HUL381/ELL457: Mind, Machines and Language

Major Test (SA)

Maximum points: 40

Form: A (**please write this Form ID on the cover of your answer book**)

All questions in this part of the exam are to be answered in your answer book. This question paper will not be collected.

Section 1. Multiple choice questions

Each question may have any number of correct choices, including zero. List all choices you believe to be correct (2 marks for each correct choice, -1 for each incorrect choice). No justification is required.

1. Why do Lakoff and Johnson think that the use of metaphor involves embodied cognition?
   (a) Most metaphors directly make use of some part or aspect of the human body.
   (b) Metaphor seems to be not just a linguistic phenomenon, but a more general cognitive tool we use to understand the world around us.
   (c) Many metaphors involve understanding a more abstract concept in terms of a more concrete one, which is often articulated in terms of our bodies and how they relate to the physical world around us.
2. In Bayesian models, one advantage is the ability to naturally incorporate prior or domain knowledge, in addition to the data being looked at. Which of the following can be incorporated *as prior knowledge* into the LDA model we discussed?
   (a) The words *watch* and *clock* have similar meanings.
   (b) Most documents tend to be focused, i.e., dominated by a single topic.
   (c) The word *play* is polysemous, i.e., has different senses in different contexts.
3. Which of the following is NOT an implication of functionalism?
   (a) Computer programs alone cannot achieve conscious mental states; only a machine whose causal structure brings about intentionality can do so.
   (b) The functions of the mind are realisable in multiple kinds of physical 'hardware'.
   (c) Mental states like emotions cannot be regarded as physical kinds, and so cannot be characterised on the basis of their physical correlates.
4. With regard to the XOR problem, which of the following are true?
   (a) Neural network models cannot compute XOR.
   (b) A single-neuron neural network cannot compute XOR.
   (c) Neural network models with intermediate representations, rather than a direct input-output mapping, can compute XOR.
   (d) A single-neuron neural network can compute XOR only if we set a very high activation threshold for that neuron.
5. Which of the following sentences contain instances of the conceptual metaphor TIME IS MONEY?
   (a) *Time flies by so quickly when you are with an interesting person.*
   (b) *What time tomorrow should we meet?*
   (c) *I will really struggle to complete all my assignments on time.*
   (d) *I've started driving to office to save some time compared to taking the Metro.*

6. Which of the following is an example of a sentence where quantification or variable-binding might pose a challenge for a connectionist model of sentence processing?

   (a) *My neighbour's baby ate a slug.*

   (b) *A bat broke my bedroom window.*

   (c) *Delhi is a better example of an Indian city than Istanbul.*

   (d) *I think you know that Deepa is aware how Amit feels about Neha.*

7. Which of the following relationships are approximately true, when looking at a large corpus of natural language text?

   (a) The frequency of occurrence of the r^{th} most frequent word is proportional to γ^{-r}, where $\gamma > 1$.

   (b) The frequency of occurrence of the r^{th} most frequent word is proportional to $r^{-\gamma}$, where $\gamma > 0$.

   (c) The frequency of occurrence of the r^{th} most frequent word is proportional to $K - \gamma \log r$, where $K > 0$, $\gamma > 0$.

Section 2. Short Answer Questions

8. The readings by Pinker, Turing, and Searle all grapple with the same questions: can machines think, and if so, what kinds of machines? In this regard, answer the following:

   (a) Why does Pinker believe that connectionist models alone may not be sufficient to explain all the abilities of the mind? Describe any two reasons he gives, with suitable examples. **[6]**

   (b) In what way did Turing anticipate the Chinese room kind of critique of the Turing test that Searle would make 30 years later? What was Turing's response to this line of criticism? **[4]**

   (c) Out of the three men, which two do you think would agree most closely with each other on these questions? Give reasons, based on the readings. **[4]**

9. Consider the following three topics learnt by an LDA model for a given corpus (for each topic, the top 10 words are listed):

   - Topic 1: *play, stage, audience, theatre, actors, drama, performance, costumes, comedy, tragedy.*
   - Topic 2: *team, game, cricket, football, tennis, player, play, ball, field, court.*
   - Topic 3: *judge, trial, case, jury, court, accused, guilty, defendant, lawyer, justice.*

   (a) Once the model has been learnt, it can be used to infer topic distributions for new documents. Suppose I have a document consisting of just one sentence: *Yesterday, in a matchfixing trial, the judge found the cricket team guilty of deliberate bad performance.* Considering only the words shown in the topics above, and assuming that these are the only words in the given topics and that each of those words has the same probability of occurring in its respective topic(s) (i.e., all these probabilities are 0.1), calculate the topic distribution for this document. Assume that all topics are equally likely *a priori*. (Hint: Use Bayes' theorem to get the posterior probability distribution for the topic of each relevant word, then average over the words to get the distribution for the whole sentence.) **[7]**

   (b) What advantage does this kind of probabilistic topic model have over a conventional semantic space model? Illustrate your answer by identifying specific instances where the above example topic model might be better than a semantic space model for the same set of words. **[3]**