

Final Exam

601.467/667 Introduction to Human Language Technology

Fall 2023

Johns Hopkins University

Co-ordinator: Philipp Koehn

14 December 2023

Complete all questions.

Use additional paper if needed.

Time: 75 minutes.

Name of student: _____

Q1. Question Answering

20 points

1. Researchers have defined various question types for QA systems. Understanding the question type is important for defining the scope of the problem and restricting the answer candidates. Please write down the name of five different question types and give an example question for each. (10 points)
2. The main components of a QA system include Question Analysis, Search, Candidate Extraction, Knowledge Sources, and Answer Ranking. Please explain the purpose of each. (5 points)
3. Here is an example of the Winograd Scheme Challenge. Explain why it is challenging for Machine Reading Comprehension (MRC) systems. (5 points)
Q: The trophy would not fit in the brown suitcase because it was too big. What was too big?
A. The trophy
B. The suitcase

Q2. Digital Humanities

20 points

1. What are some ways in which humanistic scholars resemble knowledge workers from industry, medicine, finance, etc? (10 points)
2. What are some ways in which computational approaches offer advantages or drawbacks with respect to humanistic scholarship? (10 points)

Q3. Interpretable and Explainable NLP

20 points

1. Briefly describe the main difference between black-box and white-box explanations (10 points)

2. Briefly describe how LIME works (5-6 sentences) (10 points)

Q4. Ethical Problems

20 points

1. Describe the Ethical Principle of Beneficence in AI. (10 points)

2. Describe the notion of Informed Consent in data collection. (10 points)

Q5. Large Language Models

20 points

1. Answer the following questions. (8 points)

(a) In 1-2 sentences, define “language modeling”.

(b) In no more than five sentence, describe the GPT family of models (the architecture, how they’re trained, and how they generates text at inference time.)

2. Select all the answer(s) to fill in the blank (.) in each item. (4 points)

(a) is an argument for the infeasibility of scale due to limited computing.

- ☐ Advances in computing hardware are much slower than the trends for scaling language models.
- ☐ Advances in parallel computing can support the fast pace of scaling models.
- ☐ Scaling language models continues to incur a lot of costs (monetary, carbon footprint, computing resources, etc.)
- ☐ Scaling models might reduce the overall costs: the availability of a few large models may prevent the cost of building many smaller ones.

(b) is an argument that “data” should not be a bottleneck for scaling language models.

- ☐ There size of the internet is consistently growing.
- ☐ There size of Wikipedia is consistently growing.
- ☐ One can mine data from other modalities (e.g., text data mined from videos).
- ☐ Even with limited data, we can use it more effectively to get more gains.

3. Answer the following questions in a few sentences (no more than 5 sentences for each). (8 points)

(a) Explain what long tail of problems in natural language is (provide an example).

(b) Explain how the long tail of problems in natural language poses a challenge to language models.

Q6. Computational Social Science

20 points

1. You have a data set of tweets containing terms and hashtags related to the recently released Gemini AI model. Your data has two labels (1) whether or not the post was made by an AI researcher (2) the date the tweet was posted. From glancing at the data, it seems that some people are excited about the model's capabilities, while others are critical about the lack of transparency in the development process, but you don't know much else about what people are saying.

- (a) What are two methods you might use to analyze the data? Write 2-3 sentences about how you would specifically apply the method. (10 points)

(1) Method:

Description of application:

(2) Method:

Description of application :

- (b) What are two limitations or ethical considerations of analyzing this data? (4 points)

2. What is one of the key ways in which a social scientist's approach to a problem often differs from an NLP scientist's approach? Provide an example from each discipline illustrating this difference. (6 points)