JUPITER PRIME

Consulting • Training • Advisory

جوبيتر برايم للاستشارات والتدريب

COURSE BROCHURE

Deep Learning Fundamentals Training Course

jupiterprime.uk

info@jupiterprime.uk

Registered in the United Kingdom and Saudi Arabia

Deep Learning Fundamentals Training Course

Deep Learning Fundamentals is a 1-week training course for analysts, engineers and technical managers who are new to neural networks and need to judge where layered models fit their work, ending with a deep learning use-case brief. Image, text and sensor data now drive many decisions, and organisations that grasp how networks learn choose better projects. Jupiter Prime Training delivers this introduction to Deep Learning Fundamentals.

Why Deep Learning Fundamentals Matters Now

Three forces shape the subject today: cheaper accelerated computing, the spread of pretrained transformer models, and growing scrutiny from boards and customers over how automated predictions are produced. An organisation that reads loss curves, validation splits and failure modes well can back the right use cases and stop weak ones early; one that does not risks costly models that fail silently once deployed. The course anchors oversight in ISO/IEC 42001, the international standard for an AI management system, and in the NIST AI Risk Management Framework NIST AI RMF.

Who Gains Most from Deep Learning Fundamentals

- Data analysts who prepare labelled datasets and interpret model predictions for business units
- Software and process engineers who integrate trained networks into applications or plant systems
- Technical managers who approve model experiments and allocate computing budgets
- Quality and inspection specialists who assess image or signal classification results
- Product owners who translate operational problems into prediction tasks

A cross-functional group that pairs analysts with the engineers and managers who sponsor their models gains most, leaving with one vocabulary for architectures and metrics.

Deep Learning Fundamentals Capabilities

Participants move from a single neuron to complete architectures, building judgement to question a model before production.

- Explain how weights, activations and backpropagation adjust a network during training — model sponsors
- Select between convolutional, recurrent and transformer architectures for a given data type — solution design teams
- Diagnose overfitting and underfitting from training and validation loss curves — experiment approvals
- Calculate precision, recall and F1 scores for a classification model — acceptance decisions
- Weigh transfer learning against training from scratch — computing budget holders

- Specify the data, metrics and risk controls for a deep learning proposal — steering committees

Key Deep Learning Fundamentals Terms

- Neural Network — a layered model of connected units that learns patterns by adjusting numeric weights from example data.
- Backpropagation — the method that sends prediction errors backwards through a network to calculate how each weight should change.
- Activation Function — a non-linear rule, such as ReLU, that decides how strongly a unit passes its signal onward.
- Convolutional Layer — a layer that scans small filters across an image or signal to detect local features such as edges.
- Attention Mechanism — a technique letting a model weigh which parts of an input sequence matter most for each output.
- Overfitting — the condition where a model memorises training examples and performs poorly on new, unseen data.

Deep Learning Fundamentals Day by Day

Day 1: Neural Network Building Blocks

- Perceptrons and Weighted Sums — How Single Units Form Predictions
- Activation Functions Compared — Choosing Non-Linearity for Hidden Layers
- Feedforward Layer Stacking — Depth Versus Width Trade-Offs
- Tensors and Batch Shapes — Matching Data Formats to Network Inputs
- Use-Case Screening Matrix — Deciding Where Neural Networks Add Value

Day 2: Training Networks with Backpropagation

- Loss Functions for Regression and Classification — Defining Error Correctly
- Gradient Descent and Learning Rates — Controlling Training Stability
- Optimisers Such as Adam and Momentum — Speed Versus Convergence
- Regularisation with Dropout and Weight Decay — Containing Overfitting
- Training and Validation Splits — Setting Honest Performance Baselines

Day 3: Convolutional and Sequence Architectures

- Convolution, Pooling and Feature Maps — Visual Inspection Use Cases
- Pretrained Image Backbones — When Transfer Learning Saves Effort
- Recurrent Networks and LSTM Cells — Modelling Sensor Time Series
- Embedding Layers for Text — Representing Words as Vectors
- Architecture Selection Checklist — Matching Model Families to Data Types

Day 4: Transformers and Model Evaluation

- Self-Attention and Positional Encoding — Why Transformers Scale Well
- Fine-Tuning Pretrained Language Models — Cost and Data Requirements
- Confusion Matrices, Precision and Recall — Judging Classifier Quality
- Bias and Robustness Testing — Checks Aligned with NIST AI RMF
- Explainability with Saliency Maps — Communicating Model Behaviour to Reviewers

Day 5: Practising the Deep Learning Use-Case Brief

- Problem Framing and Data Inventory — Applying the Day 1 Screening Matrix
- Baseline Experiment Plan — Reusing Loss, Optimiser and Split Choices
- Architecture Recommendation — Applying the Day 3 Selection Checklist
- Risk Register Mapped to ISO/IEC 42001 Controls — Approval Readiness
- Deep Learning Use-Case Brief — Presenting Scope, Metrics and Next Steps

Applying Deep Learning Fundamentals After the Course

- Run the use-case screening matrix on three pending analytics requests and rank them for a neural network pilot.
- Rebuild the acceptance report of an existing classifier using confusion matrices, precision and recall.

Deep Learning Fundamentals Questions FAQ

What is Deep Learning Fundamentals and why does it matter now?

Deep Learning Fundamentals covers how layered neural networks learn from data and how to judge them. It matters now because pretrained models and cheaper computing have put image, text and sensor prediction within reach of most organisations, which need staff who can question results.

Which roles gain most from Deep Learning Fundamentals training?

Analysts, engineers and technical managers new to neural networks gain most. They prepare data, integrate models or approve experiments and need shared language on architectures and metrics.

How does Deep Learning Fundamentals differ from a general machine learning course?

Deep Learning Fundamentals focuses on neural network architectures rather than the wider range of statistical learning methods. It concentrates on backpropagation, convolutional, recurrent and transformer models, leaving tree-based methods and engineering-depth experimentation to other courses.

What do you bring back from Deep Learning Fundamentals?

You bring back a structured proposal for one neural network project. It records the problem, data inventory, chosen architecture, baseline experiment, evaluation metrics and risk controls, ready for your sponsor to approve or refine.

Deep Learning Fundamentals: The Way Forward

The deep learning use-case brief gives sponsors a clear basis for the first decision: whether a candidate problem justifies a neural network experiment at all. Analysts and engineers use it to scope data collection. After the first experiment, revisit the evaluation metrics, because early results often show which errors truly matter to operations.