DevOpsSchool

Course curriculum · November 2026

LEVEL 04 OF 04 · ADVANCED

Deep Learning Advanced Training

Train deep learning models at scale, compress and deploy them for production inference and govern them responsibly.

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	48 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/deeplearning/deep-learning-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 Deep Learning course path, you will be able to:

1

Build modern neural networks

Design and train CNNs, sequence models and transformers in PyTorch for real datasets.

2

Fine-tune pretrained models efficiently

Adapt pretrained vision and language models with transfer learning and LoRA.

3

Ship models to production

Optimise, deploy and monitor deep learning models with scalable training and serving tools.

WHAT THE ADVANCED LEVEL ADDS

Distributed training, optimisation, deployment and responsible AI

Hands-on: Deploy an optimised model to production. Export a trained model to ONNX, quantize it, serve it with Triton behind an API and load-test latency and throughput.

WHO SHOULD ATTEND

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

BEFORE YOU START

- The Intermediate level, or equivalent hands-on experience.
- You can already: Apply transformers and generative models, speed up training on GPUs and tune hyperparameters systematically.
- 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

Deep Learning Advanced Training agenda

48 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: Neural network concepts, PyTorch setup, first model

Included from Foundation

1. Deep learning landscape: AI versus machine learning versus deep learning, where deep learning wins and its limits
2. Math refresher: vectors, matrices, tensors, derivatives, chain rule and probability for neural networks
3. Perceptron and multilayer networks: neurons, weights, biases, layers and forward propagation
4. Activation functions: sigmoid, tanh, ReLU, GELU and softmax and when to use each
5. Loss functions: mean squared error, cross-entropy and binary cross-entropy
6. Backpropagation and gradient descent: computing gradients, learning rate, batch, mini-batch and stochastic updates
7. Environment setup: Python, Jupyter, conda or uv, PyTorch with CUDA, Google Colab and checking the GPU
8. PyTorch tensors: creation, indexing, broadcasting, device placement and autograd
9. Building models: nn.Module, nn.Linear, nn.Sequential, optimizers and the training loop
10. Data handling: Dataset and DataLoader, train, validation and test splits, normalisation
11. Evaluation basics: accuracy, precision, recall, confusion matrix and learning curves
12. First example: classify handwritten digits on MNIST or Fashion-MNIST with a fully connected network

DAY 2 — Essential: CNNs, sequence models, regularisation and transfer learning

Included from Essential

13. Convolutional neural networks: convolutions, kernels, stride, padding, pooling and feature maps
14. CNN architectures: LeNet, VGG, ResNet skip connections and EfficientNet design ideas
15. Optimizers: SGD with momentum, RMSprop, Adam and AdamW, learning rate schedules and warm-up
16. Regularisation: dropout, weight decay, early stopping and data augmentation with torchvision transforms
17. Normalisation and initialisation: batch norm, layer norm, Xavier and He initialisation
18. Transfer learning: pretrained torchvision and timm models, freezing layers and fine-tuning heads
19. Sequence models: RNN, LSTM and GRU, vanishing gradients and sequence padding
20. Word embeddings: tokenisation, embedding layers, word2vec and GloVe concepts
21. Keras 3 overview: Sequential and functional APIs with JAX, TensorFlow or PyTorch backends
22. Experiment tracking: TensorBoard and MLflow for metrics, parameters and model artifacts
23. Saving and loading models: state_dict, checkpoints and resuming training
24. Debugging training: overfitting versus underfitting, exploding gradients and sanity checks on small batches

DAY 3 — Intermediate: Transformers, generative models and training at speed

Included from Intermediate

25. Attention and transformers: self-attention, multi-head attention, positional encoding, encoder and decoder blocks
26. Hugging Face ecosystem: Transformers, Datasets and Tokenizers, pipelines and the Trainer API
27. NLP tasks: text classification, named entity recognition and summarisation with pretrained models
28. Vision transformers: ViT patches, CLIP image-text embeddings and when to prefer them over CNNs
29. Autoencoders and VAEs: latent spaces, reconstruction loss, KL divergence and anomaly detection
30. Generative adversarial networks: generator and discriminator, training instability and DCGAN

31. Diffusion models: forward noising, denoising networks and text-to-image generation with diffusers
32. Parameter-efficient fine-tuning: LoRA and QLoRA with the PEFT library
33. Faster training: mixed precision with torch.autocast, gradient accumulation and torch.compile
34. Hyperparameter tuning: grid, random and Bayesian search with Optuna
35. Model interpretability: saliency maps, Grad-CAM and attention visualisation
36. Troubleshooting: GPU out-of-memory errors, data loader bottlenecks, NaN losses and reproducibility seeds

DAYS 4-5 **Advanced: Distributed training, optimisation, deployment and responsible AI** **New at this level**

37. Distributed training: data parallelism with DDP, FSDP sharding, multi-GPU and multi-node launch with torchrun
38. Large model techniques: gradient checkpointing, DeepSpeed ZeRO stages and memory planning
39. Model compression: post-training quantization, quantization-aware training, pruning and knowledge distillation
40. Model export: ONNX export, ONNX Runtime and TensorRT optimisation for low-latency inference
41. Serving: NVIDIA Triton Inference Server, BentoML or FastAPI, batching and GPU utilisation
42. LLM inference: vLLM, KV cache, continuous batching and serving open-weight models
43. MLOps for deep learning: data and model versioning, model registry and CI/CD for training pipelines
44. Monitoring in production: latency, throughput, data drift, prediction drift and retraining triggers
45. Responsible AI: bias evaluation, fairness metrics, explainability, model cards and the EU AI Act risk tiers
46. Security of models: adversarial examples, data poisoning, model theft and prompt injection risks
47. Cost and hardware planning: GPU selection, cloud GPU instances and cost per training run and per inference
48. Capstone: train, optimise, deploy and monitor a deep learning model end to end

PRACTICE

Hands-on labs

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

FOUNDATION **Train your first neural network**
Build, train and evaluate a PyTorch classifier on Fashion-MNIST, plot loss curves and report test accuracy.

ESSENTIAL **Fine-tune a pretrained image classifier**
Fine-tune a pretrained ResNet on a custom image dataset with augmentation and early stopping, and compare it with a CNN trained from scratch.

INTERMEDIATE **Fine-tune a transformer with LoRA**
Fine-tune a small pretrained language model with LoRA on a text classification dataset, track runs in MLflow and tune it with Optuna.

ADVANCED **Deploy an optimised model to production**
Export a trained model to ONNX, quantize it, serve it with Triton behind an API and load-test latency and throughput.

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 Deep Learning 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 DEEP LEARNING 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/deeplearning/

ENROL OR RESERVE A SEAT
Deep Learning Advanced Training

WHATSAPP · INDIA
+91 99057 40781

WHATSAPP · USA
+1 (469) 756-6329

EMAIL
contact@DevOpsSchool.com