

COURSE SPECIFICATION DOCUMENT

Academic School / Department:	Science, Innovation & Technology
Programme:	MSc Artificial Intelligence
FHEQ Level:	7
Course Title:	Natural Language Processing
Course Code:	COMP 7102
Total Hours:	200 (Lev 7) (4 US Credits)
Timetabled Hours:	39
Guided Learning Hours	21
Independent Learning Hours:	140
Credit	20 UK CATS credits 10 ECTS credits 4 US credits

Course Description:

Natural Language Processing (NLP) is a core discipline within Artificial Intelligence that focuses on the computational analysis, modelling and generation of human language. This module provides a comprehensive, postgraduate-level introduction to NLP by integrating foundational linguistic concepts with modern machine learning approaches used in contemporary language technologies. Students examine the full language-processing pipeline, including text preprocessing, vector-space representations, statistical modelling, neural sequence architectures, attention mechanisms and transformer-based models.

Prerequisites:

None

Aims and Objectives:

The aims of this module are to provide students with advanced theoretical knowledge and practical expertise in Natural Language Processing, enabling them to analyse, design and implement modern language-based AI systems.

The module develops students understanding of linguistic structures, statistical and neural modelling techniques, and transformer-based architectures, while strengthening their ability to critically evaluate NLP systems.

Programme Outcomes:

A2, A4, A5, B2, B3, C1, C3, D3.

A detailed list of the programme outcomes are found in the Programme Specification. This is located at the archive maintained by Registry and found at:

<https://www.richmond.ac.uk/programme-and-course-specifications/>

Learning Outcomes:

By the end of this course, students will be able to:

1. Demonstrate advanced understanding of linguistic structures, text representations and the computational principles underpinning modern NLP systems.
2. Critically evaluate classical, statistical, neural and transformer-based NLP approaches, identifying their strengths, limitations and appropriate applications.
3. Design and implement complete NLP pipelines for tasks such as text classification, named-entity recognition, summarisation and translation, using contemporary frameworks.
4. Apply and fine-tune transformer-based models and Large Language Models (LLMs), and evaluate their performance using appropriate metrics, robustness testing and comparative analysis.
5. Communicate technical findings clearly and professionally through structured, academically rigorous written documentation.

Indicative Content:

- Text preprocessing techniques: tokenisation, segmentation, normalisation and feature engineering
- Statistical NLP approaches: n-grams, TF-IDF, classical language models and distributional semantics
- Word embeddings and vector-space representations (e.g. word2vec, GloVe, fastText)
- Neural sequence models: RNNs, LSTMs, GRUs and their limitations
- Attention mechanisms and the transformer architecture
- Large Language Models (LLMs): prompting, fine-tuning, evaluation and practical considerations
- Core NLP tasks: text classification, named-entity recognition, question answering, summarisation and machine translation
- Model evaluation: accuracy, precision, recall, F1, BLEU, ROUGE, perplexity and error analysis
- Implementing modern NLP pipelines using spaCy, Hugging Face Transformers and related tools

Assessment:

This course conforms to the University Assessment Norms approved at Academic Board and located at: <https://www.richmond.ac.uk/university-policies/>

Teaching Methodology:

Teaching includes lectures, coding labs, implementation workshops, guided learning, model evaluation sessions and guided independent study. Students will work exclusively with tools such as PyTorch, Gymnasium and Stable-Baselines3.

Indicative Text(s):

- Alammr, J. and Grootendorst, M. (2024). *Hands-on large language models: language understanding and generation*. Farnham: O'Reilly.
- Lane, H. and Dyshel, M. (2025). *Natural language processing in action*. Second edition. Shelter Island, NY: Manning Publications.
- Rothman, D. (2024). *Transformers for natural language processing and computer vision: explore generative AI and large language models with hugging face, chatgpt, gpt-4v, and dall-e 3*. Birmingham: Packt Publishing Limited.

Journals

- *Applied Artificial Intelligence*.
- *Journal of experimental and theoretical artificial intelligence*.

Websites

- Hugging Face: <https://huggingface.co/docs>
- spaCy: Industrial-strength natural language processing: <https://spacy.io>
- Natural Language Toolkit. <https://www.nltk.org>
- Stanford Natural Language Processing Group: <https://nlp.stanford.edu>
- Association for Computational Linguistics (ACL): <https://aclanthology.org>

See syllabus for complete reading list.

Change Log for this CSD:

Nature of Change	Date Approved & Approval Body (School or AB)	Change Actioned by Registry Services
Guided Learning Hours menu updated	October 2025	
Total Hours Updated	October 2025	