

# Natural Language Processing

## Index

### Chapter 1: Introduction to NLP

- 1.1 Overview of NLP and Applications
- 1.2 History and Evolution of NLP
- 1.3 NLP Pipeline: Text Preprocessing to Modeling
- 1.4 Challenges in NLP (Ambiguity, Syntax, Semantics)
- 1.5 Tools and Libraries (NLTK, spaCy, Hugging Face)

### Chapter 2: Text Preprocessing

- 2.1 Tokenization (Word, Sentence, Subword)
- 2.2 Stopword Removal and Stemming
- 2.3 Lemmatization
- 2.4 Part-of-Speech (POS) Tagging
- 2.5 Named Entity Recognition (NER)

### Chapter 3: Text Representation

- 3.1 One-Hot Encoding
- 3.2 Bag-of-Words (BoW) Model
- 3.3 TF-IDF (Term Frequency-Inverse Document Frequency)
- 3.4 Word Embeddings (Word2Vec, GloVe, FastText)
- 3.5 Contextual Embeddings (ELMo, BERT)

### Chapter 4: Language Modeling

- 4.1 N-gram Models
- 4.2 Neural Language Models (RNNs, LSTMs)
- 4.3 Perplexity and Evaluation Metrics
- 4.4 Transformer-based Language Models
- 4.5 GPT and BERT Architectures

### Chapter 5: Syntax and Parsing

- 5.1 Context-Free Grammars (CFG)
- 5.2 Dependency Parsing
- 5.3 Constituency Parsing

- 5.4 Probabilistic Parsing
- 5.5 Applications in Question Answering

## **Chapter 6: Semantic Analysis**

- 6.1 Word Sense Disambiguation
- 6.2 Semantic Role Labeling
- 6.3 Sentiment Analysis (Lexical vs. ML Approaches)
- 6.4 Topic Modeling
- 6.5 Coreference Resolution

## **Chapter 7: Machine Translation & Seq2Seq Models**

- 7.1 Rule-Based vs. Statistical MT
- 7.2 Encoder-Decoder Model Architecture
- 7.3 Attention Mechanisms
- 7.4 Transformer Models in Translation
- 7.5 Evaluation (BLEU, METEOR)

## **Chapter 8: NLP with Deep Learning**

- 8.1 RNNs, LSTMs, and GRUs for NLP
- 8.2 Transformer Models (Self-Attention, Multi-Head Attention)
- 8.3 Transfer Learning in NLP (Fine-tuning BERT)
- 8.4 Prompt Engineering
- 8.5 Bias and Fairness in NLP Models

## **Chapter 9: Advanced NLP Applications**

- 9.1 Text Summarization (Extractive vs. Abstractive)
- 9.2 Chatbots & Dialogue Systems
- 9.3 Speech-to-Text & Text-to-Speech
- 9.4 Multilingual NLP (mBERT, XLM-R)
- 9.5 AI Ethics in NLP

## **Chapter 10: Emerging Trends & Future of NLP**

- 10.1 Large Language Models (GPT-4, Claude, Gemini)
- 10.2 Few-shot and Zero-shot Learning
- 10.3 Multimodal NLP (Text + Vision/Audio)
- 10.4 Explainability and Interpretability
- 10.5 NLP in Industry (Healthcare, Finance, Legal)