DevOpsSchool

Course curriculum · November 2026

LEVEL 04 OF 04 · ADVANCED

ML Advanced Training

Take ML models to production at scale with containers, monitoring and drift detection, while applying responsible AI and governance.

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	47 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/ml/ml-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 ML course path, you will be able to:

1

Build sound ML models

Prepare data and train, evaluate and tune classical machine learning models with scikit-learn.

2

Apply deep learning and NLP

Use TensorFlow, PyTorch and Hugging Face to solve vision, language and forecasting problems.

3

Ship ML responsibly

Deploy, monitor and govern models in production with attention to fairness, privacy and regulation.

WHAT THE ADVANCED LEVEL ADDS

Production ML, scale, responsible AI and trends

Hands-on: End-to-end production ML capstone. Build a reproducible training pipeline, deploy the model as a containerised service and add drift monitoring and a model card.

WHO SHOULD ATTEND

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

BEFORE YOU START

- The Intermediate level, or equivalent hands-on experience.
- You can already: Train deep learning and NLP models with TensorFlow, PyTorch, spaCy and Hugging Face, and package, track and serve them.
- 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

ML Advanced Training agenda

47 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: ML concepts, Python setup and first model

Included from Foundation

1. Introduction to machine learning: definition, its evolution and how it differs from rule-based programming
2. Types of machine learning: supervised, unsupervised, reinforcement and self-supervised learning
3. Potential applications of machine learning: forecasting, fraud detection, recommendations, vision and language
4. The ML workflow: problem framing, data collection, preparation, training, evaluation and deployment
5. Environment setup: Python, Jupyter, NumPy, pandas, Matplotlib and scikit-learn in a virtual environment
6. Working with data in pandas: loading CSV files, exploring, filtering and summary statistics
7. Data preprocessing and cleaning: missing values, duplicates, outliers and data types
8. Overview of supervised and unsupervised learning: labels, features, targets and clusters
9. Train and test split: why we hold out data and how data leakage happens
10. Performance metrics: accuracy, precision and recall and when accuracy misleads
11. First model: linear regression with scikit-learn fit, predict and score

DAY 2 — Essential: Supervised and unsupervised algorithms, evaluation and tuning

Included from Essential

12. Supervised learning algorithms: linear regression, logistic regression and decision trees
13. Ensemble methods: random forests and gradient boosting with XGBoost or LightGBM
14. Other classic algorithms: k-nearest neighbours, support vector machines and naive Bayes
15. Feature engineering: encoding categories, scaling, binning and creating domain features
16. scikit-learn pipelines: Pipeline and ColumnTransformer for repeatable preprocessing
17. Unsupervised learning algorithms: clustering with k-means, DBSCAN and hierarchical clustering
18. Dimensionality reduction: PCA for compression and visualisation of high-dimensional data
19. Hands-on exercises: implementing supervised and unsupervised algorithms in Python with scikit-learn
20. Model evaluation techniques: cross-validation, confusion matrix, F1, ROC-AUC and regression errors
21. Hyperparameter tuning: GridSearchCV and RandomizedSearchCV for fine-tuning models
22. Overfitting and regularisation: bias-variance trade-off, L1 and L2 penalties and learning curves
23. Imbalanced data: class weights, resampling and threshold tuning

DAY 3 — Intermediate: Deep learning, NLP and ML engineering

Included from Intermediate

24. Introduction to neural networks and deep learning: neurons, activations, loss, backpropagation and optimisers
25. Deep learning architectures: convolutional neural networks and recurrent neural networks such as LSTMs
26. Hands-on exercises: implementing deep learning in Python with TensorFlow and Keras, and an overview of PyTorch
27. Introduction to NLP and its applications: tokenisation, embeddings and text pipelines
28. NLP techniques: sentiment analysis, text classification and named entity recognition with spaCy and NLTK
29. Transformers: attention, pretrained models and fine-tuning with Hugging Face Transformers
30. Transfer learning: reusing pretrained vision and language models with limited data

31. Time series forecasting: lag features, seasonality and validation over time
32. Recommendation systems: collaborative filtering, content-based and hybrid approaches
33. Experiment tracking: MLflow runs, parameters, metrics and model registry
34. Model explainability: feature importance, SHAP values and partial dependence
35. Serving models: saving with joblib or ONNX and exposing predictions through a FastAPI endpoint

DAYS 4-5 **Advanced: Production ML, scale, responsible AI and trends** **New at this level**

36. Real-world use cases: image and speech recognition, recommendation systems and predictive modelling in production
37. Scaling training: GPUs, mixed precision, distributed data parallel and efficient data loaders
38. Advanced tuning: Bayesian optimisation with Optuna and early stopping
39. Large language models: prompt design, retrieval-augmented generation and parameter-efficient fine-tuning with LoRA
40. Deployment patterns: batch scoring, real-time APIs, Docker images and Kubernetes rollout
41. Model monitoring: data drift, concept drift and performance decay with Evidently and alerts
42. Data and model versioning: DVC, feature stores and reproducible training pipelines
43. Ethical considerations surrounding machine learning: bias, fairness metrics and privacy
44. Governance: model cards, approval workflows, audit trails and the EU AI Act risk categories
45. Future trends in machine learning: foundation models, agents, edge ML and efficient small models
46. Learning path: recap of key takeaways and additional resources for further development in ML
47. Capstone: an end-to-end ML system from raw data to a monitored, versioned production model

PRACTICE

Hands-on labs

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

FOUNDATION **Your first prediction model**

Clean a housing dataset with pandas and train, test and evaluate a linear regression model in scikit-learn.

ESSENTIAL **Churn prediction with tuned models**

Build a customer churn pipeline, compare logistic regression, random forest and gradient boosting, and tune the best model.

INTERMEDIATE **Text classifier served as an API**

Fine-tune a transformer for sentiment analysis, track runs in MLflow and serve the best model behind a FastAPI endpoint.

ADVANCED **End-to-end production ML capstone**

Build a reproducible training pipeline, deploy the model as a containerised service and add drift monitoring and a model card.

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 ML 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 ML 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/ml/

ENROL OR RESERVE A SEAT
ML Advanced Training

WHATSAPP · INDIA
+91 99057 40781

WHATSAPP · USA
+1 (469) 756-6329

EMAIL
contact@DevOpsSchool.com