Artificial Neural Networks, Machine Learning & Deep Learning (Advanced)

Course Objectives

By the end of this training, participants will be able to:

- Understand advanced ML and ANN concepts including optimization, regularization, and hyperparameter tuning
- Build, train, and evaluate complex neural network architectures (ANN, CNN, RNN/LSTM)
- Apply transfer learning and modern deep learning practices
- Use advanced training strategies to improve model accuracy and reduce overfitting
- Understand sequence models, attention mechanisms, and basics of transformer architectures
- Deploy, optimize, and monitor models using industry-standard DL tools
- Apply best practices in model evaluation, explainability, and responsible AI

Learning Outcomes

Participants will walk away with the ability to:

✓ Build ANN models using Keras/PyTorch
✓ Train CNNs for image recognition tasks
✓ Develop and tune RNN/LSTM networks for sequence data
✓ Apply transfer learning and compare model performance
✓ Use optimizers, learning rate schedules, and regularization effectively
✓ Interpret model metrics and visualize training performance
✓ Export and deploy models with scalable architectures
✓ Implement best practices in deep learning for real-world applications

---

Tools & Technologies Needed

Software

- Python 3.8+
- Jupyter Notebook / JupyterLab
- Anaconda (optional but recommended)
- VS Code / PyCharm (optional)

Libraries

- TensorFlow / Keras
- PyTorch
- NumPy, Pandas
- Matplotlib, Seaborn
- Scikit-learn
- OpenCV (optional, for CNN experiments)

Hardware

- Laptop with at least 8GB RAM
- GPU optional (CPU works for small datasets)

Cloud (Optional)

- Azure ML / Google Colab / Kaggle Notebooks for GPU access

🗓️ DAY 1 — Advanced ML & Artificial Neural Networks

Module 1: Advanced ML Concepts & Workflow (1.5 hrs)

- End-to-end ML pipeline
- Cross-validation strategies
- Regularization: L1, L2, dropout
- Feature engineering best practices
- Hyperparameter optimization (Grid, Random, Bayesian)

Module 2: Optimization Techniques in ML (1 hr)

- Gradient descent variants: SGD, Momentum, Adam
- Learning rate schedules
- Loss functions & objective formulation

Module 3: Foundations of ANN (2 hrs)

- Biological vs artificial neurons
- Activation functions
- Weight initialization (Xavier, He)

- Forward and backward propagation
- Vanishing/exploding gradients

Hands-On Lab 1 (1 hr)

Build a Multi-Layer Perceptron (MLP) classifier

- Custom layers
- Activation functions
- Regularization and tuning

📅 **DAY 2 — Deep Learning Architectures & Advanced Applications**

Module 4: Deep Neural Networks (DNN) & Training Tricks (1.5 hrs)

- Batch normalization
- Dropout
- Residual connections
- Faster convergence techniques

Module 5: Convolutional Neural Networks (CNN) (2 hrs)

- Convolution, kernels, padding, pooling
- Feature extraction concepts
- Architectures: LeNet, VGG, ResNet
- Transfer learning & fine-tuning

Hands-On Lab 2 (1 hr)

Train a CNN on MNIST/CIFAR

- Visualize training
- Apply transfer learning with MobileNet/ResNet

Module 6: RNN, LSTM, GRU & Sequence Models (1.5 hrs)

- Recurrent connections

- Long-term dependencies
- Applications in text, speech, time-series
- LSTM vs GRU
- Limitations leading to Transformers

Module 7: Introduction to Transformers & Modern DL (1.5 hrs)

- Self-attention
- Encoder/decoder structure
- Architecture basics of BERT, GPT
- How transformers outperform CNN/RNN
- Fine-tuning vs prompt engineering

Module 8: Deployment, Optimization & Best Practices (1 hr)

- Model export (SavedModel, ONNX)
- Quantization, pruning, distillation
- Monitoring model drift
- MLOps overview
- Explainability (SHAP, LIME)
- Responsible & ethical AI