

## ARTIFICIAL INTELLIGENCE & MACHINE LEARNING

**OVERALL COURSE OBJECTIVES:** Equip participants with a comprehensive understanding of Machine Learning, Artificial Intelligence, and the management of relevant projects, enabling them to cognitively apply various models for problem-solving, create AI-powered chatbots, and effectively deal with both supervised and unsupervised learning scenarios, hence fostering a holistic command over the entire ML/AI landscape.

**LEARNING OUTCOMES:** On successful completion of the course the students shall be able to:

1. Understand and implement various machine learning models using Pytorch and other open-source libraries.
2. Learn to manage and organize machine learning projects from inception to production using best practices.
3. Understand the concepts and applications of Artificial Intelligence (AI) and successfully complete a mini-project illustrating AI in action.
4. Create interactive chatbots using IBM Watson's Natural Language Processing capabilities without the need for programming.
5. Apply supervised machine learning techniques such as logistic regression, decision trees, and ensemble methods for classification and regression tasks.
6. Understand and implement unsupervised machine learning techniques using various clustering and dimension reduction algorithms.

|                                 |                                                  |
|---------------------------------|--------------------------------------------------|
| ML Foundations                  | Introduction to Machine Learning                 |
|                                 | Managing Machine Learning Projects               |
| AI & ML Algorithms – I          | Introduction to Artificial Intelligence (AI)     |
|                                 | Building AI Powered Chatbots Without Programming |
|                                 | Supervised Machine Learning: Classification      |
| AI & ML Algorithms – I          | Supervised Machine Learning: Regression          |
|                                 | Unsupervised Machine Learning                    |
| Applied ML                      | Deep Learning and Reinforcement Learning         |
|                                 | Machine Learning Capstone                        |
| ML Applications & Data Analysis | Exploratory Data Analysis for Machine Learning   |
|                                 | Applied AI with DeepLearning                     |
|                                 | Advanced Machine Learning and Signal Processing  |

### COURSE CONTENT:

**Module 1:** [Introduction to Machine Learning](#) [26 Hours]

This course will provide you with a foundational understanding of machine learning models (logistic regression, multilayer perceptrons, convolutional neural networks, natural language processing, etc.) as well as demonstrate how these models can solve complex problems in a variety of industries, from medical diagnostics to image recognition to text prediction. In addition, we have designed practice exercises that will give you hands-on experience implementing these data science models on data sets. These practice exercises will teach you how to implement machine learning algorithms with PyTorch, open source libraries used by leading tech companies in the machine learning field (e.g., Google, NVIDIA, CocaCola, eBay, Snapchat, Uber and many more).

### **Sub-Topics**

Simple Introduction to Machine Learning  
Basics of Model Learning  
Image Analysis with Convolutional Neural Networks  
Recurrent Neural Networks for Natural Language Processing  
The Transformer Network for Natural Language Processing  
Introduction to Reinforcement Learning

### **Formative Assessments:**

4 graded quizzes and 11 ungraded lab assignments.

### **Module 2: [Managing Machine Learning Projects](#) [18 Hours]**

The course in the AI Product Management Specialization that centers on managing practical aspects of machine learning (ML) projects. This course provides a thorough understanding of the integral steps of an ML project, beginning from the identification of valuable ML opportunities, right through to data collection, model building, deployment, as well as monitoring and maintaining production systems. It offers insight into the fundamental processes of data science and how to effectively organize ML projects. On successful completion of this course, participants are expected to acquire the ability to identify potential applications of ML for problem-solving and learn about important technological decisions in designing ML systems, along with leading ML projects from inception to production using best practices.

### **Sub-Topics**

Identifying Opportunities for Machine Learning  
Organizing ML Projects  
Data Considerations  
Simple Introduction to Machine Learning  
ML System Design & Technology Selection  
Model Lifecycle Management

### **Formative Assessments:**

5 graded quizzes and 1 peer review assignment.

### **Module 3: [Introduction to Artificial Intelligence](#) (AI) [8 Hours]**

In this course you will learn what Artificial Intelligence (AI) is, explore use cases and applications of AI, understand AI concepts and terms like machine learning, deep learning and neural networks. You will be exposed to various issues and concerns surrounding AI such as ethics and bias, & jobs, and get

advice from experts about learning and starting a career in AI. You will also demonstrate AI in action with a mini project.

This course does not require any programming or computer science expertise and is designed to introduce the basics of AI to anyone whether you have a technical background or not.

#### **Sub-Topics**

What is AI? Applications and Examples of AI  
AI Concepts, Terminology, and Application Areas  
AI: Issues, Concerns and Ethical Considerations  
The Future with AI, and AI in Action

#### **Formative Assessments:**

4 graded quizzes and 1 peer review assignment.

### **Module 4: [Building AI Powered Chatbots Without Programming](#) [12 Hours]**

This course will teach you how to create useful chatbots without the need to write any code. Leveraging IBM Watson's Natural Language Processing capabilities, you'll learn how to plan, implement, test, and deploy chatbots that delight your users, rather than frustrate them. True to our promise of not requiring any code, you'll learn how to visually create chatbots with Watson Assistant (formerly Watson Conversation) and how to deploy them on your own website through a handy WordPress plugin. New jobs requiring this specific skill are being added every day, consultants demand premium rates, and the interest in chatbots is quickly exploding. Gartner predicts that by 2020, 85% of customer interactions with the enterprise will be through automated means (that's chatbots and related technologies).

#### **Sub-Topics**

Introduction to Chatbots  
Working with Intents  
Entities  
Dialog  
Deployment  
Context Variables & Slots  
Digressions  
Final Exam

#### **Formative Assessments:**

8 graded quizzes.

### **Module 5: [Supervised Machine Learning: Classification](#) [25 Hours]**

This course provides an introduction to one of the major types of supervised Machine Learning models: Classification. It educates you on how to train models to classify categorical outcomes and use error metrics to compare different models. The hands-on section is geared towards utilizing best practices for classification, including train/test splits and dealing with unbalanced classes in data sets. Upon completion, you should be able to explore and apply various classification ensembles and models like logistic regression, decision trees, and other ensemble methods, and also use various error metrics to identify the best model for your data. Techniques like oversampling and undersampling to handle unbalanced classes will also be covered. This course is designed for aspiring

data scientists seeking hands-on experience with Supervised Machine Learning Classification techniques in a business context. To get the most out of this course, you should have a basic understanding of the Python programming environment, Data Cleaning, Exploratory Data Analysis, Calculus, Linear Algebra, Probability, and Statistics.

**Sub-Topics**

- Logistic Regression
- K Nearest Neighbors
- Support Vector Machines
- Decision Trees
- Ensemble Models
- Modeling Unbalanced Classes

**Formative Assessments:**

6 graded quizzes and 1 peer-review assignment.

**Module 6: [Supervised Machine Learning: Regression](#) [21 Hours]**

This course presents comprehensive knowledge of supervised Machine Learning with a focus on one of its primary modeling families: Regression. You will learn to train regression models to predict continuous outcomes, using error metrics to compare different models. The modules include best practices such as train and test splits along with regularization techniques. The skills that you should expect to gain include distinguishing applications of classification and regression, applying linear regression models, identifying the ideal linear regression model through error metrics, understanding the necessity of regularization to prevent overfitting, and utilizing regularization regressions like Ridge, LASSO, and Elastic net. It is targeted at aspiring data scientists wishing to gain practical experience with Supervised Machine Learning Regression techniques in business. Familiarity with a Python development environment and fundamental understanding of data cleaning, exploratory data analysis, calculus, linear algebra, probability and statistics will ensure maximum benefit from this course.

**Sub-Topics**

- Introduction to Supervised Machine Learning and Linear Regression
- Data Splits and Polynomial Regression
- Cross Validation
- Bias Variance Trade off and Regularization Techniques: Ridge, LASSO, and Elastic Net
- Regularization Details
- Final Project

**Formative Assessments:**

5 graded quizzes and 1 peer-review assignment.

**Module 7: [Unsupervised Machine Learning](#) [23 Hours]**

This course focuses on Unsupervised Learning, one of the major types of Machine Learning. It teaches you to derive insights from data sets without a target or labeled variable using various clustering and dimension reduction algorithms. Upon completion, you'd be able to understand the types of problems suitable for Unsupervised Learning, comprehend the concept of dimensionality curse and its impact on clustering, and apply common clustering and dimensionality reduction algorithms. You'll also be able to perform clustering where necessary, compare per-cluster models, and understand metrics important for characterizing clusters. This course is suitable for aspiring data

scientists wishing to gain practical experience in Unsupervised Machine Learning in a business context. Familiarity with the Python programming environment, Data Cleaning, Exploratory Data Analysis, Calculus, Linear Algebra, Probability, and Statistics will be advantageous.

#### **Sub-Topics**

Introduction to Unsupervised Learning and K Means  
Distance Metrics & Computational Hurdles  
Selecting a Clustering Algorithm  
Dimensionality Reduction  
Nonlinear and Distance-Based Dimensionality Reduction  
Matrix Factorization  
Final Project

#### **Formative Assessments:**

6 graded quizzes and 1 peer-review assignment.

### **Module 8: [Deep Learning and Reinforcement Learning](#) [32 Hours]**

This course will provide you a foundational understanding of machine learning models (logistic regression, multilayer perceptrons, convolutional neural networks, natural language processing, etc.) as well as demonstrate how these models can solve complex problems in a variety of industries, from medical diagnostics to image recognition to text prediction. In addition, we have designed practice exercises that will give you hands-on experience implementing these data science models on data sets. These practice exercises will teach you how to implement machine learning algorithms with PyTorch, open source libraries used by leading tech companies in the machine learning field (e.g., Google, NVIDIA, CocaCola, eBay, Snapchat, Uber and many more).

#### **Sub-Topics**

Introduction to Neural Networks  
Back Propagation Training and Keras  
Neural Network Optimizers  
Convolutional Neural Networks  
Transfer Learning  
Recurrent Neural Networks and Long-Short Term Memory Networks  
Autoencoders  
Generative Models and Applications of Deep Learning  
Reinforcement Learning

#### **Formative Assessments:**

9 graded quizzes and 1 peer-review assignment.

### **Module 9: [Machine Learning Capstone](#) [19 Hours]**

This course enables learners to build a dynamic course recommender system utilizing Python-based machine learning libraries like Pandas, scikit-learn, and Tensorflow/Keras. Participants will be involved in the analysis of course-related datasets, constructing a similarity matrix using cosine similarity, and the development of recommendation systems employing KNN, PCA, and non-negative matrix collaborative filtering. The course lets you predict course ratings by training a neural network and crafting regression and classification models. A crucial part of the course involves learners creating a Streamlit app showcasing their work, following which they share and evaluate each other's

projects.

### **Sub-Topics**

Capstone Overview  
Exploratory Data Analysis and Feature Engineering  
Unsupervised-Learning Based Recommender System  
Supervised-Learning Based Recommender Systems  
Share and Present Your Recommender Systems  
Final Submission

### **Formative Assessments:**

3 graded quizzes and 1 peer-review assignment.

### **Module 10: [Exploratory Data Analysis for Machine Learning](#) [14 Hours]**

This course begins with an introduction to Machine Learning and its constituent content, emphasizing the importance of high-quality data. The course facilitates learning about the retrieval of data from multiple sources, data cleaning, feature engineering, and preliminary analysis for hypothesis testing. Participants develop the skills to retrieve data from SQL, NoSQL databases, APIs, Cloud, handle categorical and ordinal features, missing values, and manage outliers. It also underscores the relevance of feature scaling and practices various scaling techniques. The course is ideal for aspiring data scientists interested in gaining practical experience in Machine Learning and Artificial Intelligence in a business context. To maximize the benefits, participants should be familiar with Python programming and possess a foundational understanding of Calculus, Linear Algebra, Probability, and Statistics.

### **Sub-Topics**

A Brief History of Modern AI and its Applications  
Retrieving and Cleaning Data  
Exploratory Data Analysis and Feature Engineering  
Inferential Statistics and Hypothesis Testing

### **Formative Assessments:**

4 graded quizzes and 1 peer-review assignment.

### **Module 11: [Applied Machine Learning in Python](#) [31 Hours]**

This course will introduce the learner to applied machine learning, focusing more on the techniques and methods than on the statistics behind these methods. The course will start with a discussion of how machine learning is different than descriptive statistics, and introduce the scikit learn toolkit through a tutorial. The issue of dimensionality of data will be discussed, and the task of clustering data, as well as evaluating those clusters, will be tackled. Supervised approaches for creating predictive models will be described, and learners will be able to apply the scikit learn predictive modelling methods while understanding process issues related to data generalizability (e.g. cross validation, overfitting). The course will end with a look at more advanced techniques, such as building ensembles, and practical limitations of predictive models. By the end of this course, students will be able to identify the difference between a supervised (classification) and unsupervised (clustering) technique, identify which technique they need to apply for a particular dataset and need, engineer features to meet that need, and write python code to carry out an analysis.

**Sub-Topic**

Fundamentals of Machine Learning - Intro to SciKit Learn  
Supervised Machine Learning - Part 1 & 2

**Formative Assessments:**

4 graded quizzes and 10 programming assignments/labs.

**Module 12: [Applied Social Network Analysis in Python](#) [26 Hours]**

This course will introduce the learner to network analysis through tutorials using the NetworkX library. The course begins with an understanding of what network analysis is and motivations for why we might model phenomena as networks. The second week introduces the concept of connectivity and network robustness. The third week will explore ways of measuring the importance or centrality of a node in a network. The final week will explore the evolution of networks over time and cover models of network generation and the link prediction problem.

**Sub-Topics**

Why study Networks and Basics on NetworkX  
Network Connectivity  
Influence Measures and Network Centralization  
Network Evolution

**Formative Assessments:**

4 graded quizzes and 8 programming assignments/labs.

**ASSESSMENT:**

**For summative assessments, Coursera will provide question banks for which exams can be conducted on the Coursera platform or the faculty will create their own assessments.**

*Note: If a Course or Specialization becomes unavailable prior to the end of the Term, Coursera may replace such Course or Specialization with a reasonable alternative Course or Specialization.*