NAAIC

Resource provided by

HCC | HOUSTON COMMUNITY COLLEGE

Course Outline Details

Natural Language Processing

This intermediate course explores the principles, methods, and applications of Natural Language Processing (NLP) and language understanding in artificial intelligence and robotics

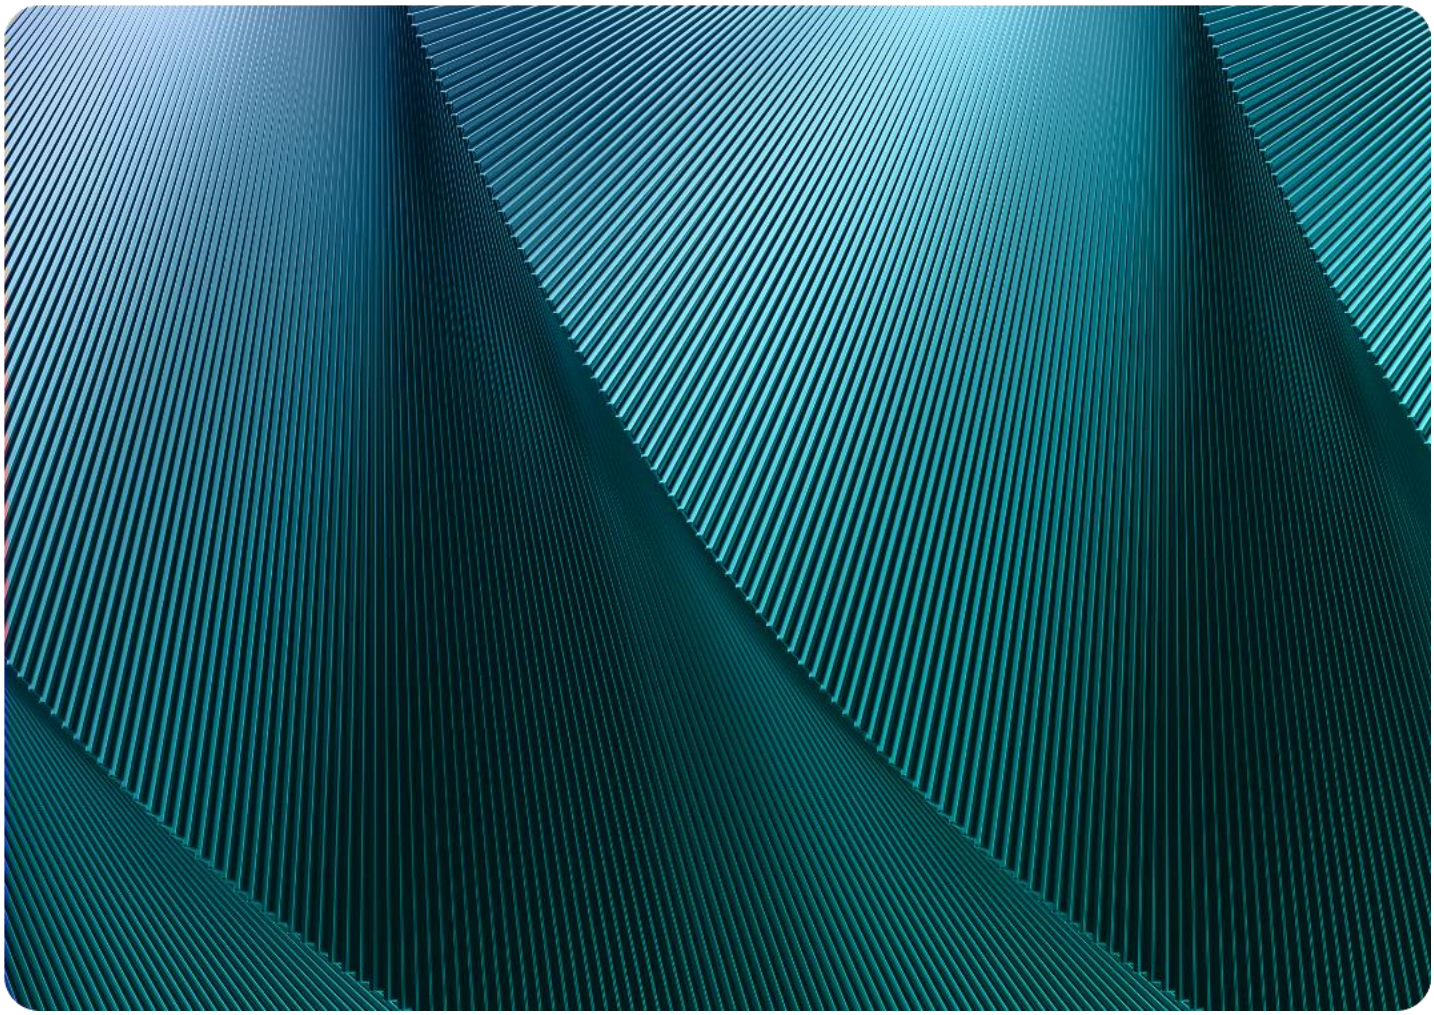

NAAIC

Course Outline

Week starting on / Module	Module Topic
Week 1	Mod 01: Introduction to NLP and NLU
Week 2	Mod 02: Text Preprocessing and Cleaning
Week 3	Mod 03: Introduction to Audio and Speech Processing
Week 4	Mod 04: Text Representation
Week 5	Mod 05: Part-of-Speech Tagging
Week 6	Mod 06: Syntax, Parsing, and Semantic Analysis
Week 7	Mod 07: Sentiment and Emotion Analysis
Week 8	Mod 08: Text Classification
Week 9	Mod 09: Named Entity Recognition (NER)
Week 10	Mod 10: Topic Modeling
Week 11	Mod 11: Language Models
Week 12	Mod 12: Machine Translation
Week 13	Mod 13: Conversational AI and NLU
Week 14	Mod 14: Summarization and Question Answering
Week 15	Mod 15: Course Review & Modern Trends, Responsible NLP
Week 16	Mod 16: Final Exam/Project Submission

NAAIC

Course Outline Details

Module 01: Introduction to NLP and NLU

- Course overview and expectations
- Key applications and industry use cases
- History and evolution of NLP and NLU
- Introduction to the NLP pipeline
- Overview of language modalities (text, speech, audio);
- Overview of NLP tools and libraries (NLTK, SpaCy)
- **Lab: Environment setup, simple tokenization and preprocessing**

Module 02: Text Preprocessing and Cleaning

- Tokenization methods
- Stop word removal
- Stemming vs. Lemmatization
- Handling punctuation, case normalization, special characters
- audio preprocessing (e.g., noise reduction, feature extraction like MFCCs) and how speech is transcribed to text for further processing
- **Lab: Text cleaning pipeline in NLTK and SpaCy**

Module 03: Introduction to Audio and Speech Processing

- Fundamentals of sound and speech
- Speech production and perception
- Acoustic features: spectrograms, MFCCs
- Audio preprocessing: noise reduction, normalization
- Overview of speech-to-text (ASR) and text-to-speech (TTS) systems
- **Lab: Loading, visualizing, and processing audio data using Librosa**

Module 04: ML Lifecycle, Overfitting, and Underfitting

- Bag of Words model
- TF-IDF representation
- Introduction to word embeddings (Word2Vec, GloVe)
- Sentence embeddings (conceptual)
- Compare text representations with audio feature representations
- **Lab: Feature extraction using Gensim and Scikit-learn**

Module 05: Part-of-Speech (POS) Tagging

- POS tagging concepts and importance
- POS taggers: rule-based, statistical
- Tagging challenges and ambiguities
- How speech recognition systems output text for POS tagging; challenges with spoken language (disfluencies, accents)
- **Lab: POS tagging using SpaCy, NLTK**

Module 06: Syntax, Parsing, and Semantic Analysis

- Introduction to syntax and parsing
- Constituency and Dependency parsing
- Introduction to semantic roles and WordNet
- Parsing spoken language; prosody and intonation as cues in speech understanding
- **Lab: Implementing dependency parsing with SpaCy, visualizing parse trees, extracting action-object pairs**

Module 07: Sentiment and Emotion Analysis

- Overview of sentiment and emotion analysis
- Techniques for sentiment analysis (Naive Bayes, SVM)
- Applications in support and robotics

- Emotion detection from speech (prosody, tone) as well as text
- **Lab: Implementing a sentiment analysis model with NLTK, using pre-trained models**

Module 08: Text Classification

- Introduction to text classification
- Feature extraction for text classification
- Common algorithms for text classification (Naive Bayes, SVM)
- Model evaluation and error analysis
- Classification of spoken utterances (intent detection in voice commands)
- **Lab: Building a text classification model with Scikit-learn, feature extraction and engineering, evaluating models**

Module 09: Named Entity Recognition (NER)

- Introduction to NER
- Common entity types
- NER in real-world and robotics applications
- Use cases of NER in extracting entities from transcribed speech (e.g., voice assistants, call center analytics)
- **Lab: Pretrained NER models and evaluation with SpaCy**

Module 10: Topic Modeling

- Introduction to topic modeling
- Latent Dirichlet Allocation (LDA)
- Applications of topic modeling
- Topic modeling on transcribed audio data (e.g., meeting recordings)
- **Lab: Implementing LDA for topic modeling with Gensim, visualizing topics**

Module 11: Language Models

- Probabilistic foundations of language modeling
- N-gram models, smoothing, evaluation metrics
- Limitations of N-gram models
- Introduction to neural language models (RNNs, LSTMs, conceptual)
- Overview of transformers and pre-trained models (BERT, GPT)

  Ethical considerations and biases in language models
- Language models in text and speech recognition/generation
- **Lab: Implement an N-gram language model in Python (briefly), explore pre-trained language models using the Hugging Face Transformers library (no fine-tuning)**

Module 12: Machine Translation

- Introduction to machine translation
- Rule-based, statistical, and neural machine translation (conceptual)
- Evaluation of machine translation systems
- Speech-to-speech translation pipeline (speech recognition → text translation → speech synthesis)
- **Lab: Using pre-trained translation models, evaluating translation quality**

Module 13: Conversational AI and NLU

- Conversational AI system architecture
- Intent detection and entity extraction
- Dialogue state tracking and context handling
- Rule-based and retrieval-based chatbots
- Introduction to Retrieval-Augmented Generation (RAG, conceptual)
- Dialogue systems for spoken and written input
- **Lab: Build a rule-based or retrieval-based chatbot for robotics/AI use case**

Module 14: Summarization and Question Answering

- Introduction to text summarization
- Extractive vs. abstractive summarization
- Introduction to question answering (QA) systems
- Summarizing spoken content (e.g., meeting or lecture recordings)
- **Lab: Implementing extractive summarization and simple QA with pre-trained models**

Module 15: Course Review and Modern Trends

- Review of all modules and key concepts
- Modern and multimodal NLP: text, image, audio (overview)
- Robotics applications of multimodal NLP
- Prompt engineering basics (overview)
- Responsible NLP: ethics, bias, privacy, and fairness (summary)
- Review activities, Q&A, project troubleshooting, portfolio

Module 16: Final Project/Final Exam

- Final project presentations or exam

Teaching Methods and Strategies

- Total Course Duration: 96 contact hours; 16 weeks
- Weekly Contact Time: 6 hours
- Weekly Structure:
    - Lecture: 2-3 hours
    - Lab: 3-4 hours