IN 4080: Additional exam questions

Dialogue

1) Turn-taking is a crucial part of conversational competence. What linguistic and extra-linguistic factors can influence how people take and release turns, and where the boundaries of these turns are likely to lie?

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.

Speech processing

1) Assume you have a user uttering the following utterance:

   Could you please take the red box and put it on the other end of the table?

   But that your speech recognition hypothesis turns out to be:

   Could you place vague red box and put it this on another end of that?

   Calculate the Word Error Rate (WER) between the ASR hypothesis and the actual utterance. Detail your calculations using an edit distance matrix.

Dialogue management

Imagine a kitchen robot whose task is to ask the user what kinds of cereals he/she wants for breakfast, wait for the user answer, and then hand out the appropriate cereal box once it knows the desired cereal. We want to design a simple dialogue system to handle the interaction with the user. A simple way to model it is via an MDP with only two states: state *UnknownCereal* where the robot doesn't know which cereal to give, and state *KnownCereal*, where the robot knows the cereal to hand out.

There are only two possible actions in the model:
• Action *AskCerealType* corresponds to the robot asking the user for the cereal box he wishes to have. The action is only available in state *UnknownCereal* and has a reward R = −1 in that state.
• Action *GiveCereal* corresponds to the robot physically giving the cereal to the user. The action is only available in state *KnownCereal* and has a reward R = +5 in that state.

When the robot executes the action *AskCerealType* in state *UnknownCereal*, it has a probability 0.8 of reaching state *KnownCereal* (if the user does answer the robot's question), and a probability 0.2 of remaining in state *UnknownCereal* (if the user ignores the question or provides an unclear answer). When the robot executes action *GiveCereal* in state *KnownCereal*, the MDP reaches a final state and finishes.

You are asked to calculate the expected cumulative reward of asking the cereal type while in the *UnknownCereal* state, i.e. *Q(s = UnknownCereal, a = AskCerealType)*. You can assume a discount factor of 0.9.
(Hint: use Bellman equation to calculate the *Q* values).

Ethics

2) Is it possible for an automated system to achieve demographic parity but remain unfair when it comes to the criteria of predictive parity or equalized odds? Illustrate your answer with a concrete (made-up) example.

3) Assume you have developed a neural model for sentiment analysis based on recurrent neural networks. You know wish to understand what your model has actually learned, and more specifically the tokens that are most important in the sentiment predictions for a given sentence. How would you use the LIME method to create such explanations? Explain your answer step-by-step.