean squared error loss function. Training was conducted with a batch size of 64 and a maximum of 50 epochs. The first 70% of observations were used for model training, the following 10% for validation and tuning, and the final 20% were reserved as a locked test set. Two callbacks were applied during training in order to increase accuracy. Early stopping halted training when the validation loss failed to improve for five consecutive epochs and restored the model weights from the epoch with the best validation performance. After prediction, the estimated sine and cosine components were converted back to wind direction using the two argument arctangent function of $\theta_t = \text{atan2}(s_t, c_t)$. The resulting angle was converted from radians to degrees and mapped to the range 0° – 360° before calculating circular forecast error and yaw misalignment^{16,23}.

Forecast performance was evaluated using circular mean absolute error and circular root mean squared error. Circular angular error was calculated as

$$e_t = |((\theta'_t - \theta_t + 180) \bmod 360^\circ) - 180^\circ|$$

which constrains the absolute directional error to the interval 0° – 180° .

Controller Strategies

The controller comparison was designed to distinguish predictive information from the information already available through the most recent observation.

The persistence controller uses the most recently observed wind direction as the target direction for the next control interval. This provides a simple forecasting and control baseline because it assumes that the wind direction will remain approximately unchanged over the next five minutes⁹.

The LSTM controller uses the one step ahead direction predicted by the neural network as its target heading. Unlike the persistence controller, it can potentially begin rotating toward a directional change before that change is observed.

Both controllers operate through the same simulated yaw actuator. Consequently, differences in performance arise from the controller's target direction rather than from giving one controller a faster or more capable actuator.

A fixed nacelle configuration was also simulated by holding the nacelle at the optimal fixed heading during controller development using developmental data, and the locked test period was not used to determine this heading. A perfect yaw reference was included as an upper bound case in which the nacelle is continuously aligned with the instantaneous wind direction without actuator limitations or yaw energy consumption.

Yaw misalignment was calculated as the smallest angular difference between the nacelle heading and the incoming wind direction, wrapped to the interval 0° – 180° . Thus, a 359° nacelle heading and a 1° wind direction correspond to a 2° misalignment rather than a 358° error.

Actuator and Energy Simulation

The final controller evaluation used the five-minute meteorological observations as the available wind measurements while simulating nacelle motion at a finer internal timestep. This separation prevented the actuator from unrealistically rotating instantaneously between consecutive meteorological observations. Wind observations therefore remained available only at their original five minute resolution, while the actuator responded internally subject to its finite rotation rate.

The actuator model was implemented at a finer internal timestep than the five-minute meteorological sampling interval. At each internal timestep, the nacelle heading was moved toward the controller target subject to the maximum yaw-rate constraint. A deadband of 8.3° prevented small corrective movements, while cumulative yaw travel and actuator energy consumption were recorded. This representation introduces finite response time and prevents instantaneous nacelle alignment, while remaining computationally tractable for comparison across controllers. More detailed turbine-specific effects, including motor inertia, acceleration limits, cable twist, hysteresis, and duty cycle restrictions, were not modeled and remain limitations of the simulation.

At each control update t , the controller is provided only with observations available before the target timestep. The most recent observed wind direction is denoted θ_{t-1} while the LSTM receives the preceding 24 observations and produces a prediction. The persistence controller uses θ_{t-1} as its target heading, whereas the LSTM controller uses the prediction. The actuator then moves from its current heading towards the selected target subject to the imposed constraints. The resulting nacelle heading is evaluated against the actual wind direction θ_t , which was not available to the controller when the command was generated. Thus the controller does not access future in-

formation and both controllers are evaluated against the same sequence.

All actuator parameters were fixed before evaluation on the locked test period. The actuator was simulated at a finer internal step, but the cumulative nacelle rotation was limited to 1.45° per five-minute control interval, while a yaw command was issued only when the target direction differed from the current nacelle heading by at least the 8.3° deadband. The same actuator constraints and control rules were applied to both controllers.

Generated turbine power was calculated using a simplified wind-speed-dependent power model rather than a normalized constant-power assumption. At each simulation timestep, the aligned turbine power $P_0(v_t)$ was determined from the corresponding wind speed v_t using the turbine power relationship implemented in the simulation. This power model is intended to provide a physically motivated estimate of the variation in available turbine power with wind speed rather than to reproduce the exact power curve of a specific commercial turbine. The effect of yaw misalignment was then represented using a multiplicative loss factor $f(\gamma_t, v_t)$, where γ_t is the angular difference between the nacelle heading and the incoming wind direction. The resulting turbine power was calculated as

$$P_t = P_0(v_t)f(\gamma_t, v_t)$$

The energy generated during each timestep was calculated from

$$E_t = P_t \Delta_t,$$

where Δ_t is the duration of the simulation timestep. Total energy production was obtained by summing the energy generated across all simulation timesteps⁵.

The yaw-loss relationship was represented using the empirical model $P_t = P_0(v_t)\cos^p(\gamma_t)$, where $P_0(v_t)$ is the aligned turbine power determined from the selected power curve, γ_t is the yaw misalignment, and p is a yaw-loss exponent. Because the value of p can depend on turbine operating conditions and atmospheric characteristics, it was treated as a modeling assumption rather than a universal physical constant. The present study did not perform a sensitivity analysis across alternative exponent values^{5,7,16}.

Consequently, the reported absolute MWh values should be interpreted as model-dependent simulated energy estimates rather than precise predictions of the annual energy production of a specific operating turbine. The primary interpretation of the energy analysis is therefore the relative comparison between controllers under identical power-curve, yaw-loss, and actuator assumptions.

Energy consumed by the yaw actuator was calculated separately and subtracted from the generated turbine energy to obtain net energy:

$$E_{net} = E_{gross} - E_{yaw}.$$

This distinction accounts for the trade-off between improved aerodynamic alignment and the electrical cost of nacelle movement. The simulation also recorded the number of yaw events and cumulative yaw travel, providing measures of mechanical activity in addition to energy production²⁴.

Controller Timing

The simulation followed a causal sequence at each five-minute controller update. First, the controller received only the wind observations available before the target timestep. Second, the controller generated a target heading using either persistence or the LSTM forecast. Third, the actuator moved the nacelle toward the target subject to its movement constraints. Finally, the resulting nacelle heading was evaluated against the physical wind direction at the corresponding simulation timestep, and turbine and actuator energy were accumulated.

At each five-minute controller update t , only wind observations available before the target timestep were provided to the controller. The latest observed wind direction was θ_{t-1} . The persistence controller selected $\theta'_t = \theta_{t-1}$, while the LSTM received the preceding 24 observations ending at $t-1$ and generated the one-step-ahead prediction θ'_t . The selected target direction was then compared with the current nacelle heading. If the angular difference exceeded the 8.3° deadband, the actuator began moving toward the target subject to its maximum yaw-rate constraint. The actuator state was updated throughout the subsequent control interval using the finer internal simulation timestep. The resulting nacelle heading was then compared with the actual wind direction θ_t , which was not available when the controller target was generated. Turbine power, yaw energy consumption, and yaw misalignment were calculated using the resulting nacelle heading and the physical wind conditions during timestep t .

Therefore, the LSTM prediction window ended at $t-1$, and neither controller received θ_t before generating its command. The resulting controller performance was evaluated against θ_t . This ensures that the comparison is causal and that the LSTM does not receive future wind-direction information.

The LSTM therefore does not receive the future wind direction as an input. Its prediction is generated solely from observations available before the target timestep. The locked test observations are not used to tune the model, controller parameters, or preprocessing decisions.

Performance Metrics

Controller performance is evaluated using generated energy, yaw energy consumption, net energy, yaw-event count, cumulative yaw travel, and angular alignment error.

The primary comparison is net energy because a controller that produces greater aerodynamic energy but consumes substantially more energy through nacelle movement does not necessarily provide a practical energy benefit.

Controller performance is also expressed relative to the fixed-nacelle and perfect-yaw references. The fixed-nacelle comparison quantifies the benefit of active yaw control, while the perfect-yaw reference provides an upper bound for the simulated system.

Because consecutive five-minute observations are temporally correlated, individual timesteps were not treated as independent statistical observations. Differences between controllers were therefore evaluated using block based resampling. Blocks preserve short term dependence while allowing uncertainty in the comparison to be estimated. Controller performance was also evaluated using cumulative yaw travel and yaw event count, since two controllers with similar energy production may impose different demands on the system.

In addition to the mean angular error, the analysis considered the upper tail quartiles, maximum observed error, the proportion of observations within selected yaw error bands, total yaw travel and yaw event count. These measures provide information about transient and extreme misalignment that cannot be captured by the mean alone.

Results

Forecasting Performance

The forecasting models were evaluated on the locked chronological test set using circular mean absolute error (MAE) and circular root mean square error (RMSE). Persistence provided a strong baseline, achieving a circular MAE of 1.29° and RMSE of 2.65° . The exponentially weighted moving average performed substantially worse, with an MAE of 3.04° . The circular autoregressive and linear sine/cosine models achieved lower errors than persistence, with MAEs of 1.03° and RMSEs of 2.26° .

The LSTM did not outperform the persistence baseline. In the finalized test evaluation, the LSTM achieved a circular MAE of 2.09° and RMSE of 3.48° , corresponding to a -62.0% skill relative to persistence. Thus, the LSTM produced less accurate one step ahead wind direction forecasts than the persistence model.

Model variability was assessed using five independent LSTM training seeds. Across these runs, test MAE was $2.02 \pm 0.24^\circ$, while test RMSE was $3.29 \pm 0.30^\circ$. This variation does not alter the qualitative conclusion that the LSTM failed to consistently outperform persistence.

The difference between the LSTM and persistence errors was further evaluated using a block bootstrap to account for temporal dependence in the five minute observations. The ob-

served difference in MAE was 0.80° in favor of persistence, with a 95% block bootstrap confidence interval of $0.63\text{--}1.00^\circ$. Because the interval lies entirely above zero, the LSTM's higher forecasting error is statistically distinguishable from the persistence baseline under the adopted block bootstrap procedure.

Stratified Forecasting Performance

To determine whether forecasting performance depended on operating conditions, the locked test was stratified by wind speed and by magnitude of the observed five minute change in wind direction. Across all wind speed categories, persistence produced lower forecasting error than the LSTM. For wind speeds below 5 m/s, persistence achieved an MAE of 2.66° while LSTM achieved 3.75° . For 5–8, 8–12 and >12 m/s, persistence MAEs were 1.26° , 0.95° and 0.72° respectively, while LSTM MAEs were 2.24° , 1.64° and 1.24° . Thus the LSTM remained less accurate even at higher wind speeds, where directional errors were generally smaller.

Performance was consistent when observations were stratified by magnitude of the five minute directional change. Persistence outperformed the LSTM in every category. For changes of $<1^\circ$, persistence achieved an MAE of 0.43° , compared with 1.44° for the LSTM. For changes of $1\text{--}2^\circ$, $2\text{--}5^\circ$, $5\text{--}10^\circ$, and $>10^\circ$, persistence MAEs were 1.42° , 3.03° , 6.72° and 17.44° respectively. The corresponding LSTM MAEs were 1.87° , 3.34° , 7.25° and 19.03° .

The distribution of wind direction changes provides context for these results. Consecutive five minute observations differed by an average of 1.91° , with a median absolute change of 0.92° . Approximately 75.5% of observations had directional changes of $<2^\circ$, and 92.7% had changes of $<5^\circ$. Large directional changes were comparatively uncommon, with only 0.327% of observations exceeding 30° , and only 0.039% exceeding 90° .

These results indicate that the Amarillo wind direction series was characterized by relatively gradual short term evolution. Under these conditions, the most recent observed direction provided a highly competitive prediction of the next five minute direction. The LSTM did not overcome this advantage including during periods with larger directional changes.

Yaw Control Performance

The forecasting models were subsequently evaluated within an actuator constrained yaw constrained simulation using the locked Amarillo test sequence. The simulation compared a fixed nacelle controller, persistence based yaw control, LSTM based predictive yaw control, and an idealized perfect yaw reference. The perfect yaw case represents the upper bound in which the turbine is consistently aligned with the incoming

wind direction while the fixed nacelle case provides the baseline without active yaw correction.

The fixed nacelle configuration generated 4,372.157 MWh over the simulated period. Perfect yaw alignment increased gross energy capture to 5,503.685 MWh, corresponding to a theoretical maximum improvement 25.88% relative to the fixed nacelle configuration.

Both active controllers recovered nearly all of this theoretical opportunity. The persistence controller generated 5,453.086 MWh before actuator energy costs and 5,447.825 after subtracting 5.261 MWh of yaw energy consumption. This represents a 24.6% increase in net energy relative to fixed nacelle case and corresponds to a 98.89% of the perfect yaw energy.

Persistence therefore achieved slightly higher net energy than the LSTM. The difference was small in absolute terms, with persistence producing a slightly higher net energy than the LSTM corresponding to approximately 0.051% of persistence's net energy. LSTM required fewer yaw events (22,510 versus 22,753 for persistence) and less total yaw travel ($182,882.75^\circ$ versus $186,437.65^\circ$). Thus, in this simulation, the more complex predictive controller did not provide an energy advantage over the simple persistence strategy.

Overall the results demonstrate that both active nacelle strategies substantially improved turbine energy capture relative to a fixed nacelle. The persistence controller marginally outperformed the LSTM in net energy, while LSTM required slightly fewer total yaw travel, consistent with the forecasting results in which persistence lowered directional prediction error.

Discussion

The principal finding of this study is that the LSTM based predictive controller did not outperform simpler forecasting strategies for short term wind direction prediction at the Amarillo site. Persistence achieved a lower locked test forecasting error than the LSTM, and the difference was consistent across wind speed and directional change regimes. The multi seed analysis demonstrated that this result was not dependent on a single LSTM initialization. However the difference in forecasting accuracy was marginal in simulated energy capture: persistence achieved 5,447.825 MWh of net energy compared with 5,445.065 MWh for the LSTM. Both controllers therefore captured approximately 99% of the idealized perfect yaw energy. This result illustrates an important distinction between forecasting performance and downstream control performance. Previous work has emphasized that wind-direction variability, actuator constraints and the temporal evolution of the wind can determine the practical value of predictive control⁹.

Table 1 Yaw Strategy Comparison

Controller	Gross Energy	Yaw Energy	Net Energy	Net vs Fixed	Perfect Yaw (%)
Fixed	4372.157	0.000	4372.157	0.00%	79.44%
Persistence	5453.086	5.261	5447.825	24.61%	98.99%
LSTM	5450.346	5.281	5445.065	24.54%	98.99%
Perfect yaw	5503.685	0.000	5503.685	25.88%	100.00%

The strength of the persistence baseline can be explained partially by the temporal characteristics of the wind direction data. Consecutive five minute observations exhibited a median absolute directional change of only 0.92° . 75.5% of changes were no greater than 2° and 92.7% were no greater than 5° . Under such conditions, the most recent observation provides a highly informative estimate of the immediate future direction. A more complex nonlinear model faces the benchmark of being able to identify predictive structure beyond an already highly informative recent observation. The importance of explicitly characterizing wind-direction variability before designing a predictive controller is consistent with recent control-oriented wind-direction research⁹.

Forecast and controller yaw error are related but not identical. The LSTM provides a target direction rather than directly setting the nacelle heading. The actuator subsequently attempts to reach that target subject to its movement limits. Precision error can be reduced, preserved or amplified by actuator dynamics. This distinction explains why the difference in forecasting accuracy between persistence and LSTM was larger than the resulting difference in simulated net energy. This distinction is consistent with recent work showing that the usefulness of wind prediction depends not only on forecast quality but also on the rotational constraints of the turbine.

These results also demonstrate that forecasting accuracy and control performance are not equivalent objectives. Though the LSTM demonstrated larger directional prediction error, its actuator constrained simulation produced almost the same energy capture as persistence. This indicates that small differences in forecast error do not necessarily produce proportionally large differences in turbine level energy output. This is more apparent when the yaw controller and actuator constraints limit the practical effect of forecast differences. Thus, model selection for predictive yaw control should also consider downstream control and energy objectives. This supports the broader view that predictive models for wind-energy control should ultimately be assessed in the context of the physical control system rather than through forecasting metrics alone^{9,25}.

Limitations

The analysis uses a single year of wind data from one location: Amarillo, Texas. The observed persistence of wind di-

rection and the relative performance of the forecasting models may therefore depend on the local meteorological regime and may not generalize to other sites or climates. The conclusion observed distribution of five minute directional changes, but additional years and locations would be required to determine whether the same relationship holds more broadly.

The energy model does not represent a specific commercial wind turbine or a turbine-specific certified power curve. Instead, it uses a simplified wind-speed-dependent power relationship and yaw-loss model to provide a comparative estimate of controller performance. Therefore, the reported MWh values should not be interpreted as precise predictions of the annual energy production of a real turbine. Their primary purpose is to quantify relative differences between controllers operating under identical wind conditions and modeling assumptions^{5,16}.

The selected yaw actuator speed represents a simulation assumption rather than a universal turbine characteristic. Different actuator speeds could alter the relative performance of persistence and LSTM control by changing how quickly forecasted or observed directional changes can be followed. Sensitivity to actuator speed should therefore be investigated in future work.

The forecasting task was limited to one step ahead prediction at five minute resolution. Longer prediction horizons may provide greater opportunity for predictive models to outperform persistence, especially if future directional changes contain an exploitable temporal structure. The present results therefore should not be interpreted as evidence that machine learning forecasting is ineffective for all predictive yaw control applications.

The yaw control evaluation was performed using a simulation instead of a real turbine. The simulated turbine, power relationship, yaw actuator constraints and actuator energy consumption represent approximations of real turbine behavior. Effects such as turbulence induced loads, wake interactions, structural dynamics, controller implementation details, and site specific turbine characteristics were not fully represented^{26,27}.

The LSTM architecture and hyperparameters represent one particular model configuration. Although the five seed analysis provides evidence that the observed result was not caused by a single random initialization, other architectures, input features, forecast horizons, or training procedures could pro-

duce different results. The conclusions therefore apply specifically to the experimental configuration evaluated in this study.

Conclusion

This study investigated whether short term machine learning based wind direction forecasting could improve wind turbine yaw control relative to simpler strategies. Using one year of five minute Amarillo wind data, a chronologically separated validation and locked test framework was used to compare persistence, statistical forecasting baselines and an LSTM model before evaluating their performance in an actuator constrained yaw simulation.

Persistence proved to be a particularly strong forecasting baseline, achieving a locked test circular MAE of 1.29° , compared with 2.09° for the LSTM. The difference remained consistent across wind speed and directional change regimes, while five independent LSTM training seeds produced a mean test MAE of $2.02 \pm 0.24^\circ$. A block bootstrap analysis further indicated that the LSTM's higher forecasting error was statistically distinguishable from persistence.

Both controllers substantially improved simulated turbine energy capture relative to a fixed nacelle. Persistence achieved 5,447.825 MWh of energy, while the LSTM achieved 5,445.065 MWh. This represents improvements of 24.61% and 24.54% respectively over the fixed nacelle baseline. The idealized perfect yaw case produced 5,503.685 MWh. Thus, both active strategies approached the simulated upper bound, while persistence marginally outperformed the more complex LSTM.

These results demonstrate that increasing forecasting model complexity does not necessarily improve downstream yaw control performance. For the short five minute prediction horizon examined here, the gradual evolution of wind direction made persistence highly competitive. This finding is consistent with recent research emphasizing that the value of wind prediction depends on wind-direction variability and the rotational constraints of the turbine⁹. Future work should evaluate longer forecast horizons, multiple sites and years, more comprehensive meteorological inputs, and realistic turbine and actuator models to determine conditions under which predictive machine learning can provide a measurable advantage over simple forecasting strategies.

References

- 1 C. Draxl, A. Clifton, B. M. Hodge, J. McCaa. The Wind Integration National Dataset (WIND) Toolkit. *Applied Energy*. Vol. 151, pg. 355–366, 2015, <https://doi.org/10.1016/j.apenergy.2015.03.121>.
- 2 A. Kusiak, W. Li, Z. Song. Control of wind turbines: past, present, and future. *Renewable Energy*. Vol. 35, pg. 2679–2688, 2010, <https://doi.org/10.1016/j.renene.2010.05.010>.
- 3 K. E. Johnson, L. Y. Pao, M. J. Balas, L. J. Fingersh. Control of variable-speed wind turbines: standard and adaptive techniques for maximizing energy capture. *IEEE Control Systems Magazine*. Vol. 26, pg. 70–81, 2006, <https://doi.org/10.1109/MCS.2006.1636312>.
- 4 F. Bernardoni, M. A. Rotea, S. Leonardi. Impact of yaw misalignment on turbine loads in the presence of wind farm blockage. *Wind Energy*. Vol. 27, pg. 535–548, 2024, <https://doi.org/10.1002/we.2899>.
- 5 P. Hulsman, C. Sucameli, V. Petrović, A. Rott, A. Gerdas, M. Kühn. Turbine power loss during yaw-misaligned free-field tests at different atmospheric conditions. *Journal of Physics: Conference Series*. Vol. 2265, pg. 032074, 2022, <https://doi.org/10.1088/1742-6596/2265/3/032074>.
- 6 A. Rott, L. Höning, P. Hulsman, L. J. Lukassen, C. Moldenhauer, M. Kühn. Wind vane correction during yaw misalignment for horizontal-axis wind turbines. *Wind Energy Science*. Vol. 8, pg. 1755–1770, 2023, <https://doi.org/10.5194/wes-8-1755-2023>.
- 7 J. Liew, A. M. Urbán, S. J. Andersen. Analytical model for the power–yaw sensitivity of wind turbines operating in full wake. *Wind Energy Science*. Vol. 5, pg. 427–437, 2020, <https://doi.org/10.5194/wes-5-427-2020>.
- 8 E. A. Bossanyi. The design of closed loop controllers for wind turbines. *Wind Energy*. Vol. 3, pg. 149–163, 2000.
- 9 S. Dallas, A. Stock, E. Hart. Control-oriented modelling of wind direction variability. *Wind Energy Science*. Vol. 9, pg. 841–858, 2024, <https://doi.org/10.5194/wes-9-841-2024>.
- 10 R. Fadaeinedjad, G. Moschopoulos, M. Moallem. Comparison of fuzzy logic and conventional PID controllers for wind turbine control. *Renewable Energy*. Vol. 34, pg. 1–10, 2009.
- 11 B. Jing, Z. Qian, Y. Pei, L. Zhang, T. Yang. Improving wind turbine efficiency through detection and calibration of yaw misalignment. *Renewable Energy*. Vol. 160, pg. 1217–1227, 2020, <https://doi.org/10.1016/j.renene.2020.07.063>.
- 12 S. Hochreiter, J. Schmidhuber. Long short-term memory. *Neural Computation*. Vol. 9, pg. 1735–1780, 1997, <https://doi.org/10.1162/neco.1997.9.8.1735>.
- 13 Z. Tang, G. Zhao, T. Ouyang. Two-phase deep learning model for short-term wind direction forecasting. *Renewable Energy*. Vol. 173, pg. 1005–1016, 2021, <https://doi.org/10.1016/j.renene.2021.04.041>.
- 14 A. P. Sari, H. Suzuki, T. Kitajima, T. Yasuno, D. A. Prasetya, A. Rabi'b. Deep convolutional long short-term memory for forecasting wind speed and direction. *SICE Journal of Control, Measurement, and System Integration*. Vol. 14, pg. 30–38, 2021, <https://doi.org/10.1080/18824889.2021.1894878>.
- 15 S. Macrí, S. Aubrun, A. Leroy, N. Girard. Experimental investigation of wind turbine wake and load dynamics during yaw maneuvers. *Wind Energy Science*. Vol. 6, pg. 585–606, 2021, <https://doi.org/10.5194/wes-6-585-2021>.
- 16 S. Tamaro, F. Campagnolo, C. L. Bottasso. On the power and control of a misaligned rotor – beyond the cosine law. *Wind Energy Science*. Vol. 9, pg. 1547–1575, 2024, <https://doi.org/10.5194/wes-9-1547-2024>.
- 17 E. J. Aju, D. Kumar, M. Leffingwell, M. A. Rotea, Y. Jin. The influence of yaw misalignment on turbine power output fluctuations and unsteady aerodynamic loads within wind farms. *Renewable Energy*. Vol. 215, pg. 118894, 2023, <https://doi.org/10.1016/j.renene.2023.06.015>.
- 18 T. Li, J. Yang, A. Ioannou. Wind forecasting-based model predictive control of generator, pitch, and yaw for output stabilisation – A 15-megawatt offshore wind turbine. *Energy Conversion and Management*. Vol. 302, pg. 118155, 2024, <https://doi.org/10.1016/j.enconman.2024.118155>.
- 19 L. P. Joseph, R. C. Deo, R. Prasad, S. Salcedo-Sanz, N. Raj, J. Soar. Near real-time wind speed forecast model with bidirectional LSTM networks.

-
- Renewable Energy. Vol. 204, pg. 39–58, 2023, <https://doi.org/10.1016/j.renene.2022.12.123>.
- 20 I. Karijadi, S.-Y. Chou, A. Dewabharata. Wind power forecasting based on hybrid CEEMDAN-EWT deep learning method. *Renewable Energy*. Vol. 218, pg. 119357, 2023.
- 21 A. Stieren, S. N. Gadde, R. J. A. M. Stevens. Modeling dynamic wind direction changes in large eddy simulations of wind farms. *Renewable Energy*. Vol. 170, pg. 1342–1352, 2021, <https://doi.org/10.1016/j.renene.2021.02.018>.
- 22 M. N. Lang, L. Schlosser, T. Hothorn, G. J. Mayr, R. Stauffer, A. Zeileis. Circular regression trees and forests with an application to probabilistic wind direction forecasting. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*. Vol. 69, pg. 1357–1374, 2020, <https://doi.org/10.1111/rssc.12437>.
- 23 R. Thedin, G. Barter, J. Jonkman, R. Mudafort, C. J. Bay, K. Shaler, J. Kreeft. Load assessment of a wind farm considering negative and positive yaw misalignment for wake steering. *Wind Energy Science*. Vol. 10, pg. 1033–1053, 2025, <https://doi.org/10.5194/wes-10-1033-2025>.
- 24 E. Kadoche, P. Bianchi, F. Carton, P. Ciblat, D. Ernst. On the importance of wind predictions in wake steering optimization. *Wind Energy Science*. Vol. 9, pg. 1577–1594, 2024, <https://doi.org/10.5194/wes-9-1577-2024>.
- 25 F. Bestehorn, F. Bürgel, C. Kirches, S. Stiller, A. M. Tillmann. Integer programming for optimal yaw control of wind farms. *Wind Energy Science*. Vol. 10, pg. 1637–1660, 2025, <https://doi.org/10.5194/wes-10-1637-2025>.
- 26 E. Simley, P. Fleming, N. Girard, L. Alloin, E. Godefroy, T. Duc. Results from a wake-steering experiment at a commercial wind plant: investigating the wind speed dependence of wake-steering performance. *Wind Energy Science*. Vol. 6, pg. 1427–1453, 2021, <https://doi.org/10.5194/wes-6-1427-2021>.
- 27 M. J. van den Broek, D. De Tavernier, P. Hulsman, D. C. van der Hoek, B. Sanderse, J.-W. van Wingerden. Free-vortex models for wind turbine wakes under yaw misalignment – a validation study on far-wake effects. *Wind Energy Science*. Vol. 8, pg. 1909–1925, 2023, <https://doi.org/10.5194/wes-8-1909-2023>.