DevOpsSchool

Course curriculum · November 2026

LEVEL 04 OF 04 · ADVANCED

Predictive Analytics Advanced Training

Deploy, monitor and govern predictive models in production, scale them on cloud platforms and move from prediction to decision-making.

LIVE ONLINE **Max 8-10 hours**	CORPORATE (DAY-WISE) **5 days · Days 1-5**	SELF-PACED VIDEO **1-year LMS access**
BATCH SIZE **Max 10 learners**	NEXT BATCHES **1st of every month · weekends: first Saturday**	CERTIFICATE **Verifiable certificate of completion**

Foundation	Essential	Intermediate	Advanced
1 day · Included	2 days · Included	3 days · Included	5 days · This curriculum

INSIDE THIS CURRICULUM

What you will achieve	Outcomes and level focus
Who should attend	Audience and prerequisites
Agenda	46 topics · Foundation to Advanced
Hands-on labs	4 labs
Formats and fees	Live, 1:1, self-paced, corporate
Trainer and reviews	Rajesh Kumar
Enrolment	Batches, payment, contacts

WHY DEVOPSSCHOOL

- Rated 4.6 on Google
- 1,500+ verifiable certificates issued in the public registry
- Trainer Rajesh Kumar: 20 years in industry
- Small batches: max 10 learners, every session recorded

Course page: www.devopsschool.com/courses/predictiveanalytics/predictive-analytics-advance-training.html
Talk to us: WhatsApp +91 99057 40781 · +1 (469) 756-6329 · contact@DevOpsSchool.com · reply within 1 hour

OUTCOMES

What you will achieve

By the end of the Predictive Analytics course path, you will be able to:

1

Build accurate predictive models

Prepare data and build, evaluate and tune regression, classification and forecasting models.

2

Explain and trust predictions

Interpret models with explainability tools and validate them against leakage and bias.

3

Turn predictions into value

Deploy, monitor and govern models and link them to measurable business decisions.

WHAT THE ADVANCED LEVEL ADDS

Deployment, MLOps, causal and prescriptive analytics at scale

Hands-on: Deploy and monitor a predictive model. Serve the churn model through an API, track versions in MLflow, monitor drift on new data and design an A/B test proving its business impact.

WHO SHOULD ATTEND

- Experienced Predictive Analytics practitioners responsible for quality, security and scale.
- Senior engineers, architects and leads who design and troubleshoot Predictive Analytics solutions.
- Teams making Predictive Analytics an organisation-wide standard.

BEFORE YOU START

- The Intermediate level, or equivalent hands-on experience.
- You can already: Build stronger models with ensembles and time series forecasting, handle imbalanced data, and explain predictions with reusable pipelines.
- Basic computer skills and a laptop with a stable internet connection for the live sessions and labs.
- This curriculum includes every earlier level, so it is complete on its own.

AGENDA · FOUNDATION TO ADVANCED

Predictive Analytics Advanced Training agenda

46 topics. Each level builds on the one before it, so this curriculum lists every earlier level in full. Days are for corporate (day-wise) delivery; live online batches cover the same agenda in max 8-10 hours.

DAY 1 — Foundation: Analytics concepts, statistics, tools and first model

Included from Foundation

1. Analytics spectrum: descriptive, diagnostic, predictive and prescriptive analytics
2. Business use cases: churn prediction, demand forecasting, credit risk, fraud and predictive maintenance
3. CRISP-DM lifecycle: business understanding, data understanding, preparation, modelling, evaluation and deployment
4. Data fundamentals: structured and unstructured data, data sources, target variables and features
5. Statistics refresher: distributions, mean, variance, percentiles and correlation vs causation
6. Probability basics: conditional probability, Bayes' theorem and expected value
7. Tooling: Python with Jupyter, pandas, scikit-learn and Matplotlib; alternatives such as R and Excel
8. Environment setup: installing Python with Anaconda or uv, creating environments and launching notebooks
9. Loading and exploring data: reading CSV and Excel files, summary statistics and simple charts
10. First predictive model: simple and multiple linear regression with scikit-learn
11. Train and test split: why we hold out data and how to evaluate with MAE and RMSE
12. Ethics and data privacy: consent, sensitive attributes and responsible use of predictions

DAY 2 — Essential: Data preparation, regression, classification and evaluation

Included from Essential

13. Data cleaning: missing values, outliers, duplicates and inconsistent categories
14. Exploratory data analysis: distributions, relationships and visual storytelling with seaborn
15. Feature engineering: encoding categories, scaling, binning, date features and interaction terms
16. Regression in depth: polynomial terms and regularised Ridge, Lasso and Elastic Net models
17. Classification models: logistic regression, decision trees, k-nearest neighbours and naive Bayes
18. Classification metrics: confusion matrix, precision, recall, F1 score and ROC-AUC
19. Choosing thresholds: cost-sensitive decisions and precision-recall trade-offs
20. Cross-validation: k-fold and stratified validation for reliable estimates
21. Overfitting and underfitting: the bias-variance trade-off and learning curves
22. Hyperparameter tuning: GridSearchCV and RandomizedSearchCV
23. Communicating results: translating model metrics into business impact for stakeholders

DAY 3 — Intermediate: Ensembles, forecasting, segmentation and explainability

Included from Intermediate

24. Ensemble methods: random forests and gradient boosting with XGBoost, LightGBM and CatBoost
25. Time series forecasting: trend, seasonality, ARIMA and exponential smoothing
26. Forecasting tools: Prophet-style models, lag features and time-based validation
27. Customer segmentation: k-means and hierarchical clustering to create predictive features
28. Imbalanced data: class weights, resampling with SMOTE and evaluating rare-event models
29. Feature selection: filter, wrapper and embedded methods and reducing dimensionality with PCA
30. Model explainability: feature importance, permutation importance and SHAP values

31. Reusable workflows: scikit-learn Pipeline and ColumnTransformer to prevent leakage
32. Data access: querying data warehouses with SQL and joining multiple sources
33. AutoML: when automated model search helps and how to validate its output
34. Troubleshooting models: data leakage, distribution shift between train and test, and unstable features

DAYS 4-5 **Advanced: Deployment, MLOps, causal and prescriptive analytics at scale** **New at this level**

35. Model deployment: batch scoring jobs and real-time prediction APIs with FastAPI
36. MLOps foundations: experiment tracking and model registry with MLflow
37. Monitoring in production: data drift, concept drift, performance decay and retraining triggers
38. Experimentation: A/B tests to prove model value and uplift modelling to target interventions
39. Causal inference basics: confounders, propensity scores and difference-in-differences
40. Prescriptive analytics: optimisation with linear programming and decision rules built on predictions
41. Scaling: distributed training and scoring with Spark MLlib
42. Cloud platforms: Amazon SageMaker, Azure Machine Learning and Google Vertex AI for predictive workloads
43. Deep learning for prediction: neural networks for tabular and sequence data and when they help
44. Responsible AI: fairness testing, bias mitigation, model documentation and regulatory expectations
45. Business adoption: ROI estimation, stakeholder buy-in and embedding models in processes
46. Capstone: deliver an end-to-end predictive solution from business question to monitored deployment

PRACTICE

Hands-on labs

You build these yourself during the course, guided by the trainer.

FOUNDATION

Predict house prices with regression

Load a housing dataset, explore it with pandas and charts, fit a linear regression, and report its error on held-out test data.

ESSENTIAL

Build a customer churn classifier

Clean and engineer features from customer data, compare logistic regression and decision trees with cross-validation, and choose a business-driven threshold.

INTERMEDIATE

Forecast product demand

Build a weekly demand forecast with seasonal features and a gradient boosting model, validate it on time-based splits and explain drivers with SHAP.

ADVANCED

Deploy and monitor a predictive model

Serve the churn model through an API, track versions in MLflow, monitor drift on new data and design an A/B test proving its business impact.

HOW YOU CAN LEARN

Formats and fees

The same agenda in four formats. Fees are in Indian rupees as of November 2026; the course page shows them in your currency.

FORMAT	DURATION	WHAT YOU GET	FEE
Live & interactive (online)	Max 8-10 hours	Instructor-led public batch of max 10 learners with demos, labs, Q&A and assignment review. Weekday batches start on the 1st of every month, weekend batches on the first Saturday of every month.	**INR 34,999**
1:1 live training (online)	Max 8-10 hours, one-to-one	The same agenda, taught one-to-one by the trainer. Starts immediately, without waiting for the next batch; sessions planned around your working hours, at your pace.	**INR 99,999**
Self-paced video	1-year LMS access	All four Predictive Analytics levels, plus 20+ more courses and 50+ tools, at your own pace.	**INR 833/month** billed yearly at INR 9,996
Corporate (day-wise)	5 days for this level	A private batch for your team, on demand. Foundation 1 day, Essential 2, Intermediate 3, Advanced 5 (cumulative).	**Quote on request**

ASSESSMENT AND CERTIFICATE

Your progress is checked through hands-on labs and assignments at every level. On completion you receive a DevOpsSchool certificate of completion with a unique credential ID that anyone can verify at devopsschool.com/certificates.

INCLUDED WITH EVERY LIVE BATCH

- Every live session is recorded and stays in your LMS for a year.
- Can't attend your batch? Move to next month's batch.
- Free demo class before you pay.
- Free 10-15 minute call with an advisor or the trainer to check the course fits you.

YOUR TRAINER

Learn from Rajesh Kumar

Principal Architect, trainer and mentor in DevOps, SRE, DevSecOps and platform engineering, with 20 years across 8+ software MNCs, including Staff DevOps/SRE Engineer for a leading Japanese mobile-payment platform and Principal DevOps Architect and Manager in Europe.

Profile: www.rajeshkumar.xyz · Videos: youtube.com/@TheDevOpsSchool

20 years in industry	**200+** corporate trainings
10,000+ engineers mentored	**70+** organisations coached

REVIEWS

What learners say

Rated 4.6 on Google · more reviews at devopsschool.com/testimonials and in our video testimonials.

> “Rajesh is very good trainer. Rajesh was able to resolve our queries and question effectively. We really liked the hands-on examples covered during this training program.”
>
> Indrayani, India

> “The training was very useful and interactive. Rajesh helped develop the confidence of all.”
>
> Abhinav Gupta, Pune

START YOUR PREDICTIVE ANALYTICS TRAINING

Join the next batch

Weekday batches start on the 1st of every month and weekend batches on the first Saturday of every month, with max 10 learners each. 1:1 training can start immediately; corporate batches run on demand. We reply within 1 hour.

Pay by UPI, credit and debit cards, net banking or PayPal. Before you pay, read our refund policy. All levels of this course: www.devopsschool.com/courses/predictiveanalytics/

ENROL OR RESERVE A SEAT
Predictive Analytics Advanced Training

WHATSAPP · INDIA
+91 99057 40781

WHATSAPP · USA
+1 (469) 756-6329

EMAIL
contact@DevOpsSchool.com