

93

Apr-25

Scoring Indicators

COURSE NAME: MACHINE LEARNING AND NEURAL NETWORKS

P-15

COURSE CODE: 4342

QID: 2102250068-B

Q. No.		Scoring Indicators	Split Score	Sub Total	Total Score
PART A					9
I	1	Machine learning is a branch of artificial intelligence that enables algorithms to uncover hidden patterns within datasets.	1	1	
I	2	Agglomerative clustering is a hierarchical clustering method that groups similar items into clusters.	1	1	
I	3	K-Means Clustering is an Unsupervised Machine Learning algorithm which groups the unlabeled dataset into different clusters.	1	1	
I	4	Hierarchical clustering, Centroids-based Clustering (Partitioning methods), Distribution-based Clustering, Density-based Clustering (Model-based methods), Fuzzy Clustering, Constraint-based (Supervised Clustering)	1	1	
I	5	Mean squared error (MSE) is a metric used to measure the average squared difference between the predicted values and the actual values in the dataset.	1	1	
I	6	Reinforcement Learning (RL) is a branch of machine learning focused on making decisions to maximize cumulative rewards in a given situation.	1	1	
I	7	Social media, marketing and sales	1	1	
I	8	An activation function is a mathematical function applied to the output of a neuron.	1	1	
I	9	Batch normalization is a deep learning technique that helps our models learn and adapt quickly.	1	1	
PART B					24
II	1	Automation and Efficiency: Machine learning algorithms can automate repetitive and time-consuming tasks, increasing efficiency and allowing humans to focus on higher value-added tasks. Improved Decision-Making: Machine learning allows companies to make accurate predictions and informed decisions based on specific data and detailed analysis. Personalization: Machine learning-based technologies, such as recommendation systems, allow personalized products and services that better meet individual users' needs and preferences.		3	

II	2	<p>Machine Learning techniques are divided mainly into the following 4 categories:</p> <ol style="list-style-type: none"> 1. Supervised Learning Supervised learning is applicable when a machine has sample data, i.e., input as well as output data with correct labels. Correct labels are used to check the correctness of the model using some labels and tags. Supervised learning technique helps us to predict future events with the help of past experience and labeled examples. 2. Unsupervised Learning In unsupervised learning, a machine is trained with some input samples or labels only, while output is not known. The training information is neither classified nor labeled; hence, a machine may not always provide correct output compared to supervised learning. 3. Reinforcement Learning Reinforcement Learning is a feedback-based machine learning technique. In such type of learning, agents (computer programs) need to explore the environment, perform actions, and on the basis of their actions, they get rewards as feedback. 4. Semi-supervised Learning Semi-supervised Learning is an intermediate technique of both supervised and unsupervised learning. It performs actions on datasets having few labels as well as unlabeled data. 		3	
II	3	<p>Data normalization is a vital pre-processing, mapping, and scaling method that helps forecasting and prediction models become more accurate. The current data range is transformed into a new, standardized range using this method. Normalization is extremely important when it comes to bringing disparate prediction and forecasting techniques into harmony. Data normalization improves the consistency and comparability of different predictive models by standardizing the range of independent variables or features within a dataset, leading to more steady and dependable results.</p> <p>Normalisation, which involves reshaping numerical columns to conform to a standard scale, is essential for datasets with different units or magnitudes across different features. Finding a common scale for the data while maintaining the intrinsic variations in value ranges is the main goal of normalization. This usually entails rescaling the features to a standard range, which is typically between 0 and 1. Alternatively, the features can be adjusted to have a mean of 0 and a standard deviation of 1.</p> <p>Z-Score Normalisation (Standardisation) and Min-Max Scaling are two commonly used normalisation techniques. In order to enable more insightful and precise analyses in a variety of predictive modelling scenarios, these techniques are essential in bringing disparate features to a comparable scale.</p>		3	
II	4	<p>Data normalization is a vital pre-processing, mapping, and scaling method that helps forecasting and prediction models become more accurate. The current data range is transformed into a new, standardized range using this method. Normalization is extremely important when it comes to bringing disparate prediction and forecasting techniques into harmony. Data normalization improves the consistency and comparability of different predictive models by</p>	3	3	

		<p>standardizing the range of independent variables or features within a dataset, leading to more steady and dependable results.</p> <p>Normalisation, which involves reshaping numerical columns to conform to a standard scale, is essential for datasets with different units or magnitudes across different features. Finding a common scale for the data while maintaining the intrinsic variations in value ranges is the main goal of normalization. This usually entails rescaling the features to a standard range, which is typically between 0 and 1. Alternatively, the features can be adjusted to have a mean of 0 and a standard deviation of 1.</p> <p>Z-Score Normalisation (Standardisation) and Min-Max Scaling are two commonly used normalisation techniques. In order to enable more insightful and precise analyses in a variety of predictive modelling scenarios, these techniques are essential in bringing disparate features to a comparable scale.</p>			
II	5	<ol style="list-style-type: none"> 1. Determine the Suitability of Data for Factor Analysis 2. Choose the Extraction Method 3. Factor Extraction 4. Determine the Number of Factors to Retain 5. Factor Rotation 6. Interpret and Label the Factors 7. Compute Factor Scores (if needed) 8. Report and Validate the Results 	3	3	
II	6	<p>The main drawbacks of partitional clustering in machine learning are: requiring the pre-defined number of clusters (k), sensitivity to initialization, potential for getting stuck in local optima, difficulty handling non-spherical clusters, and poor performance with unevenly sized clusters; essentially, it struggles when clusters have complex shapes or sizes, and the optimal number of clusters is not readily known, making it less adaptable to diverse data structures.</p> <p>Key points about partitional clustering limitations:</p> <p>Need to specify 'k':</p> <p>Unlike hierarchical clustering, partitional algorithms like K-means require users to specify the desired number of clusters beforehand, which can be challenging to determine accurately.</p> <p>Initialization dependence:</p> <p>The initial placement of cluster centroids can significantly impact the final clustering results, potentially leading to suboptimal solutions if not carefully chosen.</p> <p>Sensitivity to outliers:</p> <p>Outliers can significantly distort the clustering process, as the algorithm tries to minimize distances to these extreme points, affecting the overall cluster quality.</p> <p>Assumption of spherical clusters:</p> <p>Partitional methods typically assume clusters are roughly spherical and evenly distributed, which may not be true for real-world data with complex shapes.</p> <p>Poor performance with uneven sizes:</p> <p>If clusters have vastly different sizes, partitional clustering may struggle to accurately identify and separate them.</p>	3	3	
II	7	<p>Types of Reinforcements in RL</p> <ol style="list-style-type: none"> 1. Positive Reinforcement 	3	3	

		<p>Positive Reinforcement is defined as when an event, occurs due to a particular behavior, increases the strength and the frequency of the behavior. In other words, it has a positive effect on behavior.</p> <p>Advantages: Maximizes performance, helps sustain change over time. Disadvantages: Overuse can lead to excess states that may reduce effectiveness.</p> <p>2. Negative Reinforcement Negative Reinforcement is defined as strengthening of behavior because a negative condition is stopped or avoided.</p> <p>Advantages: Increases behavior frequency, ensures a minimum performance standard. Disadvantages: It may only encourage just enough action to avoid penalties.</p>			
II	8	<p>A loss function in machine learning is a mathematical function that measures how well a model's predictions match the actual results. It's used to guide the model's training process.</p> <p>Types of loss functions Mean absolute error (MAE): Calculates the average magnitude of errors in a set of predicted values Mean squared error (MSE): Calculates the square of the difference between predicted and actual value Huber loss: Combines the strengths of both MAE and MSE Cross entropy loss: Used when only two classes are in the dataset, such as spam or not spam Categorical cross-entropy loss: Measures the difference between classes Kullback-Leibler divergence (KL divergence): Measures how one probability distribution differs from a reference probability distribution Hinge loss: Commonly used in support vector machines (SVMs) for classification tasks</p>	3	3	
II	9	<p>1. ANN is the mathematical model which is mainly inspired by the biological neuron system in the human brain. In contrast, the biological neural network is also composed of several processing pieces known as neurons that are linked together via synapses.</p> <p>2. An artificial neural network's processing was sequential and centralized. In contrast, a biological neural network processes information in parallel and distributive.</p> <p>3. The artificial neural network is of a much smaller size than the biological neural network. In contrast, the biological neural network is large in size.</p> <p>4. The biological neural network is fault tolerant. In contrast, the artificial neural network is not fault tolerant.</p> <p>5. The processing speed of an artificial neural network is in the nanosecond range, which is faster than the biological neural network, where the cycle time associated with a neural event triggered by an external input is in the millisecond range.</p> <p>6. BNN may perform more difficult issues than artificial neural networks.</p> <p>7. The operating environment of the artificial neural network is well-defined and well-constrained. In contrast, the operating environment of the biological neural network is poorly defined and unconstrained.</p>	any 3	3	

		8.The reliability of the artificial neural network is very vulnerable. In contrast, the reliability of the biological neural network is robust.																					
II	10	<table><tr><td>Comparison Attribute</td><td>Feed-forward Neural Networks</td><td>Recurrent Neural Networks</td></tr><tr><td>Signal flow direction</td><td>Forward only</td><td>Bidirectional</td></tr><tr><td>Delay introduced</td><td>No</td><td>Yes</td></tr><tr><td>Complexity</td><td>Low</td><td>High</td></tr><tr><td>Neuron independence in the same layer</td><td>Yes</td><td>No</td></tr><tr><td>Speed</td><td>High</td><td>slow</td></tr></table>	Comparison Attribute	Feed-forward Neural Networks	Recurrent Neural Networks	Signal flow direction	Forward only	Bidirectional	Delay introduced	No	Yes	Complexity	Low	High	Neuron independence in the same layer	Yes	No	Speed	High	slow	any 3	3	
Comparison Attribute	Feed-forward Neural Networks	Recurrent Neural Networks																					
Signal flow direction	Forward only	Bidirectional																					
Delay introduced	No	Yes																					
Complexity	Low	High																					
Neuron independence in the same layer	Yes	No																					
Speed	High	slow																					
PART A						42																	
III	III	<p>1. Data Collection</p> <p>Machine learning requires training data, a lot of it. This data can either be labelled meaning Supervised Learning or not labelled meaning Unsupervised Learning.</p> <p>Accuracy of the model depends on the quality and quantity of the data. The outcome of this step is generally a representation of data which will be used for training.</p> <p>Using pre-collected data, by way of datasets from sites like Kaggle, UCI, etc. forms the basis of you Machine learning project. You may also collect data through user-surveys, analysis reports, trends, usage metrics, etc.</p> <p>2. Data Preparation</p> <p>We cannot work on raw data. Data needs to be processed by normalization, removing duplicates, errors and biases.</p> <p>Visualising data can be helpful in searching for patterns and outliers to check if the data collected is right or if it contains missing values. This can be done using libraries like seaborn, matplotlib, etc. Visualize data to help detect relevant relationships between variables or class imbalances, or perform other exploratory analysis.</p> <p>After performing data wrangling, we need to prepare the data for training. Cleaning of data is done that involves steps like removing duplicates, dealing with missing values, type conversions, correcting errors, normalizing the data, etc.</p> <p>Not all the above steps are needed to be performed as it depends entirely on the data collected. Some datasets may not require data</p>	7	7																			

		<p>preparation at all while for some data preparation step takes majority of their ML model build time.</p> <p>We can also Randomize data, which erases the effects of the particular order in which we collected and/or otherwise prepared our data. Later we can split the data into training, testing and evaluation sets</p> <p>3. Choose a Model / Algorithm The third step consists of selecting the right model. There are many models which can be used for many different purposes. Once the model is selected, it needs to meet the business goal.</p> <p>We need to have an idea about the preparation the model requires along with its accuracy and scalability. Having a complex model does not mean a better model.</p> <p>Common machine learning algorithms include Decision Trees, Random Forest, Linear Regression, Support Vector Machines (SVM), Logistic Regression, K-means, Principal Component Analysis (PCA), Naïve Bayes, and Neural Networks. Different algorithms need to be applied to different tasks, you need to choose the correct one for your use case.</p> <p>4. Training the Model Training a model forms the basis of machine learning. The goal is to use our training data and improve the predictions of our model.</p> <p>Every cycle in training a model involves updating the weights and biases in each training step. We can use labelled sample data in case supervised machine learning and unlabelled sample data for unsupervised learning.</p> <p>The goal of training is to evaluate and further improve our model accuracy and performance. Training happens in the form of iterations which is called a training step.</p> <p>5. Evaluate the Model After training the model comes evaluating the model. The larger the number of variables in the real world, the bigger the training and test data should be.</p> <p>Performance metrics are used to measure the performance of the model. These include precision, recall, accuracy, specificity, etc.</p> <p>The model is then tested against previously unseen data. The unseen data is meant to act as representative of model performance in the real world, but still helps tune the model (as opposed to test data, which does not).</p> <p>A 70/30 split, or similar, is considered a good train/eval split, which depends on things like data availability, dataset features, domain, etc.</p> <p>6. Parameter Tuning The original model parameters need to be tested after evaluating your model. By increasing the training, it can lead to better results.</p>		
--	--	---	--	--

		<p>Parameter tuning is an experimental process and hence we need to define when to stop parameter tuning otherwise it will continue to tweak the model.</p> <p>Hyperparameter tuning is an art and one that requires patience & experience. Once the model parameters are tuned it can give us better results. Some common hyperparameters include: number of training steps, learning rate, initialization values and distribution, etc.</p> <p>7. Make Predictions After the processes of collecting data, preparing the data, selecting a machine learning algorithm, training the model and evaluating the model & tuning the parameters, we need to make predictions.</p> <p>Our machine learning model can make predictions ranging from image recognition to predictive analytics to natural language processing.</p> <p>After building the model needs to be tested on a testing set to check how the model performs on unseen data. It helps to further evaluate the model and provides better approximation.</p>			
IV		<p>The integrity of data analysis is highly dependent on the quality of data preprocessing. Data preprocessing determines the usability and interpretability of data, laying the groundwork for accurate machine learning and AI models.</p> <p>Eliminating Errors Cleaning is a pivotal data preprocessing technique. It allows you to eliminate errors, impute missing values, and rectify inconsistencies. For example, a customer dataset with redundant entries due to technical mistakes would undergo cleaning to ensure each customer record is unique and accurately represented.</p> <p>Making Data Uniform Normalization is comparable to establishing a level playing field, where disparate measures are adjusted to a uniform scale, enabling equitable comparisons. For instance, normalization can help you analyze the performance of stocks from different countries despite stock prices being available in various currencies and scales. With normalization techniques such as min-max, you can convert all stock prices into a common currency, for example, USD, and then apply a min-max scaling to compare the relative performance of stocks on a uniform scale.</p> <p>Finding Hidden Patterns Diligent preprocessing can reveal concealed patterns and insights. A marketing team analyzing social media data can identify peak engagement times aligned with spam activity. However, excluding anomalies through data cleaning will allow you to pinpoint genuine peak engagement periods and optimize strategy.</p> <p>Big Data Preprocessing As datasets grow in size and complexity, preprocessing becomes even more critical. Big data has a large volume, is</p>			
	IV		7	7	

		<p>heterogeneous, and needs to be processed rapidly. Preprocessing transforms raw big data into a cleaner, more structured format, removing noise and making it easier to process.</p> <p>Similarly, advanced techniques such as parallel processing, distributed computing, and automated preprocessing pipelines are indispensable for processing big data effectively.</p>			
V		<p>A Random Forest is a collection of decision trees that work together to make predictions. In this article, we'll explain how the Random Forest algorithm works and how to use it.</p> <p>Understanding Intuition for Random Forest Algorithm Random Forest algorithm is a powerful tree learning technique in Machine Learning to make predictions and then we do voting of all the trees to make prediction. They are widely used for classification and regression task.</p> <p>It is a type of classifier that uses many decision trees to make predictions. It takes different random parts of the dataset to train each tree and then it combines the results by averaging them. This approach helps improve the accuracy of predictions. Random Forest is based on ensemble learning.</p> <p>Imagine asking a group of friends for advice on where to go for vacation. Each friend gives their recommendation based on their unique perspective and preferences (decision trees trained on different subsets of data). You then make your final decision by considering the majority opinion or averaging their suggestions (ensemble prediction).</p> <div data-bbox="343 1149 1069 1668"> </div> <p>Key Features of Random Forest Handles Missing Data: Automatically handles missing values during training, eliminating the need for manual imputation. Algorithm ranks features based on their importance in making predictions offering valuable insights for feature selection and interpretability. Scales Well with Large and Complex Data without significant performance degradation. Algorithm is versatile and can be applied to both classification tasks (e.g., predicting categories) and regression tasks (e.g., predicting continuous values).</p>			

VI	<p><i>Computer vision – Mean-Shift clustering is widely used in computer vision for object tracking, image segmentation, and feature extraction.</i></p> <p><i>Image processing – Mean-Shift clustering is used for image segmentation, which is the process of dividing an image into multiple segments based on the similarity of the pixels.</i></p> <p><i>Anomaly detection – Mean-Shift clustering can be used for detecting anomalies in data by identifying the areas with low density.</i></p> <p><i>Customer segmentation – Mean-Shift clustering can be used for customer segmentation in marketing by identifying groups of customers with similar behavior and preferences.</i></p> <p><i>Social network analysis – Mean-Shift clustering can be used for clustering users in social networks based on their interests and interactions.</i></p>	7	7	
VII	<p>Linear Discriminant Analysis (LDA) is a dimensionality reduction and classification technique commonly used in machine learning and pattern recognition. In the context of classification it aims to find a linear combination of features that best separates different classes or categories of data. It seeks to reduce the dimensionality of the feature space while preserving as much of the class-separability information as possible.</p> <p>let's walk through a simple example to understand how Linear Discriminant Analysis (LDA) works:</p> <p>Example: Iris Flower Classification</p> <p>Suppose we have a dataset of iris flowers with four features: sepal length, sepal width, petal length, and petal width. We want to classify these flowers into three species: Setosa, Versicolor, and Virginica.</p> <p>Steps:</p> <p>Data Preparation: Let's say we have 150 iris samples with four features each, and the samples are evenly distributed among the three species.</p> <p>Compute Class Statistics: Calculate the mean and covariance matrix for each feature in each class. This gives us three mean vectors and three covariance matrices (one for each class).</p> <p>Compute Between-Class and Within-Class Scatter Matrices:</p> <p>Calculate the between-class scatter matrix by computing the differences between the mean vectors of each class and the overall mean, and then summing these outer products. Calculate the within-class scatter matrix by summing the covariance matrices of each class, weighted by the number of samples in each class.</p> <p>Compute Eigenvectors and Eigenvalues: Solve the generalized eigenvalue problem using the between-class scatter matrix and the within-class scatter matrix. This gives us a set of eigenvectors and their corresponding eigenvalues.</p> <p>Select Discriminant Directions: Sort the eigenvectors by their eigenvalues in descending order. Let's say we want to reduce the dimensionality to 2, so we select the top two eigenvectors.</p> <p>Transform Data: Project the original iris data onto the two selected eigenvectors. This gives us a new two-dimensional representation of the data.</p> <p>Classification: In the reduced-dimensional space, we can use a classifier (e.g., k-nearest neighbors) to classify the iris flowers into one of the three species based on their positions in the reduced space.</p>	7	7	

VIII		<p>Types of Multidimensional Scaling</p> <p>1. Classical Multidimensional Scaling Classical Multidimensional Scaling is a technique that takes an input matrix representing dissimilarities between pairs of items and produces a coordinate matrix that minimizes the strain. The steps of a Classical MDS algorithm include setting up the squared proximity matrix, applying double centering to compute matrix B, determining the m largest eigenvalues and corresponding eigenvectors of B, and obtaining the coordinates matrix X.</p>			
	VIII	<p>2. Metric Multidimensional Scaling Metric Multidimensional Scaling generalizes the optimization procedure to various loss functions and input matrices with known distances and weights. It minimizes a cost function called "stress," often minimized using a procedure called stress majorization.</p> <p>3. Non-metric Multidimensional Scaling Non-metric Multidimensional Scaling finds a non-parametric monotonic relationship between dissimilarities and Euclidean distances between items, along with the location of each item in the low-dimensional space. It defines a "stress" function to optimize, considering a monotonically increasing function f.</p>	7	7	
IX	IX	<p>Key Components of Q-learning</p> <p>1. Q-Values or Action-Values : These represent the expected rewards for taking an action in a specific state. The agent updates these values over time using the Temporal Difference (TD) update rule.</p> <p>2. Rewards and Episodes: The agent moves through different states by taking actions and receiving rewards. The process continues until the agent reaches a terminal state which ends the episode.</p> <p>3. Temporal Difference or TD-Update: The agent updates Q-values using the formula:</p> $Q(S,A) \leftarrow Q(S,A) + \alpha(R + \gamma Q(S',A') - Q(S,A))$ $Q(S,A) \leftarrow Q(S,A) + \alpha(R + \gamma Q(S',A') - Q(S,A))$	7	7	
X	X	<p>1. Accuracy 2. Confusion Matrix 3. Precision, Recall and F1 Score 4. ROC Curve</p> <p>1. Accuracy Classification accuracy is the simplest evaluation metric. It is defined as the number of correct predictions divided by the total number of predictions multiplied by 100. The accuracy metric works great if the target variable classes in the data are approximately balanced. For example, if 60% of the classes in an animal dataset are dogs and 40% are cats, then we can say that it is a balanced dataset. It calculates the ratio of correctly predicted instances to the total instances.</p> <p>2. Confusion Matrix The confusion matrix is another way to evaluate the performance of a classifier. Here, it counts the number of times instances of class A</p>	7	7	


		<p>are classified as class B. For example, the number of times the classifier confused images of 5s with non-5s.</p> <p>This is a table that is often used to describe the performance of a classification model. It presents a summary of the predictions made by the model against the actual class labels. The confusion matrix is a matrix with four different combinations of predicted and actual classes: True Positive (TP), True Negative (TN), False Positive (FP), and False Negative (FN).</p> <p>Let's compute the confusion matrix to evaluate the performance of a classifier. We can make use of MNIST dataset to compute the confusion matrix.</p> <p>3.1 Precision: Precision provides the accuracy of the positive prediction made by the classifier. The equation is as follows: $\text{Precision} = \frac{\text{True Positive}}{\text{True Positive} + \text{False Positive}}$</p> <p>3.2 Recall: Recall is the ratio of number of true positive predictions (correctly detected by the classifier) to the total number of actual positive instances in the dataset. It measures the completeness of positive predictions. The equation is as follows: $\text{Recall} = \frac{\text{True Positive}}{\text{True Positive} + \text{False Negative}}$</p> <p>3.3 F1 Score: The F1 score is the harmonic mean of precision and recall. It favors classifiers that have similar precision and recall. Here, the classifier will only get a high F1 score if both recall and precision are high. The equation is as follows: $F1 = 2 * (\text{Precision} * \text{Recall}) / (\text{Precision} + \text{Recall})$</p> <p>4. ROC Curve</p> <p>The Receiver Operating Characteristic (ROC) curve is a graphical representation of the performance of a classification model at various thresholds. It plots the True Positive Rate (TPR) against the False Positive Rate (FPR). The Area Under the ROC Curve (AUC-ROC) is a metric to evaluate the performance of a binary classification model. AUC-ROC value lies between 0 and 1, where a higher value indicates better performance. AUC-ROC is insensitive to class distribution and gives an aggregate measure of performance across all possible classification thresholds.</p> <p>The true positive rate is calculated as: $\text{TPR} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Negatives}}$</p>			
XI	XI	<p>Key Concepts of Reinforcement Learning</p> <p>Agent: The learner or decision-maker.</p> <p>Environment: Everything the agent interacts with.</p> <p>State: A specific situation in which the agent finds itself.</p> <p>Action: All possible moves the agent can make.</p> <p>Reward: Feedback from the environment based on the action taken.</p>	7	7	
XII	XII	<p>Artificial Neural Networks and its Applications</p> <p>Last Updated : 07 Aug, 2024</p> <p>As you read this article, which organ in your body is thinking about it? It's the brain of course! But do you know how the brain works? Well, it has neurons or nerve cells that are the primary units of both the brain and the nervous system. These neurons receive sensory input from the outside world which they process and then provide the output which might act as the input to the next neuron.</p> <p>Each of these neurons is connected to other neurons in complex arrangements at synapses. Now, are you wondering how this is related to Artificial Neural Networks ? Let's check out what they are in detail and how they learn information.</p>	7	7	

Well, Artificial Neural Networks are modeled after the neurons in the human brain. If you want to gain practical skills in Artificial Neural Networks and explore their diverse applications through our interactive live data science course , perfect for aspiring data scientists.

Artificial Neural Networks

Artificial Neural Networks contain artificial neurons which are called units . These units are arranged in a series of layers that together constitute the whole Artificial Neural Network in a system. A layer can have only a dozen units or millions of units as this depends on how the complex neural networks will be required to learn the hidden patterns in the dataset. Commonly, Artificial Neural Network has an input layer, an output layer as well as hidden layers. The input layer receives data from the outside world which the neural network needs to analyze or learn about. Then this data passes through one or multiple hidden layers that transform the input into data that is valuable for the output layer. Finally, the output layer provides an output in the form of a response of the Artificial Neural Networks to input data provided.

In the majority of neural networks, units are interconnected from one layer to another. Each of these connections has weights that determine the influence of one unit on another unit. As the data transfers from one unit to another, the neural network learns more and more about the data which eventually results in an output from the output layer.



The structures and operations of human neurons serve as the basis for artificial neural networks. It is also known as neural networks or neural nets. The input layer of an artificial neural network is the first layer, and it receives input from external sources and releases it to the hidden layer, which is the second layer. In the hidden layer, each neuron receives input from the previous layer neurons, computes the weighted sum, and sends it to the neurons in the next layer. These connections are weighted means effects of the inputs from the previous layer are optimized more or less by assigning different-different weights to each input and it is adjusted during the training process by optimizing these weights for improved model performance.

Artificial neurons vs Biological neurons

		<p>The concept of artificial neural networks comes from biological neurons found in animal brains So they share a lot of similarities in structure and function wise.</p> <p>Structure : The structure of artificial neural networks is inspired by biological neurons. A biological neuron has a cell body or soma to process the impulses, dendrites to receive them, and an axon that transfers them to other neurons. The input nodes of artificial neural networks receive input signals, the hidden layer nodes compute these input signals, and the output layer nodes compute the final output by processing the hidden layer's results using activation functions.</p> <p>Synapses : Synapses are the links between biological neurons that enable the transmission of impulses from dendrites to the cell body. Synapses are the weights that join the one-layer nodes to the next-layer nodes in artificial neurons. The strength of the links is determined by the weight value.</p> <p>Learning : In biological neurons, learning happens in the cell body nucleus or soma, which has a nucleus that helps to process the impulses. An action potential is produced and travels through the axons if the impulses are powerful enough to reach the threshold. This becomes possible by synaptic plasticity, which represents the ability of synapses to become stronger or weaker over time in reaction to changes in their activity. In artificial neural networks, backpropagation is a technique used for learning, which adjusts the weights between nodes according to the error or differences between predicted and actual outcomes.</p> <p>Biological Neuron</p> <p>Artificial Neuron</p> <p>Synaptic plasticity Backpropagations</p> <p>Activation : In biological neurons, activation is the firing rate of the neuron which happens when the impulses are strong enough to reach the threshold. In artificial neural networks, A mathematical function known as an activation function maps the input to the output, and executes activations.</p>			
XIII	XIII	<p>Basic Components of Perceptron</p> <p>A Perceptron is composed of key components that work together to process information and make predictions.</p> <p>Input Features: The perceptron takes multiple input features, each representing a characteristic of the input data.</p> <p>Weights: Each input feature is assigned a weight that determines its influence on the output. These weights are adjusted during training to find the optimal values.</p> <p>Summation Function: The perceptron calculates the weighted sum of its inputs, combining them with their respective weights.</p> <p>Activation Function: The weighted sum is passed through the Heaviside step function, comparing it to a threshold to produce a binary output (0 or 1).</p> <p>Output: The final output is determined by the activation function, often used for binary classification tasks.</p> <p>Bias: The bias term helps the perceptron make adjustments independent of the input, improving its flexibility in learning.</p>	7	7	

		<p>Learning Algorithm: The perceptron adjusts its weights and bias using a learning algorithm, such as the Perceptron Learning Rule, to minimize prediction errors.</p> <p>These components enable the perceptron to learn from data and make predictions. While a single perceptron can handle simple binary classification, complex tasks require multiple perceptrons organized into layers, forming a neural network.</p>			
XI V	XIV	<p>Backpropagation (short for "Backward Propagation of Errors") is a method used to train artificial neural networks. Its goal is to reduce the difference between the model's predicted output and the actual output by adjusting the weights and biases in the network.</p> <p>Working of Backpropagation Algorithm</p> <p>The Backpropagation algorithm involves two main steps: the Forward Pass and the Backward Pass.</p> <p>How Does the Forward Pass Work?</p> <p>In the forward pass, the input data is fed into the input layer. These inputs, combined with their respective weights, are passed to hidden layers.</p> <p>For example, in a network with two hidden layers (h1 and h2 as shown in Fig. (a)), the output from h1 serves as the input to h2. Before applying an activation function, a bias is added to the weighted inputs.</p> <p>Each hidden layer applies an activation function like ReLU (Rectified Linear Unit), which returns the input if it's positive and zero otherwise. This adds non-linearity, allowing the model to learn complex relationships in the data. Finally, the outputs from the last hidden layer are passed to the output layer, where an activation function, such as softmax, converts the weighted outputs into probabilities for classification.</p> <div data-bbox="451 1319 1042 1727"> <pre> graph LR n1((n1)) -- w1j --> h1((h1)) n2((n2)) -- w2j --> h1 n3((n3)) -- w3j --> h1 n1 -- w1j --> h2((h2)) n2 -- w2j --> h2 n3 -- w3j --> h2 h1 -- w1k --> o((o)) h2 -- w2k --> o </pre> <p style="text-align: center;">Input layer Hidden layer Output layer</p> </div> <p><i>The forward pass using weights and biases</i></p> <p>How Does the Backward Pass Work?</p> <p>In the backward pass, the error (the difference between the predicted and actual output) is propagated back through the</p>	7	7	

network to adjust the weights and biases. One common method for error calculation is the **Mean Squared Error (MSE)**, given by:

$$MSE = (\text{Predicted Output} - \text{Actual Output})^2$$

$$MSE = (\text{Predicted Output} - \text{Actual Output})^2$$

Once the error is calculated, the network adjusts weights using **gradients**, which are computed with the chain rule. These gradients indicate how much each weight and bias should be adjusted to minimize the error in the next iteration. The backward pass continues layer by layer, ensuring that the network learns and improves its performance. The activation function, through its derivative, plays a crucial role in computing these gradients during backpropagation.

Example of Backpropagation in Machine Learning

Let's walk through an example of backpropagation in machine learning. Assume the neurons use the sigmoid activation function for the forward and backward pass. The target output is 0.5, and the learning rate is 1.

