

NLU+ Exam Review Guide 2024

(Some parts of this exam guide were adapted from the ANLP exam guide)

1 Exam content and format

- The final exam will cover content from the whole course.
 - That includes materials from lectures, assigned readings, tutorial and lab exercises and homework assignments.
 - The exam will include some questions that require basic mathematical understanding and fluency with concepts presented in the class, and some that require a higher level analysis of a situation or a circumstance in NLP. (For more discussion of question types, see the Revision Tips below.)
 - This course covers a lot of material. If you are struggling with that, we suggest focusing on the core material that is covered in detail in the lectures, homework, and tutorials. Some lectures also include more advanced topics where we refer to papers for details, or where we briefly mention more recent research related to the topic. While there may be questions about this on the exam, they will typically be a very small proportion of the marks, designed to identify the very strongest students.
- Questions will be worth 50 marks in total and all questions will be marked (you don't get to choose which ones to answer). The marks are divided into three questions, where each question covers mostly material that each of the lecturers taught.
- Calculators and dictionaries are not permitted.
- You will write your answers on paper in a provided exam booklet.

2 Revision tips

- **Don't wait until the last minute.** Study early and often. Research shows that spaced repetition is the most effective way to remember material.
- **Study actively.** Simply reading information is not a good way to remember or understand it. Work through examples, making sure you get the same answer that's in the text or notes, and think about what questions we might ask related to that material (see below about types of questions). Also ask yourself how different parts of the course connect to each other. What are the different themes? What are different ways to approach some of the tasks?
- **Consider the types of questions** we might ask, and how they would apply to different parts of the course. Types of questions include:

- **Bookwork** questions, which just require you to remember facts or definitions. We rarely have many marks associated with this type of question, because simply remembering a fact doesn't demonstrate that you understand how to apply it appropriately, and in a real-life situation you would typically be able to look up facts as needed. However, such questions might be a lead up for further analysis of memorised definitions or facts.
- **Applying** a method, model or linguistic concept to a particular situation.
- **Synthesising** knowledge to address a new situation. For example, given a new task related to one we've discussed, discuss how you could apply or extend methods we have seen to address it, and the problems you might expect to encounter in doing so.
- **Practice the way you will be tested.** You will need to write your answers in pen on paper, in a limited amount of time. So practice doing that, using questions from past papers, tutorial exercises, etc.
- **Do a practice exam(s).**
 - It is a good idea to do (at least one) full practice exam during revision week, after you have done most of your revision.
 - See below for more information on past papers.
 - Sit down and do the practice exam in a single 2-hour sitting, writing your answers on paper, as you will during the real exam. This may help you discover issues with timing or gaps in your knowledge.

3 Past papers

Past papers are available at this link. They cover previous years starting in 2014. Focus on more recent exams. You should *not* expect to be able to answer questions about material that was not covered in the course. If in doubt, ask the lecturers or ask on Piazza during your revision.

4 Exam-taking tips

- **Plan your time and use it wisely.** You've got 120 minutes to answer 50 marks worth of questions, so think ahead about how you want to use that time.
 - You'll probably want to spend a few minutes at the beginning to skim the questions and plan your approach, and a few minutes at the end to review your answers or come back to questions you did not fully answer.
 - After accounting for that, you'll probably have, on average, around 2 minutes per mark. However, some questions will be much quicker to answer than others. Since you don't need to answer questions in any particular order, you may want to start with the ones that you find quick and easy. You could also start with a partial or basic version of your answer, and leave some space to fill in more details later if you have time. (But see the next several tips.)
- **Keep track of your progress.** For example, you can tick off each question (or question part) when you've done it, or make a note if you've written something for that question but you want to come back to it if you have time.

- **Look for easy parts.** . Where a question has multiple parts, *often* the earlier parts are easier than the later parts, but not always!
 - Some multi-part questions have a “narrative” or natural ordering that doesn’t align well with moving from easy to hard. So do read all parts of the question: even if you can’t answer one part, you still may be able to answer the next part.
 - The number of marks on a question is only weakly correlated with its difficulty, because if all the hard questions were worth a lot, most students would do very poorly! Our goal is to achieve a spread of marks that shows how well each student has met the learning outcomes. This also means that most high-value questions have ways to achieve partial credit, and most questions where partial credit is not possible (e.g. multiple-choice questions) are not worth as much.
- **Answer the question that is asked.** Read each question carefully, and make sure you understand what it’s asking, and what type of answer is needed. For example, if the question asks you to justify your answer, you will get little or no credit for simply stating the answer, even if it is correct; you will need to provide a reason.
- **Don’t write more than needed.** There are at least two potential failure modes to avoid:
 - If a single word or phrase can clearly demonstrate you know the answer, don’t waste time writing full sentences or repeating parts of the question.
 - If (say) a single reason or a brief explanation is enough to answer the question, don’t add more just for the sake of it. And in particular, if you don’t know the answer, please don’t tell us everything you know about the topic of the question just in case you stumble on the right answer. Answers that include a lot of irrelevant information indicate to us that you do not have a strong grasp of the relevant concepts, so you are not doing yourself any favors and may be using up valuable time.

In some cases, you might be unsure how much detail is needed. Most questions have hints to help you, but if you’re still not sure, consider providing an initial answer with the most important information, and coming back to it later to add more detail if you have time (and if those extra details are actually relevant).

- **Do demonstrate your understanding clearly.** As noted above, a single word or phrase can sometimes be enough—but not always! Take note of any hints provided, and make sure you’ve expressed yourself clearly. We don’t deduct marks for minor problems with spelling or grammar, but your English does need to be good enough to convince us that you know the right answer and can explain it clearly. (Note: your job is to convince us. It is not our job to guess whether you might know the right answer given what you wrote. You do not get the benefit of the doubt.)

5 NLU+ 2024 Examinable Topics

Below is a list of concepts you should be familiar with and questions you should be able to answer if you are thoroughly familiar with the material in the course. It is safe to assume that if you have a good grasp of everything listed here, you will do well on the exam. However, we will not guarantee that only the topics mentioned here, and nothing else, will appear on the exam. In a few cases (mainly formulas) we do specify what *will not* be required for the exam, but otherwise we make no guarantees.

5.1 First block of the course

In this block (taught by Lexi) we discussed modelling language probabilistically. You should be able to describe:

- How translated data is found. How word alignments are calculated. What problems there are with treating translation as a word alignment problem.
- How to define an n-gram language model using a probabilistic model. How to estimate these probabilities.
- How to define a conditional n-gram language model which can model translation using alignment. How to decode with a conditional language model.
- How to represent n-grams with feedforward networks. One-hot encoding. How to learn model parameters and gradient descent.
- How to define a recurrent neural network and how to train it with backpropagation through time and the vanishing gradient problem.
- Long-short term memory models.
- How we model our input and output language data.
- The open vocabulary problem. Subword units and Byte Pair Encoding. BPE Dropout and Character level modelling.

5.2 Second block of the course

In this block (taught by Shay) we discussed extensively various concepts in NLP and their relationship to deep learning. You should be able to:

- Understand the idea behind sequence-to-sequence modeling, the limitations of different types of seq2seq models, how attention resolves them, and applications for such models.
- Describe and re-generate the equations for the transformer model, understand them, and understand the applications the transformer models might have.
- Understand specific models that are based on the transformer models, such as BERT, T5 and other models mentioned in the lectures.
- Understand the use and motivation behind pre-trained language models.
- Understand the different types of decoding methods that can be used with large language models.
- Understand how LSTMs and transformers can be used for parsing and what the limitations are of such models in parsing.
- Understand scaling laws and how to analyse the different parameters that drive training a large language model (including compute, training data size and model size).
- Understand basic issues about the safety of AI models.

5.3 Third block of the course

- Evaluation of generation: human, n-gram overlap metrics, embedding-based metrics, trained metrics.
- Monolingual and multilingual data for machine translation (MT), pretraining for MT, mBART and multilingual MT
- Ethics: Social impact of NLP, Types of Risks, Things think about when building a system.
- Bias in Word Embeddings, Bias in NLP systems, how to debias embeddings and NLP systems.
- Translation using LLMs, differences in translations from translation models and LLMs, Controllable MT.
- Reading Comprehension and (Open-Domain) Question Answering
- Emerging abilities of LLMs, Instruction Fine-Tuning, Reinforcement Learning from Human Feedback (RLHF)
- Retrieval-Augmented Generation (RAG)
- In-Context Learning and its limitations, and Chain-of-Thought Reasoning
- Summarisation, then and now
- Parameter-Efficient Fine-Tuning (PEFT, via Parameter Composition, Input Composition, and Function Composition)
- Unsupervised Parsing

6 Other formulas

In addition to the equations for the models listed above, you should know the formulas for the following concepts and what they may be used for and be able to apply them appropriately. Where relevant, you should be able to discuss the strengths and weaknesses of the associated method and alternatives.

- Maximum Likelihood Estimation
- Dot product, cosine similarity
- Precision, recall, and F-measure
- Feedforward equations
- Transformer equations
- Analysis of scaling laws
- RAG: Sparse and Dense retrieval
- PEFT: LoRA parameter updates and Adapters
- RLHF: REINFORCE gradient estimator
- ROUGE evaluation metrics