

Machine Learning

September 18, 2024

1 Supervised Learning

- Linear Regression
- Linear Classification
- Logistic Regression
- K-Nearest Neighbors (K-NN)
- Random Forests and Ensemble Methods
 - Bagging (Bootstrap Aggregating)
 - Boosting

2 Unsupervised Learning

- Clustering Techniques
 - K-Means Clustering
- Principal Component Analysis (PCA)

3 Deep Neural Networks

- Neurons and Perceptrons
- Multi-Layer Perceptrons (MLP)
- Backpropagation Algorithm
- Optimization
 - Stochastic Gradient Descent (SGD)
 - Mini-Batch Gradient Descent
- Regularization (Early Stopping, Dropout)

- Activation Functions (Sigmoid, ReLU, etc.)
- Weight Initialization (Xavier)
- Batch Normalization

4 Computer Vision

- Fundamentals of Convolutional Layers
- Transposed Convolution and Upsampling
- AlexNet
- ResNet
- U-Net

5 Natural Language Processing

- Introduction to Sequential Data and the Need for Word Embeddings
- Recurrent Neural Networks (RNNs)
- Long Short-Term Memory (LSTM) Networks
- Word2Vec and Word Embedding Techniques
- Transformer Architecture
- Attention Mechanisms in Vision

6 Contrastive Learning

- Introduction to Contrastive Learning
- Vision Transformers (ViT) and Image-Text Encoders in CLIP
- Applications of CLIP and Contrastive Learning