

Week 51

# A-Level Computer Science

Homework bundle for this week

本周作业合集

exercises.lol

## Contents 目录

|                                 |    |
|---------------------------------|----|
| Topic 18 quiz .....             | 2  |
| Exercise sheet 18.1 .....       | 5  |
| Past-paper questions 18.1 ..... | 12 |

## 18. Artificial Intelligence (AI)

Starter Quiz · 课前小测

A-Level Computer Science

Name: \_\_\_\_\_ Score: \_\_\_\_\_

- Machine learning means an algorithm that:
  - ☐ learns patterns from data rather than being programmed step by step
  - ☐ is written entirely in assembly
  - ☐ never needs any data
  - ☐ only follows fixed rules
- Supervised learning uses:
  - ☐ labelled data, learning to map input → label
  - ☐ unlabelled data, finding hidden structure
  - ☐ rewards from an environment
  - ☐ no data at all
- The difference between regression and classification is that regression predicts:
  - ☐ a number, while classification predicts a category
  - ☐ a category, while classification predicts a number
  - ☐ nothing
  - ☐ only true/false
- How do AI, machine learning and deep learning relate?
  - ☐ deep learning is part of machine learning, which is part of AI
  - ☐ they are three competing approaches
  - ☐ AI is part of machine learning
  - ☐ deep learning replaced AI
- Unsupervised learning is used to:
  - ☐ find structure in unlabelled data, such as clusters
  - ☐ map inputs to known labels
  - ☐ maximise a reward signal
  - ☐ compile code
- Regression predicts a continuous number (a price, a temperature), while classification predicts a category (cat vs dog).
  - ☐ True    ☐ False

7. A road network is modelled as a graph to find the shortest route. Which statements are correct? Select **all** that apply.

- ☐ each town is a node
- ☐ each road is an edge joining two nodes
- ☐ each edge carries a weight, the distance or travel time
- ☐ each town is an edge and each road a node

*Tick every correct option.*

8. In reinforcement learning the agent learns a \_\_\_\_\_ that maximises its total reward over time.

\_\_\_\_\_

9. A neural network is called "deep" when it has many hidden layers, and each neuron takes a weighted sum of its inputs (plus a bias) before applying an activation function.

- ☐ True    ☐ False

10. In gradient descent the learning rate sets how big a step each weight update takes — too large overshoots the minimum, too small makes training slow.

- ☐ True    ☐ False

## 18. Artificial Intelligence (AI)

Answers · 答案

A-Level Computer Science

1. **learns patterns from data rather than being programmed step by step**

ML learns from data; deep learning is ML using multi-layer neural networks.

2. **labelled data, learning to map input → label**

Supervised learning trains on labelled examples (e.g. images tagged cat/dog).

3. **a number, while classification predicts a category**

Regression outputs a continuous number (price, temperature); classification outputs a category.

4. **deep learning is part of machine learning, which is part of AI**

AI is the goal, machine learning is the dominant approach to it, and deep learning is machine learning with many-layered networks.

5. **find structure in unlabelled data, such as clusters**

With no labels, unsupervised learning discovers patterns/clusters in the data.

6. **True**

Number out = regression; category out = classification — the two main kinds of supervised prediction.

7. **each town is a node; each road is an edge joining two nodes; each edge carries a weight, the distance or travel time**

Places are nodes and connections are edges. Saying what the weights represent is the third mark.

8. **policy**

It learns by trial and error from rewards, with no labelled examples at all, which is why it suits sequences of decisions.

9. **True**

Stacking many hidden layers lets the network learn richer patterns — training such a network is deep learning.

10. **True**

Backpropagation finds each weight's gradient; the learning rate scales how far down that gradient the weight moves.

# 18.1 Artificial Intelligence (AI)

Name: \_\_\_\_\_ Class: \_\_\_\_\_ Date: \_\_\_\_\_

Total: 32 marks

**KEY WORDS** artificial intelligence 人工智能 • deep learning 深度学习

• machine learning 机器学习 • neural networks 神经网络 • deep neural network 深度神经网络

## Objective

Build the skills to answer exam questions on **Artificial Intelligence** 人工智能 — graphs in AI, neural networks, and the types of machine learning.

You must be able to:

- explain how **graphs** help AI and how an **artificial neural network** works
- compare **supervised**, **unsupervised** and **reinforcement** learning
- explain **deep learning** and **backpropagation**

## 1 Worked examples

Study these first. Each one shows the method for a question type used later — follow the steps and you can do the Practice and Exam-style questions yourself.



### ■ AI, ML and graphs

**Machine learning (ML)** learns patterns **from data** instead of fixed rules; **deep learning (DL)** is ML using deep neural networks. Many AI problems sit on a **graph** — **nodes** (states/places) joined by **edges** (moves, each with a cost). Two algorithms find the **shortest path**:

- **Dijkstra's algorithm** explores outward from the start, always extending the lowest-cost node, until it reaches the goal.
- **A\*** does the same but adds a **heuristic** estimate of the distance **still to go**, which steers the search **toward the goal** — so it usually checks **fewer nodes** and is faster.

### ■ Neural networks

A network has an **input** layer, one or more **hidden** layers, and an **output** layer. Each connection has a **weight**; a neuron sums its weighted inputs (plus a bias), applies an **activation function**, and passes the result on:



*How AI, machine learning and deep learning relate*

### ■ Types of machine learning

- **supervised** — data has **labels**; learns input → label (classification, regression).
- **unsupervised** — **no labels**; finds structure (e.g. **clusters**).
- **reinforcement** — an **agent** acts, gets a **reward**, and learns a **policy** that maximises total reward by trial and error.

**Training** adjusts the weights by **backpropagation**: a forward pass computes the output, a **loss function** measures the error, and the error is propagated **backwards** to update each weight (gradient descent). Deep learning needs **lots of data and compute**.

## 2 Practice

Now apply the methods above.

2.1 State what **machine learning** is. [1]

---

2.2 Name the **three** kinds of layer in a neural network. [2]

---

2.3 In a graph used for **pathfinding**, state what an **edge** represents. [1]

---

2.4 State what a **weight** is in an artificial neuron. [1]

---

## 3 Exam-style questions

3.1 Describe the difference between **supervised** and **unsupervised** learning. [2]

---

---

3.2 Explain how **reinforcement learning** works, using the terms **agent**, **reward** and **policy**. [3]

---

---

---

3.3 Explain why **deep learning** needs a large amount of **data** and **computing power**. [2]

---

---

3.4 State the role of **backpropagation** in training a neural network. [2]

---

---

3.5 State the purpose of **Dijkstra's algorithm** and **A\***, and outline **one** difference between them. [3]

---

---

---

3.6 A delivery company is building software to plan its routes and to read handwritten address labels.

(a) **Identify** one Artificial Intelligence algorithm suitable for finding the shortest distance between two depots on a road network, and **state** what the nodes and the edge weights represent. [3]

---

---

---

---

(b) **Describe** how that algorithm proceeds from the starting node. [3]

---

---

---

---

(c) The company also considers the **A\*** algorithm. **State** the one thing A\* uses that the algorithm in (a) does not, and **explain** the benefit. [3]

---

---

---

---

(d) **Describe** what is meant by **Deep Learning**, and **state** one example of where it is used. [3]

---

---

---

---

(e) **Explain** why Deep Learning suits the handwriting-reading task better than a set of hand-written rules. [3]

---

---

---

---

## Solutions

---

Each answer shows the steps, then the **result**.

### 2.1

- algorithms that **learn patterns from data** (and improve with experience) instead of being programmed with explicit rules

### 2.2

- **input** layer, **hidden** layer(s), **output** layer

### 2.3

- a **move/connection** between states (e.g. a road between two places), often with a **cost/weight**

### 2.4

- a value each input is **multiplied by**, setting how strongly that input affects the neuron's output

### 3.1

- **supervised** learning trains on **labelled** data (each input has a known correct output)
- **unsupervised** learning uses **unlabelled** data and finds structure/clusters itself

### 3.2

- an **agent** takes an **action** in an environment, which changes the state and returns a **reward**
- over many trials it learns a **policy** (strategy)
- that **maximises total reward** by trial and error

### 3.3

- deep networks have **very many weights** to learn
- so they need large datasets to learn good patterns and heavy compute (GPUs) to train in reasonable time

### 3.4

- it propagates the output **error backwards** through the layers
- to **adjust each weight** so the network's output gets closer to the target (gradient descent)

### 3.5

- both find the **shortest path** between two nodes in a weighted graph
- **A\*** uses a **heuristic** to estimate the remaining distance and head toward the goal, so it usually explores **fewer nodes / is faster**
- **Dijkstra's** has no heuristic and explores outward uniformly

### 3.6

#### (a) Dijkstra's algorithm

- the **nodes** are the depots and road junctions
- the **edge weights** are the distances, or the travel times, along each road between them

(b) it gives the start node a distance of zero and every other node infinity

- it then repeatedly takes the **unvisited node with the smallest known distance**, and for each of its neighbours checks whether going via this node is shorter than the distance recorded so far, updating it if so
- the node is then marked visited and never reconsidered; the process ends when the destination is visited

(c) A\* also uses a **heuristic** — an estimate of the remaining distance from each node to the goal, such as the straight-line distance

- it therefore explores nodes that head **towards** the destination first
- so it usually examines far fewer nodes than Dijkstra while still finding the shortest route, provided the heuristic never overestimates

(d) Deep Learning uses an **artificial neural network with many hidden layers**

- trained on large amounts of data, where each layer learns features built from the layer before it
- examples: speech recognition, image and face recognition, machine translation, self-driving vehicles

(e) handwriting varies enormously between people, so the rules that distinguish a 7 from a 1 cannot be written out completely

- a network **learns** the distinguishing features from thousands of labelled examples instead of being told them
- and it generalises to handwriting it has never seen, which a fixed rule set cannot do

Name: \_\_\_\_\_

A-Level Computer Science: **w25 31**

7 (a) Identify **one** Artificial Intelligence (AI) algorithm to find the shortest distance between two points on a graph.

.....  
..... [1]

(b) Describe Deep Learning.

.....  
.....  
.....  
.....  
.....  
.....  
.....  
.....  
.....  
..... [5]

**9 (a)** Deep Learning is a form of Machine Learning.

**(i)** State **one** example where Deep Learning is used.

.....  
 ..... [1]

**(ii)** Identify how Deep Learning can be made more effective.

.....  
 ..... [1]

**(b)** Describe the back propagation of errors method in Machine Learning.

.....  
 .....  
 .....  
 .....  
 .....  
 .....  
 .....  
 ..... [4]

**10** An exception is an error that may cause a program to halt unexpectedly.

**(a)** Describe how program termination due to an exception can be avoided.

.....  
 .....  
 .....  
 .....  
 ..... [2]

**(b)** Identify **two** possible causes of exceptions.

1 .....  
 .....  
 2 .....  
 .....  
 ..... [2]

**3** A computer program uses a digital camera to read the words on an item.

The program can read the words that are written on the item, translate the words to a chosen language and then output the words as audio. For example, when used in a supermarket, the program can output the words written on the labels on products.

The program uses Artificial Intelligence (AI).

**(a)** Explain how AI is used in the computer program described.

.....  
 .....  
 .....  
 .....  
 .....  
 ..... [4]

**(b)** State **two** social benefits of the use of AI in the computer program described.

1 .....  
 .....  
 .....  
 .....  
 .....  
 ..... [2]

**(c)** The computer program is released under a commercial software licence.

**(i)** Describe the features of a commercial software licence.

.....

.....

.....

.....

.....

.....

.....

[3]

**(ii)** Explain the reasons why an open source software licence might **not** be appropriate for the computer program described.

.....

.....

.....

.....

.....

.....

.....

[3]

**6 (a)** Describe the structure of a graph as used in an Artificial Intelligence (AI) system.

.....

.....

.....

.....

[2]

**(b)** Explain how supervised learning and unsupervised learning differ from each other.

.....

.....

.....

.....

.....

.....

.....

[4]

**10 (a)** State the purpose of the A\* and Dijkstra's algorithms.

.....  
..... [1]

**(b)** Outline the difference between the A\* and Dijkstra's algorithms.

.....  
.....  
.....  
..... [2]

**(c)** Explain how unsupervised learning takes place in machine learning.

.....  
.....  
.....  
.....  
.....  
..... [3]

**6** A car park system uses a camera to record the registration number of each car as it enters and leaves the car park.

Explain how artificial intelligence is used in the car park system to identify the car's registration number.

.....  
.....  
.....  
.....  
.....  
.....  
.....  
..... [4]