

CS 875: Natural Language Processing

NLP is one of the main tasks where AI is the key underlying technology. Processing text in main cases is the main problem to be solved specially when data is available in that form, and also when robots have to communicate with humans they may resort to it. Also, its coupling with the speech processing unit would give extra leverage to the robots to do their tasks in a more effective manner. All the universities in the comparison cohort offer this course.

Recommended Books:

1. Natural Language Processing Recipes: Unlocking Text Data with Machine Learning and Deep Learning using Python, Akshay Kulkarni, Adarsha Shivananda, Apress, 2019.
2. Deep Learning in Natural Language Processing, Li Deng, Yang Liu editors, Springer 2018.
3. Python Natural Language Processing, Jalaj Thanaki, Packt Publishing, 2017.

Credit Hours: 3 (3, 0)

Course Objectives:

- To understand vector space models and use of context to derive semantics.
- To understand modern deep learning based algorithms that are used for NLP.
- To be able to process large unstructured and unlabeled text to extract useful information.

Topics / Contents	Allocated Periods
The course covers topics such as word vectors, word senses, neural networks and matrix calculus, backpropagation and computational graphs, linguistic structure dependency and parsing, recurrent neural networks and language models, vanishing gradients and fancy RNNs, machine translation, question answering, ConvNets for NLP, subword models, generative models, safety, bias, and fairness. Other topics include words, morphology, and lexicon, noisy channel models	45

and edit distance, classification, parts of speech tags, hidden markov models, Chomsky hierarchy and natural language, context-free recognition, lexical semantics, word sense disambiguation, and interpreting social media.	
---	--