Artificial Intelligence & Machine Learning Course Syllabus

Course Overview

This course introduces the foundational concepts of Artificial Intelligence (AI) and Machine Learning (ML). It covers theory, algorithms, and real-world applications, giving students the knowledge to understand and implement AI and ML solutions. Students will learn about supervised, unsupervised, and reinforcement learning, as well as AI frameworks and tools.

Course Objectives

By the end of the course, students will:

- Understand key AI and ML concepts.
- Be familiar with supervised, unsupervised, and reinforcement learning.
- Gain hands-on experience with ML algorithms.
- Learn to evaluate model performance and tune hyperparameters.
- Understand ethical considerations in AI/ML.

Course Structure

Week	Topics
1	**Introduction to Artificial Intelligence** - History of AI - Key AI concepts and applications - Differences between AI, Machine Learning, and Deep Learning
2	**Introduction to Machine Learning** - Overview of ML types: Supervised, Unsupervised, and Reinforcement Learning - ML use cases and real-world applications
3	**Linear Algebra & Probability for Machine Learning** - Vectors, Matrices, and Operations - Probability and Statistics concepts - Python for basic math operations
4	**Supervised Learning: Regression** - Linear Regression - Polynomial Regression - Evaluation Metrics (RMSE, R^2)
5	**Supervised Learning: Classification** - Logistic Regression - Support Vector Machines (SVMs)

Week	Topics
	- k-Nearest Neighbors (k-NN) - Performance Metrics (Accuracy, Precision, Recall, F1-Score)
6	**Model Evaluation & Hyperparameter Tuning** - Cross-Validation - Grid Search, Random Search - Bias-Variance Tradeoff
7	**Unsupervised Learning: Clustering** - k-Means Clustering - Hierarchical Clustering - Gaussian Mixture Models (GMMs)
8	**Unsupervised Learning: Dimensionality Reduction** - Principal Component Analysis (PCA) - t-SNE, UMAP - Applications in Visualization
9	**Introduction to Neural Networks** - Perceptron and Multi-layer Perceptron (MLP) - Activation Functions (ReLU, Sigmoid, etc.) - Backpropagation
10	**Deep Learning & Convolutional Neural Networks (CNNs)** - CNN Architecture - Convolutional Layers, Pooling - Applications in Image Processing
11	**Recurrent Neural Networks (RNNs) & Long Short-Term Memory (LSTM)** - RNN Structure - Time-series Data - LSTMs for Sequential Data
12	**Reinforcement Learning** - Markov Decision Processes (MDPs) - Q-Learning and Deep Q-Networks (DQNs) - Policy Gradient Methods
13	**Natural Language Processing (NLP)** - Text Preprocessing (Tokenization, Lemmatization) - Bag-of-Words, TF-IDF - Word Embeddings and Transformers

Week	Topics
14	**AI Ethics and Responsible AI** - Bias in AI Models - Fairness, Accountability, and Transparency - AI in Society and Ethical Implications
15	**Capstone Project** - Choose a real-world problem - Build an AI/ML solution using the concepts learned - Final presentation and evaluation

---

Grading & Evaluation

- **Assignments (30%)**: Weekly coding exercises on ML models and AI concepts.
- **Quizzes (10%)**: Periodic quizzes to assess theoretical understanding.
- **Midterm Exam (20%)**: A written or coding exam covering the first half of the course.
- **Final Project (30%)**: A capstone project that integrates AI/ML concepts.
- **Class Participation (10%)**: Contribution to discussions and labs.

Textbooks and Resources

1. **Artificial Intelligence: A Modern Approach** by Stuart Russell and Peter Norvig
2. **Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow** by Aurélien Géron
3. **Deep Learning** by Ian Goodfellow, Yoshua Bengio, and Aaron Courville

Tools & Platforms

- Python (Anaconda, Jupyter)
- Scikit-learn, TensorFlow, Keras, PyTorch
- Google Colab for cloud-based execution