AI ML Course Details

DEEP MIND Systems

Duration: 60 Hours

Table Of Contents

AI ML Course Details

Lesson 4: Functions

- 4.1 Introduction
- 4.2 Defining Your Own Functions
- 4.3 Function Documentation
- 4.4 Parameters
- 4.5 Keyword and Optional Parameters
- 4.6 Passing Collections to a Function
- 4.7 Variable Number of Arguments
- 4.8 Scope
- 4.9 Lambda
- 4.10 Functions – "First-Class Citizens"
- 4.11 Passing Functions to a Function
- 4.12 Map
- 4.13 Mapping Functions in a Dictionary
- 4.14 Filter
- 4.15 Inner Functions
- 4.16 Closures

Lesson 5: Data Structures

- 5.1 List Comprehensions
- 5.2 Nested List Comprehensions
- 5.3 Processing Lists in Parallel
- 5.4 Dictionaries with Compound Values
- 5.5 Dictionary Comprehensions
- 5.6 Specialized Sorts
- 5.7 Time Functionality
- 5.8 Generators

Lesson 6: Generating Plots

- 6.1 Numpy – Overview
- 6.2 Setup
- 6.3 Data Types
- 6.4 Basic Operators
- 6.5 Indexing
- 6.6 Broadcasting
- 6.7 Matrix Operators
- 6.8 Matplotlib – Overview
- 6.9 Setup
- 6.10 Basic Plots
- 6.11 Customizing Plots
- 6.12 Subplots
- 6.13 3D Plots

Lesson 7: Debugging

- 7.1 Types of Errors
- 7.2 Syntax & Logical Errors
- 7.3 Syntax Error Debugging
- 7.4 Logical Error Debugging

Lesson 8: Classes

- 8.1 Principles of Object Orientation
- 8.2 Classes in Python
- 8.3 Creating Classes
- 8.4 Instance Methods
- 8.5 Special Methods
- 8.6 Class Variables
- 8.7 Inheritance
- 8.8 Polymorphism
- 8.9 Type Identification
- 8.10 Custom Exceptions
- 8.11 File Organization

Lesson 9: Regular Expressions

- 9.1 Introduction
- 9.2 Simple Character Matches
- 9.3 Special Character Matches
- 9.4 Match Objects
- 9.5 Character Classes
- 9.6 The Dot Operator
- 9.7 Quantifiers
- 9.8 Greedy Matches
- 9.9 Grouping Matches at Beginning or End
- 9.10 Compiling Regular Expressions
- 9.11 Flags
- 9.12 Substituting
- 9.13 Splitting a String

Graded Assessment

Project/Case study will be given by the Instructor at the end of the final session of Module 1

Module 2: Introduction to Artificial Intelligence

Lesson 1: Introduction to Artificial Intelligence

- 1.1 Today's AI
- 1.2 Strong AI and Weak AI
- 1.3 Artificial Intelligence Definitions
- 1.4 Machine Learning → Deep Learning → AI
- 1.5 Cognitive Science and AI
- 1.6 Cognition and the process of cognition
- 1.7 Disciplines in cognitive science
- 1.8 Multidisciplinary subject
- 1.9 Linguistics
- 1.10 Artificial Intelligence as cognitive science
- 1.11 Methods in cognitive science
- 1.12 Industry 4.0 and Industry 5.0

AI ML Course Details

DEEP MIND Systems

Lesson 2: Logical Approach to AI and Knowledge-Based Systems

- 2.1 Introduction
- 2.2 Motivation for Machine Learning
- 2.3 Applications
- 2.4 Machine Learning
- 2.5 Learning associations
- 2.6 The origin of Machine Learning
- 2.7 Uses and abuses of machine learning
- 2.8 Success cases
- 2.9 How do machines learn?
- 2.10 Abstraction and knowledge representation
- 2.11 Generalization
- 2.12 Factors to be considered
- 2.13 Assessing the success of learning
- 2.14 Supervised learning
- 2.15 Regression
- 2.16 Regression examples
- 2.17 Regression models
- 2.18 Steps in regression analysis
- 2.19 Linear regression
- 2.20 Simple linear regression
- 2.21 Least squares estimation
- 2.22 Least squares regression - Line of best fit

Lesson 3: Probabilistic Approach to AI

- 3.1 Probability Overview
- 3.2 Basic concepts
- 3.3 Probability of an event
- 3.4 Example of sample space
- 3.5 Counting rules
- 3.6 Event relations
- 3.7 Conditional probabilities
- 3.8 Defining independence
- 3.9 The law of total probability
- 3.10 Bayes' rule
- 3.11 Random variables
- 3.12 Discrete random variables
- 3.13 Probability distributions
- 3.14 Probability mass function
- 3.15 Probability density function
- 3.16 Expectations of random variables
- 3.17 Medians of random variables
- 3.18 Variance of a random variable
- 3.19 Quantiles of random variables
- 3.20 Jointly distributed random variables
- 3.21 Marginal probability distributions
- 3.22 Independence and covariance

3.23 Bayesian networks
3.24 Merits of Bayesian networks
3.25 Construction of a Bayesian network
3.26 Representation in Bayesian networks
3.27 Benefits of Bayesian networks
3.28 Why learn Bayesian networks?
3.29 Constructing Bayesian networks
3.30 Example from medical diagnostics
3.31 Software for Bayesian networks
3.32 Gaussian Bayesian networks

Lesson 4: Evolutionary Intelligence

4.1 Coefficient of determination (R-squared)
4.2 Example
4.3 Testing for significance
4.4 Testing hypothesis in simple linear regression
4.5 Illustration
4.6 Checking model adequacy
4.7 Over-fitting
4.8 Detecting over-fit models: Cross-validation
4.9 Cross-validation: The ideal procedure
4.10 Ordinary least squares estimation for multiple linear regression
4.11 Multiple linear regression model building
4.12 Partial correlation and regression model building
4.13 Interpretation of multiple linear regression coefficients - Partial regression coefficients
4.14 Standardized regression coefficients
4.15 Missing data

Lesson 5: Introduction to Machine Learning

- Data Preprocessing
- Simple Linear Regression
- Multiple Linear Regression
- Polynomial Regression
- Support Vector Regression
- Decision Tree Regression
- Random Forest Regression
- Bagging Techniques - Random Forest
- Ensemble Learning
- Boosting Techniques - An Introduction
- Boosting Techniques - AdaBoost
- Gradient Boosting Method
- Difference between AdaBoost and Gradient Boost

- Introduction to Classification Algorithms:
  - Logistic Regression
  - KNN (K-Nearest Neighbors)
  - SVM (Support Vector Machine)
  - Kernel SVM
  - Naive Bayes Algorithm
  - Decision Tree
  - Random Forest
- Introduction to Clustering Algorithms:
  - K-Means Clustering
  - Hierarchical Clustering
  - Density-based clustering methods
  - DBSCAN: Density-Based Spatial Clustering of Applications with Noise
  - When DBSCAN Does NOT Work Well
  - Python program using DBSCAN with Credit Card Dataset
  - External criteria for clustering quality
  - Different aspects of cluster validation
  - Measures of cluster validity
  - Measuring cluster validity via correlation
  - Using similarity matrix for cluster validation
  - Internal measures: SSE
  - Framework for cluster validity
  - Internal measures: Cohesion and Separation
  - Internal measures: Silhouette Coefficient

Lesson 6: Information Retrieval

6.1 Information retrieval: introduction
6.2 Information retrieval process
6.3 Information retrieval architecture
6.4 How do we represent documents?
6.5 Information retrieval models
6.6 Similarity metric
6.7 Term weighting
6.8 Retrieval in the vector space model
6.9 Constructing inverted index (word counting)
6.10 Stopwords removal
6.11 Stemming
6.12 Text document clustering:

- Agglomerative vs. divisive
- Impact of cluster distance measure
- Buckshot clustering
- Issues related to cosine similarity
- Validity of document clusters
- Text datasets
- 6.13 Experimental evaluation

AI ML Course Details

DEEP MIND Systems

Module 3: Deep Learning

Lesson 1: The Fundamentals of Deep Learning, Learning Process, and Neural Networks Model

1.1 How is deep learning different from other machine learning?
1.2 AI vs ML vs DL
1.3 Deep learning capabilities
1.4 What is special about deep learning?
1.5 Relevance of deep learning
1.6 Supervised learning
1.7 Unsupervised learning
1.8 Memory-based learning
1.9 Memory-based learning techniques
1.10 Hebbian learning
1.11 Hebbian learning modifications: Mathematical models
1.12 Competitive learning
1.13 Error-correction learning
1.14 Boltzmann learning
1.15 Memory
1.16 Adaptations
1.17 Statistical nature of the learning process
1.18 Statistical learning theory
1.19 Probably approximately correct model of learning
1.20 Adaptive filtering problems
1.21 Unconstrained optimization techniques
1.22 Linear least-squares filters
1.23 Least-mean-square algorithms
1.24 Learning curves
1.25 Learning rate annealing techniques
1.26 Perceptron
1.27 Perceptron convergence theorem
1.28 MLP concepts
1.29 Backpropagation algorithm
1.30 XOR problem
1.31 Heuristics for making backpropagation algorithm perform better
1.32 Output representation and decision rules
1.33 Feature detection
1.34 Backpropagation and Differentiation
1.35 Generalization
1.36 Approximations of functions
1.37 Cross-validations
1.38 Network pruning techniques
1.39 Virtues and limitations of backpropagation learning
1.40 Accelerated convergence of backpropagation learning
1.41 Supervised learning viewed as an optimization problem
1.42 Cover's theorem on the separability of patterns

AI ML Course Details

1.43 Interpolation problem
1.44 Regularization theory and regularization networks
1.45 Generalized radial-basis function networks
1.46 Estimation of the regularization parameter
1.47 Approximation properties of RBF networks
1.48 Comparison of RBF networks and multilayer perceptron
1.49 Kernel regression and its relation to RBF networks
1.50 Learning Strategies in RBF Networks
1.51 Simulated annealing
1.52 Boltzmann machines
1.53 Deterministic Boltzmann machine

Lesson 2: The Math Behind Neural Networks

2.1 How does a neural network look like? The matrix magic
2.2 Visualizing deep learning
2.3 The Elephant in the Room
2.4 Programmatic expression of deep learning's math constructs
2.5 Accessing and manipulation of tensors
2.6 Operations with tensors
2.7 Array broadcasting
2.8 Scalar product/Inner product of tensors
2.9 Morphing shapes of tensors
2.10 Gradient calculation
2.11 Calculation of accuracy values
2.12 Training a binary classifier

Lesson 3: Diving to the Depths of Deep Learning

3.1 Deep learning depths
3.2 Model: The molecules of DL
3.3 Loss Functions in neural networks
3.4 Optimizers in neural networks
3.5 Activation functions
3.6 Finding the perfect fit
3.7 Combating Overfitting Problem in NN
3.8 Running deep learning algorithms: The frameworks
3.9 Real examples and actual schematics of building neural nets
3.10 Data preparation and label

Lesson 4: Advanced Neural Network Architectures and Applications

4.1 Introduction to convolutional neural networks (CNNs)
4.2 Components of CNNs: Convolution layers, pooling layers, fully connected layers
4.3 Applications of CNNs in image recognition and classification
4.4 Recurrent neural networks (RNNs): Sequence prediction
4.5 Long short-term memory (LSTM) and gated recurrent unit (GRU)
4.6 Applications of RNNs in language modeling and speech recognition
4.7 Generative adversarial networks (GANs): Framework and use cases
4.8 Autoencoders: Dimensionality reduction and feature extraction

DEEP MIND Systems

4.9 Transformer architecture: Revolutionizing NLP
4.10 BERT and GPT models
4.11 Practical applications of transformer-based models
4.12 Multi-task learning in deep learning
4.13 Transfer learning and pre-trained models
4.14 Building custom architectures for specific use cases
4.15 Use of deep learning in industry verticals: Healthcare, finance, and beyond
4.16 Challenges and limitations of advanced neural networks
4.17 Case studies: Successful deployment of deep learning systems
4.18 Emerging trends and future directions in neural networks

Lesson 5: Practical Implementation and Projects

5.1 Setting up a deep learning environment
5.2 Tools and libraries for deep learning
5.3 Designing and training your first neural network
5.4 Hyperparameter tuning and model optimization
5.5 Evaluating model performance and accuracy
5.6 Implementing CNN for image classification
5.7 Implementing RNN for time-series prediction
5.8 Using GANs for image generation
5.9 Applying transfer learning for quick model deployment
5.10 Deploying a trained model to production
5.11 Real-world deep learning projects
5.12 Building a chatbot using transformer-based architecture
5.13 Developing a recommendation system
5.14 Solving industry problems using deep learning
5.15 Hands-on capstone project: From data preparation to deployment