http://www.bcs-sgai.org/seminars/2024-05-08/?section=programme

Large Language Models (LLMs):
Democratise AI with Low-Code/No-Code for Text, Images, and Audio
Content generation and content analysis in 3 lines of python code

Each exercise can be executed in 1 minute to 5 minutes using V100 GPU in Google Colab
Some of the exercises use LLMs from Google, Facebook, Microsoft and OpenAI and apply them to Biomedical, Built Environment, and e-Banking data (i.e. text, images, and sound)

Part 1: Prompting

Low-Code AI using the Huggingface pipelines
[ex1] Text generation with instruction-trained models: zero-shot prompting (basic prompts)

No-Code AI using the Huggingface pipelines (changing the prompt Not the code)
[ex2] Prompt engineering: zero-shot, few-shot, and Chain-of-Thought (CoT) prompting

---

Part 2: Low-Code AI using the Huggingface pipelines

[ex3] Generative AI with Multimodal Large Language Models (LLMs): text to video
[ex4] Generative AI with Multimodal Large Language Models (LLMs): text to image
[ex5] Sequence of words: Predicting a "masked" word
[ex6] Text Classification: zero-shot text classification
[ex7] Text Classification: Sentiment Analysis (SA)
[ex8] Token Classification: Named-entity recognition (NER)
[ex9] Text feature extraction: text similarity (comparing two words or sentences)
[ex10] Image Classification: zero-shot object detection
[ex11] Image Classification: zero-shot or open vocabulary image classification
[ex12] Image Classification: assigns a label to an image
[ex13] Image feature extraction: image similarity (comparing two images)
[ex14] Image segmentation: semantic (pixel-level), instance, and panoptic segmentation
[ex15] Audio Classification: zero-shot audio classification
[ex16] Audio Classification/Speech Recognition: Keyword Spotting (KWS); Language Identification (LID)
[ex17] Automatic Speech Recognition (ASR): converts a speech signal to text
[ex18] Machine Translation (MT): speech translation
[ex19] Audio generation: from generating speech (including singing) to generating music

---

Part 3: few lines of code without using the Huggingface pipelines

---

[ex20 extra] Text Similarity with embeddings: few lines of code
[ex21 extra] Audio generation with speaker's voice character: few lines of code
[ex22 extra] Audio generation with Massively Multilingual Speech models: few lines of code

Protein Amino Acid sequences instead of word sequences
[ex23 extra] Protein language models to predict amino acid sequences: few lines of code

Evaluation metrics to assess the performance of the models
[ex24 extra] Standard metrics (Accuracy, Precision, Recall, F1-score): few lines of code

---

A hands-on 101 Tutorial by M Arguello-Casteleiro (University of Southampton) and T Furmston (University of Manchester)