

AIIC Syllabus

1. Data Preparation

Data Processing

- Python Pandas
- Python Numpy
- Linear algebra fundamentals
- Exploratory data analysis (EDA)

Data Visualization

- Python Matplotlib
- Python Seaborn

2. Supervised Learning

Regression Models

- Linear regression (Least Squares, Gradient Descent)
- Polynomial regression

Classification Models

- Decision Trees (entropy, information gain)
- Support Vector Machines (SVM) and hyperplanes
- Naïve Bayes classification

Regularization & Overfitting

- L1 (Lasso) and L2 (Ridge) Regularization
- Cross-validation techniques

3. Unsupervised Learning

Clustering

- K-means clustering (cost function minimization)
- Hierarchical clustering
- Gaussian Mixture Models (GMM)

Dimensionality Reduction

- PCA and its application in visualization
- t-SNE for high-dimensional data representation

4. Reinforcement Learning

Markov Decision Processes (MDPs)

- States, actions, rewards
- Bellman Equations

Q-Learning

- Temporal Difference Learning
- Exploration vs. exploitation

5. AI Search

Uninformed Search

- Breadth-First Search (BFS)
- Depth-First Search (DFS)
- Uniform Cost Search

Informed Search

- A* Algorithm and heuristics
- Minimax search (game AI)

6. Logical Reasoning

Propositional & First-Order Logic

- CNF, logical connectives
- Unification in inference

Resolution Theorem Proving

- Satisfiability and SAT solvers

7. Evaluation of ML Models

Classification Metrics

- Precision, recall, F1-score
- ROC-AUC curves

Bias-Variance Tradeoff

- Overfitting detection
- K-fold cross-validation

8. Constraint Satisfaction Problems

Backtracking & Constraint Propagation

- AC-3 algorithm
- Forward checking

9. Kernel Methods

Support Vector Machines (SVM)

- Margin maximization
- Soft-margin vs. hard-margin SVM

Kernel Trick

- Transforming low-dimensional data to high-dimensional space
- Radial Basis Function (RBF) kernel

10. Recommender Systems

Collaborative Filtering

- User-based and item-based filtering
- Similarity metrics (cosine similarity, Pearson correlation)

Content-Based Filtering

- TF-IDF and word embeddings
- Feature engineering for recommendation models

Hybrid Recommender Systems

- Combining collaborative and content-based filtering
- Matrix factorization (SVD, ALS algorithms)
- Deep learning approaches (Neural Collaborative Filtering, Autoencoders for recommendations)

11. Neural Network Fundamentals

Basic Neural Network Structure

- Neurons, layers, weights, biases
- Activation functions: Sigmoid, Tanh, ReLU
- Forward propagation

Training Neural Networks

- Loss functions
- Backpropagation
- Gradient descent and optimization
- Learning rate

12. Deep Learning Optimization

Optimization Algorithms

- Stochastic Gradient Descent (SGD)
- Momentum
- Adam optimizer

Training Challenges

- Vanishing and exploding gradients
- Overfitting in deep networks
- Dropout
- Batch normalization

13. Convolutional Neural Networks (CNNs)

CNN Fundamentals

- Convolution operation
- Filters and feature maps
- Padding and stride
- Pooling layers

Applications

- Image classification
- Object detection basics
- Feature extraction

14. Recurrent Neural Networks and Sequence Models

RNN Fundamentals

- Sequential data modeling
- Hidden states
- Backpropagation through time

Advanced Sequence Models

- LSTM
- GRU
- Attention mechanism basics

15. Transformers and Modern Deep Learning

Transformer Architecture

- Self-attention
- Multi-head attention
- Positional encoding
- Encoder-decoder structure

Applications

- Large language models
- Text generation
- Embeddings
- Fine-tuning basics

16. Practical Deep Learning

Model Building

- PyTorch or TensorFlow basics
- Dataset loading
- Training loop
- Validation and testing

Model Evaluation

- Accuracy, loss curves
- Confusion matrix
- Hyperparameter tuning
- Model deployment basics