Assignment 8
Hugging Face Transformers

Part B — Transformers & LLMs

14 questions

★	**Must-know** — asked very frequently in GenAI / ML interviews (2 questions). Prepare these first.
[CODE]	**Hands-on task** — write and run the code (7 tasks).

1. What are the advantages and limitations of the Hugging Face transformers library?
2. What tasks are supported by transformers pipelines?
3. Summarize a paragraph using pipeline with the 'google-t5/t5-base' model. [CODE]
4. Do sentiment analysis using a transformers pipeline. [CODE]
5. Classify these headlines into ["entertainment", "sports", "politics", "science"] using a zero-shot classification pipeline:
   - Oscar-winning actor to play the lead role in an upcoming crime-drama film
   - Australia beat India in the final of the Cricket World Cup by six wickets
   - Election commission announces seven-phase voting for the general election
   - NASA astronaut returns to Earth after a 6-month stay on the ISS [CODE]
6. Initialize the tokenizer for distilbert-base-uncased and encode a sentence: [CODE]

```
tokenizer = AutoTokenizer.from_pretrained(model_name)
encoded = tokenizer("this is a outstanding, product love it.")
print(encoded)
# {'input_ids': [101, 2023, 2003, 1037, 5151, ...,102], 'attention_mask': [1, 1, 1, ..., 1]}
```

   (a) What are input_ids and attention_mask in this output? (b) How do you return PyTorch tensors instead?
7. Convert the sentence "This is an outstanding product. I love it." into input IDs using the BERT tokenizer. [CODE]
8. Convert the input IDs [0, 100, 101, 102, 103] back into tokens using the BERT tokenizer. What do you notice? [CODE]
9. How do you download any model from the Hugging Face Hub in PyTorch and in TensorFlow (AutoModel vs TFAutoModel)?
10. Look at this code and output:

```
model = AutoModel.from_pretrained(model_name)
output = model(**inputs)
output.last_hidden_state.shape
# torch.Size([3, 14, 768])
```

    What does each number (3, 14, 768) represent?

11. Explain the Trainer and TrainingArguments classes. ★
12. What are autoregressive models? ★
13. In output = model.generate(**inputs, num_beams=5, no_repeat_ngram_size=2, max_new_tokens=50) explain **inputs, num_beams, no_repeat_ngram_size and max_new_tokens.
14. Compute BLEU, METEOR, ROUGE and perplexity using the Hugging Face evaluate library. [CODE]