

TABLE OF CONTENTS

Chapter No.	Title	Page No.
1	Foundations of Deep Learning for NLP	01-20
	1.1 Overview of NLP tasks and challenges 1.2 Word embeddings: Word2Vec, GloVe, FastText 1.3 Neural networks: MLP, RNN, LSTM, GRU 1.4 Introduction to PyTorch/TensorFlow for NLP 1.5 Project: Text classification with word embeddings	
2	Attention and Transformers	21-40
	2.1 Attention mechanisms 2.2 The Transformer architecture 2.3 Pretraining and transfer learning in NLP 2.4 Fine-tuning BERT and other transformer models 2.5 Project: Named entity recognition (NER) using BERT	
3	Text Generation and Sequence Modeling	41-55
	3.1 Language Modeling 3.2 Sequence-to-sequence models with attention 3.3 Transformer-based generation (GPT, T5) 3.4 Evaluation metrics: perplexity, BLEU, ROUGE 3.5 Project: Text summarization or headline generation	
4	Sentiment Analysis and Social NLP	56-71
	4.1 Sentiment classification and emotion detection 4.2 Social media NLP: Twitter, Reddit, reviews 4.3 Sarcasm, hate speech, toxicity detection 4.4 Bias and fairness in NLP models 4.5 Project: Multi-class sentiment analyzer for real-world tweets	
5	Dialogue Systems and Question Answering	72-86
	5.1 Task-oriented vs open-domain chatbots 5.2 Retrieval-based vs generative QA 5.3 Use of large language models (LLMs) in QA 5.4 Memory networks and prompt engineering 5.5 Project: Build a QA system over a custom document corpus	

6	End-to-End Deployment of NLP Applications	87-102
	6.1 Model serving with FastAPI or Flask 6.2 Scalable inference with ONNX or TorchServe 6.3 Data privacy, interpretability, and ethical considerations 6.4 Continuous learning and monitoring in production 6.5 Project: End-to-end NLP system (student choice: chatbot, summarizer, sentiment engine, etc.)	