

COMP1110 COMPUTING AND DATA SCIENCE IN EVERYDAY LIFE

Quiz 2026

1. Which of the following best describes supervised learning?
 - A. Learning by trial and error through rewards and penalties
 - B. Learning from labeled examples to predict outputs for new inputs
 - C. Exhaustively searching all possible solutions without using data
 - D. Grouping unlabeled data into clusters without target labels

Solution: In supervised learning, the algorithm is trained on a labeled dataset — each training example pairs an input with its correct output (the “label”). The model learns to map inputs to outputs so that it can make predictions on unseen data. Learning by trial and error with rewards describes reinforcement learning, while grouping unlabeled data describes unsupervised (clustering) methods.

2. A messaging app reads a message and outputs “URGENT” or “NOT URGENT”. This is an example of which supervised learning task?
 - A. Classification
 - B. Regression
 - C. Reinforcement learning
 - D. Clustering

Solution: Classification assigns each input to one of a finite set of discrete categories (here: URGENT or NOT URGENT). Regression predicts a continuous numerical value; reinforcement learning involves an agent acting in an environment; clustering groups unlabeled data without predefined classes.

3. A machine learning model achieves 95% accuracy on the training set but only 60% accuracy on the test set. This is best described as:
 - A. Backfitting
 - B. Underfitting
 - C. Overfitting
 - D. Retrofitting

Solution: Overfitting occurs when a model learns the training data too well — including its noise and idiosyncrasies — so that it performs well on training data but generalises poorly to unseen test data. Underfitting is the opposite: the model is too simple to capture the patterns even in the training data. “Backfitting” and “retrofitting” are not standard ML terms in this context.

4. Which of the following is a common advantage of decision trees in machine learning?
 - A. They always produce the most accurate model
 - B. Small changes in the training data never affect the learned tree
 - C. They are relatively easy for humans to understand and interpret
 - D. They are guaranteed not to overfit the training data

Solution: Decision trees are interpretable models: the sequence of if/else splits can be read and understood by non-experts, making them useful when explainability matters. They do not always achieve the highest accuracy (ensemble methods often outperform a single tree), they are sensitive to small changes in training data (high variance), and they can overfit without pruning or depth constraints.

5. In reinforcement learning, what is the agent’s main objective?
 - A. Group similar states together without feedback
 - B. Correctly classify every example in the training set
 - C. Minimize the number of labeled training examples needed
 - D. Learn a policy that maximizes total reward

Solution: A reinforcement learning agent interacts with an environment, receiving rewards or penalties for its actions. Its goal is to discover a *policy* — a mapping from states to actions — that maximizes the cumulative (total) reward over time. Grouping states is clustering; classifying labeled examples is supervised learning; minimizing labeled data is the goal of active learning.

6. Why did neural networks become widely practical only in the last decade, despite being invented much earlier?
- A. Commercial use of neural networks was previously illegal
 - B. No training data existed before the 2010s
 - C. They require heavy computation, and earlier hardware was often too slow
 - D. The training algorithm for neural networks was discovered only in the last decade

Solution: Backpropagation, the key training algorithm for neural networks, was known since the 1980s. What changed in the last decade was the availability of powerful GPUs that can perform the massive matrix computations deep networks require, combined with large datasets (e.g., ImageNet) and advances like improved activation functions and regularisation techniques. Hardware limitations — not legal barriers, data absence, or missing algorithms — were the primary bottleneck.

7. Which of the following is a valid IPv4 address?
- A. 888.888.888.888
 - B. 88:88:88:88:88:88
 - C. 8888:8888:8888:8888:8888:8888:8888:8888
 - D. 8.8.8.8

Solution: IPv4 addresses consist of four decimal octets separated by dots, where each octet is in the range 0–255.

8 . 8 . 8 . 8 is Google’s public DNS server and is a valid IPv4 address.

888 . 888 . 888 . 888 is invalid (octets exceed 255); 88 : 88 : 88 : 88 : 88 : 88 is a MAC address format; 8888 : 8888 : . . . is an IPv6 address using colon-separated hexadecimal groups.

8. What is the primary purpose of a DNS server?
- A. To encrypt data sent between a browser and a server
 - B. To manage the physical cabling and electrical signals in a network
 - C. To translate domain names into IP addresses
 - D. To assign IP addresses to newly connected devices

Solution: The Domain Name System (DNS) acts as the Internet’s “phone book”: when you type a domain name such as `www.example.com`, the DNS server resolves it to the corresponding IP address so your device can establish a connection. Encrypting traffic is done by TLS/SSL; DHCP assigns IP addresses to new devices; physical signalling is a hardware/physical-layer concern.

9. (Cancelled) Which layer of the network model can provide end-to-end reliability between applications, for example through protocols such as TCP?
- A. Transport layer
 - B. Data link layer
 - C. Network layer
 - D. Physical layer

Solution: TCP (Transmission Control Protocol) operates at the *Transport* layer (Layer 4 in the OSI model). It provides end-to-end reliability by establishing connections, sequencing packets, acknowledging receipt, and retransmitting lost segments. The Network layer routes packets between hosts; the Data link layer handles framing on a single link; the Physical layer deals with bit transmission over a medium.

10. (Cancelled) Which of the following is *not* typically used as a transmission medium in standard computer networks?
- A. Sound waves
 - B. Electromagnetic waves
 - C. Copper wires
 - D. Optical fibers

Solution: Standard computer networks transmit data over copper wires (electrical signals), optical fibres (light pulses), and electromagnetic waves (radio/microwave for Wi-Fi and cellular). Sound waves are not a practical transmission medium for computer networks: they travel slowly through air (≈ 343 m/s) compared to light-based media and are susceptible to noise.

11. (Cancelled) What is the main advantage of UDP compared with TCP?
- A. Guaranteed delivery
 - B. Larger packet size
 - C. Lower overhead and lower latency
 - D. Better security

Solution: UDP (User Datagram Protocol) is a connectionless protocol that does not perform handshaking, acknowledgements, or retransmissions. This means it has much lower overhead and latency than TCP, making it suitable for real-time applications such as video streaming, VoIP, and online games where a small amount of packet loss is preferable to the delay caused by retransmission. UDP does *not* guarantee delivery, does not offer larger packet sizes, and provides no built-in security.

12. Why was IPv6 introduced?
- A. IPv6 is faster because its addresses are shorter
 - B. IPv4 addresses have been exhausted
 - C. IPv6 removes the need for DNS
 - D. IPv6 guarantees encryption for all traffic

Solution: IPv4 uses 32-bit addresses, providing approximately 4.3 billion unique addresses. The rapid growth of internet-connected devices exhausted the IPv4 address pool (IANA issued the last blocks in 2011). IPv6 uses 128-bit addresses, providing $2^{128} \approx 3.4 \times 10^{38}$ unique addresses. IPv6 addresses are longer (not shorter) than IPv4; DNS is still used with IPv6; and while IPv6 originally planned mandatory IPsec, encryption is not guaranteed by default in practice.

13. Which of the following is the best primary key for a “Student” table in a university database?
- A. The student’s university ID number
 - B. The student’s full name
 - C. The student’s date of birth
 - D. The student’s home address

Solution: A primary key must uniquely identify each row and should be stable over time. A university-assigned student ID number is unique per student by definition and does not change. Full names are not unique (multiple students can share the same name); dates of birth are not unique; home addresses change over time and are not unique.

14. Which of the following is a valid SQL statement?
- A. `SELECT Name FROM Student WHERE Year = 2;`
 - B. `Student.findAll(Year==2).get(Name);`
 - C. `GET /api/student/2/name`
 - D. `<Student>Name<Year>2</Year></Student>`

Solution: The first option is syntactically correct SQL: it retrieves the Name column from the Student table for all rows where Year equals 2. The second resembles an ORM (object-relational mapping) method chain, but is not standard SQL. The third is an HTTP REST API call; the fourth is XML.

15. Why is a foreign key useful in a relational database?
- A. It helps ensure that a value in one table matches an existing row in a related table
 - B. It encrypts linked data so only authorised users can read it
 - C. It stores all related data in one table so joins are never needed
 - D. It makes every row in a table unique, even if duplicate values are entered

Solution: A foreign key enforces *referential integrity*: it constrains a column’s values to match primary key values in another (or the same) table, ensuring that relationships between tables are consistent and dangling references do not exist. Foreign keys do not encrypt data, do not eliminate joins, and ensuring row uniqueness is the role of a primary key or unique constraint.

16. Why does a database management system need concurrency control?
- A. To compress data before storing it on disk
 - B. To speed up queries by caching recent results
 - C. To automatically remove missing values
 - D. To prevent data corruption when multiple users modify the same data at the same time

Solution: When multiple transactions execute concurrently without coordination, anomalies such as lost updates, dirty reads, and non-repeatable reads can occur, leaving the database in an inconsistent state. Concurrency control mechanisms (e.g., locking, timestamps, multi-version concurrency control) serialise conflicting operations to ensure isolation and data integrity.

17. In a database, what does the term “schema” refer to?

- A. A backup copy of the database
- B. A query used to retrieve data
- C. The blueprint that defines tables and their attributes
- D. The actual records stored in all tables

Solution: A database *schema* is the logical structure or blueprint of the database: it defines the tables, their columns, data types, and constraints — but not the actual data stored. The actual stored data is called the database *instance* or *state*. A schema is analogous to a class definition in OOP, while the instance is analogous to the objects created from that class.

18. What does “deterministic” mean in the context of algorithms?

- A. The algorithm always finishes within finite time
- B. The algorithm uses random guesses to speed up execution
- C. The algorithm always finds the globally optimal solution
- D. The algorithm always produces the same output for the same input

Solution: A *deterministic* algorithm is one whose behaviour is fully determined by its input: given the same input, it will always follow the same execution path and produce the same output. This is in contrast to a *randomised* algorithm, which may produce different outputs or follow different paths on the same input due to random choices.

19. Using the greedy coin-change strategy with denominations { $\$1$, $\$2$, $\$5$, $\$10$ }, how many coins does the algorithm choose to make $\$37$?

- A. 6
- B. 4
- C. 5
- D. 7

Solution: The greedy algorithm always selects the largest denomination that does not exceed the remaining amount:

- 1. $\$10 \times 3 = \30 (remaining: $\$7$, coins so far: 3)
- 2. $\$5 \times 1 = \5 (remaining: $\$2$, coins so far: 4)
- 3. $\$2 \times 1 = \2 (remaining: $\$0$, coins so far: 5)

Total: 5 coins.

20. Which statement correctly describes the binary search tree property?

- A. Every node must have exactly two children
- B. The root node always holds the largest value in the tree
- C. All values in the left subtree are smaller than the node, and all values in the right subtree are larger than the node
- D. All values in both subtrees must be smaller than the node

Solution: The BST property requires that for any node with value v : every value in the *left* subtree is less than v , and every value in the *right* subtree is greater than v . This ordering enables efficient $O(\log n)$ search, insertion, and deletion on a balanced tree. Nodes need not have two children (leaves have zero), and the root holds neither the largest nor the smallest value in general.

21. For a balanced binary search tree with N nodes, the worst-case number of nodes inspected during a search is approximately:

- A. $\frac{1}{N}$
- B. $\log_2 N$
- C. \sqrt{N}
- D. N

Solution: A balanced BST has height $\lfloor \log_2 N \rfloor$, so a search visits at most $\lfloor \log_2 N \rfloor + 1 \approx \log_2 N$ nodes before reaching the target or a null pointer. This logarithmic complexity is why balanced BSTs (e.g., AVL trees, red-black trees) are preferred over linear structures like linked lists for search operations. N would describe the worst case for an unbalanced (degenerate) tree.

22. What is a common disadvantage of greedy algorithms?
- A. They can solve only coin-related problems
 - B. They are generally very difficult to implement
 - C. They do not always produce the correct or optimal result for some problems
 - D. They are slow because they must explore all possible choices

Solution: Greedy algorithms make the locally optimal choice at each step without reconsidering previous decisions. While this makes them fast and simple, they do not always produce a globally optimal solution. A classic counter-example is the coin-change problem with non-standard denominations (e.g., $\{1, 3, 4\}$ coins for amount 6: greedy picks $4+1+1=3$ coins, but the optimal is $3+3=2$ coins).

23. In the theory of computation, the Halting Problem shows that:
- A. Not every computational problem can be solved by an algorithm
 - B. Infinite loops occur only in poorly written software
 - C. A computer will eventually freeze if a program enters an infinite loop
 - D. Every program will eventually halt if given enough memory

Solution: Alan Turing proved in 1936 that no general algorithm can determine, for an arbitrary program and input, whether that program will halt or run forever. This is the *Halting Problem*, and it establishes that there exist *undecidable* problems — problems for which no algorithmic solution can exist, independent of hardware speed or memory.

24. Which of the following is an example of a *non-functional* requirement in software engineering?
- A. Users can reset their password by email
 - B. The web page must load in under 2 seconds
 - C. Customers can place vegetarian food orders through the app
 - D. Students can search for courses by instructor

Solution: *Functional* requirements describe *what* a system does (specific behaviours and features). *Non-functional* requirements describe *how well* the system performs — qualities such as performance, reliability, security, and usability. “Load in under 2 seconds” is a *performance* non-functional requirement. The other three options all describe specific user-facing features (functional requirements).

25. In software engineering, what is the main purpose of a penetration test?
- A. To test whether users can navigate deep menu structures efficiently
 - B. To identify security vulnerabilities before attackers can exploit them
 - C. To measure how thoroughly automated tests cover all branches of a program
 - D. To measure how well Wi-Fi signals pass through walls in an office building

Solution: A penetration test (“pen test”) involves authorised security professionals attempting to exploit a system’s vulnerabilities in the same way a malicious attacker would, in order to discover and remediate weaknesses before real attackers do. Usability testing evaluates user navigation; code coverage measures test completeness; Wi-Fi signal measurement is a network infrastructure task.

26. What is the main purpose of version control in software development?
- A. To track changes to code over time and support collaboration

- B. To automatically find and fix bugs in source code
- C. To convert design documents directly into working software
- D. To improve the visual appearance of the user interface automatically

Solution: Version control systems (e.g., Git) maintain a history of every change made to a codebase. This allows developers to review past changes, revert to earlier versions, branch for experimental features, and merge contributions from multiple collaborators without overwriting each other's work.

27. In a Gantt chart, what is a dependency between two tasks?
- A. Two tasks grouped under the same milestone
 - B. A relationship in which one task cannot start until another task has finished
 - C. Two tasks assigned to the same developer
 - D. Two tasks with the same deadline

Solution: A task dependency (“finish-to-start”) means that a successor task cannot begin until its predecessor has completed. Gantt charts visualise these dependencies with arrows connecting tasks, helping project managers identify the *critical path* — the longest chain of dependent tasks that determines the minimum project duration.

28. Which of the following best describes a common trade-off of pair programming?
- A. It splits one task into two independent coding tasks, which speeds development but makes integration harder
 - B. It uses two developers on one task, which may increase effort upfront but can improve code quality and reduce defects
 - C. It requires writing production code and test code at the same time, which improves coverage but doubles the codebase
 - D. It runs the same application on two identical servers, which improves reliability but doubles hosting cost

Solution: In pair programming, two developers work together at one workstation: one writes code (the “driver”) while the other reviews each line (the “navigator”). The immediate trade-off is that two people are needed for what one person could do alone. However, research shows that pair programming reduces defect rates, improves design quality, and facilitates knowledge sharing.

29. In data analysis, what is an outlier?
- A. A data point that lies unusually far from the others
 - B. The target variable being predicted
 - C. A missing value recorded as blank
 - D. A duplicate record

Solution: An *outlier* is an observation that deviates markedly from the overall pattern of the dataset. Outliers may be caused by measurement errors, data entry mistakes, or they may represent genuine but rare phenomena. They are important to detect because they can disproportionately affect summary statistics (e.g., mean) and distort model training.

30. What is the main purpose of exploratory data analysis (EDA)?
- A. To train and evaluate the final predictive model
 - B. To understand a dataset by summarising patterns, spotting anomalies, and checking data quality
 - C. To encrypt and secure sensitive information before storage
 - D. To gather new raw data from external sources

Solution: EDA is the process of examining a dataset before formal modelling — using summary statistics, visualisations (histograms, box plots, scatter plots, etc.), and correlation analysis to understand the data's structure, distribution, and relationships. It helps identify missing values, outliers, data quality issues, and promising features before committing to a model.

31. Why are precision and recall often reported in addition to accuracy for classification tasks?
- A. Accuracy is used only for training data, while precision and recall are used only for test data
 - B. Precision and recall are always larger than accuracy
 - C. Accuracy alone can be misleading, especially when the classes are imbalanced
 - D. Precision and recall automatically remove extreme data points before evaluation

Solution: Consider a dataset where 99% of emails are not spam. A classifier that predicts “not spam” for every email achieves 99% accuracy yet is entirely useless. Precision (fraction of predicted positives that are truly positive) and recall (fraction of actual positives that are correctly identified) capture the classifier’s performance on the minority class.

32. In data science, which of the following statements about causation is true?
- A. If two variables are correlated, then one must cause the other
 - B. Causation means that changing one variable directly affects another variable
 - C. If two variables move in opposite directions, they cannot be causally related
 - D. Causation can always be proven just by observing enough data points

Solution: Causation means that a direct mechanism exists whereby changing variable *A* produces a change in variable *B*. Correlation alone does not imply causation: two variables may be correlated due to a common cause (confounding variable) or by coincidence. Establishing causation requires controlled experiments (e.g., randomised controlled trials) or careful causal inference methods, not merely more observational data.

33. An investor studies only highly successful traders and notices that many of them took large risks. He ignores traders who took large risks and failed, then concludes that taking large risks is a reliable way to become rich. What statistical misconception does this best illustrate?
- A. No misconception; the conclusion is justified
 - B. Base rate fallacy
 - C. Simpson’s paradox
 - D. Survivorship bias

Solution: *Survivorship bias* occurs when analysis focuses exclusively on the subjects that “survived” a selection process, ignoring those that did not. Here, only successful risk-takers are observed; the many who took equal risks and failed are invisible. Correcting for survivorship bias requires explicitly including the full population, not just the “survivors”.

34. An active sensor emits its own energy and measures the reflection. Which of the following is an active sensor?
- A. Radar
 - B. Thermometer
 - C. Optical camera
 - D. Microphone

Solution: An *active* sensor emits energy into the environment and measures the return signal. Radar emits radio waves and detects their reflections from objects, which is why it can operate in darkness, rain, or fog. A thermometer measures ambient temperature passively; an optical camera captures ambient or reflected light it did not emit; a microphone passively captures sound waves.

35. Which of the following tasks is LiDAR especially suitable for?
- A. Building a precise 3D map of the surroundings
 - B. Reading QR codes printed on a poster
 - C. Recording a sunset with accurate colour detail
 - D. Identifying the colour of a traffic light

Solution: LiDAR (Light Detection And Ranging) emits rapid pulses of laser light and measures the time-of-flight of each return pulse, generating a dense *point cloud* that encodes the precise distance to thousands of surfaces per second. This makes LiDAR ideal for high-resolution 3D mapping — used in autonomous vehicles, surveying, and robotics. LiDAR does not capture colour or texture information, so it is not suited to reading QR codes, photographing scenes, or distinguishing traffic light colours.

36. A robot arm must reach a specific point in space. You know the desired end-effector position and need to determine the required joint angles. What is this problem called?
- A. Inverse kinematics B. Forward kinematics C. Delayed kinematics D. Fuzzy kinematics

Solution: *Forward kinematics* computes the end-effector position given known joint angles. *Inverse kinematics* (IK) solves the reverse problem: given a desired end-effector position, find the joint angles that achieve it. IK is generally harder than forward kinematics and may have multiple solutions, no solution, or require iterative numerical methods.

37. Which of the following is primarily used as an actuator?
- A. GPS B. Servo motor C. QR code reader D. LiDAR

Solution: An *actuator* converts energy into mechanical motion, causing a robot to interact physically with its environment. A *servo motor* is a powered actuator that can rotate to a commanded angle with feedback control, making it the standard actuator in robotic joints and RC mechanisms. GPS, QR code readers, and LiDAR are all *sensors* — they gather information about the environment but do not produce motion.

38. What is the plaintext of the ciphertext “ABC” under ROT13, that is, a Caesar cipher with shift 13?
- A. NOP B. CBA C. XYZ D. 123

Solution: ROT13 shifts each letter forward by 13 positions in the alphabet (wrapping around at Z). Since $13 + 13 = 26$ (one full cycle), encryption and decryption are the same operation:

$$A \xrightarrow{+13} N, \quad B \xrightarrow{+13} O, \quad C \xrightarrow{+13} P$$

Therefore the plaintext of ciphertext “ABC” under ROT13 is **NOP**.

39. Why is a monoalphabetic substitution cipher rarely used in modern cryptography?
- A. It requires specialised computer hardware for encryption
- B. It is too slow because of its mathematical complexity
- C. It can often be broken by frequency analysis if the ciphertext is long enough
- D. It works only for messages written in English

Solution: In a monoalphabetic substitution cipher each plaintext letter is always replaced by the same ciphertext letter. Since natural language has a well-known letter frequency distribution (e.g., ‘E’ is most common in English), an attacker who observes which ciphertext symbol appears most frequently can guess the key iteratively. This *frequency analysis* attack, used since the 9th century by Al-Kindī, renders monoalphabetic ciphers insecure for any sufficiently long message.

40. In a public-key cryptosystem such as RSA, which of the following must be kept secret?
- A. Nothing; everything is public in a public-key cryptosystem
- B. The private key
- C. The public key
- D. The ciphertext

Solution: In public-key (asymmetric) cryptography, each party has a *key pair*: a **public key** that is freely distributed and a **private key** that must be kept secret. Anyone can encrypt a message using the recipient’s public key, but only the recipient holding the corresponding private key can decrypt it. Exposing the private key would allow any attacker to decrypt messages or forge digital signatures.

41. What does the security of the RSA cryptosystem primarily rely on?
- A. The difficulty of integer factorisation
- B. The difficulty of recognising human faces

- C. The difficulty of dividing by zero
- D. The difficulty of sorting very large arrays

Solution: RSA key generation involves selecting two large prime numbers p and q and publishing their product $n = pq$ as part of the public key. Breaking RSA requires factoring n back into p and q . Integer factorisation of large numbers (e.g., 2048-bit) is believed to be computationally infeasible with current classical computers. Note that sufficiently large quantum computers running Shor’s algorithm could factor large integers efficiently, which is why post-quantum cryptography is an active research area.

42. Which type of malware encrypts a victim’s files and demands payment for the decryption key?
- A. Zombie process
 - B. Ransomware
 - C. Adware
 - D. Honeypot

Solution: *Ransomware* is a class of malware that encrypts the victim’s files (or locks their device) and demands a ransom — typically in cryptocurrency — in exchange for the decryption key. High-profile examples include WannaCry (2017) and Ryuk. Adware displays unwanted advertisements; a honeypot is a security decoy (not malware); a zombie process is a Unix/Linux concept for a terminated child process not yet reaped by its parent.

43. How does a DDoS (Distributed Denial of Service) attack typically work?
- A. By intercepting communications between two parties
 - B. By installing hidden software that monitors user activity
 - C. By injecting malicious code into database queries
 - D. By overwhelming servers with traffic and making them unavailable

Solution: A DDoS attack floods the target server or network with an enormous volume of requests, exhausting its bandwidth, CPU, or memory so that legitimate users cannot access the service. The “distributed” aspect means traffic originates from many compromised machines (a botnet), making it difficult to block a single source IP. Intercepting communications describes a man-in-the-middle attack; monitoring user activity describes spyware; injecting malicious code into queries describes SQL injection.

44. What does “authentication” mean in cybersecurity?
- A. Encrypting data so that no one can read it during transmission
 - B. Proving that a user is really who they claim to be
 - C. Making a backup copy of a system in case of failure
 - D. Deciding what actions a user is allowed to perform after logging in

Solution: *Authentication* answers the question “Who are you?” — verifying a claimed identity, typically via a password, biometric, hardware token, or cryptographic certificate. It is distinct from *authorisation* (determining what an authenticated user is permitted to do) and *confidentiality* (keeping data secret via encryption).

45. Which language is most commonly used to implement logic running inside a web browser?
- A. C++
 - B. Perl
 - C. JavaScript
 - D. Kotlin

Solution: JavaScript is the only programming language natively supported by all major web browsers. It runs client-side, enabling dynamic page updates, user interaction handling, and communication with servers (via APIs and fetch/AJAX). C++ can target browsers via WebAssembly but is not the primary browser language; Perl is a server-side scripting language; Kotlin is primarily used for Android and JVM development.

46. What does the “front end” of a web application refer to?
- A. The database server that stores persistent records
 - B. The backup system that archives logs
 - C. The router that forwards Internet packets
 - D. The part that runs in the user’s browser and handles what the user sees and interacts with

Solution: The *front end* (client side) encompasses everything the user directly sees and interacts with: HTML for structure, CSS for styling, and JavaScript for behaviour, all executed in the browser. The *back end* (server side) handles business logic, data storage, and authentication on a server. Databases, backup systems, and network routers are all back-end or infrastructure components.

47. What is the main purpose of a Content Delivery Network (CDN)?

- A. To deliver static content from servers closer to users, often improving load time and reducing latency
- B. To ensure that websites never experience server failures
- C. To replace the database by storing all application data in the browser
- D. To encrypt all traffic between users and servers

Solution: A CDN is a geographically distributed network of servers (“edge nodes” or “points of presence”) that cache and serve static assets — images, CSS, JavaScript files, videos — from the location nearest to the user. By reducing the physical distance between the server and the client, a CDN lowers round-trip latency and improves page load times. CDNs do not prevent all failures, cannot replace a database, and do not inherently encrypt traffic (that is TLS’s role).

48. What does “cross-platform mobile development” mean?

- A. Writing one codebase that can be used to build apps for both Android and iOS
- B. Building an app that supports multiple spoken languages, such as English and Chinese
- C. Ensuring an app’s data is automatically backed up across multiple cloud servers
- D. Designing a user interface that automatically fits both smartphones and tablets

Solution: Cross-platform mobile development allows developers to write a single codebase that compiles or runs on both Android and iOS. Frameworks such as Flutter (Dart), React Native (JavaScript), and Xamarin (.NET) achieve this. Supporting multiple languages is called *localisation*; multi-server backups are a cloud storage concern; fitting different screen sizes is *responsive design*.

49. What is JavaScript minification?

- A. Reducing JavaScript file size using general-purpose compression methods such as gzip
- B. Translating JavaScript into a low-level language such as Assembly for faster execution
- C. Splitting one JavaScript file into many smaller files for modularity
- D. Reducing JavaScript file size by removing unnecessary characters and shortening identifiers while preserving functionality

Solution: Minification removes whitespace, comments, and newlines, and replaces long variable and function names with short ones (e.g., `calculateTotalPrice` → `a`), without changing the program’s behaviour. This is distinct from *compression* (gzip applies lossless byte-level compression and is a separate step) and *code splitting* (dividing a bundle into smaller chunks for lazy loading).

50. What does “edge computing” refer to?

- A. Processing data close to where it is generated instead of sending everything to a cloud server
- B. Compressing files at the edge of a network by removing metadata
- C. A mobile design style that places buttons near the border of the screen
- D. Running graphics calculations only near the edge of a screen

Solution: Edge computing moves computation and data processing closer to the source of data generation (e.g., IoT sensors, smart cameras, factory equipment) rather than routing all data to a centralised cloud data centre. This reduces latency (important for real-time decisions), lowers bandwidth usage, and improves privacy by keeping sensitive data local.