

| COURSE BROCHURE · 2026

Data Science

3 months · live + recorded · Statistics, ML, deep learning, NLP, Gen AI — all shipped end-to-end. Four deployed projects. The path to Rs 10-35 LPA Data Scientist roles.

DURATION

3 mo

MODULES

2

ROLES UNLOCKED

6+

SALARY BAND

4-40 LPA

| ABOUT THIS PROGRAMME

Data Scientists in 2026 are the rare profile that bridges statistics, ML, and business impact — and the ones who can ship end-to-end command Rs 10-35 LPA at product companies. This three-month programme is built for engineers, analysts, and graduates targeting Data Scientist roles. You'll cover statistics, EDA, classical ML, deep learning, NLP, time series, MLOps, and Gen AI integration. By the end you'll have four deployed portfolio projects spanning regression, classification, NLP, and a Gen AI-powered application.

| BUILT FOR

- Engineers + analysts transitioning into data science roles
- Final-year graduates targeting Data Scientist positions
- Career switchers building 4 deployable portfolio projects

OUR COMMITMENT

Job-focused training. Real instructors. Lifetime access. We don't sell certificates — we get you hired.

WHY NOW · 2026 MARKET SIGNAL

Data Scientists who can pair classical ML with Gen AI are the most-hired profile at Indian product startups in 2026. The shift from 'analyst with notebook' to 'engineer with pipeline' is complete.

01

Data Scientist postings up 165% YoY
(LinkedIn India)

02

Median DS offer at product cos: Rs
16 LPA

03

Gen AI + production ML combo
earns Rs 5-8 LPA premium

TECHNOLOGY STACK

Technologies you'll master

The exact stack senior engineers + consultants are billing for in 2026 — nothing taught for show, everything taught because it lands on a job posting.

- Python — NumPy, Pandas, Matplotlib, Seaborn, scikit-learn
- Statistics + probability + hypothesis testing fundamentals
- Classical ML — regression, classification, clustering, ensembles
- Deep learning — TensorFlow, PyTorch, Keras
- Natural Language Processing — transformers, BERT, Hugging Face
- Time-series — ARIMA, Prophet, LSTM, Exponential Smoothing
- A/B testing + experimental design
- SQL for data exploration + feature engineering
- MLflow + FastAPI for model lifecycle + deployment
- Vector databases + semantic search basics

WHY THIS PROGRAMME

What makes this different

- ✓ Four deployed portfolio projects: regression, classification, NLP, Gen AI app
- ✓ Time series forecasting with classical and deep learning approaches
- ✓ Statistics deep-dive — hypothesis testing, A/B, Bayesian fundamentals
- ✓ MLOps lifecycle — experiment tracking, deployment, monitoring
- ✓ Classical ML and deep learning with PyTorch and scikit-learn
- ✓ Real-world datasets from BFSI, retail, and healthcare
- ✓ NLP with transformers — fine-tuning and RAG patterns
- ✓ Interview prep — 60+ data science questions and case studies

CAREER OUTCOMES

By the end, you will be able to

- ✓ Run rigorous EDA and statistical tests on enterprise-scale data
- ✓ Build and tune ML models with proper validation discipline
- ✓ Fine-tune transformer models for classification and sequence tasks
- ✓ Deploy ML models behind APIs with monitoring and drift detection
- ✓ Design A/B tests with proper power calculations and sequential analysis
- ✓ Ship Gen AI applications — RAG, embeddings, LLM evaluation
- ✓ Pass Data Scientist interviews at product cos paying Rs 10-30 LPA
- ✓ Translate ambiguous business problems into rigorous ML hypotheses

SALARY OPPORTUNITIES

Job roles + earning potential

Indian-market salary bands for 2026, sourced from active listings on Naukri, LinkedIn, and Instahyre. Bands assume the certifications and projects this programme delivers.

	JOB ROLE	EXPERIENCE	SALARY
01	Junior Data Scientist	Fresher to 2+ Years	4–8 LPA
02	Data Scientist (Entry/Mid)	Fresher to 3+ Years	6–12 LPA
03	Machine Learning Engineer (Junior)	1 to 3+ Years	6–12 LPA
04	Business Analyst / BI Analyst	2 to 4+ Years	7–12 LPA
05	Data Scientist (Mid-Level)	3 to 5+ Years	12–20 LPA
06	Machine Learning Engineer (Mid-Level)	3 to 5+ Years	12–22 LPA
07	Senior Data Analyst	4 to 6+ Years	10–18 LPA
08	Senior Data Scientist	4 to 7+ Years	20–35 LPA
09	Applied Scientist (ML/NLP/CV)	5 to 8+ Years	20–40 LPA
10	Data Engineer (ML/AI Focused)	5 to 8+ Years	12–25 LPA

| CAREER PATHS

Roles you can target

Roles ordered from entry → senior. The same role titles you'll see on Naukri, LinkedIn, and Instahyre filtered to the tech stack you'll master.

- Data Scientist
- Machine Learning Engineer
- Applied Scientist
- AI Engineer
- Research Scientist
- Senior Data Scientist

Emerging — high-demand roles for 2026

Roles that didn't exist five years ago, paying premium because the talent pool is small.

- LLM Engineer / Researcher
- Causal Inference Specialist
- MLOps Engineer
- AI Product Scientist

| CAREER TRAJECTORY

Where this programme leads, in 5 years

Median 2026 salary bands for the role progression this curriculum unlocks. Numbers are conservative — top performers cross the upper band earlier.

YEAR 1

Data Scientist (Associate)

Rs 8-14 LPA

Ship first production model. Own one ML initiative.

YEAR 3

Senior Data Scientist

Rs 16-28 LPA

Lead initiative end-to-end from problem to production.

YEAR 5

Lead DS / Applied Scientist

Rs 30-55 LPA

Shape company ML strategy. Mentor 5+ scientists.

| COURSE CURRICULUM

What you'll build, module by module

Every module is taught hands-on with real datasets, instances, or codebases — not slides. By the end, you'll have built, not just watched.

FOUNDATIONS OF DATA SCIENCE

- Python syntax essentials
- Variables, data types
- Control flow (if/else), loops
- Functions, modules, virtual environments
- Introduction to OOPS in Python
- Jupyter Notebook setup
- Basic scripting
- Working with `os`, `json`, `requests`
- Small exercises (temperature converter, file reader,...)
- Arrays, vectorization
- DataFrames, indexing, grouping
- Plotting foundations (line, bar, scatter, histograms)
- NumPy exercises (array math, reshaping)
- Pandas exercises (filtering, groupby, joins)
- Exploratory plots with Matplotlib/Seaborn
- Mini-project: Clean + Explore a CSV dataset
- Descriptive stats: mean, variance, skew, kurtosis
- Probability basics
- Distributions (Normal, Bernoulli, Binomial)
- Correlation vs causation
- Confidence intervals, hypothesis testing
- Distribution visualizations
- t-test, chi-square test on sample data
- Correlation heatmaps and interpretation
- Data cleaning workflow
- Handling missing values & outliers
- Feature scaling & encoding
- Train-test splits
- Leakage and good ML practice
- Handling NaNs in Pandas
- Encoding categorical variables (OneHot, Label)
- Normalization & standardization
- Pipeline building in scikit-learn
- Mini-project: full cleaning workflow
- ML building blocks
- Bias-variance tradeoff
- Cross-validation
- Regularization (Lasso, Ridge)
- Algorithms:
- Linear regression
- Logistic regression
- KNN
- Naïve Bayes
- Decision trees
- Random forest, Gradient Boosting (intro)
- Build regression + classification models in sklearn
- Model evaluation metrics: RMSE, MAE, Accuracy, P...
- Hyperparameter tuning with GridSearchCV
- Mini-project: Prediction task (housing prices or churn)
- Clustering concepts
- K-Means algorithm
- Hierarchical clustering
- Dimensionality reduction: PCA theory & intuition
- Elbow method, silhouette score
- Apply K-Means on a customer segmentation dataset
- Hierarchical clustering dendrogram
- PCA decomposition & visualization

ADVANCED DATA SCIENCE & DEEP LEARNING

- Neuron, weights, biases
- Activation functions
- Loss functions
- Gradient descent & backpropagation (conceptual)
- Overfitting & regularization in neural nets
- Build a small neural net from scratch in NumPy (for ...)
- Tensors
- Autograd
- Datasets/Dataloaders
- Optimizers
- Build a classifier for MNIST or a tabular dataset
- Custom training loop (forward → loss → backward → ...)
- Saving/loading models
- CNNs — convolution, filters, pooling
- RNNs, LSTMs, GRUs (overview only)
- Regularization in deep nets: dropout, batch norm
- CNN for image classification (CIFAR-10 or fashion-...)
- Compare performance with data augmentation
- Modern vision models
- Transfer learning
- Feature extraction vs fine-tuning
- ResNet, EfficientNet (high-level)
- Fine-tune a pretrained ResNet on a small custom da...
- Implement augmentations using torchvision
- Why RNNs fail
- Self-attention mechanism
- Encoder-decoder design
- Transformer block anatomy
- Positional embeddings
- Overview of BERT, GPT, Vision Transformers
- Attention mechanism in PyTorch
- Tokenization, subword embeddings
- Pretraining vs fine-tuning
- Task types: classification, NER, QA
- Fine-tune BERT (HuggingFace) for text classification
- Evaluate & interpret embeddings
- Export model for inference
- Intro to LLMs beyond transformers
- Vector databases & retrieval
- MLOps foundations
- Model interpretability (SHAP/LIME)
- Reinforcement learning (very high level)

I INSTRUCTED BY

Trained by working practitioners

Not academics. Not pre-recorded YouTubers. Engineers and consultants who still ship production code at top companies — the exact people you want to learn from.

MK

Mr. Karan E

Industry Trainer

A senior Data Science & Artificial Intelligence practitioner with 19+ years of experience in designing, deploying, and scaling real-world ML and AI systems. With strong expertise in statistics, software engineering, and applied research, he has led large-scale machine learning, deep learning, and generative AI initiatives across multiple industries. He also contributes to the AI community through publications, patents, and mentorship in AI strategy and scalable system design. Approved trainer by Raj Cloud Technologies.

APPROVED TRAINER · RAJ CLOUD TECHNOLOGIES

| WHAT'S INCLUDED

Everything you get beyond the curriculum

The sessions are the tip. Here's the full value stack a single fee unlocks — mentorship, community, career support, and lifetime updates.

- + Lifetime access to recordings — re-watch any session, forever
- + Private community access — get unblocked in minutes, not days
- + 1:1 mentor reviews on every portfolio project you build
- + Weekly doubt-clearing sessions led by the instructor
- + Resume + LinkedIn rewrite by a working hiring manager
- + Two full mock interviews with real-time feedback
- + Interview question bank — 50+ scenario-based questions
- + Real-world datasets + playground accounts you keep
- + Verifiable course completion certificate
- + Lifetime curriculum updates — when modules refresh, you get them

| ENROLMENT

Pick your tier and enrol

Two tiers — pick what fits your budget. Both include lifetime access to recordings and a verifiable certificate.

LIVE + RECORDED

Rs 29,999

- ✓ All live sessions + Q&A
- ✓ Lifetime recordings
- ✓ Doubt-clearing sessions
- ✓ Resume + interview prep
- ✓ Certificate of completion

MOST POPULAR

| HOW TO ENROL

Three steps to your seat

01 Book a free demo

30-min walk-through with the instructor — ask anything before you commit.

02 Pick your tier + enrol

Live + Recorded or Recorded-only. UPI, card, EMI — all accepted.

03 Start on day one

Onboarding kit, course portal access, and your first live session calendar invite.

I HOW WE STACK UP

Compared to your alternatives

Free tutorials and generic bootcamps both have their place. Here's exactly where they fall short — and where the fee here pays for itself.

WHAT YOU GET	FREE YOUTUBE TUTORIALS	GENERIC BOOTCAMPS	RAJ CLOUD
Live practitioner instructors	×	×	✓
Real datasets + playground accounts	×	partial	✓
1:1 mentor reviews on your portfolio	×	×	✓
Mock interviews with feedback	×	partial	✓
Resume + LinkedIn rewrite	×	×	✓
Job referral network access	×	×	✓
Lifetime updates + recordings	partial	×	✓
Verifiable certificate	×	✓	✓

I QUESTIONS ANSWERED

Common questions, answered upfront

Q. Do I need prior experience to start?

Most students begin with little to no experience. We assume only basic computer literacy. The first two weeks ramp you up to the level the rest of the curriculum needs — no one is left behind.

Q. What if I miss a live session?

Recordings are uploaded within 24 hours of every session. Weekly doubt-clearing sessions catch you up on missed material. Your community is active 24/7. Nothing gets lost, only delayed.

Q. Is there a job placement guarantee?

We don't promise placement — we engineer outcomes. Resume + LinkedIn rewrite, two mock interviews, a referral network, and portfolio reviews are all included. The work happens; the offer is yours to earn.

Q. Can I pay in instalments?

Yes. UPI EMI, credit-card EMI, and a 3-month no-cost split-pay are all supported. Email admissions and we'll set up a plan that fits your cash flow.

Q. Does the certificate hold weight in interviews?

Yes. It's dated, verifiable via our website, and signed by the trainer. Alumni reference it on LinkedIn and inside offer-stage interview rounds — but the portfolio you build matters more, and that's the point.

Talk to admissions

EMAIL

info@rajcloudtech.com

PHONE

[+91 93804 87284](tel:+919380487284)

WEB

rajcloudtech.com

OUR COMMITMENT

Job-focused training. Real instructors. Lifetime access. We don't sell certificates — we get you hired.