DevOpsSchool

Course curriculum · November 2026

LEVEL 04 OF 04 · ADVANCED

Artificial Intelligence Advanced Training

Fine-tune and evaluate models, build tool-using AI agents and run AI in production with MLOps, security 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/ai/artificial-intelligence-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 Artificial Intelligence course path, you will be able to:

1

Build machine learning models

Prepare data and train, tune and evaluate classical and deep learning models in Python.

2

Apply LLMs and RAG

Use prompt engineering, embeddings and retrieval to build useful generative AI applications.

3

Run AI in production

Deploy, monitor, secure and govern AI systems with MLOps practices.

WHAT THE ADVANCED LEVEL ADDS

Fine-tuning, agents, MLOps and AI governance

Hands-on: Ship a monitored AI agent. Build a tool-using agent, evaluate it against a test set, deploy it on Kubernetes and add monitoring and guardrails.

WHO SHOULD ATTEND

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

BEFORE YOU START

- The Intermediate level, or equivalent hands-on experience.
- You can already: Apply transformers and LLMs with prompt engineering and RAG, track experiments and deploy models as APIs.
- 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

Artificial Intelligence 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: AI concepts, Python setup and first model

Included from Foundation

1. What AI is: narrow and general AI, machine learning, deep learning and generative AI in context
2. Classical AI: search algorithms such as BFS, DFS and A*, rule-based and expert systems
3. Types of learning: supervised, unsupervised, self-supervised and reinforcement learning
4. Real-world use cases: vision, language, recommendation, forecasting and automation
5. Python environment: Anaconda or venv, pip, Jupyter Notebook and VS Code setup
6. Python data stack: NumPy arrays, pandas DataFrames and Matplotlib plots
7. Math for AI: vectors and matrices, probability, statistics and the intuition of gradients
8. Data basics: datasets, features, labels, train/test splits and data quality
9. First model: training a scikit-learn classifier on the Iris dataset and making predictions
10. Evaluation basics: accuracy, confusion matrix and why a held-out test set matters
11. Responsible AI basics: bias, fairness, privacy and transparency from day one

DAY 2 Essential: Core machine learning and neural networks

Included from Essential

12. Data preparation: cleaning, missing values, encoding categories, scaling and pandas pipelines
13. Exploratory data analysis: distributions, correlations and visualisation with Seaborn
14. Regression: linear, polynomial and regularised regression with Ridge and Lasso
15. Classification: logistic regression, k-NN, decision trees, random forests and gradient boosting
16. Clustering and dimensionality reduction: k-means, DBSCAN, PCA and t-SNE
17. Model evaluation: precision, recall, F1, ROC-AUC, MAE and RMSE for the right problem
18. Cross-validation and tuning: k-fold, GridSearchCV, RandomizedSearchCV and avoiding leakage
19. Feature engineering: feature creation, selection, interaction terms and scikit-learn Pipelines
20. Neural network basics: perceptrons, activation functions, loss, backpropagation and optimisers
21. Deep learning frameworks: building and training a network with PyTorch or TensorFlow Keras
22. NLP basics: tokenisation, bag of words, TF-IDF and text classification
23. Computer vision basics: images as tensors, convolution and a simple image classifier

DAY 3 Intermediate: Deep learning, LLMs, RAG and deployment

Included from Intermediate

24. Convolutional networks: CNN architectures, data augmentation and transfer learning with ResNet
25. Sequence models and transformers: RNN and LSTM limits, attention and transformer architecture
26. Hugging Face ecosystem: Transformers pipelines, model hub, tokenizers and datasets
27. Large language models: how LLMs are trained, context windows, tokens and calling model APIs
28. Prompt engineering: system prompts, few-shot examples, structured output and evaluation
29. Embeddings and vector databases: semantic search with FAISS, Chroma or pgvector
30. Retrieval-augmented generation: chunking, retrieval, grounding and citing sources

31. Experiment tracking: MLflow runs, parameters, metrics, artifacts and the model registry
32. Model serving: wrapping models in FastAPI, containerising with Docker and testing endpoints
33. Explainability: SHAP and LIME to explain predictions to stakeholders
34. Troubleshooting models: overfitting, underfitting, class imbalance and data drift signals
35. GPU and compute basics: CUDA, Colab and cloud GPU instances, batch size and memory limits

DAYS 4-5 **Advanced: Fine-tuning, agents, MLOps and AI governance** **New at this level**

36. Fine-tuning LLMs: supervised fine-tuning, LoRA and QLoRA with PEFT, and when not to fine-tune
37. AI agents: tool calling, planning loops, memory and frameworks such as LangGraph
38. LLM evaluation: test sets, LLM-as-judge, hallucination and groundedness metrics
39. MLOps pipelines: automated training, validation and deployment with CI/CD and Kubeflow or Airflow
40. Scalable serving: Kubernetes, KServe and vLLM, autoscaling and GPU scheduling
41. Distributed training: data and model parallelism, mixed precision and checkpointing
42. Production monitoring: data and concept drift, latency, cost and quality dashboards
43. AI security: prompt injection, data leakage and the OWASP Top 10 for LLM applications
44. AI governance: NIST AI RMF, EU AI Act risk categories, model cards and audit trails
45. Cost and performance optimisation: quantisation, distillation, caching and model selection
46. Multimodal AI: vision-language models, speech-to-text and image generation in applications
47. Capstone: design, build and deploy a governed AI application end to end

PRACTICE

Hands-on labs

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

FOUNDATION **Train your first classifier**
Set up Jupyter with pandas and scikit-learn, explore a small dataset and train and evaluate a first classification model.

ESSENTIAL **Build an end-to-end prediction pipeline**
Clean a real dataset, compare three algorithms with cross-validation and train a small neural network for the same task.

INTERMEDIATE **Build a RAG question-answering service**
Index a document set with embeddings, build a RAG pipeline over an LLM and serve it as a FastAPI endpoint in Docker.

ADVANCED **Ship a monitored AI agent**
Build a tool-using agent, evaluate it against a test set, deploy it on Kubernetes and add monitoring and guardrails.

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 Artificial Intelligence 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 ARTIFICIAL INTELLIGENCE 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/ai/

ENROL OR RESERVE A SEAT
Artificial Intelligence Advanced Training

WHATSAPP · INDIA
+91 99057 40781

WHATSAPP · USA
+1 (469) 756-6329

EMAIL
contact@DevOpsSchool.com