
Graduate Certificate in AI for Renewable Energy Forecasting

AI in Wind Energy Predictions

Artificial Intelligence (AI): The simulation of human intelligence in machines that are programmed to think like humans and mimic their actions. The term may also be applied to any machine that exhibits traits associated with a human mind such as learning and problem-solving.

Renewable Energy Forecasting: The use of statistical and machine learning techniques to predict the availability of renewable energy sources, such as wind and solar, in the future. Accurate renewable energy forecasting is crucial for the effective integration of renewable energy into the power grid.

Wind Energy Predictions: The use of AI and other techniques to predict the amount of energy that can be generated by a wind turbine or wind farm at a given point in time. Wind energy predictions can help energy operators to more effectively manage the power grid and ensure that the right amount of power is available to meet demand.

Artificial Neural Networks (ANNs): A type of machine learning model inspired by the structure and function of the human brain. ANNs are composed of interconnected nodes, or "neurons," that work together to learn and make predictions.

Deep Learning: A subset of machine learning that is based on the use of artificial neural networks with multiple layers. Deep learning models are able to learn and make predictions based on large amounts of data, making them well-suited for tasks such as wind energy predictions.

Long Short-Term Memory (LSTM): A type of recurrent neural network (RNN) that is capable of learning and making predictions based on data that changes over time. LSTMs are well-suited for tasks such as wind energy predictions, as they are able to take into account historical data as well as current conditions.

Convolutional Neural Networks (CNNs): A type of neural network that is designed to process data that has a grid-like structure, such as an image. CNNs are composed of convolutional layers, which apply a series of filters to the input data to extract features, and pooling layers, which reduce the dimensionality of the data.

Support Vector Machines (SVMs): A type of supervised learning algorithm that can be used for both classification and regression tasks. SVMs work by finding the optimal boundary, or "hyperplane," that separates data points into different classes.

Random Forests: A type of ensemble learning algorithm that combines the predictions of multiple decision trees to make a final prediction. Random forests are robust to overfitting and can handle large amounts of data with high dimensionality.

Gradient Boosting: A type of ensemble learning algorithm that combines the predictions of multiple weak models to make a final prediction. Gradient boosting algorithms work by iteratively adding new models to the ensemble, with each new model focusing on correcting the errors of the previous models.

Feature Engineering: The process of selecting and transforming the input data, or "features," in a machine learning model in order to improve its performance. Feature engineering can involve tasks such as scaling and normalization, dimensionality reduction, and the creation of new features.

Data Preprocessing: The process of cleaning and transforming raw data into a format that can be used by a machine learning model. Data preprocessing may include tasks such as handling missing values, outlier detection and removal, and the conversion of categorical data into a numerical format.

Overfitting: A common problem in machine learning where a model learns the training data too well and performs poorly on new, unseen data. Overfitting can be caused by a model that is too complex, a lack of data, or a lack of regularization.

Regularization: A technique used to prevent overfitting in machine learning models by adding a penalty term to the loss function. Regularization encourages the model to learn simpler patterns in the data, which can improve its performance on new, unseen data.

Cross-Validation: A technique used to evaluate the performance of a machine learning model by splitting the data into multiple folds and training and testing the model on each fold. Cross-validation can help to reduce the risk of overfitting and provide a more accurate estimate of the model's performance.

Hyperparameter Tuning: The process of selecting the optimal values for the hyperparameters of a machine learning model in order to improve its performance. Hyperparameters are parameters that are set before training and cannot be learned from the data, such as the learning rate or the number of layers in a neural network.

Bias-Variance Tradeoff: The balance that must be struck between the bias, or assumptions, of a machine learning model and the variance, or complexity, of the model. A model with high bias will underfit the data, while a model with high variance will overfit the data.

Supervised Learning: A type of machine learning in which the model is trained on labeled data, meaning that the correct answer, or "label," is provided for each example in the training data. Supervised learning algorithms are typically used for tasks such as classification and regression.

Unsupervised Learning: A type of machine learning in which the model is trained on unlabeled data, meaning that the correct answer is not provided for each example in the training data. Unsupervised learning algorithms are typically used for tasks such as clustering and dimensionality reduction.

Reinforcement Learning: A type of machine learning in which the model, or "agent," learns to make decisions by interacting with an environment and receiving rewards or penalties for its actions.

Reinforcement learning algorithms are typically used for tasks such as game playing and robot control.

Markov Property: The property of a stochastic process where the future state depends only on the current state and not on the past states. The Markov property is the basis for many types of time series models, including hidden Markov models (HMMs) and Markov chains.

Hidden Markov Models (HMMs): A type of statistical model that is used to represent a system that is governed by a set of unobserved, or "hidden," states. HMMs are well-suited for tasks such as wind energy predictions, as they can take into account the underlying physical processes that govern the behavior of the wind.

Markov Chains: A type of stochastic process that is used to model a system that transitions between a finite number of states. Markov chains are often used in wind energy predictions to model the transitions between different wind speed regimes.

Autoregressive Integrated Moving Average (ARIMA): A type of time series model that is used to forecast future values based on past values and errors. ARIMA models are well-suited for tasks such as wind energy predictions, as they can take into account the autocorrelation and trend in the data.

Prophet: An open-source software package for forecasting time series data that is based on decomposable time series models and additive regression models. Prophet is well-suited for tasks such as wind energy predictions, as it can handle multiple seasonalities and trends in the data.

Wind Rose: A graphical representation of the wind speed and direction at a particular location. Wind roses are often used in wind energy predictions to visualize the distribution of wind speeds and directions at a wind farm.

Wind Shear: The variation in wind speed and direction with height above the ground. Wind shear can have a significant impact on the performance of wind turbines, as it can affect the amount of power that can be generated.

Turbulence: The chaotic motion of the wind, which can be caused by a variety of factors such as atmospheric stability, surface roughness, and the presence of obstacles. Turbulence can have a significant impact on the performance of wind turbines, as it can affect the loads on the turbine blades and the amount of power that can be generated.

Wake Effect: The reduction in wind speed and increase in turbulence that occurs downstream of a wind turbine. The wake effect can have a significant impact on the performance of downstream turbines, as it can reduce the amount of power that can be generated.

Power Curve: A graph that shows the relationship between the wind speed and the power output of a wind turbine. The power curve is an important parameter in wind energy predictions, as it determines the amount of power that can be generated at different wind speeds.

Cut-In Speed: The minimum wind speed at which a wind turbine begins to generate power. The cut-in speed is an important parameter in wind energy predictions, as it determines the minimum wind