

Deep Learning for Natural Language Processing

Deep Learning is a subfield of machine learning based on the use of artificial neural networks that, through a hierarchy of layers with non-linear processing units, learn high-level abstractions for data. In recent years, these representations have enabled outstanding performance in various fields of artificial intelligence (AI) such as: computer vision; reinforcement learning; and, as addressed in this course, natural language processing or NLP.

NLP is a crucial field of AI that studies the interactions between computers and human language. The goal is making computers "process" or "understand" natural language (as opposed to programming languages), allowing them to perform useful tasks. Examples of these tasks include sentiment analysis, automatic translation, automatic text summarization, or the search for answers to questions posed by humans in natural language. This course will explore the main Deep Learning technologies for NLP and how they can be used to solve these problems.