

**Natural Language Processing
Final Exam**

September 11th, 2024

1. **[2 points]** With reference to linguistic theory, answer the following questions.
 - (a) Introduce the notions of root and affix, providing some examples.
 - (b) What is the distinction between inflectional morphology and derivational morphology? Again, provide some examples.
2. **[6 points]** With reference to the use of recurrent neural networks for language modelling, answer the following questions.
 - (a) Provide and explain the mathematical equations defining the inference process.
 - (b) Discuss the training process and the associated objective function.
 - (c) What are the advantages of the use of recurrent neural networks as opposed to feedforward neural networks for language modelling?
3. **[2 points]** With reference to the task of part of speech tagging, define the notion of open class and closed class tags, providing some examples.
4. **[6 points]** In the context of transition-based parsing, answer the following questions.
 - (a) Define the notion of spurious ambiguity, as we have introduced it in the course lectures.
 - (b) State two different sequences of transitions that make an arc-standard parser produce the projective dependency tree consisting of the following unlabeled dependency relations

head	w_3	w_1	$\langle \text{ROOT} \rangle$	w_5	w_3
dependent	w_1	w_2	w_3	w_4	w_5

(see next page)

5. **[5 points]** Consider the neural network architecture called sentence-BERT (SBERT), and answer the following questions.
 - (a) Discuss the tasks in which the model is usually applied.
 - (b) Explain how training is carried out in sentence-BERT.
 - (c) Explain how inference is carried out in sentence-BERT.
6. **[5 points]** To make practical use of pre-trained LLMs, we need to interface these with downstream applications.
 - (a) Briefly explain the notion of fine-tuning, providing a simple example.
 - (b) Introduce the techniques for parameter efficient fine-tuning that have been discussed in the course lectures.
7. **[5 points]** Consider the task of span-based machine reading, and answer the following questions.
 - (a) Present and discuss the neural approach using contextual embeddings produced by BERT.
 - (b) For many datasets the passages have length larger than the maximum number of input tokens allowed by BERT. How is this problem solved in the above approach?
8. **[2 points]** With reference to linguistic theory, outline the distinction between the field of study of general semantics and the field of study of pragmatics.