

STUDY PLAN

Week	Module	Date	Day	Delivery Type	Session Type	Session Title	Duration	Sub topics
0	Founder's Session	20-Jun-2026	Saturday					
0	Orientation Session	21-Jun-2026	Saturday					
1	Python Fundamentals	27-Jun-2026	Saturday	Live Class		Control Structures – If Else - Loops - UDF		if, elif, else statements Nested conditions for loops with range() while loops and loop control: break, continue, pass else clause with loops User-Defined Functions
1	Python Fundamentals	28-Jun-2026	Sunday	Live Class		Strings, Lists, Tuples		String objects. Creating, accessing, slicing lists List comprehension Tuples: declaration, immutability, use cases Tuple unpacking and zip() basics
1	Python Fundamentals	30-Jun-2026	Tuesday		DC + Focus			
2	Python Fundamentals	4-Jul-2026	Saturday	Live Class		Sets & Dictionaries		Creating sets, uniqueness Set operations: union(), intersection(), difference(), symmetric_difference() in keyword and set membership Creating and accessing key-value pairs get(), items(), keys(), values() Updating and deleting entries in operator, nested dictionaries
2	Python Fundamentals	5-Jul-2026	Sunday	Live Class		NumPy Arrays		Need for efficient numeric computation Python lists vs NumPy arrays (memory, speed) Creating Arrays Basic Array Properties Simple Indexing and Slicing Simple Array Operations Advanced Array Creation Deep Dive into Indexing
2	Python Fundamentals	7-Jul-2026	Tuesday		DC + Focus			
3	Python Fundamentals	11-Jul-2026	Saturday	Live Class		Vectorized Operations Intro to Pandas Series Object		Vectorized Operations Mathematical and Statistical Functions Reshaping and Flattening What is Pandas and why use it? Pandas Series Object Creating a Series Accessing elements in a Series Operations on Series
3	Python Fundamentals	12-Jul-2026	Sunday	Live Class		Creating and Exploring DataFrames, Data Selection, Filtering		What is a DataFrame Creating DataFrames Exploring DataFrames - .head(), .tail(), .shape, .info(), .describe() Exploring Dataframes - .columns, .index, .dtypes Accessing Columns and Rows Basic Attributes Column Selection Row Selection Filtering Data (Boolean Masking)
3	Python Fundamentals	14-Jul-2026	Tuesday		DC + Focus			
4	Python Fundamentals	18-Jul-2026	Saturday	Live Class		Sorting, Data Cleaning and Handling Missing Values		Sorting Data Setting and Resetting Index Handling Duplicates Renaming Columns Adding New Columns Deleting Columns Reordering Columns Cleaning Column Names Handling Missing values
4	Python Fundamentals	19-Jul-2026	Sunday	Live Class		Aggregation & Grouping in Pandas Advanced DataFrame Operations		groupby() with mean(), sum(), count() Using agg() for multi-aggregation Multi-index results Pivot tables with pivot_table() Merging/joining DataFrames Reshaping Data Date and Time Handling Exporting Data

Week	Module	Date	Day	Delivery Type	Session Type	Session Title	Duration	Sub topics
4	Python Fundamentals	21-Jul-2026	Tuesday		DC + Focus			
5	Python Fundamentals	25-Jul-2026	Saturday	Live Class		Data Visualization with Matplotlib & Seaborn		Line plot, bar chart, histogram, pie chart using Pandas .plot() Customizing plots: labels, titles, colors Plotting directly from groupby results Intro to matplotlib.pyplot: plot(), bar(), scatter() Styling and subplot basics
1	Understanding SQL	26-Jul-2026	Sunday	Live Class		Understanding SQL, Creating Tables		Creating tables using SQL CREATE TABLE statements
1	Python Fundamentals & SQL	28-Jul-2026	Tuesday		DC			
2	Understanding SQL	1-Aug-2026	Saturday	Live Class		Populating Tables in SQLite Querying Data with SELECT		Understanding constraints like PRIMARY KEY, NOT NULL, DEFAULT Inserting single and multiple rows using INSERT INTO Verifying table structure and inserted data Combining conditions using AND, OR, and NOT Sorting results using ORDER BY Limiting output using LIMIT Using column aliases with AS for readability
2	Understanding SQL	2-Aug-2026	Sunday	Live Class		Updating, Deleting, and Altering Data Aggregations and Grouping		Adding and removing columns from an existing table Understanding the risks and best practices in modifying data Using aggregate functions such as COUNT, SUM, AVG, MIN, and MAX Grouping data using GROUP BY to summarize results Filtering grouped results using HAVING Combining GROUP BY with ORDER BY to sort aggregated output Applying aggregation in real-world analysis scenarios
2	Understanding SQL	4-Aug-2026	Tuesday		DC + Focus			
3	Understanding SQL	8-Aug-2026	Saturday	Live Class		Joins and Subqueries		Writing subqueries in SELECT, WHERE, and FROM clauses Combining joins and subqueries for complex data retrieval Practical examples involving multiple-table queries
1	Statistics & Mathematics	9-Aug-2026	Sunday	Live Class		Descriptive Statistics & Probability		Using Python to compute statistical summaries Visualizing data distributions using histograms and boxplots with libraries like Matplotlib and Seaborn Fundamental probability concepts and rules Independent and dependent events Conditional probability and Bayes' theorem (conceptual) Addition and multiplication rules
1	SQL & Statistics	11-Aug-2026	Tuesday		DC + Focus			
2	Statistics & Mathematics	15-Aug-2026	Saturday	Live Class		Probability & Probability Distributions		Calculating probabilities from frequency tables and simulations Difference between discrete and continuous distributions Binomial distribution and its applications Normal distribution and the bell curve Properties of the normal distribution: symmetry, standard deviation, z-scores Using Python (NumPy, SciPy) to generate and visualize distributions Plotting probability distribution curves using Seaborn and Matplotlib
2	Statistics & Mathematics	16-Aug-2026	Sunday	Live Class		Inferential Statistics & Intro to Hypothesis Testing		Using Python (NumPy, SciPy, statsmodels) to simulate sampling and compute confidence intervals Concept of statistical hypothesis: null and alternative hypotheses Understanding Type I and Type II errors Significance level (alpha) and interpretation of p-values Steps in conducting a hypothesis test

Week	Module	Date	Day	Delivery Type	Session Type	Session Title	Duration	Sub topics
2	Statistics & Mathematics	18-Aug-2026	Tuesday		DC + Focus			
	Assessment	19-Aug-2026	Wednesday		ET1	Employability Test 1		
3	Statistics & Mathematics	22-Aug-2026	Saturday	Live Class		Hypothesis Testing Contd..	3	Introduction to one-sample z-test and t-test Performing basic hypothesis tests using Python (SciPy, statsmodels) Conducting two-sample hypothesis tests (independent and paired) Choosing between z-test and t-test based on sample conditions Interpreting test results in the context of real-world problems Performing two-sample t-tests using Python (SciPy, statsmodels) Visualizing group differences and test outcomes using Seaborn and Matplotlib
3	Statistics & Mathematics	23-Aug-2026	Sunday	Live Class		Linear Algebra Basics & Calculus Essentials	3	Understanding scalars, vectors, and matrices Matrix operations: addition, subtraction, multiplication Dot product and its interpretation Transpose and identity matrices Using NumPy for linear algebra computations Applications of linear algebra in machine learning contexts (conceptual) Understanding functions, limits, and derivatives (conceptual overview) Concept of gradients and their role in optimization Introduction to partial derivatives
3	Statistics & Mathematics	25-Aug-2026	Tuesday		DC + Focus		3	
4	Statistics & Mathematics	29-Aug-2026	Saturday	Live Class		Calculus Contd.. & Correlation & Covariance	3	Visualizing functions and slopes using Python (Matplotlib, NumPy) Understanding the role of calculus in training machine learning models (gradient descent) Understanding covariance and how it measures joint variability Difference between covariance and correlation Pearson correlation coefficient and its interpretation Visualizing relationships using scatter plots and heatmaps Using Python (Pandas, NumPy, Seaborn) to compute and visualize correlations Identifying spurious correlations and understanding limitations
4	Supervised Learning	30-Aug-2026	Sunday	Live Class		Linear Regression	2	Assumptions of linear regression Loss function and least squares method Implementing simple linear regression using scikit-learn Multiple predictors and the idea of feature contribution Multicollinearity and correlation among predictors Implementing multiple linear regression in Python
4	Statistics & Mathematics	1-Sep-2026	Tuesday		DC		2	
1	Supervised Learning	5-Sep-2026	Saturday	Live Class		Linear Regression	2.5	Interpreting coefficients and residuals Exploring feature importance in linear models Dataset walkthrough: features, target variable, and objectives Data preprocessing: handling missing values, encoding, scaling Model building and evaluation on the dataset Interpreting outputs and drawing business insights Discussing limitations and possible model improvements
1	Supervised Learning	6-Sep-2026	Sunday	Live Class		Logistic Regression	2.5	Sigmoid function and probability output Decision boundary and thresholding Implementing binary logistic regression using scikit-learn Evaluating classification performance: accuracy, confusion matrix, precision, recall, F1-score Introduction to Customer Churn dataset and target variable Data exploration and preprocessing for classification
1	Supervised Learning	8-Sep-2026	Tuesday		DC + Focus		3	

Week	Module	Date	Day	Delivery Type	Session Type	Session Title	Duration	Sub topics
2	Supervised Learning	12-Sep-2026	Saturday	Live Class		Logistic Regression Contd., K-Nearest Neighbour	3	(2H) Fitting logistic regression to predict churn Interpreting coefficients and model predictions Comparing predictions against actual churn outcomes Insights and business implications of churn predictions Discussion on model limitations and threshold adjustment (1H) Concept of instance-based learning and non-parametric models Distance metrics: Euclidean, Manhattan How KNN works: selecting 'k', voting mechanism Bias-variance tradeoff and impact of different values of 'k'
2	Supervised Learning	13-Sep-2026	Sunday	Live Class		K-Nearest Neighbour	3	Implementing KNN using scikit-learn Evaluating model performance with accuracy and confusion matrix Data preparation for KNN (scaling importance) Model training and prediction on a prepared dataset Effect of distance metric and k-value on accuracy Comparison with logistic regression results (optional) Visualizing decision boundaries (if dataset is suitable) Key considerations for using KNN in production
2	Supervised Learning	15-Sep-2026	Tuesday		DC + Focus		3	
3	Supervised Learning	19-Sep-2026	Saturday	Live Class		Naïve Bayes Classification	1	using scikit-learn Evaluating predictions using accuracy, precision, recall, F1-score
3	Supervised Learning	20-Sep-2026	Sunday	Live Class		Naïve Bayes Classification Intro to Decision Tree	2	Preparing textual data using CountVectorizer or TfidfVectorizer Training a Naïve Bayes classifier on text (e.g., spam detection or sentiment analysis) Model performance interpretation and common pitfalls Why Naïve Bayes performs well on text data Real-world considerations for deploying NB-based text classifiers
3	Supervised Learning	22-Sep-2026	Tuesday		DC + Focus		3	
4	Supervised Learning	26-Sep-2026	Saturday	Live Class		Decision Tree	3	Building a basic decision tree using scikit-learn Visualizing the decision tree structure Controlling tree complexity with max_depth, min_samples_split, etc. Pruning techniques and early stopping Evaluating model performance on unseen data Feature importance and model interpretability Comparing with previous classifiers (logistic, KNN) Preparing the Credit Risk dataset for modeling
4	Supervised Learning	27-Sep-2026	Sunday	Live Class		Decision Tree (Case Study) Random Forests	3	2H Exploratory analysis of credit risk dataset Data preprocessing: encoding, handling missing values Fitting decision tree classifier and evaluating results Visualizing and interpreting feature splits Drawing business insights and decision rules from the model Discussion on model limitations and edge cases 1H Concept of ensemble learning and bagging How Random Forest improves over a single decision tree Key hyperparameters: n_estimators, max_features, max_depth Out-of-bag error and model stability
4	Supervised Learning	29-Sep-2026	Tuesday		DC + Focus		3	
5	Supervised Learning	3-Oct-2026	Saturday	Live Class		Random Forests	3	Building and evaluating a Random Forest model using scikit-learn Understanding feature importance Revisiting the Credit Risk dataset Modeling with Random Forest and comparing with Decision Tree Using feature importance to support credit scoring decisions Evaluating performance and explaining results Discussion on model reliability, interpretability, and real-world usage
5	Supervised Learning	4-Oct-2026	Sunday	Live Class		Case Study	2.5	Cross-Validation & Hyperparameter Tuning: Understanding overfitting and the need for validation K-Fold cross-validation and Stratified K-Fold GridSearchCV and RandomizedSearchCV with scikit-learn Using cross-validation to tune Logistic Regression, KNN, and Decision Trees Model Comparison & Selection: Review of evaluation metrics across models (classification & regression)

Week	Module	Date	Day	Delivery Type	Session Type	Session Title	Duration	Sub topics
5	Supervised Learning	6-Oct-2026	Tuesday		DC		2	
6	Supervised Learning	10-Oct-2026	Saturday	Live Class		Case Study	3	"Contd from previous session if required Case Study Wrap-Up: Integrated case study using a new or blended dataset Students apply all models (or select best fit), tune them, and justify decisions Peer discussion or guided walk-through Reflection on business insights, model effectiveness, deployment readiness"
1	Unsupervised Learning	11-Oct-2026	Sunday	Live Class		Dimensionality Reduction with PCA	3	Understanding high-dimensional data and the curse of dimensionality Concept of feature extraction vs feature selection What is PCA and when to use it Understanding variance and covariance matrix Explained variance ratio and selecting number of components Hands-on PCA with scikit-learn using synthetic datasets Visualizing data in 2D/3D using PCA Analyzing principal components and feature loadings Integrating PCA with preprocessing pipelines Impact of standardization on PCA Applying PCA before clustering Using PCA to prepare features for Customer Segmentation
1	Supervised Learning & Unsupervised	13-Oct-2026	Tuesday		DC		2	
2	Unsupervised Learning	17-Oct-2026	Saturday	Live Class		Recap/Pending of PCA(Optional) K-Means Clustering	3	Recap of PCA if required or pending. What is clustering and how K-Means works Concept of centroids and iterative updates Distance metrics: Euclidean distance Choosing the optimal number of clusters using Elbow Method Limitations of K-Means: initialization sensitivity, non-convex clusters Implementing K-Means in scikit-learn with synthetic data Data preprocessing for clustering: scaling and feature selection
2	Unsupervised Learning	18-Oct-2026	Sunday	Live Class		K-Means Case study Contd(Optional) Intro to Hierarchical Clustering	3	Applying K-Means to a customer dataset Interpreting cluster outputs and labeling segments Visualizing clusters using PCA-reduced data Evaluating clustering results using Silhouette Score Drawing business insights from segmented customer groups Understanding the concept of hierarchy in clustering Agglomerative vs divisive clustering approaches Linkage criteria: single, complete, average Distance matrix and dendrogram construction
2	Unsupervised Learning	20-Oct-2026	Tuesday		DC + Focus		3	
3	Unsupervised Learning	24-Oct-2026	Saturday	Live Class		Hierarchical Clustering	3	Hands-on implementation using scipy.cluster.hierarchy Visual interpretation of dendrograms and how to choose cluster cut-off Applying hierarchical clustering to the Customer Segmentation dataset Visual comparison of K-Means vs hierarchical clustering outputs Advantages and limitations of each technique When to prefer hierarchical methods in business scenarios Wrap-up of unsupervised learning strategies and evaluation techniques
1	Neural Networks	25-Oct-2026	Sunday	Live Class		Introduction to Artificial Neurons & Network Architecture Forward Propagation, Loss Functions, and Cost Minimization	1.5	Gradient descent and backpropagation algorithm Hyperparameters: learning rate, batch size, epochs Overfitting and regularization: dropout, L2 penalty Hands-on: tuning ANN hyperparameters in code

Week	Module	Date	Day	Delivery Type	Session Type	Session Title	Duration	Sub topics
1	Unsupervised Learning & Neural Networks	27-Oct-2026	Tuesday		DC		2	
2	Neural Networks	31-Oct-2026	Saturday	Live Class		Backpropagation, Hyperparameters, and Overfitting	2.5	Gradient descent and backpropagation algorithm Hyperparameters: learning rate, batch size, epochs Overfitting and regularization: dropout, L2 penalty Hands-on: tuning ANN hyperparameters in code
2	Neural Networks	1-Nov-2026	Sunday	Live Class		Deep Learning vs Traditional ML, and Mathematical Foundations Optimization Algorithms and Training Strategies	3	Deep learning vs traditional ML: when and why Vectors, matrices, tensors – review and importance in DL Computational graphs and matrix-based operations Hands-on: working with tensors in TensorFlow Gradient descent variants: SGD, Adam, RMSProp Cost function intuition: MSE and cross-entropy
2	Neural Networks	3-Nov-2026	Tuesday		DC + Focus		3	
3	Neural Networks	7-Nov-2026	Saturday	Live Class		Optimization Algorithms and Training Strategies(Contd) Deep Learning Challenges & Model Tuning with MNIST	3	Training strategies: epochs, mini-batch gradient descent Learning rate tuning and scheduling techniques Vanishing and exploding gradients – the curse of depth Solutions: ReLU, batch normalization, careful initialization Regularization: dropout, L2 regularization Advanced strategies: early stopping, batch normalization
3	Neural Networks	8-Nov-2026	Sunday	Live Class		Introduction to Image Data & Convolutions CNN Architecture	2	Understanding image data: pixels, channels, resolution Image representation as tensors in NumPy and TensorFlow Preprocessing image data: resizing, normalization, color conversion Concept of convolution operation and feature extraction Understanding filters/kernels, stride, and padding (conceptually and visually) Visualizing feature maps using simple filters (edge detection)
3	Neural Networks	10-Nov-2026	Tuesday		DC		2	
4	Neural Networks	14-Nov-2026	Saturday	Live Class		CNN on Image Classification	3	Building a simple CNN using TensorFlow/Keras Training and validation process for image classification Understanding model.summary() and layer shapes Recap of CNN pipeline and key hyperparameters Dataset walkthrough: CIFAR-10 or MNIST, Model training with evaluation metrics: accuracy, confusion matrix Tuning epochs, batch size, and dropout rate, Visualizing learning curves and prediction outputs Interpreting misclassifications and refining model

Week	Module	Date	Day	Delivery Type	Session Type	Session Title	Duration	Sub topics
1	NLP	15-Nov-2026	Sunday	Live Class		Normalization, Vectorization using Bag-of-Words and TF-IDF	3	Text normalization pipeline design Hands-on using NLTK and SpaCy for preprocessing Impact of preprocessing choices on model performance Need for numerical representation of text Bag-of-Words model: concept, advantages, limitations TF-IDF: intuition behind term frequency and inverse document frequency Vectorizer tools in Python: CountVectorizer, TfidfVectorizer (from scikit-learn) Building and visualizing sparse matrices
1	Neural Networks & NLP	17-Nov-2026	Tuesday		DC		3	
2	NLP	21-Nov-2026	Saturday	Live Class		POS Tagging, NER, and Text Classification Introduction to Word Embeddings	3	Part-of-Speech (POS) tagging and syntactic roles Named Entity Recognition (NER): types of entities and use cases Hands-on with SpaCy for POS tagging and NER Text classification workflow using labeled data Model evaluation: accuracy, confusion matrix, precision/recall Limitations of BoW and TF-IDF (sparsity, lack of semantic meaning) Introduction to dense word vectors and distributed representations Word2Vec: CBOW vs Skip-gram (conceptual)
2	NLP	22-Nov-2026	Sunday	Live Class		Word Embeddings (Word2Vec, GloVe) Language Models and Embeddings Evolution	3	GloVe: global co-occurrence matrix-based embeddings Hands-on using pre-trained Word2Vec/GloVe embeddings with Gensim Visualizing embeddings with t-SNE or PCA Comparison of embeddings with sparse methods on simple tasks What is a language model and how it predicts sequences Evolution from count-based models to neural language models Contextual embeddings vs static embeddings Overview of ELMo and BERT as contextual models Use case: contextual word usage in different sentences
2	NLP	24-Nov-2026	Tuesday		DC + Sprint		3	
3	NLP	28-Nov-2026	Saturday	Live Class		Transformers and Self-Attention BERT, GPT & Fine-Tuning	2	Positional encoding and its role in sequence modeling Visualizing attention weights (tools or conceptual diagrams) Key differences between BERT and GPT architectures (encoder vs decoder) Pretraining vs fine-tuning paradigm in LLMs Common downstream tasks: sentiment analysis, question answering Hands-on: using Hugging Face pipeline for zero-shot or fine-tuned inference
3	NLP	29-Nov-2026	Sunday	Live Class		Fine-Tuning Contd.. Applications and Ethics in LLMs	0.5	Bias and fairness concerns in LLMs Risks: hallucination, misinformation, toxic outputs Mitigation techniques: instruction tuning, safety filters Hands-on: testing edge cases using prompts with GPT models Discussion on responsible deployment of generative NLP models
3	NLP	1-Dec-2026	Tuesday		DC + Sprint		3	
	Assessment	19-Dec-2026	Saturday		ET2	Employability Test 2		