Syllabus by Week

I. ML Fundamentals and Core Concepts

Week 1. ML Foundations and Classical Models I

Week 2. Model Selection and Classical Models II

Week 3. Classical Models II and Important Model Concepts

Week 4. Important Model Concepts and Learning Paradigms

II. Advanced/Modern Models

Week 5. Neural Nets & Modern Architectures

III. ML on Specialized Data Modalities

Week 6. Tabular & Graph ML Models

Week 7. Temporal, Text & Image ML Models

Detailed Syllabus

I. ML Fundamentals and Core Concepts

ML Intro

Definition, Types of ML, Applications.
The Machine Learning Workflow.
Data Preprocessing and Feature Engineering.

Classical Models I: Regression & Classification

Linear Regression & Logistic Regression

Model Evaluation, Selection, Generalization

Bias-Variance Trade-off, Overfitting, and Underfitting.
Metrics, Train-Test Split, Cross-Validation.
Hyper-parameter Optimization (HPO).

Classical Models II: Decision Trees & Ensemble Models

RFs, Boosted DTs: Gradient Boosting (e.g., XGBoost, LightGBM)

Important Model Concepts

Uncertainty Est. & Calibration

Model Interpretability

Built-in vs Post-hoc (SHAP, LIME)

Robustness (Adversarial learning, Shortcut learning, etc.), Fairness

Learning Paradigms

Supervised, Semi-supervised

Self-supervised Learning - Losses

Unsupervised: Clustering, Anomaly Detection, Dim. Reduction

Other

- Active Learning
- Reinforcement Learning
- Meta learning
- Continual / Lifelong Learning
- Federated Learning
- Transfer Learning

II. Advanced/Modern Models

Neural Nets (tabular), CNN, RNN,

LSTM, Attention, Transformer

III. ML on Specialized Data Modalities

ML on Tabular Data

Foundation models: TabPFN, FoM0D

Outlier detection: kNN, IsolationForest, DeepSVDD, AEs

Clustering: k-means, DBSCAN

ML on Graphs

Graph/Node embeddings (for e.g. Recommender systems)

Graph neural networks / transformers (for e.g. Fraud detection)

Heterogenous graphs

Graph Partitioning (for e.g. Community Detection)

ML on Temporal/Seq. Data

Models: ARIMA (classic), LSTM (modern), TFT (advanced).
Tasks: Forecasting, anomaly detection, classification.

ML on Text & Image

LLMs: text generation / classification
Image generation / classification / object detection / segmentation

Multimodal Learning (Vision-Language models)

CLIP

BLIP-2: image captioning and visual question answering (VQA)