

**Natural Language Processing
Final Exam**

July 9th, 2026

1. **[5 points]** In the context of subword tokenization, answer the following questions.
 - (a) Define the byte-pair encoding (BPE) learner algorithm.
 - (b) Define the variant of BPE called byte-level byte-pair encoding (BBPE).
 - (c) Explain the advantages of BBPE over BPE.
2. **[5 points]** In the context of static word embeddings, answer the following questions.
 - (a) Specify the notation and the algorithm called skip-gram with negative sampling (SGNS).
 - (b) Introduce and discuss the objective function used by the SGNS algorithm.
3. **[5 points]** In the context of neural language modeling (NLM), answer the following questions.
 - (a) Introduce the NLM based on recurrent neural networks that has been presented in the course lectures, and report the model equations as used in the inference phase.
 - (b) Define and explain the objective function used by this model in the training phase.
4. **[2 points]** Define and explain the intrinsic evaluation measure of language models called perplexity.

(see next page)

5. **[5 points]** Introduce the neural model called sentence-BERT and its applications. Discuss the three objective functions for sentence-BERT that have been presented in the course lectures.
6. **[4 points]** Introduce the techniques for parameter efficient fine-tuning that are usually exploited in post-training and that have been discussed in the course lectures.
7. **[5 points]** Consider the Latin sentence with part-of-speech annotation ‘Parva/ADJ puella/NOUN rosam/NOUN pulchram/ADJ in/ADP horto/NOUN videt/VERB ./PUNCT’, which translates into the English sentence ‘The small girl sees a beautiful rose in the garden.’ Assume the projective dependency tree consisting of the following unlabeled dependency relations.

head	puella	videt	videt	rosam	horto	videt	⟨ROOT⟩	videt
dependent	Parva	puella	rosam	pulchram	in	horto	videt	.

Answer the following questions.

- (a) Draw a graphical representation of the dependency tree above, with arcs directed from the head to the dependent.
 - (b) Apply to the above tree the oracle presented in the syntactic parsing lecture, to construct a sequence of training instances for the arc-standard parser.
8. **[2 points]** Introduce the task of text completion, and provide a few examples of how to cast traditional NLP tasks as text completion.