

# ELL457/HSL622: End-Term Examination

April 30, 2024

Maximum Marks: 32

## Instructions:

- Please clearly indicate the question number, and part number if applicable, at the start of each response.
- Please read all questions carefully.
- Please ensure that your responses are to-the-point and that you write only what is asked for on the answer script you submit.
- While the exam is open-notes, all your answers must be written entirely in your own words, without any copying from anywhere.
- Please try to be clear and careful with all formal/mathematical notation, so that there is no ambiguity in the expressions/formulae you write down. Try to stick to the kind of notation used in class as far as possible.

## Multiple-Choice Questions

[Each question may have any number of correct choices. List, in clear handwriting, the index letters corresponding to all options you believe to be correct for each question. Nothing else should be written on your answer script for these questions (for any rough work, please use the back of the script). +3 for each correct choice, -1 for each wrong choice.]

1. How many inputs and outputs does the transition function of a Turing machine's finite-state machine have?
  - (a) 2 inputs, 2 outputs
  - (b) 3 inputs, 4 outputs
  - (c) 2 inputs, 3 outputs
  - (d) 3 inputs, 2 outputs
2. Are the Dreyfus and Searle attacks on Strong AI dependent on whether one is considering symbolic or connectionist AI?
  - (a) No, both critiques are independent of the computational framework
  - (b) Searle's critique applies primarily to connectionist models, but Dreyfus' is independent of the framework
  - (c) Both critiques apply primarily to symbolic models
  - (d) Dreyfus' critique applies primarily to symbolic models, but Searle's is independent of the framework
3. Consider a single perceptron or linear threshold unit with 2 binary inputs,  $x_1$  and  $x_2$ , and binary output  $y$ . Let the synaptic weights for the two inputs be set to  $w_1 = 1$  and  $w_2 = -1$ . Which of the following Boolean functions is it possible for this unit to implement, for *some* value of the threshold  $\theta$ ?
  - (a)  $y = x_1 \text{ AND } x_2$
  - (b)  $y = x_1 \text{ AND } (\text{NOT } x_2)$
  - (c)  $y = (\text{NOT } x_1) \text{ OR } x_2$
  - (d)  $y = (\text{NOT } x_1) \text{ AND } x_2$
  - (e)  $y = x_1 \text{ OR } (\text{NOT } x_2)$
  - (f)  $y = x_1 \text{ OR } x_2$
  - (g)  $y = \text{NOT } x_2$

4. Which of the following are true with regard to the role of representations in symbolic and connectionist models?
  - (a) Representations can act as a bridge between the frameworks, because connectionist models allow for easy hardcoding of representations, and symbolic models allow for easy finetuning of such representations
  - (b) Connectionist modellers usually regard representations as unimportant for AI
  - (c) Connectionism makes representation a more central aspect of AI than symbolism, because connectionist models gain their power from the ability to learn good representations
  - (d) While representations are important in both approaches, they go more naturally with symbolic approaches because the latter better facilitate representation learning
  - (e) Because connectionist models can be seen as operating at the subsymbolic level of learning representations from data, and symbolic models can then be used for higher-level (e.g., logical) computation with such representations, we can see the notion of representation as providing a natural bridge between the two approaches
  
5. Why do sparse connectivity and weight sharing in CNNs naturally complement each other?
  - (a) Sparse connectivity allows the model to process localised information from a small part of the input image, and weight sharing allows multiple types of features to be extracted from such small regions
  - (b) Sparse connectivity reduces the number of parameters in the model, making it less complex; and weight sharing allows these parameters to more flexibly capture a diverse range of features
  - (c) Sparse connectivity allows the model to process localised information from a small part of the input image, and weight sharing allows such localised features to be uniformly computed over the entire image
  - (d) Sparse connectivity reduces the number of parameters in the model, making it less complex; and weight sharing allows these parameters to be specialised to local receptive fields in the larger input image

## Longer-Answer Questions

6. With regard to modelling semantic memory for words like BANK which have multiple distinct meanings, answer the following as precisely as possible.
  - (a) Why does a semantic network model, where nodes are words and links between nodes denote semantic relationships, not work very well in such cases? [1]
  - (b) Why does a semantic space model, where each word is a point in a vector space such that words with related meanings are close to each other, not work very well in such cases? [1]
  - (c) How exactly does the Bayesian topic modelling approach solve these issues? Don't try to write out all details of the approach; focus only on the specific aspects which are relevant to dealing with this kind of case, and try to illustrate how they work via clear examples. [3]
  
7. The year is 2035. ChatGPT-17 has just been released. It is the first-ever AI system to rigorously pass the Turing test: thousands of human participants who have engaged in detailed conversations with it all agree that its responses are indistinguishable from human ones. Leading AI researcher Nam Altmas has triumphantly declared: "The Chinese Room argument is finally resolved! Strong AI is possible after all!".
  - (a) Is he right? Explain precisely why or why not. [2.5]
  - (b) What would be the respective views of Daniel Dennett and John Searle on whether ChatGPT-17 has intentionality or not? Give clear and precise reasons for your answers, based on the readings/discussions from the course. [4]
  - (c) How would Alan Turing respond to those who claim that ChatGPT-17 doesn't have real understanding or intelligence, because it lacks consciousness and qualia? Again, give a precise answer based on the course materials. [2.5]