Machine Learning Course Content

1. Introduction to Machine Learning
2. Python for Machine Learning (Hands on)
3. Mathematics and Statistics for ML
4. Data Preprocessing and Feature Engineering (Hands on)
5. Supervised Learning: Regression (Hands on)
6. Supervised Learning: Classification (Hands on)
7. Ensemble Learning and Gradient Boosting (Hands on)
8. Unsupervised Learning (Hands on)
9. Model Evaluation, Tuning and Explainability (Hands on)
10. Time Series and Recommender Systems
11. MLOps and Deployment (Hands on)
12. Capstone Projects

Introduction to Machine Learning

- Types of ML: Supervised, Unsupervised, Reinforcement
- ML Lifecycle, Problem Framing, CRISP-DM, Roles in ML Teams
- Bias-Variance Tradeoff, Overfitting, Train, Validation and Test Sets

Python for Machine Learning (Hands on)

- Python Basics, Functions, OOP, Virtual Environments
- NumPy, Pandas, Polars Overview, Data Wrangling
- Jupyter Notebooks, Git and GitHub

Mathematics and Statistics for ML

- Linear Algebra and Calculus Basics for ML
- Probability, Distributions, Sampling
- Hypothesis Testing, Confidence Intervals, Correlation

Data Preprocessing and Feature Engineering (Hands on)

- Missing Values, Outliers, Encoding, Scaling
- Feature Creation, Feature Selection, Handling Imbalanced Data
- scikit-learn Pipelines, ColumnTransformer, Avoiding Data Leakage

Supervised Learning: Regression (Hands on)

- Linear, Polynomial, Ridge, Lasso and ElasticNet Regression
- Decision Tree and Random Forest Regressors
- Metrics: MAE, RMSE, R-squared, Residual Analysis

Supervised Learning: Classification (Hands on)

- Logistic Regression, KNN, Naive Bayes, SVM, Decision Trees
- Class Weights, SMOTE, Threshold Selection
- Metrics: Precision, Recall, F1, ROC-AUC, Confusion Matrix

Ensemble Learning and Gradient Boosting (Hands on)

- Bagging, Random Forest, AdaBoost, Stacking
- XGBoost, LightGBM and CatBoost
- Why Boosted Trees Lead on Tabular Data

Unsupervised Learning (Hands on)

- K-Means, Hierarchical Clustering, DBSCAN
- PCA, t-SNE and UMAP for Dimensionality Reduction
- Anomaly Detection with Isolation Forest, Customer Segmentation

Model Evaluation, Tuning and Explainability (Hands on)

- Cross Validation, Grid Search, Random Search, Optuna
- SHAP and LIME for Explaining Predictions
- Fairness, Bias and Responsible ML

Time Series and Recommender Systems

- Trend, Seasonality, ARIMA, Prophet, Forecasting Metrics
- ML-Based Forecasting with Lag Features
- Collaborative and Content-Based Recommenders

MLOps and Deployment (Hands on)

- Experiment Tracking and Model Registry with MLflow
- FastAPI Model Serving, Docker, CI/CD with GitHub Actions
- Monitoring, Data Drift, Cloud ML Overview: AWS SageMaker, Azure ML

Capstone Projects

- Credit Risk, Customer Churn, Demand Forecasting and Recommendation Projects
- GitHub Portfolio, Resume Building and Interview Preparation