CCBP TECH 4.0

AI-ML Engineer

AI/ML is a cutting-edge technology that is revolutionizing the way people think, learn and work across various fields. In this program, you'll learn several applications of AI, how to build machine learning systems, how the algorithms behind them work and how to build those algorithms!

Fundamentals of ML

In this course, you will learn the widely used terminology in AI and ML. You'll also understand the essential algorithms behind some supervised learning applications. Most importantly, you will build solutions for some real-world problems.

Introduction to AI and ML

- What is AI, ML?
- Applications of ML
- ML Components
- Supervised Learning
- Classification, Regression
- Python for ML

Linear Regression

- Hypothesis function
- Cost function
- Gradient Descent
- Normal Equation
- Vectorization
- Polynomial Regression

K - Nearest Neighbours

- k-NN for Classification
- Choice of k
- KD Trees
- LSH and Inverted Indices
- Instance Based Learning

Logistic Regression

- Sigmoid function
- Cost function
- Maximum Likelihood Estimate
- GD for Logistic Regression
- Multiclass Classification

Data preprocessing

- Handling missing values
- One-hot Encoding
- Feature Scaling
- Feature Selection

Regularization

- Overfitting
- Underfitting
- Regularized Linear Regression
- Regularized Logistic Regression

Course Project

COVID-19 Survival Prediction

Build an ML model that predicts the survival of a COVID-19 infected patient.

Supervised Machine Learning Algorithms

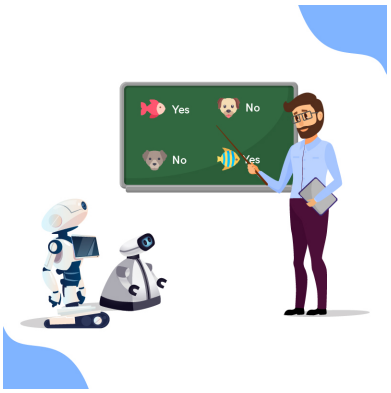

In this course, you'll learn advanced supervised learning algorithms like SVM, decision trees, naive bayes etc. You will also learn ensembling techniques which will boost the performance of ML models. During the course, you will get to apply the algorithms learnt to solve real-world problems.

SVM

- Maximum Margin Classifier
- Support Vectors
- Handling Outliers
- Kernels
- SVM as Large Margin Classifier
- Constrained Optimization

Decision Trees

- Choosing the best attribute
  - Entropy
  - Information Gain
  - Gini Index
- ID3, C4.5
- Handling missing values
- Pruning
  - Reduced-error pruning
  - Rule post-pruning
- Limitations of Decision Trees

Naive Bayes

- Conditional Probability & Bayes rule
- Naive Bayes Classifier
- Naive Bayes Algorithm
- Laplace Smoothing
- NB for Text Classification
- Handling Real-world problems

Algorithms for Regression

- Regression Trees
  - Choice of Threshold & Attribute
  - Overfitting
  - Cost Complexity Pruning
- k-NN for Regression
- Locally Weighted Regression

Learning Theory

- Evaluation of Hypothesis function
- Tuning Hyperparameters
- Bias, Variance and Noise Decomposition
- Error Analysis
- Performance Optimization

Bagging

- Bootstrapping
- Out-of-Bag error
- Bagged Decision Trees
- Random Forests

Boosting

- Gradient Boosting
- AdaBoost
- Stochastic Gradient Boosting
- ECOC
- Stacking

Course Project

Detecting Water Pollution

Develop an ML model that predicts the pollution levels of water from the chemical composition of the contaminated water.

Competitive Machine Learning

Does Competitive Programming excite you? This course will take that excitement to the next level. You'll get the right toolbox to take on your competition by showing your expertise in data pre-processing, feature engineering, modelling and maximizing the performance of the model.

Introduction

- Introduction to Competitive ML
- Exploratory Data Analysis
  - Introduction to Pandas
  - Visualizations
  - Dataset Cleaning
- Introduction to scikit-learn

Advanced Pre-processing & Feature Engineering

- Review of previously covered techniques
- Missing Values Imputation
- Label Encoding
- Target Encoding
- Mean Target Encoding
- Feature Interactions

Validation

- K-fold Cross-Validation
- Validation Strategies
- Data-splitting Strategies
- Pros and Cons of Validation

Modelling

- Recap of ML algorithms covered so far
- Hyperparameter Tuning
- Grid Search
- Advanced Model Ensembling
- Recap of Metrics covered
- Metrics Optimization

Course Project
Kaggle Expert

Get to the Kaggle Expert level by winning atleast 2 medals in public competitions.

Introduction to Deep Learning

Deep learning has taken the world by a surprise enabling rapid progress in diverse fields. In this course, you'll understand the fundamentals of deep learning algorithms and get practical experience in building deep neural networks.

Introduction to Neural Networks

- Understanding how the Brain works
- Model representation
- Vectorizing
- Activation Functions

Backpropagation

- Intuition and the Algorithm
- Gradient Checking
- Initialization

Introduction to TensorFlow

- "Hello World" from Neural networks
- Loading Training Data
- Coding Computer Vision NN
- Using callbacks

Deep Neural Networks

- Forward propagation
- Building blocks
- Forward and Backward propagation
- Hyperparameters

Practical considerations for Applying Neural Networks

- Normalizing Inputs
- Bias, Variance
- Initialization
- Regularization
- Dropouts
- Gradient Checking

Hyperparameter Tuning

- Using appropriate scale
- Normalizing activations
- Batch Normalization
- Softmax regression

Optimization

- Mini-batch gradient descent
- Exponentially weighted averages
- Gradient descent with momentum
- RMSprop
- Learning rate decay

Practical Advice for applying Deep Learning

- Orthogonalization
- Working with Train/Test datasets
- Setting proper evaluation metrics
- Data augmentation
- Improving model performance
- Error analysis
- Working with different data
- Cleaning bad data
- Transfer learning
- Multi-task learning

ML Engineering

- Dynamic Training and Inference
- Data dependencies
- Fairness

CNNs and Computer Vision

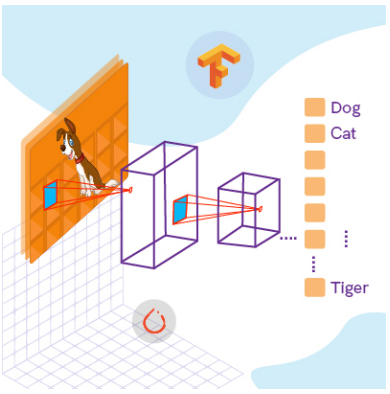

Recent developments in deep learning have advanced the performance of Computer Vision based systems significantly. In this course, you'll gain knowledge of cutting-edge research in Convolutional Neural Networks and apply it to build visual detection and recognition systems.

CNN Basics

- Edge Detection
- Strided Convolutions
- Convolutions over volume
- Pooling Layers

Deep CNNs

- ResNets
- Inception Network
- Transfer Learning
- Computer Vision
- Object Detection
- Face Recognition

Intro to CNNs in TensorFlow

- Implementing Convolutional layers
- Implementing pooling layers
- Walking through CNN of real-world images (dev, training)
- Testing accuracy

Advanced CNNs in TensorFlow

- Working with Larger Datasets
- Avoiding overfitting with
- Transfer Learning
- Multi-class Classifications

Course Project
Face Recognition

Build a face recognition software using CNNs.

RNNs and Natural Language Processing

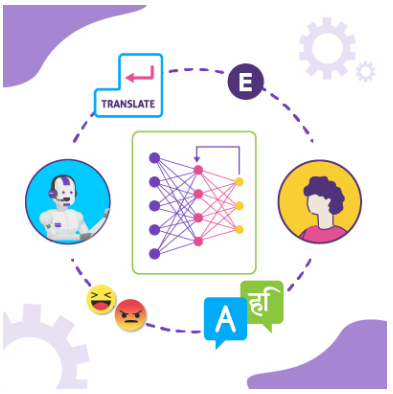

Deep learning approaches achieve very high performance across many different NLP tasks. In this course, you will gain all the necessary skills to design and implement Recurrent Neural Networks and its variants. You'll apply them to build interesting NLP based applications.

Intro to RNNs

- Sequence models
- Backpropagation through time
- Sampling
- GRU
- LSTM
- Bidirectional RNNs

Natural Language Processing

- Word Representation
- Word Embeddings
- Word2Vec
- GloVe vectors
- Sentiment classification

Speech Recognition

- Beam Search
- Bleu Score
- Attention Model
- Trigger Word Detection

NLP in TensorFlow

- Word-based encodings
- Tokenizer
- Word embeddings with IMDb dataset
- Implementing LSTMs
- Working with sequence models

Advanced RNNs in TensorFlow

- Working with Time series problems
- Forecasting problem
- Lambda layers
- Dynamically adjusting learning rates
- Using Bi-directional LSTMs on real-world time series data

Course Project
Sentiment Analysis

Build an ML model that detects the emotions/sentiments expressed in twitter trends.

Unsupervised Learning Algorithms

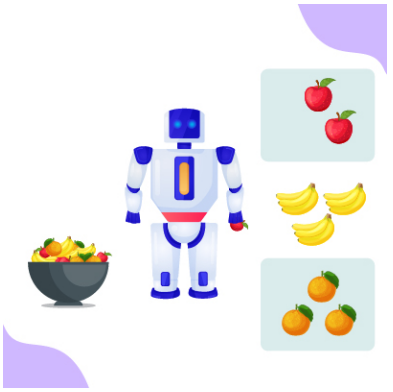

Unsupervised learning is an important area of machine learning which infers the unknown patterns in unlabelled data. In this course, you'll learn clustering and dimensionality reduction techniques in depth. You'll also understand the applications of unsupervised learning in a wide range of domains like medical imaging, market research, e-commerce, etc.

Clustering

- Introduction to unsupervised learning
- K-Means clustering
- Hierarchical clustering
- Density based clustering
- Gaussian Mixture models

Dimensionality Reduction Techniques

- Principal Component Analysis (PCA)
- Independent Component Analysis (ICA)
- PCA vs ICA
- Singular Value Decomposition (SVD)
- Non-negative Matrix Factorization

Semi-Supervised Learning

- Introduction to Semi-Supervised Learning
- Self Training
- Generative Models
  - Mixture of Gaussian Distributions (GMM)
  - Mixture of Multinomial Distributions
  - Hidden Markov Models (HMM)
- Graph-Based Algorithms
  - Mincut
  - Harmonic
  - Local and Global consistency
  - Manifold Regularization
- S3VMs
- Multiview Algorithms
  - Co-training
  - Variants of Co-training
  - Multiview Learning
- Pros and Cons of Semi-Supervised Learning algorithms

Course Project

Customer Segmentation

Identify the target customers segment of an e-commerce website.

Search-based AI

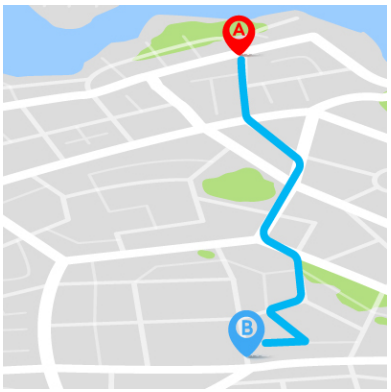

Ever wondered how Google Maps suggests the best route among multiple possible routes? Search-based AI algorithms help us choose the one with the best possible outcome. This course will cover search based AI algorithms like Alpha-beta pruning, A* search, Hill-climbing, etc. You'll also apply these algorithms to build exciting games and applications.

Uninformed Search

- Breadth-first search
- Depth-first Search
- Depth Limited Search
- Iterative Deepening Search
- Uniform Cost Search
- Bidirectional Search

Informed Search

- Heuristic evaluation functions
- Greedy Best-First Search
- A* Search
- Memory Bounded Search
- Iterative Deepening A* Memory Bounded Search
- RBFS

Non-classic Search

- Hill-climbing Search
- Local Beam Search
- Simulated Annealing
- Search in non-deterministic actions
- Search with partial observations

Adversarial Search

- Minimax algorithm
- Alpha-Beta Pruning
- AI for Stochastic Games
- AI for Partially observable Games

Course Project
Board Games Bots

Create unbeatable bots for board games like chess, checkers, etc. using adversarial search based algorithms.

Reinforcement Learning

Heard about AlphaGo? It has learned to play the board game 'Go' on it's own through reinforcement learning. In this course, you will get to learn RL algorithms like temporal difference learning, Monte Carlo, Q-learning, etc. By the end of the course, you will be able to apply RL to various domains like traffic control systems, robotics, advertising, stock market, etc.

Markov decision process

- Agent & Environment Interface
- Goals & Rewards
- Markov Decision processes
- Value functions
- Optimal Value functions

Planning Algorithms

- Bellman equations
- Dynamic Programming
- Policy Evaluation
- Policy Improvement
- Policy Iteration
- Value Iteration

Monte Carlo methods

- Monte Carlo Prediction
- Monte Carlo Estimation of Action Values
- Monte Carlo Control
- Off-policy Prediction
- Off-Policy Monte Carlo Control

Temporal Difference learning

- TD(0) algorithm
- TD(λ) algorithm
- Q-learning algorithm
- SARSA
- Games

Value-Based methods

- Deep Q-Learning network (DQN)
- Double DQN
- Noisy DQN
- Prioritized Replay
- Deep Reinforcement Learning for Robotics

Policy-Gradient methods

- Policy Gradient Theorem
- REINFORCE
- Generalized Advantage Estimation (GAE)
- Natural Policy Gradient (NPG)
- Trust-Region Policy Optimization (TRPO)
- Proximal Policy Optimization (PPO)

Course Project
Undying Flappy Bird

Build an RL model for the flappy bird game such that the bird will learn to fly forever by itself.

Recommendation Systems

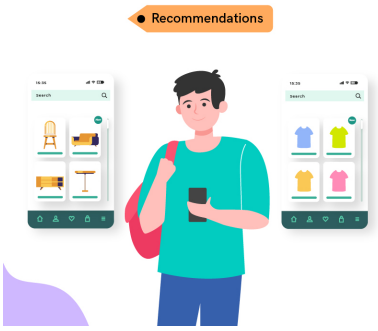

Did you ever wonder how amazon knows the products that you wish to buy? Were you surprised how Netflix gives personalized movie suggestions? Recommendation systems are revolutionizing the way an application communicates to the user. In this course, you will master the art of building powerful recommendation systems using advanced machine learning techniques. By the end of this course, you will be able to build an end-to-end workflow for a recommendation engine.

Introduction to Recommendation systems

- Definitions
- Applications of recommendation systems
- Real-world use cases
- Types of recommendation systems

Content-based recommendation systems

- Components of content-based systems
- Preprocessing and Feature Engineering
- Bayes Classifiers
- Rule-based Classifiers
- Regression-Based Models

Collaborative filtering recommendation systems

- Types of user feedback data
- Decision and Regression Trees
- Matrix Factorization Methods
  - ALS algorithm
  - Principal Component Analysis (PCA)
  - Singular Factor Decomposition (SVD)
- Rule-Based collaborative filtering
- Naive Bayes collaborative filtering

Neighborhood-based collaborative filtering

- Similarity Metrics
- User-Based collaborative filtering
- Item-Based collaborative filtering
- Combination of techniques

Other Recommendation systems

- Context-based
- Location sensitive
- Time-Based
- Hybrid & Ensemble

Deep Learning for recommendation systems

- Context-Aware algorithms
- Restricted-Boltzmann Machines (RBM)
  - Introduction
  - Evaluating the RBM recommender
  - Tuning the RBM recommender
- Modelling using Context-Aware algorithms
- Clickstream recommendations with RNN's

Evaluation of recommendation systems

- Validation
- Evaluation Metrics
- Evaluation Design
- Accuracy and Coverage
- Novelty and Diversity
- Robustness, Stability and Scalability
- Limitations of recommender systems

Course Project

Movie Recommendation System

Build a movie recommendation system and compete in a competition similar to the famous 'Netflix Prize'.

Contact us

support@nxtwave.tech

www.ccbp.in