



LIVE ONLINE TRAINING · NEXT BATCH DATE TO BE CONFIRMED

AI & Machine Learning with Python

From Python basics to a deployed ML app.

Learn Python and essential SQL, build and deploy machine-learning models, and create a simple Generative AI application — in 72 live one-hour classes. Every class is concept and demo first, then guided practice inside the same hour, with one customer-churn capstone carried through data prep, modelling, deployment and documentation. All tools and datasets are free; no paid API or GPU is needed.

9

Modules

72

Live classes

1 hr

Per live class

72 hrs

Live training

NEXT BATCH

Collecting interest
Date to be confirmed

Live online on Zoom ·
join from anywhere

DAYS AND TIMING

To be confirmed

Tell us what suits you

FEE

To be confirmed

Shared with the batch details
before you decide



Mentor: confirmed together with the batch dates · this session plan is prepared by the mentor

S1-10

Python & Git

S11-20

Data & SQL

S21-25

Statistics

S26-43

Machine
Learning

S44-58

Deploy · Deep
Learning · Text

S59-72

GenAI · RAG ·
Capstone

Course at a glance

Everything you need to know about the course: format, practice and support.

Course name	AI & Machine Learning with Python
Total duration	72 live classes · days and start date to be confirmed
Each class	1 hour, live and instructor-led · last 5–10 minutes reserved for doubts and wrap-up · timing confirmed with the batch
Total live hours	72 hours of live, instructor-led classes
Mode	Live online on Zoom · join from anywhere · recordings 1 year from full payment
Teaching style	Concept and demo, then guided practice inside the same hour; practical and project classes are student-built with the mentor guiding
Tools & datasets	All free and open-source — no paid API subscription or GPU purchase needed
Project work	Customer-churn capstone (sessions 43–69): data prep, modelling, Streamlit deployment, GitHub documentation; plus image, text and document Q&A mini-projects
Assessment	Assignment after most classes, capstone presentation, code review, final assessment and a mock interview
Certificate	JavitaTech Certificate of Completion (not a vendor certification)
Mentor	Confirmed together with the batch dates · session plan prepared by the mentor

Who this course is for

A complete beginner and a working analyst can both sit in the same class and get value.

Freshers and graduates

Session 1 installs Python with you. No coding or maths background needed.

Professionals moving into data and ML

From spreadsheets to pandas, SQL, trained models and a deployed app you can show.

Analysts, Excel and SQL users

Two full modules on data analysis and SQL before the modelling starts.

Developers who want GenAI skills

LLMs, prompt engineering, embeddings, vector stores and a working RAG Q&A app.

— What you'll learn

Twelve outcomes in nine modules — each one is taught live, practised in the same hour and reused in the capstone.

- ✓ Python: collections, functions, OOP, files, CSV/JSON
- ✓ NumPy, pandas, Matplotlib & Seaborn EDA
- ✓ Applied statistics: probability, confidence intervals, hypothesis tests
- ✓ Pipelines, evaluation metrics, cross-validation & tuning
- ✓ Deploying an ML app with Streamlit Community Cloud
- ✓ Text applications: TF-IDF, sentiment, sentence embeddings
- ✓ Git & GitHub for every exercise
- ✓ Essential SQL on SQLite: filters, aggregations, joins
- ✓ Machine learning: regression, classification, clustering, PCA
- ✓ Imbalanced data, leakage and error analysis
- ✓ Deep learning: Keras networks, CNNs, transfer learning
- ✓ Generative AI: LLMs, prompt engineering, embeddings, RAG

— Learning & career support

What comes with every seat in the course.

- ✓ Live classes, one hour each
- ✓ Assignment after most classes, reviewed by the mentor
- ✓ Capstone app deployed publicly and documented on GitHub
- ✓ Final assessment and mock interview
- ✓ JavitaTech Certificate of Completion
- ✓ Hands-on practice inside every hour
- ✓ Recordings · 1 year from full payment
- ✓ Capstone presentation and code review
- ✓ Interview preparation on common ML/AI questions

YOUR MENTOR

To be confirmed

The mentor for the next batch is confirmed together with the batch dates. Every class fits one hour with time for your doubts.

NEXT BATCH

Collecting interest

TIME

To be confirmed

START DATE

To be confirmed

CLASS LENGTH

1 hour per class

Course roadmap — 9 modules, 72 sessions

Python and data first, then statistics and machine learning, then deployment, deep learning, text and Generative AI — finishing with the capstone, code review and interview preparation.

M1 M2 M3 M4 M5 M6 M7 M8 M9			
Session 1		Bar width = hours per module	Session 72
#	MODULE	SESSIONS	WHAT YOU BUILD
M1	Python Foundations & Git	S1–S10 · 10 hrs	Contact-book app, OOP classes, file/CSV/JSON scripts, your first GitHub repo
M2	Data Analysis & Essential SQL	S11–S20 · 10 hrs	Titanic and Superstore EDA in pandas, SQL on a SQLite sales database, mini EDA project
M3	Applied Statistics	S21–S25 · 5 hrs	Descriptive stats, probability simulations, confidence intervals and a t-test on real data
M4	Machine Learning	S26–S43 · 18 hrs	Ten model types on real datasets, a full churn pipeline with tuning and error analysis, capstone kickoff
M5	ML Application & Deployment	S44–S48 · 5 hrs	A Streamlit churn-prediction app, validated, deployed publicly and documented on GitHub
M6	Deep Learning & Image Applications	S49–S54 · 6 hrs	A Keras neural network, a trained CNN and a MobileNetV2 transfer-learning image classifier
M7	Text Applications	S55–S58 · 4 hrs	A sentiment classifier with TF-IDF, then with sentence embeddings
M8	Generative AI & Document Q&A	S59–S65 · 7 hrs	A document Q&A app with embeddings, a vector store and RAG, tested for reliability
M9	Capstone, Review & Interview Preparation	S66–S72 · 7 hrs	Final churn model, capstone app deployed, presentation, code review, assessment and mock interview

By the end you can

- Write clean Python, keep it in Git and manage environments and packages
- Clean, merge, analyse and visualise real datasets with pandas, and query them with SQL
- Choose, train, evaluate and tune ML models, and explain their errors
- Ship a model as a public Streamlit app with validation, logging and documentation
- Build and train neural networks and CNNs, and fine-tune a pretrained image model
- Build a RAG document Q&A app with embeddings and a vector store, and test it for reliability

Session-wise plan · sessions 1–10

Every session fits in one hour — concept and demo, then you build it yourself. Assignments and doubt discussions stay inside the hour.

MODULE 1 OF 9

Python Foundations & Git

Sessions 1–10 · 10 hours

SESSION	TOPIC & SUBTOPICS	PRACTICAL IN CLASS
S1	Course Kickoff & Environment Setup Course roadmap and learning outcomes · Installing Python, VS Code and Google Colab · Variables and data types (int, float, str, bool)	You build: Set up local Python environment and a Google Colab account; run a first Python script.
S2	Operators, I/O & Conditional Statements Arithmetic, comparison and logical operators · input()/print() formatting · if / elif / else	You build: Build a simple grade-calculator script using conditionals.
S3	Loops & Practice Problems for and while loops · break / continue · Nested loops	You build: Solve 5 loop-based problems (pattern printing, number games).
S4	Python Collections Lists and list methods · Tuples · Sets · Dictionaries and dict methods	You build: Build a contact-book program using a dictionary of lists.
S5	Functions & Comprehensions Function definition, default/keyword args, *args/**kwargs · Return values and scope · Lambda, map/filter · List and dict comprehensions	You build: Refactor Session 4's contact book into reusable functions.
S6	Object-Oriented Programming Basics Classes and objects · __init__ and attributes · Methods · Basic inheritance	You build: Build a Student class, then an inherited GraduateStudent class.
S7	File Handling & Exceptions Reading/writing text files · try / except / finally · Custom exceptions	You build: Build a script that logs user input to a file and handles invalid input gracefully.
S8	Working with CSV and JSON csv module · json module · Reading/writing structured data · Nested JSON parsing	You build: Parse a sample CSV of student records, convert to JSON and back.
S9	Debugging, Packages & Environments Debugging with print/breakpoints · pip and virtual environments (venv) · Installing/importing third-party packages	You build: Create a project venv, install numpy/pandas, debug a broken script.
S10 PRACTICAL	Git/GitHub Basics & Module 1 Review git init/add/commit/push/pull · Branches and basic merging · Creating a GitHub repo and README	You build: Push all Week 1–2 exercises to a personal GitHub repository.

Session-wise plan · sessions 11–20

Every session fits in one hour — concept and demo, then you build it yourself. Assignments and doubt discussions stay inside the hour.

MODULE 2 OF 9

Data Analysis & Essential SQL

Sessions 11–20 · 10 hours

SESSION	TOPIC & SUBTOPICS	PRACTICAL IN CLASS
S11	NumPy Arrays & Operations ndarray creation · Indexing and slicing · Broadcasting · Vectorized operations	You build: Perform array-based calculations on a sample sales-numbers dataset.
S12	Pandas Series & DataFrame Basics Series vs DataFrame · Reading CSV/Excel · Selecting rows/columns (loc/iloc) · Basic DataFrame operations	You build: Load the Titanic dataset and explore its structure.
S13	Data Cleaning — Missing Values & Duplicates Detecting missing values · fillna/dropna strategies · Removing duplicates · Outlier detection basics	You build: Clean the Titanic dataset end to end.
S14	Grouping & Merging DataFrames groupby and aggregation · Pivot tables · merge / join / concat	You build: Merge orders and customers tables; produce grouped summaries.
S15	Data Visualization with Matplotlib Line, bar, scatter, histogram plots · Subplots · Labels, legends and styling	You build: Visualize sales trends from the Superstore dataset.
S16	Visualization with Seaborn & EDA Distribution plots, boxplots, heatmaps, pairplots · Structuring an EDA workflow	You build: Produce a full EDA report on the Titanic dataset.
S17	Basic SQL — Filtering & Selection SELECT, WHERE, ORDER BY, LIMIT · Comparison/logical operators in SQL · DISTINCT	You build: Write 8 SQL queries against a sample sales database.
S18	SQL Aggregations & Grouping COUNT, SUM, AVG, MIN/MAX · GROUP BY and HAVING · Subqueries (intro)	You build: Answer 6 business questions using aggregate SQL queries.
S19	SQL Joins INNER / LEFT / RIGHT joins · Joining 3+ tables · Common join pitfalls	You build: Join customers, orders and products tables into a combined report.
S20 PRACTICAL	SQL-to-Pandas & Module 2 Mini EDA Project Running SQL from Pandas (sqlite3 + read_sql) · Comparing SQL vs Pandas approaches	You build: Recreate 5 SQL queries in Pandas; consolidate the Module 2 mini EDA project.

Session-wise plan · sessions 21–30

Every session fits in one hour — concept and demo, then you build it yourself. Assignments and doubt discussions stay inside the hour.

MODULE 3 OF 9

Applied Statistics

Sessions 21–25 · 5 hours

SESSION	TOPIC & SUBTOPICS	PRACTICAL IN CLASS
S21	Central Tendency & Variation Mean, median, mode · Variance and standard deviation · Types of data (categorical/numerical)	You build: Compute descriptive statistics column-by-column on a real dataset.
S22	Probability Basics Basic probability rules · Independent vs dependent events · Intro to normal & binomial distributions	You build: Simulate coin-toss / dice probability experiments in Python.
S23	Sampling & Confidence Intervals Population vs sample · Sampling methods · Confidence intervals: concept and calculation	You build: Compute a confidence interval for a sample mean from a dataset.
S24	Correlation vs Causation Pearson correlation coefficient · Correlation matrix/heatmap · Why correlation $\neq$ causation	You build: Analyze correlations in a dataset and flag spurious-correlation examples.
S25	Hypothesis Testing Basics Null vs alternative hypothesis · p-values and significance · Simple two-sample t-test	You build: Run a t-test comparing two groups (e.g., survival rates) in a dataset.

MODULE 4 OF 9

Machine Learning

Sessions 26–43 · 18 hours

SESSION	TOPIC & SUBTOPICS	PRACTICAL IN CLASS
S26	ML Fundamentals & Workflow Supervised vs unsupervised learning · The ML workflow: problem → data → model → evaluation · Intro to scikit-learn	You build: Load a dataset and walk through the ML workflow skeleton (no modeling yet).
S27	Linear Regression Simple & multiple linear regression · Cost function and gradient descent (concept) · Fitting with scikit-learn	You build: Predict housing prices using linear regression.
S28	Logistic Regression Logistic regression for classification · Sigmoid function · Decision boundary	You build: Predict Titanic survival with logistic regression.
S29	Decision Trees Split criteria: Gini / entropy · Overfitting in trees · Visualizing a tree	You build: Train a decision tree on the Titanic dataset and visualize it.
S30	Random Forests Bagging concept · Random forest algorithm · Feature importance	You build: Train a random forest and compare to a single decision tree.

Session-wise plan · sessions 31–42

Every session fits in one hour — concept and demo, then you build it yourself. Assignments and doubt discussions stay inside the hour.

MODULE 4 OF 9 · CONTINUED

Machine Learning

Sessions 26–43 · 18 hours

SESSION	TOPIC & SUBTOPICS	PRACTICAL IN CLASS
S31	Boosting (Gradient Boosting / XGBoost Intro) Boosting concept · Gradient Boosting basics · Intro to XGBoost/LightGBM	You build: Train a gradient boosting model and compare to the random forest.
S32	K-Nearest Neighbors (KNN) KNN algorithm and distance metrics · Choosing k · Why scaling matters for KNN	You build: Train a KNN classifier, tune k, plot accuracy vs k.
S33	Support Vector Machines — Introductory SVM concept and margin · Linear vs kernel SVM (intro)	You build: Train an SVM classifier and compare with previous models.
S34	K-Means Clustering K-Means algorithm · Choosing k: elbow method · Interpreting clusters	You build: Cluster customers by spending behaviour (unsupervised).
S35	PCA & Dimensionality Reduction Why reduce dimensionality · PCA concept and explained variance · Visualizing PCA components	You build: Apply PCA to a high-dimensional dataset and visualize in 2D.
S36	Preprocessing & Train/Validation/Test Split Encoding categorical variables · Feature scaling: standardization/normalization · Train/val/test split strategy	You build: Preprocess the customer-churn dataset end to end.
S37	Building ML Pipelines scikit-learn Pipeline and ColumnTransformer · Combining preprocessing + model in one object	You build: Wrap Session 36's preprocessing and a classifier into a single pipeline.
S38	Model Evaluation Metrics Accuracy, precision, recall, F1-score · Confusion matrix · ROC-AUC curve	You build: Evaluate the churn pipeline with all key metrics.
S39	Cross-Validation & Hyperparameter Tuning k-fold cross-validation · GridSearchCV vs RandomizedSearchCV	You build: Tune the churn model's hyperparameters.
S40	Overfitting, Underfitting & Regularization Bias-variance tradeoff · L1/L2 regularization · Learning curves	You build: Plot learning curves for the churn model; diagnose over/underfitting.
S41	Imbalanced Data Handling The class-imbalance problem · Class weighting · SMOTE oversampling	You build: Apply class weighting/SMOTE to the churn dataset and re-evaluate.
S42	Data Leakage & Error Analysis Common data-leakage scenarios · Systematic error analysis on misclassified samples	You build: Audit the churn pipeline for leakage; analyze misclassified customers.

Session-wise plan · sessions 43–50

Every session fits in one hour — concept and demo, then you build it yourself. Assignments and doubt discussions stay inside the hour.

MODULE 4 OF 9 · CONTINUED

Machine Learning

Sessions 26–43 · 18 hours

SESSION	TOPIC & SUBTOPICS	PRACTICAL IN CLASS
S43 PROJECT START	Capstone Kickoff — Churn Project (Problem Framing & EDA) Defining the churn-prediction problem statement · Capstone EDA plan	You build: Begin the customer-churn capstone: problem framing + full EDA (student-built).

MODULE 5 OF 9

ML Application & Deployment

Sessions 44–48 · 5 hours

SESSION	TOPIC & SUBTOPICS	PRACTICAL IN CLASS
S44	Saving & Loading the Complete Pipeline joblib / pickle for model persistence · Saving the full preprocessing + model pipeline · Versioning basics	You build: Save the churn pipeline and reload it in a fresh script.
S45	Building a Streamlit Prediction App Streamlit basics: widgets and layout · Loading the saved pipeline · Building a prediction form	You build: Build a Streamlit app that takes customer inputs and predicts churn (student-built).
S46	Input Validation & Application Testing Validating user inputs · Handling edge cases/errors in the app · Basic testing of the prediction function	You build: Add validation and test the churn app with 5 edge cases.
S47	Deployment & GitHub Documentation Preparing requirements.txt · Deploying to Streamlit Community Cloud (free) · Writing a project README	You build: Deploy the churn app publicly and document it on GitHub.
S48	Model Versioning & Basic Error Tracking Simple model-version naming conventions · Logging predictions/errors · When to retrain (concept)	You build: Add basic logging to the deployed app; simulate a model-update scenario.

MODULE 6 OF 9

Deep Learning & Image Applications

Sessions 49–54 · 6 hours

SESSION	TOPIC & SUBTOPICS	PRACTICAL IN CLASS
S49	Neural Network Fundamentals Perceptron · Multi-layer perceptron (MLP) · Activation functions: ReLU, sigmoid, softmax	You build: Build a simple MLP forward pass by hand with NumPy for a toy problem.
S50	Backpropagation, Gradient Descent & Loss Functions Backpropagation concept · Gradient descent variants: SGD, Adam (intro) · Common loss functions	You build: Train a basic neural network with TensorFlow/Keras on a small tabular/digits dataset.

Session-wise plan · sessions 51–58

Every session fits in one hour — concept and demo, then you build it yourself. Assignments and doubt discussions stay inside the hour.

MODULE 6 OF 9 · CONTINUED

Deep Learning & Image Applications

Sessions 49–54 · 6 hours

SESSION	TOPIC & SUBTOPICS	PRACTICAL IN CLASS
S51	CNN Fundamentals Convolution operation · Pooling layers · CNN architecture overview	You build: Build a CNN architecture layer-by-layer in Keras (no training yet).
S52	Building & Training a CNN Compiling and training a CNN · Validation split · Monitoring accuracy/loss	You build: Train a CNN on Fashion-MNIST/CIFAR-10 (subset).
S53	Transfer Learning with a Pretrained Model Why transfer learning · Using a pretrained CNN (e.g., MobileNetV2) · Fine-tuning basics	You build: Fine-tune a pretrained MobileNetV2 on a small custom image set.
S54 PRACTICAL	Guided Image Classification Project Consolidating CNN + transfer learning into one project · Evaluation and confusion matrix for images	You build: Complete an image-classification mini-project (student-built, instructor-guided).

MODULE 7 OF 9

Text Applications

Sessions 55–58 · 4 hours

SESSION	TOPIC & SUBTOPICS	PRACTICAL IN CLASS
S55	Text Cleaning & Tokenization Text normalization: lowercasing, punctuation removal · Tokenization · Stemming vs lemmatization	You build: Clean and tokenize a sample product-reviews dataset.
S56	Bag of Words & TF-IDF Bag-of-Words representation · TF-IDF concept and computation · scikit-learn vectorizers	You build: Vectorize the reviews dataset with BoW and TF-IDF and compare.
S57	Text / Sentiment Classification Building a sentiment classifier on TF-IDF features · Evaluating text classifiers	You build: Train a sentiment classifier on the reviews dataset.
S58 PRACTICAL	Sentence Embeddings & Mini Project Intro to sentence embeddings (Sentence-Transformers, free/local) · Embeddings vs TF-IDF performance	You build: Re-run sentiment classification using embeddings; consolidate into a mini sentiment-analysis project.

Session-wise plan · sessions 59–67

Every session fits in one hour — concept and demo, then you build it yourself. Assignments and doubt discussions stay inside the hour.

MODULE 8 OF 9

Generative AI & Document Q&A

Sessions 59–65 · 7 hours

SESSION	TOPIC & SUBTOPICS	PRACTICAL IN CLASS
S59	LLM Fundamentals & Limitations What is an LLM: tokens, context windows · Strengths and limitations: hallucination, knowledge cutoff	You build: Explore an LLM via a free playground/API and observe hallucination examples.
S60	Prompt Engineering Techniques Zero-shot vs few-shot prompting · System/role prompts · Prompt-structuring best practices	You build: Design and test 5 prompts of increasing sophistication for a given task.
S61	Document Loading & Splitting Loading PDFs/text files programmatically · Chunking strategies: size and overlap · Why chunking matters for retrieval	You build: Load a sample document set and split it into chunks.
S62	Embeddings & Vector Stores What are embeddings · Generating embeddings with a free local model (sentence-transformers) · Storing embeddings in FAISS/Chroma	You build: Generate embeddings for Session 61's chunks and store them in FAISS/Chroma.
S63	Retrieval-Augmented Generation (RAG) Concepts Retriever + generator architecture · Similarity search · Combining retrieved context with an LLM prompt	You build: Perform similarity search against the vector store; manually construct a RAG prompt.
S64 PRACTICAL	Build a Document Q&A Application Wiring loader → splitter → embeddings → retriever → LLM · Adding source references to answers	You build: Build a working document Q&A app (student-built, instructor-guided) with source citations.
S65	Testing the Q&A App — Reliability, Privacy & Cost Testing for incorrect/unsupported answers · Privacy considerations with uploaded documents · Understanding free-tier API costs/limits	You build: Stress-test the Q&A app with tricky/out-of-scope questions; document failure cases.

MODULE 9 OF 9

Capstone, Review & Interview Preparation

Sessions 66–72 · 7 hours

SESSION	TOPIC & SUBTOPICS	PRACTICAL IN CLASS
S66	Capstone — Model Training & Pipeline Finalizing the feature set from Session 43's EDA · Training and comparing final candidate models	You build: Train and select the final churn-prediction model (student-built).
S67	Capstone — Streamlit App Integration Integrating the final model into the Module 5 Streamlit template	You build: Wire the finalized churn model into a deployable Streamlit app (student-built).

Session-wise plan · sessions 68–72

Every session fits in one hour — concept and demo, then you build it yourself. Assignments and doubt discussions stay inside the hour.

MODULE 9 OF 9 · CONTINUED

Capstone, Review & Interview Preparation

Sessions 66–72 · 7 hours

SESSION	TOPIC & SUBTOPICS	PRACTICAL IN CLASS
S68	Capstone — Testing & Refinement End-to-end testing · Refining based on evaluation metrics and edge cases	You build: Test and refine the capstone app (student-built).
S69	Capstone — Documentation & Deployment Writing project documentation/README · Deploying the final app publicly	You build: Deploy the finished capstone app and finalize documentation (student-built).
S70	Presentation Preparation Structuring a project presentation · Explaining technical decisions to a non-technical audience	You build: Prepare and rehearse a 5–7 minute capstone presentation.
S71 REVIEW	Code Review & Individual Understanding Check Peer/instructor code review of the capstone project · Individual Q&A on code and design choices	You build: Present code for review; answer instructor questions on design decisions.
S72 REVIEW	Final Assessment & Interview Preparation Final skills assessment · Common ML/AI interview questions · Mock interview based on the capstone project	You build: Complete the final assessment and a mock interview based on the capstone project.

Your capstone: customer-churn prediction, deployed

One running project through the second half of the course, so you leave with a live app and a documented GitHub repository.

Telco Customer Churn · sessions 43 to 69

Started in Module 4 right after you have learned the full ML workflow, finished in Module 9, then presented, reviewed and used as the base for your mock interview.

S43

Problem framing and full EDA on the churn dataset

S44–48

Save the pipeline, build a Streamlit prediction app, validate, deploy, log

S66–67

Train and select the final model; wire it into the app

S68–69

Test, refine, deploy publicly, write the README

S70–72

Present, code review, final assessment and mock interview

How a one-hour class runs

Concept classes

Concept & demo 35 min · guided practice 15 min · doubts & wrap-up 10 min

Practical & project classes

Task briefing 5 min · hands-on build 45 min · doubts & wrap-up 10 min

Review classes

Review / assessment 45 min · feedback 10 min · doubts & wrap-up 5 min

Tools you'll work with

All free and open-source. No paid API subscription and no GPU purchase — Colab's free GPU covers the CNN sessions.

Python 3.x

VS Code

Google Colab (free GPU)

pip & venv

Git & GitHub

NumPy

pandas

Matplotlib

Seaborn

SQLite & DB Browser

SciPy & statsmodels

scikit-learn

XGBoost

imbalanced-learn

joblib

Streamlit

Streamlit Community Cloud

TensorFlow / Keras

MobileNetV2

NLTK / spaCy

sentence-transformers

FAISS / ChromaDB

LangChain

Free-tier LLM access (Google AI Studio / Groq)

Datasets you'll work on

California Housing

Iris

Wine

Digits

Titanic

Superstore orders & customers

Mall Customers

Telco Customer Churn (capstone)

IMDB / Amazon reviews

Fashion-MNIST / CIFAR-10

Kaggle image dataset

Sample SQLite sales.db

Sample PDF / text documents

Free accounts you'll create

- Google Colab — notebooks and the free GPU tier for the CNN sessions
- GitHub — every exercise, mini-project and the capstone are pushed here
- Streamlit Community Cloud — public deployment of your ML app
- Kaggle — public datasets used through the course
- Google AI Studio or Groq free tier — LLM access for Module 8, no paid API

Assessment structure

- Assignment after most classes, submitted and reviewed
- Module consolidation in sessions 10 and 20; mini-projects in sessions 54, 58 and 64
- Capstone build in sessions 43 and 66–69, presentation in 70
- Code review and individual understanding check in session 71
- Final assessment and mock interview in session 72

— Your next step: register your interest

Tell us the days and times that suit you, and we'll contact you about suitable batch options. Registering interest does not reserve a seat or confirm a batch.

Collecting interest — next batch date to be confirmed.

Register your interest on [javitatech.com](https://www.javitatech.com). Your demo registration is confirmed only when you receive a separate confirmation for that session.

Register your interest → www.javitatech.com

WHATSAPP

+91 87122 87999

EMAIL

info@javitatech.com

WEBSITE

javitatech.com

In the first 3 classes

- Session 1 · Course kickoff and environment setup: Python, VS Code, Google Colab and your first script
- Session 2 · Operators, input/output and conditionals: build a grade-calculator script
- Session 3 · Loops and practice problems: solve five loop-based problems in a notebook

Good to know

- Live online course, on Zoom; sessions in IST
- 1 hour per class · days and timing confirmed with the batch
- Recordings · 1 year from full fee payment
- Any recent laptop works — no GPU purchase and no paid software
- Fee confirmed with the batch · refund policy at javitatech.com/cancellation-and-refund

From Python basics to a deployed ML app.

— Team JavitaTech · www.javitatech.com