



UNIVERSITY OF RUHUNA

Faculty of Engineering

End-Semester 8 Examination in Engineering: December 2025

Module Number: EE8210

Module Name: Intelligent Systems Design (C-18)

Part - A


[30 Minutes]

[This paper consists of 4 pages.]

Instructions for candidates

- Write your index number at the top of every page.
- The question paper contains 20 multiple choice questions (MCQs).
- Answer all questions. Each question has only one correct answer.
- For each question, put an X mark on the letter: (a), (b), (c), or (d) which corresponds to the correct answer, by using a black or blue pen.
- The total marks allocated for the 20 MCQ questions is 10 marks.

1. Consider the following Neural Network.



Let $a^1 = x \in \mathbb{R}^{n+1}$ denotes the input with ($a_0^{(1)} = 1$), How do you compute $a^{(2)}$?

- (a) $a^{(2)} = \emptyset^{(1)} a^{(1)}$
- (b) $z^{(2)} = \emptyset^{(2)} a^{(1)}$; $a^{(2)} = g(z^{(2)})$
- (c) $z^{(2)} = \emptyset^{(1)} a^{(1)}$; $a^{(2)} = g(z^{(2)})$
- (d) $z^{(2)} = \emptyset^{(2)} g(a^{(1)})$; $a^{(2)} = g(z^{(2)})$

2. In the context of image-based learning, what does the term “curse of dimensionality” refer to?

- (a) Difficulty in visualizing high-dimensional data.
- (b) Overfitting due to small datasets.
- (c) Rapid growth of feature space when quadratic features are used.
- (d) Vanishing gradients in deep networks.

Index number:

3. In evolutionary neural networks, evolutionary algorithms are mainly used to:
 - (a) Replace neural network training.
 - (b) Optimize weights, architecture, or learning parameters.
 - (c) Generate fuzzy rules directly.
 - (d) Perform symbolic reasoning.
4. Which of the following correctly describes the purpose of backpropagation?
 - (a) To compute the output of the neural network.
 - (b) To initialize weights randomly.
 - (c) To calculate weight updates in order to minimize the cost function.
 - (d) To normalize the input data.
5. Which of the following best explains why Genetic Algorithms work effectively?
 - (a) They perform an exhaustive search over the solution space.
 - (b) They rely only on deterministic optimization.
 - (c) They combine directed and stochastic search with exploration and exploitation.
 - (d) They always converge to the global optimum.
6. Why are Fully Connected Neural Networks (FCNNs) not well-suited for image processing tasks?
 - (a) They cannot learn non-linear relationships in image data.
 - (b) They fail to preserve spatial information and involve a large number of parameters.
 - (c) They cannot be trained using backpropagation.
 - (d) They are unable to perform classification tasks.
7. Which of the following statements correctly compares max pooling, min pooling, and average pooling in Convolutional Neural Networks (CNN)?
 - (a) Max pooling captures the most prominent features, average pooling smooths feature maps, while min pooling emphasizes low-intensity or background features.
 - (b) Max pooling reduces overfitting, min pooling increases spatial resolution, and average pooling increases the number of trainable parameters.
 - (c) All three pooling methods preserve exact spatial information while reducing computational complexity.
 - (d) Average pooling is primarily used for edge detection, whereas max and min pooling are employed mainly for noise reduction.
8. In a CNN, which of the following correctly describes the typical progression of feature extraction through the layers?
 - (a) Early layers detect high-level concepts, middle layers detect edges and textures, and final layers detect basic shapes.
 - (b) Early layers detect edges and textures, middle layers detect shapes and patterns, and final layers detect high-level concepts.

Index number:

- (c) All layers detect only low-level features such as edges and corners.
 - (d) Early layers extract simple features, and subsequent layers gradually combine them to produce class predictions.
9. Applying a large penalty term in L1 or L2 regularization primarily results in:
- (a) A highly wiggly curve that memorizes noise.
 - (b) A simpler, smoother model that may underfit if the penalty is too high.
 - (c) No change in the model.
 - (d) Increased number of complex terms in the model.
10. A convolutional layer is applied to an input image of size 64×64 using a 5×5 filter with a stride of 1. To maintain the same spatial dimensions in the output feature map as the input, which type of padding should be used?
- (a) Valid padding (no padding)
 - (b) Reflective padding
 - (c) Same padding
 - (d) Zero padding is not required
11. Which member of the expert system development team is specifically responsible for interviewing the domain expert and representing their knowledge in the system?
- (a) The Programmer
 - (b) The Project Manager
 - (c) The Knowledge Engineer
 - (d) The End-user
12. "Data-driven" reasoning, which starts from known facts and proceeds forward to reach a conclusion by adding new facts to the database, is known as:
- (a) Forward Chaining
 - (b) Backward Chaining
 - (c) Conflict Resolution
 - (d) Knowledge Engineering
13. Why is the Explanation Facility considered a vital component of an expert system?
- (a) It translates computer code into human languages like English.
 - (b) It allows the user to understand how the system reached a conclusion.
 - (c) It automatically updates the knowledge base with new rules.
 - (d) It speeds up the inference engine by skipping unnecessary rules.
14. In the context of a recruitment platform, why is algorithmic bias considered a major ethical concern?
- (a) It significantly reduces the processing speed of the candidate screening system.
 - (b) It causes the AI to perform more slowly when evaluating unseen data.

Index number:

- (c) It can amplify existing societal inequalities in hiring practices.
 - (d) It prevents the system from being able to retrain itself without human intervention.
15. Which of the following remains a significant privacy risk in Federated Learning, even though raw data never leaves the client's device?
- (a) Excessive use of the local device battery.
 - (b) The inability of the central server to aggregate updates.
 - (c) Algorithmic bias in the final global model.
 - (d) The possibility of reconstructing private client data from model updates.
16. Which principle is most violated when companies use data without approval to train AI models?
- (a) Open-source Compliance
 - (b) Consent and Data Ownership
 - (c) Accuracy of Algorithms
 - (d) Fairness
17. In a Recurrent Neural Network (RNN), what does the hidden state represent?
- (a) The final output of the network.
 - (b) The relationship between inputs and outputs.
 - (c) The learning rate of the network.
 - (d) A summary of past inputs in the sequence.
18. What is the primary characteristic that distinguishes RNNs from feedforward neural networks?
- (a) Use of convolutional layers
 - (b) Ability to process inputs in parallel
 - (c) Presence of feedback connections
 - (d) Requirement of large labeled datasets
19. Which type of data are RNNs best suited to handle?
- (a) Static images
 - (b) Tabular data
 - (c) Sequential data
 - (d) Graph-structured data
20. Which problem commonly affects standard RNNs when learning long sequences?
- (a) Overfitting
 - (b) Vanishing or exploding gradients
 - (c) High computational cost
 - (d) Underfitting