CUET UG 2026 Computer Science / Information Practices May 24 Shift 1 Question Paper with Solutions PDF

Question Paper with Solutions

Conducted by the National Testing Agency (NTA)

Time Allowed: 60 Minutes **Maximum Marks: 250**

General Instructions

(i) This is the Computer Science / Information Practices question paper of CUET (UG) 2026, conducted by the National Testing Agency (NTA) in Computer Based Test (CBT) mode.

(ii) The paper contains 85 multiple choice questions in 3 sections. Section A Common is for every candidate; attempt either Section B1: Computer Science or Section B2: Informatics Practices, as instructed at the start of each section.

(iii) The duration of this test paper is 60 minutes. The maximum score is 250 marks.

(iv) Each question has four options. Select the one correct answer for each question.

(v) Marking scheme: $+5$ marks for a correct answer, -1 mark for an incorrect answer, and 0 marks for an unanswered question.

(vi) For questions based on a passage or table, read the given material carefully and answer the questions that follow.

Section A Common (75 marks)

1. _____ helps to enforce integrity and identify relationships among tables in a database. [5 marks]

(A) Row

(B) Tuple

(C) Key

(D) Column

Correct Answer: (C) Key

Solution:

Concept:

- A key is an attribute or a set of attributes used to identify records or establish links between relations.
- Primary-key constraints enforce unique, non-null identification. Foreign-key constraints maintain referential integrity between related tables.

Step 1: Identify the function described The statement combines two roles: preserving the validity of related data and expressing relationships among tables. Keys supply these roles. For example, a department identifier can be a primary key in a department table and a foreign key in a student table. The foreign key refers to an existing department and connects the two relations.

Step 2: Eliminate the other terms A row, also called a tuple, is one record of a relation; it is not itself the database mechanism that enforces a relationship. A column is an attribute or field. Although a key is formed from one or more columns, an arbitrary column need not be unique or refer to another table. The general term covering the required constraint mechanism is therefore key.

Step 3: Conclusion Keys help enforce integrity and identify relationships among tables. Hence, the correct option is (C).

Quick Tip to Solve: Link uniqueness with a primary key and valid cross-table references with a foreign key; both belong to the broader category of keys.

2. Which of the following can uniquely identify the tuples of a relation in a database?

(A) Foreign Key

(B) Primary Key

(C) Candidate Key

(D) Composite Primary Key

Choose the correct answer from the options given below: [5 marks]

(A) (A), (B) and (D) only

(B) (A), (C) and (D) only

(C) (A), (B), (C) and (D)

(D) (B), (C) and (D) only

Correct Answer: (D) (B), (C) and (D) only

Solution:

Concept:

- A candidate key is a minimal set of attributes that uniquely identifies each tuple. One candidate key is selected as the primary key.
- A composite primary key uses several attributes together. A foreign key refers to another relation and need not be unique in the referencing relation.

Step 1: Identify the keys that guarantee uniqueness Statement (B) is true because a primary key uniquely identifies every row. Statement (C) is true because each candidate key is capable of unique identification, whether or not it is the chosen primary key. Statement (D) is also true: a composite primary key consists of multiple attributes, but their combined value is unique. Its individual attributes need not be unique separately.

Step 2: Test the foreign-key statement and select the combination Statement (A) is false as a general claim. Many students may have the same department identifier as a foreign key, for example, so that value alone cannot distinguish the student records. A foreign key may happen to be unique under an additional constraint, but foreign-key status does not guarantee it. Therefore reject every combination containing (A) and retain the combination (B), (C) and (D) only.

Step 3: Conclusion Primary Key, Candidate Key and Composite Primary Key can uniquely identify tuples; a foreign key alone does not guarantee uniqueness. Hence, the correct option is (D).

Quick Tip to Solve: Test whether uniqueness is guaranteed by the key definition; a foreign key guarantees a reference, not unique identification in its own table.

3. Rama stores the data in a specific format or structure in a file. If the structure or format itself is changed, all the existing application programs accessing that file also need to be changed. Otherwise, the programs may not work correctly. What is this concept called? [5 marks]

(A) Data Sharing

(B) Data Inconsistency

(C) Data Dependency

(D) Data Isolation

Correct Answer: (C) Data Dependency

Solution:

Concept:

- Program-data dependency occurs when application code depends on the physical structure or format of stored data.
- Data independence means that a change in data organisation can be accommodated without corresponding changes to every application.

Step 1: Trace the effect of changing the file Rama changes the storage format or structure, and the application programs must then be changed to continue accessing the file correctly. This means those programs have assumptions about the old record layout built into them. A change to the data representation therefore propagates to the programs: this is data dependency.

Step 2: Distinguish the alternative concepts Data sharing concerns access to the same data by different users or applications. Data inconsistency concerns conflicting values for the same fact. Data isolation concerns data being scattered across separate files or formats and consequently being difficult to combine or access. None names the explicit need to rewrite programs when the stored format changes.

Step 3: Conclusion The programs depend on the file structure, so the concept is Data Dependency. Hence, the correct option is (C).

Quick Tip to Solve: When the wording says that changing a file format forces changes in programs, identify program-data dependency.

4. _____ is known as design of database. [5 marks]

(A) Constraints

(B) Keys

(C) Schema

(D) Tuples

Correct Answer: (C) Schema

Solution:

Concept:

- A database schema describes the logical design or structure of a database, including relations and their attributes.

- The schema is the blueprint; the actual records stored at a particular time form a database instance.

Step 1: Identify the database blueprint The question asks for the name of the database design. Its schema specifies how data is organised: the tables, their columns and data types, and the structural relationships or constraints. For instance, defining a STUDENT table with Name, Class and Marks describes part of a schema, independently of which students are currently stored.

Step 2: Rule out components and records Constraints are rules restricting acceptable data, such as a uniqueness requirement; they are elements of a design rather than its overall name. Keys identify records or relate tables, so they too are components of the structure. Tuples are the individual rows containing current data. They may change without changing the design itself.

Step 3: Conclusion The design of a database is its schema. Hence, the correct option is (C).

Quick Tip to Solve: Use schema for the database blueprint and instance for the data present at a particular moment.

5. STUDENT

Name	Class	DOB	Gender	City	Marks
Rohan	XI	2008	M	Delhi	85
Priya	XII	2006	F	Mumbai	92
Ankit	XI	2007	M	Delhi	88
Nidhi	XII	2006	F	Mumbai	90
Aarav	IX	2009	M	Kolkata	78
Simran	XI	2007	F	Chandigarh	94
Rahul	XII	2006	M	Pune	80

From the above table 'STUDENT', which query will return all the students from class 'XI'? [5 marks]

(A) SELECT * FROM STUDENT WHERE Class ='XI';

(B) SELECT * FROM STUDENT WHERE LIKE 'XI';

(C) SELECT Name, DOB FROM STUDENT WHERE Class >'XI';

(D) SELECT * FROM STUDENT WHERE Class IN ='XI';

Correct Answer: (A) SELECT * FROM STUDENT WHERE Class ='XI';

Solution:

Concept:

- SELECT chooses output columns, FROM identifies the relation, and WHERE filters individual rows.
- The equality operator tests an exact value; a text literal such as XI is enclosed in single quotation marks.

Step 1: Translate the requirement into a predicate All columns of the matching students are requested, so SELECT * is appropriate. The source is STUDENT and the condition is Class = 'XI'. Applying it to the supplied table selects Rohan, Ankit and Simran; the students in XII and IX do not satisfy the equality test.

Step 2: Compare the query syntax Option (A) places the exact comparison after WHERE and is valid. Option (B) has LIKE without a column to compare. Option (C) uses a greater-than comparison and selects only Name and DOB, so it does not express the required class equality. Option (D) incorrectly combines IN with an equals sign; IN requires a parenthesised list or suitable subquery. These alternatives cannot replace the exact equality query.

Step 3: Conclusion The required query is `SELECT * FROM STUDENT WHERE Class = 'XI';` Hence, the correct option is (A).

Quick Tip to Solve: For one exact class value use WHERE Class = followed by the quoted value; use SELECT * when every column is required.

---

6. Which Sql Statement correctly retrieves records from a 'Commodity' table where the Price is either above 1000 or the product Category is 'Sales'? [5 marks]

(A) SELECT * FROM Commodity WHERE Price>1000 OR Category ='Sales';

(B) SELECT * FROM Commodity WHERE Price>1000 AND Category ='Sales';

(C) SELECT * FROM Commodity IF Price>1000 OR Category ='Sales';

(D) SELECT * FROM Commodity WHERE Price>1000 || Category =Sales ;

Correct Answer: (A) SELECT * FROM Commodity WHERE Price>1000 OR Category ='Sales';

Solution:

Concept:

- A WHERE clause selects rows for which its search condition is true.
- Logical OR accepts a row when at least one predicate is true; AND requires both predicates to be true.

Step 1: Express the two permitted cases The first predicate is Price > 1000; the second is that Category equals the text value Sales. A product priced above 1000 qualifies even outside Sales. Conversely, a Sales product qualifies even at a lower price. Combining the predicates with OR expresses exactly this inclusive choice, including a row satisfying both.

Step 2: Check the alternatives Option (A) correctly uses WHERE, OR and a quoted text literal. Option (B) uses AND and wrongly excludes records meeting just one condition. Option (C) uses IF in place of the required WHERE clause. In option (D), Sales is unquoted, so it is treated as an identifier rather than the requested string value; the meaning of the double vertical bar is also SQL-dialect dependent. The standard explicit form in (A) avoids these problems.

Step 3: Conclusion Use `WHERE Price>1000 OR Category = 'Sales'`. Hence, the correct option is (A).

Quick Tip to Solve: Translate either qualifying condition into OR, and remember that text values such as Sales require quotes.

---

7. Which query correctly finds a category in the 'Commodity' table that has more than 5 products with a Price greater than 600? [5 marks]

(A) SELECT category FROM Commodity GROUP BY category HAVING COUNT(*)>5 AND Price >600;

(B) SELECT category FROM Commodity WHERE Price>600 GROUP BY category HAVING COUNT(*)>5;

(C) SELECT category, COUNT(*) FROM Commodity GROUP BY Category HAVING Price>600 AND COUNT(*)>5;

(D) SELECT category FROM Commodity GROUP BY HAVING COUNT(Price>600)>5;

Correct Answer: (B) SELECT category FROM Commodity WHERE Price>600 GROUP BY category HAVING COUNT(*)>5;

Solution:

Concept:

- WHERE filters individual records before grouping; GROUP BY collects the remaining records by the grouping attribute.
- HAVING filters groups using aggregate results such as COUNT(*). The count must be taken after the price filter.

Step 1: Apply the price condition before counting First retain only rows with Price > 600. Next group these qualifying products by category. COUNT(*) in each group now measures the number of products in that category whose price exceeds 600, rather than the total number of all products in the category.

Step 2: Apply the group-size condition and compare queries The required groups satisfy $\mathrm{COUNT}(*) > 5$, so place that aggregate test in HAVING. Option (B) has the sequence WHERE, GROUP BY, HAVING and implements both stages correctly. Options (A) and (C) leave Price as an ungrouped, nonaggregate condition after grouping and fail to restrict the counted input rows correctly. Option (D) omits the grouping column after GROUP BY; counting a Boolean expression is also not the required row-filtering operation.

Step 3: Conclusion Filter prices with WHERE, group by category, and keep groups using HAVING COUNT(*)>5. Hence, the correct option is (B).

Quick Tip to Solve: For a count of qualifying products, filter product rows with WHERE first, then test the group count with HAVING.

8. Match List-I with List-II:

List-I	List-II
(A) Retrieve all rows with duplicates removed	(I) GROUP BY
(B) Filter rows based on condition	(II) ORDER BY
(C) Aggregate data into groups	(III) DISTINCT
(D) Sort rows in ascending or descending order	(IV) WHERE

Choose the correct answer from the options given below: [5 marks]

(A) (A) - (IV), (B) - (III), (C) - (I), (D) - (II)

(B) (A) - (I), (B) - (II), (C) - (III), (D) - (IV)

(C) (A) - (III), (B) - (IV), (C) - (I), (D) - (II)

(D) (A) - (III), (B) - (IV), (C) - (II), (D) - (I)

Correct Answer: (C) (A) - (III), (B) - (IV), (C) - (I), (D) - (II)

Solution:

Concept:

- DISTINCT removes duplicate result rows; WHERE selects rows satisfying a condition.
- GROUP BY forms groups for aggregation, whereas ORDER BY sorts the result.

Step 1: Match row selection and duplicate removal The description in List-I (A) requires duplicate rows to be removed. SQL SELECT DISTINCT performs this task, so its match is (III). List-I (B) asks for filtering on a condition; WHERE is the row-filtering clause, so its match is (IV). Filtering determines which source rows participate in the result; it does not put them in sorted order.

Step 2: Match grouping and sorting List-I (C) describes forming groups of rows, which is GROUP BY, listed as (I). List-I (D) asks for ascending or descending arrangement, which is ORDER BY, listed as (II). Combining these four matches gives A–III, B–IV, C–I, D–II. The first option interchanges filtering and duplicate removal; the second misassigns all four operations. The fourth interchanges grouping and sorting. Only the third combination preserves every match.

Step 3: Conclusion The matching is A–III, B–IV, C–I, D–II. Hence, the correct option is (C).

Quick Tip to Solve: Remember the four verbs: DISTINCT removes duplicates, WHERE filters, GROUP BY groups, and ORDER BY sorts.

9. Which of the following are multiple row functions?

(A) Sum

(B) Count

(C) Left

(D) Max

Choose the correct answer from the options given below: [5 marks]

(A) (A), (B) and (C) only

(B) (A) and (D) only

(C) (A), (B) and (D) only

(D) (A) and (B) only

Correct Answer: (C) (A), (B) and (D) only

Solution:

Concept:

- A multiple-row or aggregate function computes one result from a collection of rows.
- SUM, COUNT and MAX aggregate values; LEFT is a string function applied to an individual value.

Step 1: Classify the numeric and counting functions SUM adds the non-null numeric values in a group. COUNT counts rows or non-null expression values, depending on its argument. MAX returns the greatest non-null value in a group. Thus statements (A), (B) and (D) each name an aggregate function. Their defining feature is that they summarise a collection rather than merely transform one input string.

Step 2: Exclude the string function and compare the sets LEFT extracts a specified number of characters from the left end of a string. It operates on a single string value and is therefore a single-row string function, so (C) must be excluded. The first choice wrongly includes LEFT; the second leaves out COUNT, and the fourth leaves out MAX. The third choice contains exactly SUM, COUNT and MAX, so it includes every required function and no single-row function.

Step 3: Conclusion SUM, COUNT and MAX are the multiple-row functions. Hence, the correct option is (C).

Quick Tip to Solve: Classify what a function does: summarising a group indicates aggregation; extracting characters indicates a string function.

10. Arrange the evolution of networking as per the development timeline:

(A) The Berners - Lee at CERN developed HTML and URL, thus giving birth to World Wide Web (www).

(B) TCP/IP introduced as standard protocol on ARPANET.

(C) Roy Tomlinson develops network messaging or mail symbol@ comes to mean "at".

(D) ARPANET became functional by connecting UCLA and SRI.

Choose the correct answer from the options given below: [5 marks]

(A) (A), (B), (C), (D)

(B) (D), (C), (B), (A)

(C) (B), (A), (D), (C)

(D) (C), (B), (D), (A)

Correct Answer: (B) (D), (C), (B), (A)

Solution:

Concept:

- A chronological sequence places the earliest networking milestone first.
- ARPANET connectivity preceded network email, standard TCP/IP adoption and the World Wide Web.

Step 1: Place the early networking milestones ARPANET first connected UCLA and SRI in 1969, so event (D) is earliest. Ray Tomlinson developed network email and introduced the use of the @ symbol for addressing in 1971, placing event (C) next. Email already operated on early networks before the later transition to the TCP/IP protocol suite; it was not an invention of the World Wide Web.

Step 2: Place the protocol transition and the Web ARPANET adopted TCP/IP as its standard protocol on 1 January 1983, corresponding to event (B). Tim Berners-Lee developed the Web at CERN later, with the proposal in 1989 and the initial implementation in 1990, giving event (A). The resulting order is D, C, B, A. The first option starts with the latest event, the third puts the Web before the first ARPANET connection, and the fourth incorrectly puts email before ARPANET connectivity.

Step 3: Conclusion The chronological order is D, C, B, A. Hence, the correct option is (B).

Quick Tip to Solve: Anchor the timeline at ARPANET in 1969 and TCP/IP in 1983; email lies between them and the Web comes later.

11. Which of the following statement is CORRECT about GROUP BY clause in SQL? [5 marks]

(A) The GROUP BY clause is used to sort data in ascending or descending order.

(B) GROUP BY allows you to aggregate data, but it cannot be used with aggregate functions like SUM, AVG or COUNT.

(C) The GROUP BY clause is used to arrange identical data into groups often used with aggregate functions such as SUM, AVG or COUNT.

(D) The GROUP BY clause merges the contents of two or more tables.

Correct Answer: (C) The GROUP BY clause is used to arrange identical data into groups often used with aggregate functions such as SUM, AVG or COUNT.

Solution:

Concept:

- GROUP BY partitions rows according to equal values in the grouping columns.
- Aggregate functions such as SUM, AVG and COUNT can then return a result for each group.

Step 1: Identify the purpose of grouping Suppose a table contains employees from several departments. GROUP BY Department places all records having the same Department value into the same group. An aggregate such as SUM(Salary) can then return a separate salary total for each department. This is the operation described by the third statement: identical grouping-column values are arranged into groups for aggregation.

Step 2: Eliminate the competing SQL operations Sorting in ascending or descending order is performed by ORDER BY, so the first statement describes another clause. The second statement is false because aggregate functions are commonly used with GROUP BY; claiming that they cannot be used reverses its principal purpose. Combining data from multiple tables requires operations such as joins or unions, depending on the intended result; grouping does not itself merge tables. Therefore the fourth statement also fails.

Step 3: Conclusion GROUP BY arranges rows with identical grouping values into groups, often for aggregate functions. Hence, the correct option is (C).

Quick Tip to Solve: Use GROUP BY for one summary per category and ORDER BY for the order in which result rows are displayed.

12. A _____ is an access point between an organization's network and the outside world of the internet. [5 marks]

(A) Switch

(B) Gateway

(C) Repeater

(D) Modem

Correct Answer: (B) Gateway

Solution:

Concept:

- A gateway provides a point of entry or exit between a network and another network.
- Network devices are distinguished by their primary role: forwarding, regeneration, modulation or interconnection.

Step 1: Identify the required network boundary The blank describes the access point connecting an organisation's network with the external Internet. This is a gateway: traffic destined outside the local network passes through this network boundary. In an IP network, a router often serves as the default gateway, receiving packets for remote destinations and forwarding them towards the appropriate outside network.

Step 2: Distinguish the other devices A switch chiefly connects devices within a local network and forwards frames towards the appropriate ports; this alone does not define an Internet access gateway. A repeater regenerates weakened signals to extend their reach, without choosing an external network route. A modem performs modulation and demodulation for a communication link. A device may combine several functions, but the role named in the question is specifically the gateway function.

Step 3: Conclusion The required access point is a gateway. Hence, the correct option is (B).

Quick Tip to Solve: When a question describes entry to or exit from a local network, look for gateway rather than a signal-regenerating device.

13. An analog device which regenerates the weakened signal appearing on the cable and puts them back is known as _____. [5 marks]

(A) Switch

(B) Modem

(C) Bridge

(D) Repeater

Correct Answer: (D) Repeater

Solution:

Concept:

- Transmission signals weaken as they travel through a physical medium.
- A repeater regenerates a weakened signal so that communication can continue over a longer link.

Step 1: Identify the signal-restoration function The decisive words are that the device regenerates the weakened signal on a cable and puts it back onto the link. That task belongs to a repeater. The device receives the signal and retransmits a restored version, helping overcome the distance limitation imposed by attenuation. It is concerned with signal transmission rather than selecting a destination network or inspecting the application data.

Step 2: Separate regeneration from other device functions A switch forwards frames towards selected devices in a local network, while a bridge connects and filters traffic between LAN segments. Their distinguishing role is frame forwarding, rather than merely restoring the physical signal. A modem converts between digital information and a signal suitable for its communication medium through modulation and demodulation. None of those descriptions is as direct a match as repeater for the regeneration function specified here.

Step 3: Conclusion The device is a repeater. Hence, the correct option is (D).

Quick Tip to Solve: Associate attenuation and signal regeneration with a repeater; associate modulation and demodulation with a modem.

14. Match List-I with List-II:

List-I (Network Devices)	List-II (Description)
(A) Switch	(I) Sends signals only to the selected devices instead of sending to all.
(B) Repeater	(II) An analog device that works with signals on the cables to which it is connected.
(C) Hub	(III) A network device used to connect different devices through wires.
(D) Ethernet Card	(IV) A network adapter acts as an interface between the computer and the network.

Choose the correct answer from the options given below: [5 marks]

(A) (A) - (I), (B) - (II), (C) - (III), (D) - (IV)

(B) (A) - (II), (B) - (I), (C) - (III), (D) - (IV)

(C) (A) - (I), (B) - (II), (C) - (IV), (D) - (III)

(D) (A) - (IV), (B) - (III), (C) - (II), (D) - (I)

Correct Answer: (A) (A) - (I), (B) - (II), (C) - (III), (D) - (IV)

Solution:

Concept:

- A switch forwards frames towards selected ports, while a hub provides a shared wired connection point.
- A repeater restores signals and an Ethernet card interfaces a computer with an Ethernet network.

Step 1: Match forwarding and signal regeneration A switch learns where devices are reachable and normally sends a received unicast frame towards the appropriate port instead of broadcasting it to every device. This matches description (I). A repeater operates on signals along the connected

medium and regenerates them, matching description (II). These functions distinguish selective traffic forwarding from the restoration of signal strength or form.

Step 2: Match the connection point and network adapter A hub connects multiple devices through wired ports and repeats incoming signals to its other ports, matching description (III). An Ethernet card is the network interface adapter installed in or attached to a computer, so it matches description (IV). The complete mapping is A–I, B–II, C–III, D–IV. The second choice swaps the switch and repeater roles; the third swaps the hub and Ethernet card roles. The fourth misassigns these functions more broadly. Only the first choice fits every description.

Step 3: Conclusion The mapping is A–I, B–II, C–III, D–IV. Hence, the correct option is (A).

Quick Tip to Solve: Recall each device by its key role: switch selects, repeater regenerates, hub connects wired ports, and Ethernet card interfaces the computer.

15. The RJ in RJ-45 connector stands for : [5 marks]

(A) Registered Joint

(B) Routing Jack

(C) Registered Jack

(D) Ring Joint

Correct Answer: (C) Registered Jack

Solution:

Concept:

- RJ is the abbreviation for Registered Jack. It names a standardised telecommunications connector arrangement.
- RJ-45 is commonly associated with the modular connector used to connect twisted-pair Ethernet network cables.

Step 1: Expand the abbreviation Read RJ as two initial letters: R stands for Registered and J stands for Jack. Therefore the complete expansion required by the question is Registered Jack. The question asks what the letters mean, rather than what a cable does or how packets are forwarded.

Step 2: Compare the alternatives Registered Joint has the correct first word but replaces Jack with Joint, so it is not the expansion. Routing Jack substitutes Routing for Registered; routing is a networking function but is unrelated to the expansion of RJ. Ring Joint describes neither the standard abbreviation nor the Ethernet connector. Registered Jack is consequently the only matching pair of words.

Step 3: Conclusion The RJ in RJ-45 stands for Registered Jack. Hence, the correct option is (C).

Quick Tip to Solve: Remember RJ as Registered Jack; do not replace Jack with Joint merely because both suggest a connection.

Section B1: Computer Science (175 marks)

Attempt either Section B1: Computer Science or Section B2: Informatics Practices, answering all its 35 questions.

16. Arrange the following statements in the correct sequence in order to write multiple strings to a file.

(A) myobj.close()

(B) lines= ["Hello everyone \n", "writing multiple strings into file"]

(C) myobj.writelines(lines)

(D) myobj = open ("myfile.txt", 'w')

Choose the correct answer from the options given below: [5 marks]

(A) (D), (B), (C), (A)

(B) (C), (D), (B), (A)

(C) (B), (A), (D), (C)

(D) (C), (B), (D), (A)

Correct Answer: (A) (D), (B), (C), (A)

Solution:

Concept:

- A file must be opened before methods are called on its file object. Opening in mode w prepares it for writing.
- writelines accepts a sequence of strings and writes them in sequence; close releases the file after the writing is complete.

Step 1: Establish the prerequisites Statement D creates myobj by opening myfile.txt in write mode. Statement B then assigns the two strings to the list lines. Before statement C can call myobj.writelines(lines), both the file object and the list must exist. Calling writelines before either required object exists would not accomplish the task.

Step 2: Complete and compare the sequence After creating myobj and lines, execute C to write the strings and A to close the file. This gives D, B, C, A, which is the first option. The second and fourth options start with C before the required setup. The third closes myobj before the opening statement and also places the actual write after the close. Although the independent list could be prepared before opening a file in another valid program, the first listed sequence is the valid offered one.

Step 3: Conclusion Open the file, prepare the list, write its strings, and then close the file. Hence, the correct option is (A).

Quick Tip to Solve: Check dependencies in ordering questions: create the file object and data before writelines, and close only after writing.

17. Stack follows the _____ principle. [5 marks]

(A) Last-In-First-out

(B) First-In-First-out

(C) Enqueue - Dequeue

(D) Last-In-Last-out

Correct Answer: (A) Last-In-First-out

Solution:

Concept:

- A stack inserts and removes elements at the same accessible end, called its top.
- LIFO means Last-In-First-Out: the most recently pushed element is the first element removed by pop.

Step 1: Trace a short stack Suppose an empty stack receives 7 and then 5. The second element, 5, lies at the top above 7. A pop removes 5 first, although 7 entered earlier. This simple trace demonstrates that entry order is reversed during removal, provided no intervening operations change the stack.

Step 2: Distinguish the other descriptions Last-In-First-out describes precisely the observed behaviour. First-In-First-out describes a queue, where the oldest element leaves first. Enqueue and dequeue are names of queue operations rather than the principle of a stack. Last-In-Last-out would make the most recent element leave last, contradicting removal from the top.

Step 3: Conclusion A stack follows the Last-In-First-Out principle. Hence, the correct option is (A).

Quick Tip to Solve: Think of plates placed on a pile: the last plate placed on top is the first one you can take away.

18. The 'raise' keyword in Python is used ____. [5 marks]

(A) to throw an exception

(B) to avoid an exception

(C) to handle a keyword

(D) to find an exception

Correct Answer: (A) to throw an exception

Solution:

Concept:

- An exception signals an error or exceptional condition during the execution of a Python program.
- The raise statement explicitly raises an exception; a matching except clause handles an exception when control reaches it.

Step 1: Identify the action of raise The keyword raise is used when a program deliberately signals an exception. For example, a program can raise ValueError when an input violates a required condition. Execution then transfers according to the exception-handling rules; the statement does not itself provide the recovery action.

Step 2: Eliminate the remaining meanings The description to throw an exception is another way of saying to raise one, and is therefore correct. Raising does not avoid an exception: it initiates one. It does not handle a keyword, since keywords are reserved elements of Python syntax. It also does not search for or find an exception; handling is associated with try and except, whereas raise actively signals it.

Step 3: Conclusion The raise keyword is used to throw an exception. Hence, the correct option is (A).

Quick Tip to Solve: Associate raise with signalling an error and except with handling it; they perform different roles.

19. Which of the following statement is correct with respect to exception handling in python?

[5 marks]

(A) There can be multiple except clause associated with a single try block.

(B) The finally clause cannot be used if there are more than one except clause present with a try block.

(C) The raise statement is mandatory in a try block.

(D) An except clause can exist without any try block.

Correct Answer: (A) There can be multiple except clause associated with a single try block.

Solution:

Concept:

- A try suite can be followed by several except clauses, allowing different exception classes to receive different handlers.
- A finally suite can accompany exception handlers and runs for cleanup; an explicit raise is not required inside every try suite.

Step 1: Check the valid relationship A single try block may have one except clause for ValueError and another for ZeroDivisionError, for example. When an exception occurs, Python searches the associated handlers for a matching type and executes the first suitable handler. Thus multiple except clauses associated with one try block are legal.

Step 2: Test the three false statements The presence of more than one except clause does not prohibit finally; cleanup can still be placed in a finally suite. A try suite may contain ordinary operations that raise exceptions automatically, so writing raise explicitly is not mandatory. Finally, an except clause belongs to a try statement and cannot stand alone without the associated try block. These checks leave only the first statement correct.

Step 3: Conclusion There can be multiple except clauses associated with a single try block. Hence, the correct option is (A).

Quick Tip to Solve: For exception questions, separate the roles: try contains risky work, except selects a handler, and finally performs cleanup.

20. Which of the following methods is used to position a file object to a specific position in a file?

[5 marks]

(A) tell()

(B) peep()

(C) move()

(D) seek()

Correct Answer: (D) seek()

Solution:

Concept:

- A file object maintains a current position that determines where the next read or write operation takes place.
- The seek method changes that position, whereas tell reports the current position without moving it.

Step 1: Select the positioning operation To move the file object to a specific permitted position, call its seek method with the required offset, such as myobj.seek(0) to return to the beginning. Later file operations use the new position. In text files only supported offsets and positions should be used; the key point here is that seek performs the repositioning.

Step 2: Compare with the distractors tell() answers the question of where the file position currently is; it does not set a new location. peep() and move() are not the standard Python file-object methods for positioning the file pointer. Therefore only seek() supplies the action requested in the question, while tell() supplies information about that action's current result.

Step 3: Conclusion The method used to position a file object at a specific location is seek(). Hence, the correct option is (D).

Quick Tip to Solve: Remember the contrast: tell tells you the current position; seek sets the file position.

21. Consider an empty stack. What will the stack look like (starting from the top) after executing the following operations in sequence?

1. push(7)

2. push(5)

3. push(9)

4. push(11)

5. pop()

6. push (22)

(A) 22

(B) 9

(C) 5

(D) 7

Choose the correct answer from the options given below: [5 marks]

(A) (A), (B), (C), (D)

(B) (D), (C), (B), (A)

(C) (B), (A), (D), (C)

(D) (C), (B), (D), (A)

Correct Answer: (A) (A), (B), (C), (D)

Solution:

Concept:

- A push adds an element to the top of a stack, and pop removes its current top element.
- The final stack must be reported starting from the top, so reverse the bottom-to-top insertion record after accounting for removals.

Step 1: Trace the pushes and pop Write the stack from bottom to top while processing each operation. After push(7), it is [7]. After push(5), it is [7, 5]. After push(9), it is [7, 5, 9]. After push(11), it is [7, 5, 9, 11]. The pop operation removes 11, the most recently inserted element, leaving [7, 5, 9].

Step 2: Add the final element and read from the top Pushing 22 produces [7, 5, 9, 22] from bottom to top. Therefore, from top to bottom the values are 22, 9, 5, 7. The stated labels for these values are A, B, C, D respectively. The other sequences either start at the bottom or reorder the interior elements, and none describes removal from the current stack top correctly.

Step 3: Conclusion The final stack from the top is 22, 9, 5, 7, matching the first sequence. Hence, the correct option is (A).

Quick Tip to Solve: Track the stack in one fixed direction throughout the operations, then reverse it only if the question asks from the top.

22. In the below stack operation, what will be the output value for the variable 'result'?

result = 0

stack = [10,20,30]

stack.append(40)

result +=stack.pop()

result +=stack.pop() [5 marks]

(A) 70

(B) 60

(C) 90

(D) 50

Correct Answer: (A) 70

Solution:

Concept:

- A stack follows last-in, first-out (LIFO): the most recently pushed element is removed first.
- For a Python list used as a stack, append adds at the right end and pop without an index removes and returns the rightmost element.

Step 1: Trace the push The initial stack, written from bottom to top, is [10, 20, 30]. The statement stack.append(40) adds 40 above 30, so the stack becomes [10, 20, 30, 40]. The variable result is still zero because append changes the list rather than adding to result.

Step 2: Trace the two removals The first stack.pop() returns 40 and leaves [10, 20, 30]. Thus result becomes 40. The second stack.pop() returns 30 and leaves [10, 20]. Consequently, result becomes $40 + 30 = 70$. The two returned values are added; the remaining values 10 and 20 are not included. A result of 60, 90 or 50 does not follow these two LIFO removals.

Step 3: Conclusion The final value stored in result is 70. Hence, the correct option is (A).

Quick Tip to Solve: Write the stack top at the right and cross out one rightmost element for each pop before updating the accumulator.

23. What will be the postfix equivalent of the expression (A + B)* (C - D)? [5 marks]

(A) A B + C D - *

(B) A B * C D + -

(C) A B - C D * +

(D) A B C * + D -

Correct Answer: (A) A B + C D - *

Solution:

Concept:

- In postfix notation, an operator is placed after both of its operands; parentheses are unnecessary once evaluation order is encoded.

- A parenthesised subexpression must be converted as a whole before combining it with another subexpression.

Step 1: Convert the two parenthesised operands The infix sum (A + B) becomes A B + because the addition operator follows A and B. Similarly, the difference (C - D) becomes C D -. The order C then D must be preserved, since subtraction is not commutative.

Step 2: Place the outer multiplication last The complete infix expression multiplies the sum by the difference. Write the postfix form of the first operand, then the postfix form of the second operand, and finally *. This gives A B + C D - *. In option B the early * multiplies A and B instead; option C subtracts A and B; option D multiplies B and C before addition. Each changes the original grouping.

Step 3: Conclusion The postfix equivalent is A B + C D - *. Hence, the correct option is (A).

Quick Tip to Solve: Treat each parenthesised expression as one operand, convert it first, and append the operator that joins the two groups.

24. Which of the following elements will be present in a queue after the end of the following operation sequence? Assume the queue is empty initially.

enqueue(10)
enqueue(20)
enqueue(30)
dequeue()
enqueue(40)
dequeue()
enqueue(50)
enqueue(60)
(A) 60
(B) 30
(C) 20
(D) 50

Choose the correct answer from the options given below: [5 marks]

(A) (A), (B) and (D) only
(B) (A), (B) and (C) only
(C) (A), (B), (C) and (D)
(D) (B), (C) and (D) only

Correct Answer: (A) (A), (B) and (D) only

Solution:

Concept:

- A queue follows first-in, first-out (FIFO): enqueue adds at the rear, while dequeue removes the element at the front.

- To determine final membership, record each removal in sequence; an element does not disappear merely because later elements are inserted.

Step 1: Process the first removal Starting with an empty queue, the first three insertions produce [10, 20, 30], with the front on the left. The first dequeue removes 10, leaving [20, 30]. Inserting 40 then gives [20, 30, 40].

Step 2: Process the remaining operations The second dequeue removes 20, the oldest element still present, leaving [30, 40]. Inserting 50 and then 60 produces the final queue [30, 40, 50, 60]. Of the listed elements, statement (A), 60, is present; statement (B), 30, is present; statement (C), 20, is absent; and statement (D), 50, is present. Thus the required combination is (A), (B) and (D) only. Every other offered combination incorrectly includes the removed element 20.

Step 3: Conclusion The final queue is [30, 40, 50, 60], so among the listed elements 60, 30 and 50 remain, while 20 has been removed. Hence, the correct option is (A).

Quick Tip to Solve: For a queue trace, add to the right and remove from the left; check the statement labels only after completing every operation.

25. Which of the following is a linear data structure?

(A) Deque

(B) Queue

(C) Stack

(D) Tree

Choose the correct answer from the options given below: [5 marks]

(A) (A), (B) and (C) only

(B) (A), (B) and (D) only

(C) (A), (B), (C) and (D)

(D) (B), (C) and (D) only

Correct Answer: (A) (A), (B) and (C) only

Solution:

Concept:

- A linear data structure organises elements in one sequential order, even when insertion and deletion are restricted to particular positions.
- A nonlinear structure allows hierarchical or branching relationships rather than placing every element along one single sequence.

Step 1: Classify deque, queue and stack A deque is a double-ended queue whose elements form a single ordered sequence. A queue also has a single sequence, with insertion at the rear and removal at the front. A stack has a single sequence with its accessible top at one end. Hence statements (A), (B) and (C) all name linear data structures. Different access rules do not make a structure nonlinear.

Step 2: Exclude the tree A tree represents a hierarchy: a node can have multiple children and

branching paths lead to different descendants. It is therefore nonlinear, so statement (D) must be excluded. Select the combination containing deque, queue and stack, and no tree. The other combinations all include (D) and therefore cannot be correct, regardless of which linear structure they omit.

Step 3: Conclusion Deque, queue and stack are linear data structures; a tree is nonlinear. Hence, the correct option is (A).

Quick Tip to Solve: Classify by the relationship between elements: one sequence is linear, while parent-child branching indicates a nonlinear structure.

26. Which of the following statements is true in the context of a Deque? [5 marks]

(A) Insertion is done from front and deletion is done from rear.

(B) Insertion and deletion can be done at both ends.

(C) Insertion is done from front and deletion is done from front.

(D) Insertion is done from rear and deletion is done from rear.

Correct Answer: (B) Insertion and deletion can be done at both ends.

Solution:

Concept:

- A deque means double-ended queue: it maintains a linear sequence with both a front and a rear.
- In an unrestricted deque, both insertion and deletion are permitted independently at either end.

Step 1: Identify the defining operations A deque supports insertion at the front, insertion at the rear, deletion at the front and deletion at the rear. This flexibility is its defining feature. Choosing one particular pair of operations does not describe the full capability of the structure.

Step 2: Compare the four descriptions Option B states the general rule that insertion and deletion can be done at both ends, so it is the complete definition. Option A mentions only front insertion and rear deletion, while C mentions only front insertion and front deletion. Option D mentions only rear insertion and rear deletion. These are individual operation patterns that a deque can perform, but they omit other supported operations and do not state its distinguishing property. The question therefore selects B as the general description.

Step 3: Conclusion Insertion and deletion can both be performed at either end of a deque. Hence, the correct option is (B).

Quick Tip to Solve: Expand deque as double-ended queue, then check that the answer permits both kinds of operation at both ends.

27. Which of the following are the applications of binary search?

(A) Telephone Directory

(B) Finding the element with a minimum or maximum value in a sorted list

(C) Implementing routing tables

(D) Indexing in database

Choose the correct answer from the options given below: [5 marks]

(A) (A), (B) and (D) only

(B) (A), (B) and (C) only

(C) (A), (B), (C) and (D)

(D) (B), (C) and (D) only

Correct Answer: (C) (A), (B), (C) and (D)

Solution:

Concept:

- Binary search uses ordered keys: comparing a target with a middle key eliminates one half of the remaining search interval.
- Its applications include ordered lookup tables and indexes, where the relevant records or search boundaries can be examined systematically.

Step 1: Check the ordered-data applications A telephone directory orders names, so repeated comparison with a middle entry can narrow the range containing a requested name. Database indexing likewise supports fast lookup over ordered keys. The endpoints of a sorted list identify its minimum and maximum; binary search techniques can also locate ordered boundaries or extreme qualifying values. Thus the sorted-list setting in (B) is compatible with ordered searching, although merely reading a known endpoint is faster when only the unrestricted minimum or maximum is needed.

Step 2: Check routing tables and select the combination Routing-table lookup can be implemented using ordered lookup structures and binary-search-based techniques to narrow candidate routes; the statement does not claim that every routing table must use this technique. Thus (C) is also an application context. All four listed contexts can employ ordered search, so select the combination including (A), (B), (C) and (D). Each other option omits at least one applicable context.

Step 3: Conclusion All four listed application contexts can employ ordered searching. Hence, the correct option is (C).

Quick Tip to Solve: Look for an ordering of keys or boundaries that lets a comparison discard half the candidates; distinguish a possible application from the fastest special-case operation.

28. What happens even if a key is not present in the list during a linear search? [5 marks]

(A) The algorithm stops mid way.

(B) The algorithm skips elements.

(C) The algorithm gives an error.

(D) The algorithm searches through entire list.

Correct Answer: (D) The algorithm searches through entire list.

Solution:

Concept:

- Linear search compares the search key with successive elements, usually from the beginning of the list.
- An unsuccessful search can be declared only after every candidate element has been checked and none matches the key.

Step 1: Apply the stopping condition At each position, linear search compares the current element with the required key. A match allows an immediate successful return. If the key is absent, that condition never becomes true, so traversal continues to the next element until the list is exhausted.

Step 2: Evaluate the alternatives The algorithm therefore searches through the entire list, as stated in option D. It cannot safely stop midway because an unchecked element might contain the key. It does not skip elements: skipping would invalidate the guarantee of finding an arbitrary matching value. Absence is also an ordinary search outcome, not an algorithmic error; the implementation typically reports a failure value such as -1 or False after completing the traversal. For a list of n elements, an unsuccessful ordinary linear search makes n key comparisons.

Step 3: Conclusion When the key is absent, linear search checks the entire list before reporting failure. Hence, the correct option is (D).

Quick Tip to Solve: For an absent key, count one comparison per list element; early stopping is possible only after an actual match.

29. If a linear search is performed on 10 elements, how many comparisons will be needed in worst case? [5 marks]

(A) 1

(B) 10

(C) 9

(D) 11

Correct Answer: (B) 10

Solution:

Concept:

- Linear search examines elements sequentially, comparing the search key with one element at a time.

- The worst case occurs when the key is absent or is stored at the last position. For a list of n elements, this requires n key comparisons.

Step 1: Identify the worst-case position There are 10 elements. If the first nine elements fail to match, the algorithm must still inspect the tenth. Stopping after nine comparisons would leave the final element unchecked, so the search could not correctly report failure.

Step 2: Count the comparisons and reject the alternatives Each element contributes one comparison with the key, giving $1 + 1 + \cdots + 1 = 10$ comparisons. One comparison is the best case, when the first element matches. Nine comparisons do not cover the entire list. Eleven would count an extra operation rather than an additional element-to-key comparison; loop-control tests are not what this search-comparison question counts.

Step 3: Conclusion The worst case needs 10 comparisons. Hence, the correct option is (B).

Quick Tip to Solve: For worst-case linear search, imagine the key is absent: every element must be checked once.

30. How many passes are required to sort a list of 6 elements using bubble sort? [5 marks]

(A) 5

(B) 6

(C) 4

(D) 3

Correct Answer: (A) 5

Solution:

Concept:

- Bubble sort compares adjacent elements and swaps a pair when it is in the wrong order.
- After one full pass, the largest remaining element reaches its final position at the end of the unsorted portion.

Step 1: Apply the standard pass count For a list of n elements, the standard bubble-sort procedure needs at most $n - 1$ passes to guarantee a sorted result. After fixing the first $n - 1$ final positions, the remaining element is automatically in its correct place; it does not need a separate sorting pass.

Step 2: Substitute the list size Here $n = 6$, so the required standard pass count is $6 - 1 = 5$. These passes make 5, 4, 3, 2 and 1 adjacent comparisons respectively in the usual shrinking-bound implementation. Six passes include a redundant pass, while four or three passes cannot guarantee sorting arbitrary input. An early-exit implementation can stop sooner on an already sorted or favourable list, but no such special input is supplied here.

Step 3: Conclusion The standard bubble-sort count for six elements is 5 passes. Hence, the correct option is (A).

Quick Tip to Solve: Count passes separately from comparisons: bubble sort needs at most n minus 1 passes, with fewer comparisons on successive passes.

31. Which sorting algorithm involves repeatedly selecting the smallest element and placing it in its correct position? [5 marks]

(A) Bubble Sort

(B) Selection Sort

(C) Insertion Sort

(D) Quick Sort

Correct Answer: (B) Selection Sort

Solution:

Concept:

- Selection sort maintains a sorted prefix and an unsorted remainder of the list.
- On each pass it selects the smallest element of the unsorted remainder and moves it to the next position of the sorted prefix.

Step 1: Match the stated operation The defining phrase is repeatedly selecting the smallest element. In ascending selection sort, the first pass finds the global minimum and places it at the first position. The next pass finds the minimum among the remaining elements and places it second. This repeats until the list is sorted.

Step 2: Distinguish the competing algorithms Bubble sort works by successive comparisons and swaps of neighbouring elements, rather than selecting one minimum over the unsorted region. Insertion sort takes the next element and inserts it into an already sorted prefix. Quick sort chooses a pivot and partitions the list around it, then sorts the resulting portions. Only selection sort directly follows the minimum-selection procedure described in the question.

Step 3: Conclusion The described algorithm is Selection Sort. Hence, the correct option is (B).

Quick Tip to Solve: The word selection is literal: look for a pass that selects the minimum or maximum from the unsorted part.

32. A loop within a loop will have the time complexity as _____. [5 marks]

(A) n^2

(B) log(n)

(C) n

(D) n^3

Correct Answer: (A) n^2

Solution:

Concept:

- For two nested loops with n iterations each and constant-time work in the inner body, total work is the product of their iteration counts.
- Time complexity describes growth with input size, ignoring constant multiplicative factors and lower-order terms.

Step 1: Count the inner work for one outer iteration Under the standard interpretation of this question, the outer loop runs n times and the inner loop also runs n times for each outer iteration. One outer iteration therefore performs n executions of the basic inner operation.

Step 2: Multiply across the outer loop The total basic-operation count is $n \times n = n^2$, giving quadratic time. Linear growth would describe a single such loop, while logarithmic growth typically involves repeatedly reducing the remaining problem by a fixed factor. Cubic growth would require an additional independent factor of n. Nesting alone is not a universal guarantee: fixed or logarithmic loop bounds change the result. The listed choice intended for two full input-sized loops is quadratic.

Step 3: Conclusion For the usual pair of n-iteration nested loops, the time complexity is $O(n^2)$. Hence, the correct option is (A).

Quick Tip to Solve: Write each loop bound beside the loop, then multiply the counts for nested loops; do not infer complexity from indentation alone.

33. Which of the following are examples of unstructured data?

(A) Writing tweets

(B) Recording a video

(C) Marking attendance by the teacher

(D) Filling a form for opening an account

Choose the correct answer from the options given below: [5 marks]

(A) (A) and (B) only

(B) (A), (B) and (C) only

(C) (C) and (D) only

(D) (B), (C) and (D) only

Correct Answer: (A) (A) and (B) only

Solution:

Concept:

- Structured data follows predefined fields or a regular tabular schema, such as attendance records and account-opening details.
- Unstructured data includes free-form natural-language content and multimedia, whose main content is not arranged as fixed database fields.

Step 1: Classify the free-form content Writing tweets produces free-form text, so statement (A) is an example of unstructured content. Recording a video produces audiovisual content rather than a fixed set of tabular field values, so statement (B) is also unstructured. A platform may attach structured metadata, such as a timestamp, without making the main text or video itself structured.

Step 2: Classify the records and select the combination Marking attendance normally records predefined values against student names or identifiers and dates, so statement (C) is structured. Filling a form for opening an account enters values into predefined fields, such as name, address and account details; statement (D) is structured too. Thus the required combination includes (A) and

(B) and excludes both (C) and (D). The first option expresses exactly this combination; choices that include attendance or account-form records do not fit.

Step 3: Conclusion Writing tweets and recording a video are the unstructured examples. Hence, the correct option is (A).

Quick Tip to Solve: Ask whether the main content fits fixed fields: free-form text and video differ from attendance tables and completed forms.

34. Calculate the standard deviation of the given temperature below recorded for the 7 days of a week:

24, 24, 21, 20, 26, 27, 30 [5 marks]

(A) 24

(B) 3.2

(C) 24.57

(D) 11.95

Correct Answer: (B) 3.2

Solution:

Concept:

- Population standard deviation measures spread about the mean: $\sigma = \sqrt{\frac{\sum x^2}{N} - \bar{x}^2}$.
- For all seven recorded days, use $N = 7$. The square root returns the spread to the same unit as the original temperatures.

Step 1: Find the mean and squared sum The total is $24 + 24 + 21 + 20 + 26 + 27 + 30 = 172$, so $\bar{x} = 172/7 \approx 24.5714$. The sum of squared temperatures is

$$24^2 + 24^2 + 21^2 + 20^2 + 26^2 + 27^2 + 30^2 = 4298.$$

The mean is a location measure; it is not the standard deviation.

Step 2: Calculate the variance and take its square root Using the seven-day data as the population,

$$\sigma^2 = \frac{4298}{7} - \left(\frac{172}{7}\right)^2 = \frac{502}{49} \approx 10.2449.$$

Therefore

$$\sigma = \frac{\sqrt{502}}{7} \approx 3.2008 \approx 3.2.$$

The options 24 and 24.57 represent typical temperatures or the mean, not the spread. The value 11.95 is the sample variance obtained with a divisor of six; it is neither the population standard deviation nor a square-rooted spread.

Step 3: Conclusion The standard deviation is approximately 3.2 in the temperature unit used. Hence, the correct option is (B).

Quick Tip to Solve: Keep full precision for the mean and variance; round only after taking the square root to avoid a wrong spread.

35. Which measure of central tendency is better suited for dealing with an outlier? [5 marks]

(A) Median

(B) Mean

(C) Mode

(D) Range

Correct Answer: (A) Median

Solution:

Concept:

- The median is the middle value after ordering the observations, or the average of the two middle values for an even number of observations.
- An outlier is an unusually extreme observation. A robust measure of centre changes little when an extreme value is made still more extreme.

Step 1: Explain why the median resists outliers The median depends mainly on rank and the central position in the ordered list. Increasing a value already at the extreme end does not normally alter the middle observation. For example, the median of 2, 3, 4, 5, 100 is 4; replacing 100 by 1000 leaves the median equal to 4.

Step 2: Compare the other measures The mean uses every numerical value in its sum, so a very large or small observation pulls it toward that extreme. The mode identifies the most frequent value; it can be absent or nonunique and does not generally provide the most suitable numerical centre for skewed data with an outlier. The range is the difference between the largest and smallest values, making it a measure of dispersion rather than central tendency and particularly sensitive to outliers. The median is therefore the standard robust choice.

Step 3: Conclusion The median is better suited to describing the centre when an outlier is present. Hence, the correct option is (A).

Quick Tip to Solve: For a skewed numerical dataset with an extreme value, sort the data and check its middle position before choosing the mean.

36. **Match List-I with List-II**

List-I	List-II
KEY CONCEPTS IN DATABASE	DESCRIPTION
(A) Database Schema	(I) Snapshot of the database
(B) Database Instance	(II) Restriction on the type of data
(C) Data Dictionary	(III) Skeleton of the database
(D) Data Constraint	(IV) Database catalog

Choose the correct answer from the options given below: [5 marks]

(A) (A) - (III), (B) - (I), (C) - (IV), (D) - (II)

(B) (A) - (I), (B) - (III), (C) - (II), (D) - (IV)

(C) (A) - (I), (B) - (II), (C) - (IV), (D) - (III)

(D) (A) - (III), (B) - (IV), (C) - (I), (D) - (II)

Correct Answer: (A) (A) - (III), (B) - (I), (C) - (IV), (D) - (II)

Solution:

Concept:

- A database schema is its structural design: the relations, attributes, data types and relationships that define its skeleton.
- A database instance is the data stored at a particular moment. A data dictionary stores metadata, while constraints restrict permitted data.

Step 1: Match the design and the current contents The schema describes how the database is organised, so (A) matches (III), skeleton of the database. An instance changes as records are inserted, deleted or updated; it is a snapshot at a particular time, so (B) matches (I). A schema and an instance therefore describe different aspects of the same database.

Step 2: Match metadata and restrictions The data dictionary is the database catalog containing descriptions of objects such as tables and columns, so (C) matches (IV). A data constraint is a rule that restricts acceptable values or states, so (D) matches (II). For example, a domain restricts the type of values an attribute may contain.

Step 3: Compare the combinations The complete matching is A–III, B–I, C–IV, D–II. The first choice gives all four pairs. The other combinations confuse the structural skeleton with the current data, or confuse metadata with a restriction; matching only one familiar pair is insufficient.

Step 4: Conclusion Schema matches skeleton, instance matches snapshot, dictionary matches catalog, and constraint matches restriction. Hence, the correct option is (A).

Quick Tip to Solve: Remember schema as structure, instance as an instant snapshot, dictionary as metadata, and constraint as a restriction.

37. Sometimes, on calling a customer service center, the Interactive Voice Response System (IVRS) tells us to wait till a support person is available. Which data structure this application is of? [5 marks]

(A) Enqueue

(B) Binary Tree

(C) Stack

(D) Queue

Correct Answer: (D) Queue

Solution:

Concept:

- A queue is a linear data structure that ordinarily serves elements in first-in, first-out (FIFO) order.
- Enqueue inserts an element at the rear; dequeue removes an element from the front. These are operations on the queue.

Step 1: Identify the waiting-line behaviour When all support persons are busy, incoming callers wait for service. The earliest waiting caller is ordinarily connected when a support person becomes available, while a newly arriving caller joins behind callers already waiting. This is the familiar first-come, first-served service discipline.

Step 2: Relate the behaviour to a data structure A queue represents this waiting line. A new call is enqueued at its rear, and the call at the front is dequeued for service. The question asks for the structure modelling the application, rather than the name of only the insertion operation.

Step 3: Reject the other choices Enqueue is an operation, so it is not the name of the data structure. A stack uses last-in, first-out order, which would serve the newest caller first. A binary tree represents a branching hierarchy and does not directly express the stated waiting-line discipline.

Step 4: Conclusion The IVRS waiting line is an application of a queue. Hence, the correct option is (D).

Quick Tip to Solve: For a service waiting line, identify FIFO first; choose queue rather than its insertion operation, enqueue.

38. Calculate the degree and cardinality of the given relation below:

ID	Name	Email	Address	Phone No.
001	Ajay	ajay@abc.com	A1-Shakti Nagar, Bangalore	55218936
002	Amit	amit@abc.com	C32- Indira Nagar, Lucknow	45321789
003	Ashu	ashu@abc.com	117/82 Patel Nagar, Delhi	98567894
004	Ashok	ashok@abc.com	249-250, Motinagar, Gujrat	97698541

[5 marks]

(A) Degree = 4, Cardinality = 5

(B) Degree = 5, Cardinality = 4

(C) Degree = 5, Cardinality = 1

(D) Degree = 4, Cardinality = 1

Correct Answer: (B) Degree = 5, Cardinality = 4

Solution:

Concept:

- The degree of a relation is its number of attributes, represented by the columns of the table.
- The cardinality of a relation is its number of tuples, represented by the data rows; the heading row is excluded.

Step 1: Count the attributes The relation has the attributes ID, Name, Email, Address and Phone No. Each heading denotes one attribute, regardless of how long the values below it are. Therefore the degree is 5. In particular, an address containing several words still occupies one field and does not create extra attributes.

Step 2: Count the tuples The displayed data contains one row each for Ajay, Amit, Ashu and Ashok, with IDs 001, 002, 003 and 004. Thus there are four tuples and the cardinality is 4. The row containing the column names describes the relation and is not an additional tuple.

Step 3: Compare both numbers with the options The required ordered pair is degree 5 and cardinality 4. The first choice reverses these counts. The third and fourth choices give cardinality 1, although four records are shown; the fourth also miscounts the attributes. Only the second choice agrees with both definitions.

Step 4: Conclusion The relation has degree 5 and cardinality 4. Hence, the correct option is (B).

Quick Tip to Solve: Write degree = columns and cardinality = data rows before counting; never include the heading row as a tuple.

39. Suppose you want to select distinct salaries working in department 101 from the 'Employee' relation. What will be the correct order of writing the query?

(A) FROM Employee

(B) SELECT

(C) WHERE

(D) DISTINCT Salary

(E) Dept_ID= 101

Choose the correct answer from the options given below: [5 marks]

(A) (A), (B), (C), (D), (E)

(B) (A), (C), (B), (D), (E)

(C) (B), (D), (A), (C), (E)

(D) (B), (E), (C), (D), (A)

Correct Answer: (C) (B), (D), (A), (C), (E)

Solution:

Concept:

- A SELECT query is written with its projection first, followed by FROM and the relation name, then WHERE and its condition.
- DISTINCT appears immediately after SELECT to remove duplicate values from the selected result. Written clause order differs from logical evaluation order.

Step 1: Assemble the intended query To obtain different salary values among employees in department 101, the SQL statement is `SELECT DISTINCT Salary FROM Employee WHERE Dept_ID = 101;`. SELECT introduces the result list, DISTINCT modifies it, FROM supplies the table, and WHERE introduces the department restriction.

Step 2: Map each fragment to its position The given fragment labels are B for SELECT, D for DISTINCT Salary, A for FROM Employee, C for WHERE, and E for the condition. Therefore the required fragment order is B, D, A, C, E. This order uses every fragment once and produces a complete syntactically valid statement.

Step 3: Check the printed answer choices The third choice gives B, D, A, C, E and produces the required query. The first choice begins with FROM instead of SELECT. The second begins with FROM WHERE, putting clauses in the wrong order. The fourth puts the condition immediately after SELECT and places the relation last. Thus only the third choice has the correct written SQL syntax.

Step 4: Conclusion The correct sequence is B, D, A, C, E. Hence, the correct option is (C).

Quick Tip to Solve: For written SQL order, place SELECT DISTINCT before the selected column, then FROM the table and WHERE the condition.

40. _____ is a search technique that makes use of the ordering of elements in the list to quickly search for a key. [5 marks]

(A) Linear Search

(B) Binary Search

(C) Hashing

(D) Bubble Search

Correct Answer: (B) Binary Search

Solution:

Concept:

- Binary search locates a key in a sorted sequence by comparing it with the middle element and eliminating one half of the remaining search interval.
- The ordering of elements makes that elimination valid: values on one side cannot contain the key after the comparison.

Step 1: Identify the clue in the question The decisive clue is that the technique makes use of the ordering of elements. In an ascending list, if a key is smaller than the middle element, it can only occur in the left half; if it is greater, it can only occur in the right half. An equality comparison ends the search successfully.

Step 2: Name the method and explain its speed This repeated halving is binary search. The search interval becomes about half as large after each unsuccessful comparison, yielding logarithmic search time, $O(\log n)$, for a sorted random-access list of n elements. The method therefore exploits sorted order rather than merely inspecting every entry.

Step 3: Eliminate the alternatives Linear search checks entries successively and does not require an ordered list. Hashing computes a position using a hash function rather than relying on sorted order. Bubble sort is a sorting procedure; the listed bubble search is not the standard search technique described by repeated comparisons in an ordered list.

Step 4: Conclusion The search technique is binary search. Hence, the correct option is (B).

Quick Tip to Solve: When a search question stresses sorted order and discarding half the candidates, recognise binary search.

41. Sales

ID	customer_id	order_date	total_amount
1	C101	2022-01-15	1000
2	C102	2022-02-20	500
3	C103	2022-03-10	800
4	C101	2022-04-05	600

What is the result of the following query with respect to given relation 'Sales'?
SELECT COUNT(DISTINCT customer_id) FROM Sales; [5 marks]

(A) 1
(B) 2
(C) 3
(D) 4

Correct Answer: (C) 3

Solution:

Concept:

- COUNT(DISTINCT column) counts the different non-NULL values occurring in that column, rather than counting every row.
- The DISTINCT keyword removes repetitions before the aggregate count is determined.

Step 1: Read the values used by the query The query uses only `customer_id`, whose four displayed entries are C101, C102, C103 and C101. Neither the different dates nor the different total amounts changes the customer identifiers being counted. There is no WHERE clause to remove any of these rows.

Step 2: Remove duplicate identifiers The identifier C101 appears twice, but DISTINCT retains it only once. The distinct customer identifiers are therefore C101, C102 and C103. All three are non-NULL values, so each contributes one to the aggregate count.

Step 3: Apply COUNT and compare the choices Counting the three unique identifiers gives 3. A count of 4 would be the row count and would incorrectly count C101 twice. A count of 1 would count only the repeated customer or the number of result rows, while 2 would omit one of the distinct customers. The aggregate returns one result row containing the value 3.

Step 4: Conclusion The query returns 3. Hence, the correct option is (C).

Quick Tip to Solve: For COUNT(DISTINCT ...), list unique values in the named column first; repeated purchases by one customer count once.

42. Which protocol is used to format and transmit the information from the World Wide Web?

[5 marks]

(A) FTP

(B) TCP

(C) HTTP

(D) DNS

Correct Answer: (C) HTTP

Solution:

Concept:

- HTTP stands for Hypertext Transfer Protocol. It defines the application-level request and response messages used to transfer web resources.
- A browser requests resources from a web server, which responds with the requested representation and status information.

Step 1: Connect the protocol with web access The World Wide Web operates through clients such as browsers and servers that exchange resources. HTTP specifies how these requests and responses are structured and communicated. For example, the browser can send a GET request for a resource, and the server returns a response with headers and content.

Step 2: Distinguish the application protocol from other roles The wording about formatting and transmitting web information refers to the protocol governing the web exchange. HTTP serves that role. HTML is a document markup language, whereas HTTP is the communication protocol; the question lists protocol choices and asks for the latter.

Step 3: Eliminate the remaining protocols FTP is used for file transfer rather than the standard browser-server web exchange. TCP supplies reliable transport services to applications but does not define the web request and response format. DNS translates domain names to addresses and does not transfer the requested web page. HTTP is therefore the matching application protocol.

Step 4: Conclusion The protocol is HTTP, Hypertext Transfer Protocol. Hence, the correct option is (C).

Quick Tip to Solve: Associate HTTP with web requests and responses, FTP with file transfer, DNS with name lookup, and TCP with reliable transport.

43. Match List-I with List-II

List-I	List-II
Topology	Characteristics
(A) Star	(I) Single backbone cable
(B) Bus	(II) Central hub or switch
(C) Ring	(III) Unidirectional link
(D) Hybrid	(IV) Multiple branches

Choose the correct answer from the options given below: [5 marks]

(A) (A) - (I), (B) - (II), (C) - (III), (D) - (IV)

(B) (A) - (I), (B) - (III), (C) - (II), (D) - (IV)

(C) (A) - (I), (B) - (II), (C) - (IV), (D) - (III)

(D) (A) - (II), (B) - (I), (C) - (III), (D) - (IV)

Correct Answer: (D) (A) - (II), (B) - (I), (C) - (III), (D) - (IV)

Solution:

Concept:

- A network topology describes the arrangement of devices and the connections between them.
- The defining structural feature distinguishes star, bus, ring and hybrid topologies; the matching uses these features rather than their names alone.

Step 1: Match star and bus to their connections In a star topology, each device connects to a central hub or switch, so (A) matches (II). A bus topology uses a single backbone cable to which its nodes connect, so (B) matches (I). These two matches already distinguish the correct alternative: reversing them incorrectly assigns the shared cable to a star and the central device to a bus. A central hub is not the defining structure of a bus.

Step 2: Complete the ring and hybrid matches A ring connects devices around a closed loop; the listed unidirectional link characteristic is the conventional ring description used here. Thus (C) matches (III). A hybrid combines arrangements and can have multiple branches, giving (D) with (IV). The completed mapping is therefore A-II, B-I, C-III, D-IV. The first alternative interchanges the star and bus features, the second also misidentifies ring, and the third swaps the ring and hybrid characteristics. Only the fourth alternative preserves all four pairings.

Step 3: Conclusion The matching is A-II, B-I, C-III, D-IV. Hence, the correct option is (D).

Quick Tip to Solve: Start with the unmistakable pair: star has a central switch, while bus has one backbone cable.

44. What describes the work of a repeater? [5 marks]

(A) Regenerates weak signals

(B) Resolves domain to IP address

(C) Unique hardware address

(D) Circuit switching

Correct Answer: (A) Regenerates weak signals

Solution:

Concept:

- A repeater is a physical-layer device that receives a weakened signal and regenerates it before transmitting it onward.
- Signal attenuation increases with distance. Regeneration restores the signal shape and strength so that it can travel over a further network segment.

Step 1: Identify the problem a repeater solves As a signal travels along a transmission medium, its strength decreases and its pulses can become distorted. A receiver must still be able to distinguish the original data reliably. A repeater is inserted along the path to receive this weakened signal, reconstruct it and transmit a refreshed version. Thus the device extends the usable transmission distance; it does not interpret an application request or decide a destination using a domain name.

Step 2: Compare the four descriptions Regenerating weak signals is exactly the repeater function, so the first description applies. Resolving a domain name into an IP address is a DNS service. A unique hardware address is a MAC address identifying a network interface, rather than the operation of a repeater. Circuit switching establishes a dedicated communication path between endpoints and is a switching technique. These three descriptions therefore refer to different networking concepts.

Step 3: Conclusion A repeater regenerates weak signals. Hence, the correct option is (A).

Quick Tip to Solve: Associate repeater with restoring a signal over distance; associate domain-name resolution with DNS.

45. What is the measurment unit of bandwidth? [5 marks]

(A) Hertz

(B) Kilobytes

(C) Milliseconds

(D) Megabytes

Correct Answer: (A) Hertz

Solution:

Concept:

- For a communication channel, frequency bandwidth is the difference between its upper and lower limiting frequencies.
- Frequency is measured in hertz (Hz), so the difference between two frequencies is also measured

in hertz.

Step 1: Identify the quantity being measured The question offers hertz, two units of data storage and a unit of time. Here bandwidth refers to the width of the frequency band available for transmission. If a channel passes frequencies from a lower limit to an upper limit, subtracting these limits gives its bandwidth. Because both limits are frequencies, their difference has a frequency unit. The unit hertz means cycles per second.

Step 2: Eliminate units for other quantities Hertz is therefore the appropriate choice. Kilobytes and megabytes count amounts of digital data; without a time denominator they do not measure a transmission rate either. Milliseconds measure time and may be used for delay or latency, not frequency bandwidth. In everyday networking, bandwidth can also describe data-transfer capacity in bits per second, but that unit is absent from the choices. Among the printed options, hertz is the precise unit for frequency bandwidth.

Step 3: Conclusion Frequency bandwidth is measured in hertz. Hence, the correct option is (A).

Quick Tip to Solve: Check whether bandwidth means a frequency interval or a data rate; bare kilobytes and megabytes measure neither.

46. What is the range of frequency for radio waves? [5 marks]

(A) 3 KHz to 300 GHz

(B) 1 GHz to 10 Ghz

(C) 300 GHz to 400 Thz

(D) 3 KHz to 1 GHz

Correct Answer: (D) 3 KHz to 1 GHz

Solution:

Concept:

- In the networking classification of unguided transmission media, radio waves occupy the low-frequency range from 3 kHz to 1 GHz.
- Microwaves and infrared are treated as separate transmission-media bands in this classification, even though broader electromagnetic terminology can group microwaves with radio.

Step 1: Use the communication-media classification For computer-network transmission media, the radio-wave range is taken as 3 kHz to 1 GHz. Radio waves in this range can be used for wireless communication and broadcasting. The question must be read using this classification, which distinguishes radio-wave transmission from microwave transmission. Its lower endpoint is in kilohertz and its upper endpoint is in gigahertz.

Step 2: Compare the boundaries in the alternatives The fourth option reproduces both endpoints, 3 kHz and 1 GHz. The interval 1 GHz to 10 GHz lies in the microwave band of this classification. The interval 300 GHz to 400 THz corresponds to the infrared range. The first interval, 3 kHz to 300 GHz, includes the microwave band as well and therefore is broader than the separately classified radio-wave transmission range used here. Although that broader usage occurs in electromagnetic-spectrum

terminology, the networking categorisation makes the fourth alternative the intended answer.

Step 3: Conclusion In the networking classification, radio waves range from 3 kHz to 1 GHz. Hence, the correct option is (D).

Quick Tip to Solve: For networking questions, keep the radio band ending at 1 GHz distinct from the separately listed microwave band.

---

47. Consider the following numList:

8	7	13	1	-9	4

Sequence the following output obtained after each pass of comparison using bubble sort.

(A) 7, 8, 1, -9, 4, 13

(B) 7, 1, -9, 4, 8, 13

(C) -9, 1, 4, 7, 8, 13

(D) 1, -9, 4, 7, 8, 13

Choose the correct answer from the options given below: [5 marks]

(A) (B), (D), (C), (A)

(B) (A), (B), (C), (D)

(C) (A), (B), (D), (C)

(D) (C), (B), (D), (A)

Correct Answer: (C) (A), (B), (D), (C)

Solution:

Concept:

- Ascending bubble sort compares neighbouring elements from left to right and swaps a pair whenever the left element is larger.
- Each complete pass moves the largest remaining unsorted element to the end of the unsorted region; later passes keep the sorted suffix unchanged.

Step 1: Trace the first two passes Begin with 8, 7, 13, 1, -9, 4. On the first pass, swap 8 and 7, then let 13 move right by swapping it successively with 1, -9 and 4. The result is 7, 8, 1, -9, 4, 13, which is arrangement (A). On the second pass, 8 moves right past 1, -9 and 4; the final 13 remains fixed. This produces 7, 1, -9, 4, 8, 13, arrangement (B).

Step 2: Trace the remaining passes and select the order For the third pass, 7 is compared with 1, -9 and 4 and moves to its correct position just before 8. The resulting list is 1, -9, 4, 7, 8, 13, arrangement (D). During the fourth pass, 1 and -9 are swapped, giving -9, 1, 4, 7, 8, 13, arrangement (C). Therefore the chronological sequence is (A), (B), (D), (C). The other alternatives either start with a later pass or place the completely sorted arrangement before an unsorted one.

Step 3: Conclusion The pass order is A, B, D, C. Hence, the correct option is (C).

Quick Tip to Solve: Track the growing sorted suffix: 13, then 8 and 13, then 7, 8 and 13. This fixes the pass order quickly.

---

48. Match List-I with List-II:

List-I (Malware)	List-II (Threat)
(A) Keylogger	(I) It generates revenue using "pay per click" mechanism to charge its clients against the number of clicks on their displayed ads.
(B) Adware	(II) Makes the logs of daily keyboard usage and may send it to an external entity as well.
(C) Spyware	(III) Standalone programs that are capable of replicating on their own and can spread to other computers through the network.
(D) Worm	(IV) Usually tracks internet usage data and sells it to advertisers.

Choose the correct answer from the options given below: [5 marks]

(A) (A) - (II), (B) - (IV), (C) - (I), (D) - (III)

(B) (A) - (II), (B) - (I), (C) - (IV), (D) - (III)

(C) (A) - (I), (B) - (II), (C) - (IV), (D) - (III)

(D) (A) - (III), (B) - (I), (C) - (IV), (D) - (II)

Correct Answer: (B) (A) - (II), (B) - (I), (C) - (IV), (D) - (III)

Solution:

Concept:

- Malware types are identified by their principal behaviour, such as recording input, displaying advertisements, monitoring activity or self-replication.
- A matching question requires every malware description to agree; one correct pairing alone does not establish the correct combination.

Step 1: Identify keylogger and adware A keylogger records keyboard usage, including potentially sensitive text such as passwords, and may send these records to another entity. This is description (II), so (A) matches (II). Adware is associated with advertising and the revenue generated through displayed advertisements, including pay-per-click schemes. This matches description (I), giving (B) with (I). The descriptions identify the primary behaviour rather than claiming that all advertising software necessarily steals data.

Step 2: Identify spyware and worm, then check alternatives Spyware monitors user activity; the description about tracking internet usage and selling that data to advertisers is (IV). Thus (C) matches (IV). A worm is a standalone program capable of replicating itself and spreading through a network, which is description (III), so (D) matches (III). The full mapping is A-II, B-I, C-IV, D-III. The first alternative swaps adware and spyware. The third gives the keylogger the advertising behaviour. The fourth incorrectly gives self-replication to the keylogger and keyboard logging to the worm. Only the second alternative satisfies all four matches.

Step 3: Conclusion The matching is A-II, B-I, C-IV, D-III. Hence, the correct option is (B).

Quick Tip to Solve: Use the names as anchors: keylogger logs keys, adware concerns ads, spyware monitors activity and a worm replicates.

49. Which statement is true about a Trojan? [5 marks]

(A) It is not a malware.

(B) It does not self-replicate or infect other files.

(C) It cannot damage files.

(D) It does not spreads through user interaction.

Correct Answer: (B) It does not self-replicate or infect other files.

Solution:

Concept:

- A Trojan horse is malicious software disguised as a useful or legitimate program so that a user is persuaded to run it.
- Unlike a virus or worm, a Trojan does not characteristically propagate by self-replication or by inserting copies of itself into other files.

Step 1: Identify the distinguishing property A Trojan gains entry by deception: the user believes the program is useful, opens an attachment or installs a seemingly legitimate application. Once executed, its malicious payload may steal information, change files or give an attacker access. Its defining propagation method is not autonomous copying. This distinguishes it from a worm, which can replicate and spread, and a virus, which infects other files.

Step 2: Test the statements against that property The statement that a Trojan does not self-replicate or infect other files describes this distinction and is correct. Saying that it is not malware is false, because a Trojan is a type of malicious software. Saying that it cannot damage files is also false: its payload can delete, alter or expose files. Finally, user interaction is a common means of introducing a Trojan, so the claim that it does not spread through user interaction is incorrect. A Trojan may deliver other malware, but this does not make self-replication its own defining mechanism.

Step 3: Conclusion A Trojan does not self-replicate or infect other files as its propagation mechanism. Hence, the correct option is (B).

Quick Tip to Solve: Differentiate a Trojan by deceptive installation, a virus by infected files and a worm by independent replication.

50. Which malware type disguises itself as legitimate software? [5 marks]

(A) Spyware

(B) Adware

(C) Virus

(D) Trojan

Correct Answer: (D) Trojan

Solution:

Concept:

- A Trojan is malicious software presented as a useful or legitimate program, so the user is persuaded to install or run it.
- Malware names describe different behaviours: a virus replicates through a host, spyware gathers information secretly, and adware displays advertisements.

Step 1: Identify the defining clue The decisive phrase is "disguises itself as legitimate software". A Trojan uses a misleading appearance or claimed purpose, such as a useful utility, to gain the user's trust. Once executed, it can perform a harmful action. Its defining feature is this deception; self-replication is not required.

Step 2: Eliminate the other malware categories Spyware is classified by covert monitoring and collection of user information. Adware is classified by unwanted advertising. A virus is classified by attaching to other files or programs and replicating when the infected host runs. These behaviours do not by themselves define the disguise described in the question. Trojan therefore matches the specific clue most directly.

Step 3: Conclusion The malware that disguises itself as legitimate software is a Trojan. Hence, the correct option is (D).

Quick Tip to Solve: Associate Trojan with a deceptive disguise; distinguish virus by replication and spyware by secret monitoring.

Section B2: Informatics Practices (175 marks)

Attempt either Section B1: Computer Science or Section B2: Informatics Practices, answering all its 35 questions.

51. Which of the following operations can be performed with the UNION clause in SQL? [5 marks]

(A) Fetch only common rows between two tables

(B) Combine rows from two tables without duplicates

(C) Combine rows from two tables allowing duplicates

(D) Combine columns with duplicates

Correct Answer: (B) Combine rows from two tables without duplicates

Solution:

Concept:

- UNION combines the result rows of two compatible SELECT queries and removes duplicate rows.
- UNION ALL retains duplicate rows; a union requires the same number of selected columns with compatible data types.

Step 1: Apply the row-combination rule Suppose one SELECT returns names Ana and Ben, while another returns Ben and Cara. UNION gives Ana, Ben and Cara: Ben appears once even though both input results contain it. This illustrates the operation asked for: combining rows and then eliminating duplicate result rows.

Step 2: Distinguish the alternatives Fetching only common rows is intersection, not union. Combining rows while retaining repeated entries is UNION ALL, so that description is not correct for plain UNION. Combining columns is associated with joins rather than a set union of compatible query results. Therefore, among the listed descriptions, combining rows without duplicates is the correct one.

Step 3: Conclusion UNION combines rows from two tables or query results without duplicates. Hence, the correct option is (B).

Quick Tip to Solve: Check whether the keyword ALL is present: plain UNION removes duplicate rows, while UNION ALL preserves them.

52. Which clause is used to combine rows from two tables that meet a condition? [5 marks]

(A) UNION

(B) INTERSECT

(C) JOIN

(D) MINUS

Correct Answer: (C) JOIN

Solution:

Concept:

- A JOIN combines related rows from two tables according to a join condition, commonly specified using ON.

- A set operator such as UNION, INTERSECT or MINUS operates on compatible query results rather than pairing columns from related rows.

Step 1: Identify the operation required The question asks to combine rows from two tables that meet a condition. For example, a student row and a class row can be paired when their class identifiers are equal. An SQL JOIN expresses this relationship and returns columns from the matched rows in a combined result.

Step 2: Compare the remaining clauses UNION appends compatible query results and normally removes duplicate rows. INTERSECT returns rows common to compatible result sets. MINUS returns rows in the first result that do not appear in the second, in SQL systems that support that name. None of these is the general clause for combining related rows under a join condition. JOIN directly provides the required operation.

Step 3: Conclusion The clause used to combine rows from two tables that meet a condition is JOIN. Hence, the correct option is (C).

Quick Tip to Solve: When two tables are linked by a condition such as matching identifiers, look for JOIN and its ON condition.

53. What is the degree of the resulting table when performing a Cartesian product of two tables with degrees 3 and 4? [5 marks]

(A) 3

(B) 4

(C) 7

(D) 12

Correct Answer: (C) 7

Solution:

Concept:

- The degree of a relation is its number of attributes or columns; its cardinality is its number of rows.
- A Cartesian product contains every column of both input relations, so the resulting degree is the sum of their degrees.

Step 1: Count the attributes in the product Let the first table have attributes A1, A2 and A3. Let the second have B1, B2, B3 and B4. Each row of the Cartesian product contains the first table's three values followed by the second table's four values. Thus every result row has seven attribute positions.

Step 2: Calculate and reject the row-count distractor The degree is therefore

$$3 + 4 = 7.$$

The values 3 and 4 each retain the degree of only one input table and omit the other table's columns. The value 12 comes from multiplying 3 and 4, but degrees are added. It is the cardinalities, if known, that are multiplied to obtain the number of result rows in a Cartesian product.

Step 3: Conclusion The resulting table has degree 7. Hence, the correct option is (C).

Quick Tip to Solve: For a Cartesian product, add column counts to find degree; multiply row counts to find cardinality.

54. Sequence the following steps involved in Project Based Learning:

(A) Fixing of a time frame and processing

(B) Identification of a project

(C) Defining a plan

(D) Providing guidance and monitoring a project

(E) Outcome of a project

Choose the correct answer from the options given below: [5 marks]

(A) (A), (B), (C), (E), (D)

(B) (A), (B), (C), (D), (E)

(C) (B), (A), (D), (C), (E)

(D) (B), (C), (A), (D), (E)

Correct Answer: (D) (B), (C), (A), (D), (E)

Solution:

Concept:

- Project based learning proceeds from choosing a project to planning it, scheduling and processing it, monitoring its progress, and obtaining its outcome.
- Planning establishes what will be done; scheduling assigns a time frame to that plan, and monitoring accompanies its implementation.

Step 1: Order identification, planning and scheduling First identify the project, which is statement (B): the team must know the problem or task it will address. Next define a plan, statement (C), to set out the activities and approach. Then fix a time frame and carry out processing, statement (A), so the planned work can be organised and executed.

Step 2: Place guidance before the outcome Providing guidance and monitoring the project, statement (D), helps check progress and address difficulties during implementation. The outcome of the project, statement (E), follows the work. The complete sequence is therefore (B), (C), (A), (D), (E). The alternatives that begin with a time frame before identification or put the outcome before monitoring do not follow this progression.

Step 3: Conclusion The sequence is identification, plan, time frame and processing, guidance and monitoring, then outcome. Hence, the correct option is (D).

Quick Tip to Solve: Choose the project before planning it, then schedule and monitor the work; place the outcome last.

55. Match List-I with List-II

List-I: AGGREGATE FUNCTIONS	List-II: DESCRIPTION
(A) COUNT()	(I) Returns the average of the values in a specified column
(B) AVG()	(II) Returns the number of records in a table
(C) COUNT(*)	(III) Truncate numerical values to a specified position
(D) ROUND()	(IV) Returns the number of values in a specified column, ignoring NULL values

Choose the correct answer from the options given below: [5 marks]

(A) (A) - (IV), (B) - (III), (C) - (II), (D) - (I)

(B) (A) - (II), (B) - (I), (C) - (IV), (D) - (III)

(C) (A) - (IV), (B) - (I), (C) - (II), (D) - (III)

(D) (A) - (IV), (B) - (I), (C) - (III), (D) - (II)

Correct Answer: (C) (A) - (IV), (B) - (I), (C) - (II), (D) - (III)

Solution:

Concept:

- COUNT(column) counts non-NULL entries in that column, whereas COUNT(*) counts rows whether individual fields are NULL or not.
- AVG(column) returns the arithmetic mean of non-NULL numeric values. The question's ROUND() description refers to controlling numerical precision.

Step 1: Match the three aggregate functions COUNT() in list entry (A) denotes counting values in a specified column, ignoring NULL values, so it matches (IV). AVG(), entry (B), returns an average, so it matches (I). COUNT(*), entry (C), counts the number of records in a table, including rows that contain NULL fields, so it matches (II). These three matches already identify a unique option.

Step 2: Complete the printed match carefully The remaining list entry is ROUND(), entry (D), paired in the printed table with description (III). Strictly, ROUND rounds to a specified precision; truncation discards digits without rounding and is normally performed by TRUNCATE. The table uses imprecise wording here, but the only option that correctly matches all three unambiguous aggregate functions is the third option, with (D) paired to the remaining description.

Step 3: Conclusion The intended matching is COUNT()–IV, AVG()–I, COUNT(*)–II and ROUND()–III, with the wording caveat explained above. Hence, the correct option is (C).

Quick Tip to Solve: Use COUNT(column) versus COUNT(*) to separate non-NULL values from all rows; ROUND and TRUNCATE are distinct functions.

56. Which of the following are the components of successful teamwork?

(A) Teams listen to each other and follow the steps of their own

(B) Share ideas with others

(C) Do not have respect for team members

(D) Participate in completing the project

Choose the correct answer from the options given below: [5 marks]

(A) (B) and (D) only

(B) (A), (C) and (D) only

(C) (A), (B) and (D) only

(D) (B), (C) and (D) only

Correct Answer: (A) (B) and (D) only

Solution:

Concept:

- Successful teamwork requires sharing ideas and contributing to completing the common project.
- Listening must support coordinated action and mutual respect; working independently of agreed team steps undermines cooperation.

Step 1: Test the positive components Sharing ideas with others, statement (B), lets members contribute knowledge and improve the joint work. Participating in completing the project, statement (D), ensures that members contribute to the shared outcome. Both are clear components of successful teamwork.

Step 2: Read both rejected statements in full Statement (A) begins with listening, which is helpful, but adds that teams follow steps of their own. In this contrast, independently following one's own steps rather than coordinating the team's plan is not a sound teamwork component. Statement (C) explicitly says members do not have respect for team members, which is incompatible with successful cooperation. Therefore only (B) and (D) are accepted; combinations containing (A) or (C) fail the complete-statement test.

Step 3: Conclusion Sharing ideas and participating in completing the project are the accepted components. Hence, the correct option is (A).

Quick Tip to Solve: Read an entire statement before selecting it: a positive opening does not rescue an uncooperative second clause.

57. What will be the output of df.isnull().any()? [5 marks]

(A) True if all columns contain null values

(B) False if no missing values are present in any column

(C) True or false for each column depending whether it contains missing values or not

(D) A count for missing values in all columns

Correct Answer: (C) True or false for each column depending whether it contains missing values or not

Solution:

Concept:

- The isnull() method makes a Boolean DataFrame of the same shape as the original: each missing value becomes True and each present value becomes False.
- For a DataFrame, any() uses axis=0 by default. It tests each column separately and returns a Series indexed by the column names.

Step 1: Apply the missing-value test First evaluate `df.isnull()`. A column containing at least one missing entry has at least one True in its Boolean mask. A column containing no missing entry has only False values. The operation detects missing values; it does not count them or combine the entire table into a single Boolean.

Step 2: Reduce each column Next apply `.any()` to that mask. Its result is True for a column if any entry of that column is missing, and False otherwise. Thus the output contains one Boolean per column. The first choice incorrectly requires all columns to contain nulls; the second describes only a special case and omits the per-column result. A count would require a counting reduction such as sum() on the Boolean mask, rather than any().

Step 3: Conclusion The result gives True or False for each column according to whether that column contains missing values. Hence, the correct option is (C).

Quick Tip to Solve: Read a chain of Pandas methods from left to right: isnull() creates the mask, then any() reduces each column.

58. Identification of a project leader generally falls under which of the following steps of Project Based Learning? [5 marks]

(A) Defining a plan

(B) Fixing of a time frame and processing

(C) Identification of a project

(D) Providing guidance and monitoring a project

Correct Answer: (A) Defining a plan

Solution:

Concept:

- A project plan translates the selected project into an organised course of action: tasks, roles, responsibilities and the resources needed.

- A project leader coordinates team members and their responsibilities, so selecting that person belongs to organising the plan before implementation.

Step 1: Locate the decision in the project process The decision here concerns who will lead the group, rather than which problem the group should investigate. Once the project is identified, the team must arrange how the work will be carried out. Naming its leader is part of that arrangement and therefore fits defining a plan.

Step 2: Distinguish the remaining stages Identification of a project chooses the subject or problem itself; it does not primarily assign leadership. Fixing a time frame and processing concerns scheduling and executing the planned work. Providing guidance and monitoring concerns supervising progress and helping resolve difficulties during the project. These activities may involve the leader after appointment, but they are not the stage in which the leadership role is normally assigned.

Step 3: Conclusion Identifying the project leader generally falls under defining a plan. Hence, the correct option is (A).

Quick Tip to Solve: Separate choosing the project from choosing the people and roles: the latter is a planning decision.

59. In data handling, pandas support _____ index values. [5 marks]

(A) Unique

(B) Non-unique

(C) Homogeneous

(D) NaN

Correct Answer: (B) Non-unique

Solution:

Concept:

- A Pandas index supplies labels for the positions along an axis. Labels identify data for selection and alignment; they are not required to behave like a relational database primary key.
- Pandas permits repeated labels in an index. Such an index is called non-unique, even though the data still occupy distinct physical positions.

Step 1: Recognise the supported index property For example, a Series may use the index labels `[x, x, y]`. The first two entries have the same label but remain separate entries. Label-based selection for x can therefore return more than one value. This illustrates the support for non-unique index values asked about here.

Step 2: Compare the alternatives Unique labels are possible and often convenient, but uniqueness is not a compulsory restriction; the distinctive supported feature in this question is non-uniqueness. Homogeneous describes sameness of type and is not the property at issue. NaN denotes a missing-value marker rather than the repeated-label property. Therefore the option referring to non-unique labels best completes the statement.

Step 3: Conclusion Pandas supports non-unique index values. Hence, the correct option is (B).

Quick Tip to Solve: Do not treat a Pandas index as a mandatory unique key: repeated labels can be valid, so selection may return several entries.

60. Match List-I with List-II

List-I	List-II
(A) pip install pandas	(I) one-dimensional array containing a sequence of values
(B) Numpy array	(II) can have different data types
(C) Pandas DataFrame	(III) requires homogeneous data
(D) Series	(IV) install Pandas library

Choose the correct answer from the options given below: [5 marks]

(A) (A) - (IV), (B) - (III), (C) - (II), (D) - (I)

(B) (A) - (I), (B) - (II), (C) - (III), (D) - (IV)

(C) (A) - (I), (B) - (II), (C) - (IV), (D) - (III)

(D) (A) - (III), (B) - (IV), (C) - (I), (D) - (II)

Correct Answer: (A) (A) - (IV), (B) - (III), (C) - (II), (D) - (I)

Solution:

Concept:

- pip installs Python packages. The command pip install pandas installs the Pandas library into the selected Python environment.
- A standard NumPy array holds homogeneous data; a DataFrame can hold columns of different types; a Series is a one-dimensional labelled sequence.

Step 1: Match the installation command and NumPy array The installation command is not itself an array or a DataFrame: it performs the library installation, so it matches statement IV. A NumPy array uses a common element dtype for its stored data, matching the requirement for homogeneous data in statement III. Thus the first two links are fixed by the meanings of the command and array.

Step 2: Match the Pandas structures and check the code A DataFrame is two-dimensional and may have one numeric column alongside another text column, so it matches statement II, which allows different data types. A Series is the one-dimensional structure containing a sequence of values, matching statement I. The resulting correspondence agrees with the first offered combination. The other combinations assign installation or dimensional descriptions to the wrong objects and fail one or more of these direct definitions.

Step 3: Conclusion The correct correspondence is installation to IV, NumPy array to III, DataFrame to II and Series to I. Hence, the correct option is (A).

Quick Tip to Solve: Fix the unmistakable match pip install pandas to installation first; then use array homogeneity and Series dimensionality.

61. While appending two dataframes, if we want to raise error in case row labels are duplicate, we shall use which parameter of append() method? [5 marks]

(A) ignore_index = False

(B) row_index = False

(C) in_place = True

(D) verify_integrity = True

Correct Answer: (D) verify_integrity = True

Solution:

Concept:

- Appending DataFrames combines rows and normally retains their row labels unless a fresh index is requested.
- The verify_integrity parameter requests a check that the resulting index contains no duplicates. Setting it to True raises an error if that check fails.

Step 1: Apply the requested integrity check The question explicitly wants an error when row labels are duplicated, not automatic relabelling. The appropriate argument in the append() interface described is `verify_integrity=True`. For example, if both inputs contribute a row labelled 0 and their labels are retained, this check detects the collision instead of silently accepting it.

Step 2: Eliminate the other parameters `ignore_index=False` preserves the existing labels but does not itself demand that they be unique. It therefore cannot guarantee the requested error. `row_index=False` is not the duplicate-index verification parameter of append(). Likewise, `in_place=True` is not the requested append integrity control; it neither names nor enables the check. The relevant switch specifically verifies the combined index.

Step 3: Conclusion Use `verify_integrity=True` to raise an error for duplicate row labels. Hence, the correct option is (D).

Quick Tip to Solve: Preserving labels and validating labels are different tasks: ignore_index controls relabelling; verify_integrity controls duplicate detection.

62. Order the following steps to for creation of series from NumPy Arrays.

(A) >>> **import numpy as np**

(B) >>> **array1 = np.array([1,2,3,4])**

(C) >>> **import pandas as pd**

(D) >>> **series3 = pd.Series(array1)**

(E) >>> **print(series3)**

Choose the correct answer from the options given below: [5 marks]

(A) (A), (D), (C), (B), (E)

(B) (A), (C), (B), (D), (E)

(C) (B), (A), (D), (C), (E)

(D) (C), (B), (D), (A), (E)

Correct Answer: (B) (A), (C), (B), (D), (E)

Solution:

Concept:

- A Python name must be defined before an expression uses it: np requires importing NumPy, and pd requires importing Pandas.
- To construct a Series from array1, first create array1, then call pd.Series(array1), and finally print the resulting Series.

Step 1: List the dependencies of each statement The statement `array1 = np.array([1,2,3,4])` requires the NumPy import. The statement `series3 = pd.Series(array1)` requires both the Pandas import and the completed array. Printing series3 must occur after series3 is assigned. The two import statements may be interchanged provided these dependencies are respected.

Step 2: Test the offered sequences The second sequence imports NumPy, imports Pandas, creates the array, constructs the Series, and prints it. Every required name is therefore available when used. The first sequence constructs the Series before the Pandas import and before array creation. The third starts by creating the array before importing NumPy. The fourth also creates the array before importing NumPy. Those sequences therefore fail with undefined names when executed in a fresh session.

Step 3: Conclusion The valid offered order is the NumPy import, Pandas import, array creation, Series creation and printing. Hence, the correct option is (B).

Quick Tip to Solve: For ordering code, draw dependencies mentally: imports before aliases, array before Series, and Series before print.

63. What will be the output of the command df.count(axis = 1)? [5 marks]

(A) The total number of row in a dataframe.

(B) The total number of values in each row.

(C) The total number of columns in each dataframe

(D) The total number of NaN values in each row

Correct Answer: (B) The total number of values in each row.

Solution:

Concept:

- DataFrame.count() counts non-missing observations. It excludes missing values such as NaN rather than counting them as ordinary entries.
- With axis=1, the method works across columns separately for each row. Its output is a Series giving a count for every row.

Step 1: Interpret the axis argument In `df.count(axis=1)`, axis=1 makes each row the unit whose available entries are counted across the columns. It does not calculate the total number of rows in the table. If one row has four non-missing entries, its output count is 4; if another row has only three non-missing entries, its count is 3.

Step 2: Select the intended row-wise description The choice describing the number of values in each row is the intended answer, with the essential qualification that count() includes only non-missing values. The total number of rows would come from the row dimension of the DataFrame. The number of columns is a table dimension, not a separate observed-value count for each row. Counting NaN entries would instead require a missing-value mask followed by a row-wise sum.

Step 3: Conclusion The command returns the count of non-missing values in each row, corresponding to the description of values in each row. Hence, the correct option is (B).

Quick Tip to Solve: Remember both parts of count(axis=1): axis=1 means row-wise, and count excludes missing values.

64. What does pivot_table( ) do in pandas that pivot( ) cannot? [5 marks]

(A) It aggregates the values from rows with duplicate entries for the specified columns.

(B) It does not allow duplicate data to be reshaped.

(C) It merges multiple dataframes into one.

(D) It converts NaN values to zero.

Correct Answer: (A) It aggregates the values from rows with duplicate entries for the specified columns.

Solution:

Concept:

- A pivot rearranges data into a table whose rows and columns are identified by selected keys.
- The pivot_table method also performs an aggregation, such as mean, sum or count, when multiple observations belong to the same output cell.

Step 1: Identify the duplicate-key problem Suppose two records have the same chosen row key and column key. A plain pivot cannot decide which of their values should occupy that single cell and therefore cannot reshape such duplicate combinations directly.

Step 2: Use aggregation to form one cell The pivot_table method combines the values associated with each key combination using its aggregation function. For example, two sales records for the same product and month can be summed into one total. Thus the first statement describes its additional capability. The second states the restriction of pivot instead; merging frames is a separate operation, and replacing missing values by zero is not the defining difference.

Step 3: Conclusion pivot_table aggregates duplicate key combinations into a single value for each output cell. Hence, the correct option is (A).

Quick Tip to Solve: If several records map to one pivot cell, look for an aggregation function such as sum or mean.

65. _____ will display the total number of values for each column or row of a DataFrame. [5 marks]

(A) count()

(B) tail(n)

(C) head(n)

(D) max(n)

Correct Answer: (A) count()

Solution:

Concept:

- A DataFrame stores values in rows and columns, and count computes the number of non-missing entries.
- The axis argument selects the direction: axis=0 counts down columns; axis=1 counts across each row.

Step 1: Identify the required operation The question asks for the number of values rather than for selected records or an extreme value. The count method performs this counting operation separately along the chosen axis, returning one count for each column or row.

Step 2: Distinguish the offered methods Calling df.count() gives column-wise counts by default, whereas df.count(axis=1) gives row-wise counts. Missing entries such as NaN are excluded, so a row containing a missing entry can have a count smaller than the number of columns. The tail(n) and head(n) methods display the last or first n rows. The max method finds maximum values rather than the number of entries; its displayed n argument does not make it a counting method.

Step 3: Conclusion count() is the method that displays counts of non-missing values along the selected axis. Hence, the correct option is (A).

Quick Tip to Solve: Use count for present values; use axis=1 to obtain a separate count for every row.

66. Which of the following functions are used to split the data into groups based on some criteria?

[5 marks]

(A) split()

(B) GROUP BY()

(C) first()

(D) reset_index()

Correct Answer: (B) GROUP BY()

Solution:

Concept:

- Grouping divides observations into sets according to one or more keys before an operation is applied.
- In pandas, the groupby method implements the split-apply-combine approach to grouped data analysis.

Step 1: Identify the grouping operation Records sharing the same criterion, such as the same department or class, belong in the same group. This is the grouping operation represented by GROUP BY in the options. In actual pandas code, its method spelling is lower-case groupby, for example df.groupby('Class').

Step 2: Eliminate unrelated operations A string split method separates text around a delimiter; it is not the DataFrame operation for forming groups of records by a criterion. first selects an initial entry or entries in its applicable context and does not itself specify the grouping keys. reset_index changes the index representation, often turning index levels into columns, rather than splitting the records into groups. Therefore the intended grouping option is the second one.

Step 3: Conclusion GROUP BY() denotes the intended grouping operation, implemented as groupby() in pandas. Hence, the correct option is (B).

Quick Tip to Solve: Associate grouping by a column with df.groupby(column); apply count, sum or mean after forming the groups.

67. Arrange the following steps to calculate the total number of values for each row of a DataFrame.

(A) Load the DataFrame as 'df'

(B) Use the df.count() method

(C) Print the result

(D) import pandas as pd

Choose the correct answer from the options given below: [5 marks]

(A) (D), (C), (B), (A)

(B) (D), (A), (B), (C)

(C) (D), (C), (A), (B)

(D) (C), (B), (D), (A)

Correct Answer: (B) (D), (A), (B), (C)

Solution:

Concept:

- A Python program must import pandas before referring to its pd alias to construct or load a DataFrame.
- Counting requires an existing DataFrame, and the result must be computed before it is printed.

Step 1: Resolve the dependencies Step (D) imports pandas as pd, making the library available. Step (A) then loads the DataFrame as df. Once df exists, step (B) can call its count method. Finally step (C) displays the computed result. Each stage supplies something needed by the next, giving the dependency order (D), (A), (B), (C).

Step 2: Select the valid sequence The second option is the only sequence that imports the library, obtains the data, computes the counts and then prints the result. The first option prints before counting and loads df last. The third also prints before loading the DataFrame and computing its counts. The fourth prints before the calculation and attempts counting before loading df. For the row-wise counts requested in the question, the actual call must specify df.count(axis=1); the printed generic df.count() name denotes the counting step, but its default axis would instead give column counts. The axis detail does not change the order of the four program stages.

Step 3: Conclusion The required order is import, load, count, print. Hence, the correct option is (B).

Quick Tip to Solve: Order program statements by their dependencies; for counts across each row, use axis=1 in the counting call.

68. Extracting a part of a series can be done through _____ . [5 marks]

(A) Indexing

(B) Subtracting

(C) Slicing

(D) Partitioning

Correct Answer: (C) Slicing

Solution:

Concept:

- A Series is a one-dimensional labelled data structure.
- Slicing selects a range or portion of a sequence or Series using boundaries, rather than retrieving just one scalar entry.

Step 1: Identify selection of a portion The phrase extracting a part refers to selecting a contiguous range of entries. A slice expresses such a range through start and stop positions or labels, with an optional step where supported. In positional selection, s.iloc[1:4] retrieves positions 1, 2 and 3.

Step 2: Choose the precise term Indexing is the broader process of selecting data using positions or labels and can retrieve an individual item. Slicing is the specific form that extracts a portion, so it is the most precise answer among the alternatives. Subtracting is an arithmetic operation and changes values rather than selecting a range. Partitioning is not the standard term used here for Series range extraction. Label-based slices and positional slices can have different endpoint rules, so their semantics should not be confused.

Step 3: Conclusion Extracting a portion of a Series is called slicing. Hence, the correct option is (C).

Quick Tip to Solve: Use iloc for a positional slice and remember that its stop position is excluded.

69. Which type of chart is most suitable for statistical summary, including quartiles? [5 marks]

(A) Line Plot

(B) Box Plot

(C) Pie chart

(D) Scatter Plot

Correct Answer: (B) Box Plot

Solution:

Concept:

- Quartiles divide ordered observations into four parts; the median is the second quartile.
- A box plot summarises location and spread using the median, lower and upper quartiles, and whiskers, with possible outliers shown separately.

Step 1: Match quartiles to the chart The box extends from the first quartile to the third quartile, so its length represents the interquartile range. A line inside the box marks the median. This directly displays the quartile-based statistical summary requested by the question.

Step 2: Compare the alternatives A line plot is generally used to reveal changes or trends across ordered observations, often time. A pie chart displays how categories contribute to a whole. A scatter plot shows pairs of numerical observations and helps identify relationships between two variables. None of those charts provides the quartiles as its standard visual structure. A box plot is therefore the suitable choice when the aim is to compare medians and spreads across distributions. Its whisker convention should be checked before interpreting extremes.

Step 3: Conclusion A box plot presents a statistical summary including quartiles and the median. Hence, the correct option is (B).

Quick Tip to Solve: For a quartile or interquartile-range question, look for the box boundaries and median line of a box plot.

70. What does the explode property in the pie chart do? [5 marks]

(A) Adds border to the pie chart

(B) Changes the chart color

(C) Specifies the fraction of the radius with which to expand each slot

(D) Rotates the chart clockwise

Correct Answer: (C) Specifies the fraction of the radius with which to expand each slot

Solution:

Concept:

- A pie chart divides a circle into wedges representing proportions of a total.
- The explode parameter supplies a radial offset for each wedge, measured as a fraction of the pie radius.

Step 1: Interpret the offset values A zero value leaves a wedge in its ordinary position. A positive value moves that wedge outward from the centre along its radial direction, separating it visually from the remaining pie. For example, an offset of 0.1 displaces the chosen wedge by one tenth of the radius.

Step 2: Distinguish other chart properties The third option describes this fractional-radius displacement, even though it calls the wedge a slot. Exploding a wedge highlights a category without changing the underlying quantity or its angular share. Adding a border concerns edge styling, changing colours concerns the colour settings, and changing the starting angle or direction concerns rotation or orientation. These are separate properties and are not controlled by explode. An explode sequence normally gives one offset for each wedge, allowing selected slices to be emphasised independently.

Step 3: Conclusion The explode property specifies the fractional-radius outward displacement of each pie wedge. Hence, the correct option is (C).

Quick Tip to Solve: Set all explode entries to zero except the slice you wish to highlight; 0.1 moves it outward by 10 percent of the radius.

71. Which statements are true regarding the scatter plot?

(A) Shows how two variables are correlated

(B) Summarizes categorical data

(C) Displays proportions as a whole

(D) Uses dots to represent data

Choose the correct answer from the options given below: [5 marks]

(A) (A), (B) and (D) only

(B) (A), (B) and (C) only

(C) (A), (B), (C) and (D)

(D) (A) and (D) only

Correct Answer: (D) (A) and (D) only

Solution:

Concept:

- A scatter plot displays paired numerical observations as points on two axes. Each dot represents one pair of values for the two variables.
- The pattern of dots helps identify positive correlation, negative correlation or no clear relationship; it does not by itself prove causation.

Step 1: Test the true statements Statement (A) is true: comparing the positions of the points shows how two quantitative variables are related. An upward trend suggests positive association, whereas a downward trend suggests negative association. Statement (D) is also true because individual observations are represented by dots rather than by slices or category bars.

Step 2: Exclude the other chart descriptions Statement (B) describes a common use of a bar chart, which compares categories; it is not the defining purpose of a scatter plot. Statement (C) describes a part-to-whole display such as a pie chart. Thus both (B) and (C) must be excluded. The first three combinations each include at least one of these false statements, while the fourth contains precisely the two true statements.

Step 3: Conclusion Only statements A and D are true. Hence, the correct option is (D).

Quick Tip to Solve: For scatter-plot statements, look for paired numerical values and dots; categories suggest bars and proportions suggest a pie chart.

72. Match List-I with List-II

List-I	List-II
(A) xticks, yticks	(I) Sets the user defined markings for x and y-axis
(B) kind	(II) Displayed by default associating the colours with the plotted data
(C) legend	(III) A symbol representing data value
(D) marker	(IV) Customizing different plot types

Choose the correct answer from the options given below: [5 marks]

(A) (A) - (I), (B) - (II), (C) - (III), (D) - (IV)

(B) (A) - (I), (B) - (III), (C) - (II), (D) - (IV)

(C) (A) - (I), (B) - (IV), (C) - (II), (D) - (III)

(D) (A) - (III), (B) - (IV), (C) - (I), (D) - (II)

Correct Answer: (C) (A) - (I), (B) - (IV), (C) - (II), (D) - (III)

Solution:

Concept:

- Axis ticks mark positions on an axis, while a legend identifies plotted series by colours or symbols.
- The kind parameter selects a plot type, whereas a marker specifies the symbol used for an individual plotted value.

Step 1: Match the axis and plot-type controls The functions xticks and yticks control the user-defined tick positions or markings on the horizontal and vertical axes. Therefore (A) matches (I). The kind parameter selects among plot types, such as bar, line or histogram in a plotting interface that supports it, so (B) matches (IV). It is not the name for a symbol representing one observation.

Step 2: Match the explanatory key and data symbol A legend associates colours or other visual styles with the plotted data series, so (C) matches (II). A marker is the symbol, such as a dot, square or triangle, used to represent a data value, so (D) matches (III). Combining the four independent matches gives the third option. The first two options incorrectly pair kind, and the fourth incorrectly treats axis ticks as a data symbol.

Step 3: Conclusion Axis ticks match I, kind matches IV, legend matches II, and marker matches III. Hence, the correct option is (C).

Quick Tip to Solve: Match the distinctive terms first: kind selects the plot kind, and marker is the point symbol; then assign ticks and legend.

73. How can you save a matplotlib figure to a file? [5 marks]

(A) plt.savefile()

(B) plt.savefig(filename)

(C) plt.storefig(filename)

(D) plt.save(filename)

Correct Answer: (B) plt.savefig(filename)

Solution:

Concept:

- Matplotlib pyplot provides savefig to export the current figure to an image or another supported file format.
- The filename argument supplies the destination; its extension can indicate a format such as PNG or PDF.

Step 1: Identify the figure-export function With pyplot imported as plt, the command is `plt.savefig(filename)`. It writes the current figure to the specified file. For example, `plt.savefig("chart.png")` saves a PNG image of the plot. The filename is an argument, so the destination is passed inside the parentheses rather than being inferred from the displayed chart title.

Step 2: Eliminate the similar-looking names The names `plt.savefile()` and `plt.storefig(filename)` do not identify the required pyplot function. The fourth choice, `plt.save(filename)`, is also incorrect. They sound plausible because they contain save or figure-related words, but the actual function name combines save and fig. This operation exports the plotted figure, including its axes and graphical elements, to a file that can be opened or shared later.

Step 3: Conclusion Use `plt.savefig(filename)` to save the figure. Hence, the correct option is (B).

Quick Tip to Solve: Remember the exact pyplot verb: savefig saves the figure; give it a filename such as chart.png.

74. Which of the following determines the height of each column in a histogram? [5 marks]

(A) The number of intervals

(B) The frequency of datapoints within the range

(C) The x-axis values

(D) The y-axis values

Correct Answer: (B) The frequency of datapoints within the range

Solution:

Concept:

- A basic frequency histogram groups numerical observations into contiguous class intervals, called bins.
- For equal-width bins in an ordinary frequency histogram, the height of a column is the number of

observations falling in that bin.

Step 1: Interpret one histogram column Choose a bin, such as a range of examination scores, and count the data points that fall within it. That count is the frequency represented by the height of its column. A bin containing more observations produces a taller column than another equal-width bin containing fewer observations. Thus the count within the range determines the height in the standard frequency histogram intended here.

Step 2: Distinguish the axes from the determining quantity The number of intervals determines how many bins are drawn, rather than the frequency of an individual bin. The x-axis values locate the numerical ranges. Although the vertical axis displays the frequency scale and can be used to read the height, saying y-axis values does not identify the data quantity that determines it. The explicit frequency description in the second option is therefore the precise answer.

Step 3: Conclusion The frequency of data points within the range determines the column height. Hence, the correct option is (B).

Quick Tip to Solve: In a standard frequency histogram, read ranges horizontally and counts vertically; a taller bar means more observations in that bin.

---

75. Arrange the following steps in sequence to transmit a message over the internet.

(A) The message is displayed on a web page

(B) Packets are reassembled at the destination

(C) The message is disassembled as packets

(D) Packets are transmitted over the network

Choose the correct answer from the options given below: [5 marks]

(A) (A), (B), (C), (D)

(B) (A), (C), (B), (D)

(C) (B), (A), (D), (C)

(D) (C), (D), (B), (A)

Correct Answer: (D) (C), (D), (B), (A)

Solution:

Concept:

- Internet transmission divides a message into packets so that smaller units can travel over the network.
- The destination reassembles the received packets before the application can present the complete message.

Step 1: Place the sender-side operations first Before a message can travel as packets, it must be disassembled into packets. Statement (C) is therefore the first step. Those packets are then transmitted over the network, which is statement (D). Sending the pieces must occur before the receiver can put them together; neither display nor reassembly can be the first operation in this sequence.

Step 2: Complete the receiver-side operations After transmission, the packets are reassembled at

the destination, giving statement (B). The complete message can then be displayed on a web page, giving statement (A). The causal order is therefore (C), (D), (B), (A). The other three sequences begin with display or reassembly, placing a receiver-side result before the transmission and packet-creation operations needed to produce it.

Step 3: Conclusion The order is split into packets, transmit, reassemble, then display. Hence, the correct option is (D).

Quick Tip to Solve: Use the dependency chain: split, send, reassemble, display. Each step needs the result of the preceding step.

76. Which of the browser settings contains cookies and site data? [5 marks]

(A) General Panel

(B) Privacy and Security Panel

(C) Home panel

(D) Search Panel

Correct Answer: (B) Privacy and Security Panel

Solution:

Concept:

- Browser cookies are small pieces of site-related data stored by the browser to remember information about visits.
- Controls for stored site data, tracking and related permissions belong with privacy and security settings.

Step 1: Identify the relevant settings category Cookies and site data affect what websites can store about a user and their browsing sessions. The panel concerned with these controls is Privacy and Security. This category lets users inspect or manage stored data and related privacy choices. The question asks for the settings category, rather than the name of a cookie or the action of deleting it.

Step 2: Eliminate unrelated panels The General panel usually groups broad browser preferences, such as startup or download behaviour. The Home panel concerns the home page or new-window presentation. The Search panel concerns search engines and search preferences. None of those names specifically identifies the controls for cookies and site data. Privacy and Security is the category that directly matches both the nature of the data and the user control involved.

Step 3: Conclusion Cookies and site data are found in the Privacy and Security Panel. Hence, the correct option is (B).

Quick Tip to Solve: Associate cookies and site data with privacy; associate the Search panel with search-engine choices.

77. Which topology has a backbone wire that is shared among all nodes? [5 marks]

(A) Ring

(B) Bus

(C) Star

(D) Mesh

Correct Answer: (B) Bus

Solution:

Concept:

- Network topology describes the arrangement of nodes and their communication links.
- A bus topology connects devices to a common backbone cable that serves as the shared communication medium.

Step 1: Match the shared-backbone description In a bus topology, all nodes connect to the same main cable. A transmission travels along this shared backbone, and the intended recipient recognises the message addressed to it. The key feature in the question is that the backbone wire is shared among all nodes, which directly identifies a bus arrangement.

Step 2: Compare the alternative arrangements A ring connects nodes in a closed loop, with each node linked along the ring rather than attached to one common straight backbone. A star connects every node through a central device using separate links. A mesh uses multiple interconnections between nodes, potentially a link for each pair in a full mesh. These structures do not fit the single shared-backbone description. The common cable is therefore the deciding clue, independent of the number of connected computers.

Step 3: Conclusion The topology is Bus. Hence, the correct option is (B).

Quick Tip to Solve: Remember bus as the shared backbone; star uses a central device, ring forms a loop, and mesh has many links.

78. Which of the following can be viewed as Project Specification?

(A) Item Category is Book, then the Item Type can be either Academic or Non-Academic, and Item Id should be prefixed with a 'B'.

(B) Teenagers spend a good amount of time on social networking sites, and it plays a vital role in their behaviour.

(C) Total population, total male population and total female population aged 10 to 24 in India.

(D) Realising the importance of Reduce, Reuse and Recycle, the Bookworm club every year organises a Book Donation Camp.

Choose the correct answer from the options given below: [5 marks]

(A) (A), (C) and (D) only

(B) (B), (C) and (D) only

(C) (A) and (C) only

(D) (B) and (D) only

Correct Answer: (C) (A) and (C) only

Solution:

Concept:

- A project specification states the concrete requirements, data items or rules that a proposed project must implement.
- A motivation or a general observation describes the context; it does not by itself specify the project data or processing rules.

Step 1: Identify the concrete requirements Statement (A) specifies permitted values for Item Type when Item Category is Book, together with a precise prefix rule for Item Id. These are directly implementable validation and classification requirements. Statement (C) identifies the population measures, age group and country whose data are to be handled. It therefore gives a definite data requirement for a population project.

Step 2: Separate background from specifications Statement (B) is a broad observation about teenagers and social networking; it gives no concrete data fields or implementation rule. Statement (D) describes a club activity and its environmental motivation, without defining the proposed software or data-analysis requirements. Thus (B) and (D) are contextual statements, while (A) and (C) identify project requirements. Options containing (D), or replacing (A) with (B), do not give the required pair.

Step 3: Conclusion The project specifications are statements A and C only. Hence, the correct option is (C).

Quick Tip to Solve: Look for an implementable rule or named data requirement; a useful social activity alone is not a project specification.

79. The physical organization of devices and connecting media in a network is called _____.

[5 marks]

(A) Web Server

(B) Web page

(C) WAN

(D) Topology

Correct Answer: (D) Topology

Solution:

Concept:

- Network topology describes the arrangement of network devices and the links that connect them.
- Physical topology concerns the actual placement and interconnection of devices and transmission media, such as a star, bus or ring.

Step 1: Recognise the defining phrase The question asks about the physical organisation of both devices and connecting media. This is precisely what physical network topology describes. For example, devices linked to a central switch form a star topology; devices sharing a common backbone form a bus topology. The term therefore identifies the structure of the connections rather than a particular service or document.

Step 2: Eliminate the other terms A web server is hardware or software that responds to web requests and supplies website resources. A web page is a document that a browser displays. Neither names the arrangement of network links. WAN means Wide Area Network and classifies a network by its geographical extent; networks within that category can still have different physical topologies. Only Topology matches the definition in the question.

Step 3: Conclusion The physical organisation of devices and connecting media is called topology. Hence, the correct option is (D).

Quick Tip to Solve: For questions about arrangement or connection pattern choose topology; for geographical coverage distinguish LAN, MAN and WAN.

80. _____ attribute name is used to display a NumPy ndarray having all the values in the DataFrame, without the axes labels. [5 marks]

(A) DataFrame.columns

(B) DataFrame.size

(C) DataFrame.shape

(D) DataFrame.values

Correct Answer: (D) DataFrame.values

Solution:

Concept:

- A DataFrame stores tabular data with row labels (an index), column labels and the underlying cell values.
- The `values` attribute returns a NumPy ndarray of the DataFrame values, without carrying its row and column labels.

Step 1: Select the attribute that returns cell data The required result is an ndarray containing all cell values rather than labels or metadata. Accessing `DataFrame.values` obtains this underlying array. For example, a frame with two rows and two numeric columns yields a two-dimensional array containing those four numbers; the names assigned to its index and columns do not accompany those numbers in the array.

Step 2: Compare the remaining attributes `DataFrame.columns` gives the column labels, which the question explicitly excludes. `DataFrame.size` gives the total number of elements as a scalar, rather than the elements themselves. `DataFrame.shape` gives a tuple containing the number of rows and columns. Neither size nor shape returns all cell values. Consequently, values is the only attribute among the choices that matches both the content and the requested return type.

Step 3: Conclusion The attribute is `DataFrame.values`. Hence, the correct option is (D).

Quick Tip to Solve: Remember: columns gives labels, shape gives dimensions, size gives the element count, and values gives the data array.

81. Match List-I with List-II

List-I	List-II
ETIQUETTE	DESCRIPTION
(A) Credibility	(I) Choose your password wisely.
(B) Security	(II) Cautious while making comments.
(C) Reliability	(III) No copyright violation.
(D) Ethical	(IV) Think before you upload anything.

Choose the correct answer from the options given below: [5 marks]

(A) (A) - (I), (B) - (II), (C) - (III), (D) - (IV)

(B) (A) - (I), (B) - (III), (C) - (II), (D) - (IV)

(C) (A) - (II), (B) - (I), (C) - (IV), (D) - (III)

(D) (A) - (III), (B) - (IV), (C) - (I), (D) - (II)

Correct Answer: (C) (A) - (II), (B) - (I), (C) - (IV), (D) - (III)

Solution:

Concept:

- Cyber etiquette includes responsible comments, secure access practices, dependable sharing and respect for the rights of creators.
- Match each description to its most direct category, then check the complete four-part mapping rather than selecting from a single familiar pair.

Step 1: Fix the security and ethical pairs Security is directly related to choosing a password wisely, so (B) matches (I). Ethical conduct includes avoiding copyright violation, since using another creator's work must respect the rights attached to it, so (D) matches (III). These two unambiguous matches already identify the third combination among the options.

Step 2: Complete and verify the mapping Credibility concerns the trustworthiness of a person's online behaviour, so being cautious while making comments matches (A) with (II). Reliability of shared information requires thinking before uploading anything, so (C) matches (IV). The full mapping is A-II, B-I, C-IV and D-III. The first and second options wrongly connect credibility to password choice, while the fourth connects security to thinking before uploads and reliability to passwords; these do not give the direct category matches.

Step 3: Conclusion The correct mapping is A-II, B-I, C-IV, D-III. Hence, the correct option is (C).

Quick Tip to Solve: Start with the clearest anchors: passwords mean security, and respecting copyright means ethical behaviour.

82. A permission agreement between the creator and the party who wants to use the creator's work is called _____. [5 marks]

(A) Trademark

(B) Copyright

(C) Patent

(D) Licensing

Correct Answer: (D) Licensing

Solution:

Concept:

- A licence is permission granted by a rights holder to another party to use protected work under specified terms.
- A licensing agreement records that permission and may specify the permitted uses, duration, fees and restrictions.

Step 1: Identify the relationship described The question describes a creator and another party who wishes to use the creator's work. Their permission agreement concerns the authorised use of that work. This is licensing: the owner grants defined rights of use to a licensee. Permission to use a work does not necessarily transfer its ownership; the agreement controls what the user may do.

Step 2: Distinguish protection from permission Copyright is the legal protection attached to an original creative work, rather than the agreement through which someone else receives permission to use it. A trademark protects identifying signs such as brand names or logos. A patent protects an invention subject to the relevant legal requirements. Although an intellectual property right can itself be licensed, none of these names the permission agreement requested here. The agreement is therefore licensing.

Step 3: Conclusion A permission agreement for the use of the creator's work is licensing. Hence, the correct option is (D).

Quick Tip to Solve: Separate the right from permission to use it: copyright protects a work, while a licence states the authorised uses.

83. Using another person's work, without permission or without paying for it, if it is being sold is known as_____. [5 marks]

(A) Cyber crime

(B) Fair use

(C) Copyright infringement

(D) Identity theft

Correct Answer: (C) Copyright infringement

Solution:

Concept:

- Copyright protects the rights of creators over the use and reproduction of their original works.
- Copyright infringement occurs when a protected work is used in a way that violates those rights, such as unauthorised copying or use without the required permission or payment.

Step 1: Apply the description in the question The defining action is using another person's work without the necessary permission or, where payment is required for the intended use, without paying for it. The question is concerned specifically with a creator's work and the rights governing its use. The named violation is copyright infringement. Merely finding a work online does not grant unrestricted permission to reproduce or use it.

Step 2: Reject the alternatives carefully Cyber crime is a broad category of unlawful activities involving computers or networks; it is not the specific intellectual-property term requested. Fair use refers to certain limited uses permitted by law, so it cannot label every unauthorised use described here. Identity theft means using another person's identifying information to impersonate them or commit fraud. This question describes misuse of creative work, rather than theft of a person's identity.

Step 3: Conclusion The described unauthorised use is copyright infringement. Hence, the correct option is (C).

Quick Tip to Solve: When the issue concerns another creator's work, check copyright; when it concerns pretending to be another person, check identity theft.

84. Which of the following is an example of cyberbullying? [5 marks]

(A) Sending phising links to known friends or aquaintance.

(B) Creating fake profiles of your friends and uploading embarrassing pictures.

(C) Blocking your friends from any communication with you.

(D) Hacking in to your friends system and read his mails

Correct Answer: (B) Creating fake profiles of your friends and uploading embarrassing pictures.

Solution:

Concept:

- Cyberbullying uses digital communication or online platforms to harass, embarrass, threaten or humiliate another person.

- Other harmful online actions, such as phishing and hacking, have different defining aims even though several kinds of misconduct can occur together.

Step 1: Identify the action aimed at humiliation Creating fake profiles of friends and uploading embarrassing pictures uses an online identity and digital content to humiliate or harm those people. This directly fits cyberbullying: the action targets another person's dignity and social reputation. The use of fake profiles also provides a means to misrepresent the victims and spread embarrassing material.

Step 2: Distinguish the remaining choices Sending phishing links is primarily an attempt to deceive users into disclosing information or visiting a malicious site. Hacking into a friend's system and reading mail is unauthorised access and a privacy violation. Blocking a person from communication can be a normal account control or personal boundary and does not by itself establish harassment. The fake-profile and embarrassing-picture option is the clearest direct example of cyberbullying among the four choices.

Step 3: Conclusion Creating fake profiles of friends and uploading embarrassing pictures is an example of cyberbullying. Hence, the correct option is (B).

Quick Tip to Solve: Find the purpose of the action: humiliation points to cyberbullying, credential deception to phishing, and unauthorised entry to hacking.

85. The 3 "R" of e-waste management leading to efficient disposal are: [5 marks]

(A) Reduce, Reuse, Return

(B) Retire, Reuse, Recycle

(C) Recycle, Reuse, Reduce

(D) Rescue, Reuse, Return

Correct Answer: (C) Recycle, Reuse, Reduce

Solution:

Concept:

- The three Rs of waste management are Reduce, Reuse and Recycle. They help conserve resources and reduce the amount of discarded material requiring final disposal.
- Electronic waste includes discarded electrical and electronic equipment. Applying the three Rs reduces avoidable purchases, extends the useful life of devices and recovers materials from equipment that can no longer be used.

Step 1: Recall the meaning of each R Reduce means limiting the generation of e-waste, for example by buying equipment only when necessary and choosing durable products. Reuse means using a working device again, passing it to another user or extending its service life through suitable repair. Recycle means processing unusable electronic equipment through appropriate recycling facilities to recover useful materials. These are distinct actions: continuing to use a working computer is reuse, whereas recovering metals from an unusable computer is recycling.

Step 2: Compare the three-word combinations Option (C) contains Recycle, Reuse and Reduce,

so it includes all three required terms. Although the familiar phrase usually lists Reduce first, the question asks for the three Rs, and changing their order does not change the set of terms. Option (A) substitutes Return for Recycle. Option (B) substitutes Retire for Reduce. Option (D) includes Rescue and Return and omits both Reduce and Recycle. Returning a product may support a collection scheme, but Return is not one of the three terms in the standard 3R principle.

Step 3: Conclusion The three Rs are Reduce, Reuse and Recycle. The combination containing all three is Recycle, Reuse, Reduce. Hence, the correct option is (C).

Quick Tip to Solve: Check that a 3R option contains all three exact terms: Reduce, Reuse and Recycle. Their order may vary; words such as Return or Retire do not replace them.