

DEEP LEARNING AND NATURAL LANGUAGE PROCESSING

OVERALL COURSE OBJECTIVES:

The participant will gain a thorough understanding of deep learning and natural language processing principles and applications. They will be skilled in designing, implementing, and optimizing diverse types of neural network models, resulting in the ability to build unique solutions for a variety of real-world problems. The learner will gain competence in using cutting-edge technologies and methodologies to effectively handle complex tasks such as sentiment analysis, language translation, text summarizing, and question answering.

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

1. Construct and train neural networks, applying deep learning knowledge to generate actionable insights from complex datasets.
2. Implement various types of neural network models including recurrent neural networks, convolutional neural networks, and sequence models for tasks such as sentiment analysis, language translation, and named entity recognition.
3. Utilize advanced techniques like vector space models, Word2Vec, and GLoVe word embeddings to handle language processing tasks and generate creative solutions such as auto-correct and automatic text generation.
4. Apply Transformer and attention models to summarize text, answer questions, and build chatbots, demonstrating an understanding of how to leverage these models for complex NLP problems.
5. Utilize technology trends effectively, such as TensorFlow and other machine learning frameworks, to enhance the efficiency and performance of deep learning applications.
6. Evaluate and enhance model performance and reliability through techniques such as hyperparameter tuning, model optimization, and regularization.

Deep Learning - I	Neural Networks and Deep Learning
	Improving Deep Neural Networks: Hyperparameter Tuning, Regularization and Optimization
	Structuring Machine Learning Projects
Deep Learning - II	Convolutional Neural Networks
	Sequence Models
Natural Language Processing - I	Natural Language Processing with Classification and Vector Spaces
	Natural Language Processing with Probabilistic Models
Natural Language Processing - II	Natural Language Processing with Sequence Models
Natural Language Processing - III	Natural Language Processing with Attention Models

COURSE CONTENT:

Module 1: [Neural Networks and Deep Learning](#) [25 Hours]

In this course, you will study the foundational concept of neural networks and deep learning. By the end, you will be familiar with the significant technological trends driving the rise of deep learning; build, train, and apply fully connected deep neural networks; implement efficient (vectorized) neural networks; identify key parameters in a neural network's architecture; and apply deep learning to your own applications.

Sub-Topics:

Deep Neural Networks
Introduction to Deep Learning
Neural Networks Basics
Shallow Neural Networks

Formative Assessments:

4 Programming Assignments & 4 Graded Quizzes

Module 2: [Improving Deep Neural Networks: Hyperparameter Tuning, Regularization and Optimization](#) [24 Hours]

In this course, you will open the deep learning black box to understand the processes that drive performance and generate good results systematically. By the end, you will learn the best practices to train and develop test sets and analyze bias/variance for building deep learning applications; be able to use standard neural network techniques such as initialization, L2 and dropout regularization, hyperparameter tuning, batch normalization, and gradient checking; implement and apply a variety of optimization algorithms, such as mini-batch gradient descent, Momentum, RMSprop and Adam, and check for their convergence; and implement a neural network in TensorFlow.

Sub-Topics:

Hyperparameter Tuning, Batch Normalization and Programming Frameworks
Optimization Algorithms
Practical Aspects of Deep Learning

Formative Assessments:

5 Programming Assignments & 3 Graded Quizzes

Module 3: [Structuring Machine Learning Projects](#) [7 Hours]

In this course, you will learn how to build a successful machine learning project and get to practice decision-making as a machine learning project leader. By the end, you will be able to diagnose errors in a machine learning system; prioritize strategies for reducing errors; understand complex ML settings, such as mismatched training/test sets, and comparing to and/or surpassing human-level performance; and apply end-to-end learning, transfer learning, and multi-task learning.

Sub-Topics

ML Strategy

Multi-task Learning

Training and Testing on Different Distributions

Formative Assessments:

2 Graded Quizzes

Module 4: [Convolutional Neural Networks](#) [36 Hours]

In this course, you will understand how computer vision has evolved and become familiar with its exciting applications such as autonomous driving, face recognition, reading radiology images, and more. By the end, you will be able to build a convolutional neural network, including recent variations such as residual networks; apply convolutional networks to visual detection and recognition tasks; and use neural style transfer to generate art and apply these algorithms to a variety of image, video, and other 2D or 3D data.

Sub-Topics:

Deep Convolutional Models: Case Studies

Foundations of Convolutional Neural Networks

Object Detection

Special Applications: Face recognition & Neural Style Transfer

Formative Assessments:

8 Programming Assignments & 4 Graded Quizzes

Module 5: [Sequence Models](#) [37 Hours]

In this course, you will become familiar with sequence models and their exciting applications such as speech recognition, music synthesis, chatbots, machine translation, natural language processing (NLP), and more. By the end, you will be able to build and train Recurrent Neural Networks (RNNs) and commonly-used variants such as GRUs and LSTMs; apply RNNs to Character-level Language Modeling; gain experience with natural language processing and Word Embeddings; and use HuggingFace tokenizers and transformer models to solve different NLP tasks such as NER and Question Answering.

Sub-Topics:

Natural Language Processing & Word Embeddings

Recurrent Neural Networks

Sequence Models & Attention Mechanism

Transformer Network

Formative Assessments:

8 Programming Assignments & 4 Graded Quizzes

Module 6: [Natural Language Processing with Classification and Vector Spaces](#) [34 Hours]

In this Course, you will:

- a) Perform sentiment analysis of tweets using logistic regression and then naïve Bayes,
- b) Use vector space models to discover relationships between words and use PCA to reduce the dimensionality of the vector space and visualize those relationships, and
- c) Write a simple English to French translation algorithm using pre-computed word embeddings and locality-sensitive hashing to relate words via approximate k-nearest neighbor search.

Sub-Topics

Machine Translation and Document Search
Sentiment Analysis with Logistic Regression
Sentiment Analysis with Naïve Bayes
Vector Space Models

Formative Assessments:

4 Programming Assignments

Module 7: [Natural Language Processing with Probabilistic Models](#) [31 Hours]

In this course, you will:

- a) Create a simple auto-correct algorithm using minimum edit distance and dynamic programming,
- b) Apply the Viterbi Algorithm for part-of-speech (POS) tagging, which is vital for computational linguistics,
- c) Write a better auto-complete algorithm using an N-gram language model, and
- d) Write your own Word2Vec model that uses a neural network to compute word embeddings using a continuous bag-of-words model.

Sub-Topics

Autocomplete and Language Models
Autocorrect
Part of Speech Tagging and Hidden Markov Models
Word embeddings with neural networks

Formative Assessments:

4 Programming Assignments

Module 8: [Natural Language Processing with Sequence Models](#) [45 Hours]

In this Course, you will:

- a) Train a neural network with GLoVe word embeddings to perform sentiment analysis of tweets,

- b) Generate synthetic Shakespeare text using a Gated Recurrent Unit (GRU) language model,
- c) Train a recurrent neural network to perform named entity recognition (NER) using LSTMs with linear layers, and
- d) Use so-called 'Siamese' LSTM models to compare questions in a corpus and identify those that are worded differently but have the same meaning.

Sub-Topics

LSTMs and Named Entity Recognition
Recurrent Neural Networks for Language Modeling
Siamese Networks

Formative Assessments:

3 Programming Assignments

Module 9: [Natural Language Processing with Attention Models](#) [62 Hours]

In this Course, you will:

- a) Translate complete English sentences into German using an encoder-decoder attention model,
- b) Build a Transformer model to summarize text,
- c) Use T5 and BERT models to perform question-answering, and
- d) Build a chatbot using a Reformer model.

Sub-Topics

Neural Machine Translation
Question Answering
Text Summarization

Formative Assessments:

3 Programming Assignments

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.