

Course Code: W7108GW

Course Title: What is Natural Language Processing (NLP)?

Description:

This course introduces key concepts and methods in Natural Language Processing (NLP), the subfield of data science and artificial intelligence that deals with computer interaction with human language. The course covers 1) The definition of NLP, and its relationship to data science, artificial intelligence, and other subfields therein, 2) Historical thought and development in NLP as well as recent breakthroughs in NLP enabled by artificial intelligence, and 3) Select methods and use cases for NLP, focused mainly on business contexts. The course will also briefly introduce learners to award-winning IBM Watson NLP tools, which make the discussed technologies accessible to non-technical users. Upon completion of the course, learners will feel comfortable describing what NLP is, how it has evolved from historical precedents to contemporary developments, and several specific methods and use cases for NLP technologies in organizational settings. Learners will leave this course well-prepared for more applied and technical courses on NLP.

Objectives:

• Defining NLP • NLP History • NLP Applications • Working With Text • From Models to Systems

Prerequisites:

No prerequisite knowledge is assumed, however a basic understanding of artificial intelligence and machine learning (i.e. what they are, what they do) will be beneficial.

Duration:

4.8 Hrs

Topics:

Defining NLP • Defining natural languages, and contrasting them from constructed and formal languages • Examining the ways that computers handle natural language as numbers (data) • Recognizing simple applications for computer processing of natural language data NLP History • Identifying early thinkers who speculated about computation, computers, and human language • Analyzing successes and failures in early NLP experiments in machine translation and chatbots • Examining key enabling technologies that supported the advancement of NLP over the past few decades NLP Applications • Relating early NLP experiments to modern applications for individuals and organizations • Describing specific sources of value brought by NLP applications to businesses, such as in customer engagement and sentiment analysis • Considering the future of common NLP applications Working With Text • Identifying computer processes (algorithms) useful for handling and transforming text • Recognizing common challenges in working with textual data, such as issues related to data quality • Analyzing components of NLP frameworks, and linking them to common NLP methods such as named-entity recognition (NER) From Models to Systems • Differentiating text-level NLP methods from NLP systems • Relating system-level NLP to NLP applications • Recognizing the intuitions behind two complex NLP systems: Information Extraction and Conversational AI

Audience:

The course is intended primarily for learners with a focus on business and/or technology, and specifically those seeking to better understand what NLP is, where it came from, and how it can benefit their organizations. Although the course does not cover the technical aspects of building NLP systems, Data Scientists, Machine Learning Engineers, and other technically-inclined learners may still find concepts and best practices introduced in the subject matter to be useful.