

Exercises IN 4080

(The following questions are all part of previous exams)

Question 1

Here is a short example of dialogue:

Person 1: did you manage to finish the third obligatory assignment due yesterday?

Person 2: I sent it last week already!

Person 1: Impressive!

Could you send me your solutions?

Person 2: Sure, I'll send you the PDF

Answer the following questions based on this dialogue:

- 1) What are the speech acts (according to Searle's taxonomy) associated with each of those 5 utterances? (2 points)
- 2) Do you observe some conversational implicatures? Briefly explain. (2 points)
- 3) Does the common ground of those two persons evolve in the course of this short interaction? Explain in 2-3 sentences. (2 points)

Question 2

Take the following utterance:

robot please look at the ball no sorry the box

Analyse the disfluent part based on Shriberg's disfluency model.

Question 3

Assume you wish to develop a chatbot to answer questions about the current time in various locations around the world. The chatbot should for instance be able to respond to queries such as "What is the current time in Buenos Aires?", "What is the time difference between Boston and Oslo?" or "If it is 6:00 AM in Oslo, what time is it in Tokyo?".

Your first task is to decide what kind of chatbot development strategy you wish to follow. We have covered four alternative approaches during the course: handcrafted chatbots, IR-based chatbots, sequence-to-sequence chatbots and NLU-based chatbots.

Answer the questions below:

- 1) Which type of approach (among the four approaches above) do you think is most suitable for this task? Motivate your choice in one or two paragraphs.
- 2) Which data will you need to collect to train/develop your chatbot, based on the approach chosen above? What type of annotations would you need to add to this data, if any? Explain in a few sentences.

- 3) Which system modules (machine learning models or rule-based components) would you need to integrate in your chatbot? Describe in one or two paragraphs the general processing pipeline of your chatbot.
- 4) How would you evaluate the performance of the resulting chatbot? Describe in one or two paragraphs the evaluation procedure you would follow.

Question 4:

Calculate the Word Error Rate (WER) between this utterance:

could you go to my office and pick up my NLP book

And the hypothesis generated by a speech recogniser:

could you got you my office and pickup my NLB book

Show your calculations using an edit distance matrix. You can assume that insertions, deletions and substitutions all have a cost of 1.

Question 5:

- 1) Is it possible to have a Word Error Rate (WER) that is larger than 100%? Explain.
- 2) During the lecture on speech processing, we mentioned that one challenge faced when training speech recognition models was the lack of explicit alignments between the speech inputs and the output transcriptions. What did we mean by that? Explain in 2-3 sentences.

Question 6

What is the fundamental frequency (F0) in acoustics? And why might it be useful for a spoken dialogue system to measure it?