GATE Data Science and Artificial Intelligence

Sample Paper – 2

Duration: 128 Minutes **Maximum Marks: 72**

Instructions

- This paper contains **46 questions** covering the core **Data Science and Artificial Intelligence** subject of the GATE paper conducted by the IITs and IISc (General Aptitude and Engineering Mathematics are provided as separate papers).
- **Q1–Q20 carry 1 mark each** and **Q21–Q46 carry 2 marks each**, for a total of **72 marks**.
- Each question carries **negative marking of one-third of its allotted marks**: a wrong 1-mark question costs **0.33** and a wrong 2-mark question costs **0.67**. Unattempted questions carry no penalty.
- Every question here is a single-correct **MCQ**. The real GATE paper also has Multiple Select (MSQ) and Numerical Answer Type (NAT) questions, and **those carry no negative marking**.
- Attempt this paper in one timed sitting of about **128 minutes**, the share this core subject earns at GATE's overall pace of 180 minutes for 65 questions. All logarithms in information-theoretic quantities (entropy, information gain) are to base 2 and expressed in bits unless stated otherwise.

Probability & Statistics

Q1. A box contains 4 red and 6 blue balls. Two balls are drawn one after another *without replacement*. The probability that *both* balls drawn are red is: **[1 mark]**

(A) $\frac{4}{25}$

(B) $\frac{2}{9}$

(C) $\frac{2}{15}$

(D) $\frac{1}{5}$

Q2. For two events A and B the probabilities are shown on the Venn diagram below (the region outside both circles has probability 0.30). The probability $P(A \cup B)$ is: **[1 mark]**

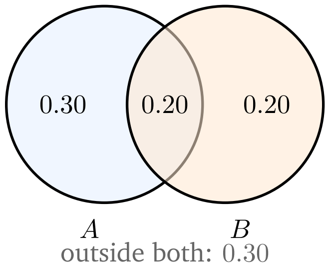

(A) 0.20

(B) 0.90

(C) 0.50

(D) 0.70

Q3. A screening test for a rare condition is modelled by the probability tree below. The condition is present (C) in 1% of people. The test is positive ($+$) with probability 0.90 when the condition is present and with probability 0.05 when it is absent ($\bar{C}$). A randomly chosen person tests positive. The probability that the person actually has the condition is nearest to: **[1 mark]**

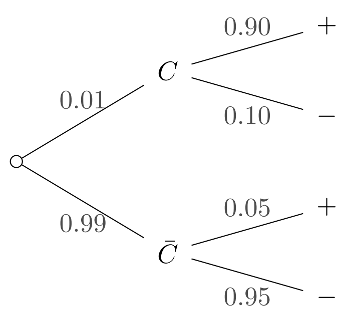

(A) 0.154

(B) 0.900

(C) 0.010

(D) 0.500

Q4. A feature X is normally distributed with mean $\mu = 70$ and standard deviation $\sigma = 10$. Using the shaded right tail of the standard normal curve shown, the probability $P(X > 90)$ is nearest to (take $P(Z \leq 2) = 0.9772$): **[1 mark]**

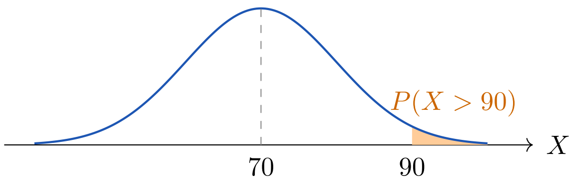

(A) 0.9772

(B) 0.0228

(C) 0.4772

(D) 0.1587

Q5. Emails arrive at an inbox following a Poisson process with an average of 2 emails per hour. The probability of receiving *exactly* 3 emails in a given one-hour window is nearest to (take $e^{-2} = 0.135$): **[1 mark]**

(A) 0.135

(B) 0.270

(C) 0.090

(D) 0.180

Q6. A discrete random variable X is uniformly distributed over the values $1, 2, 3, 4$ (each with probability $\frac{1}{4}$). The variance of X is: **[1 mark]**

(A) 2.50

(B) 1.25

(C) 2.00

(D) 0.25

Q7. A sample of size $n = 64$ is drawn from a population with known standard deviation $\sigma = 16$. For a 95% confidence interval for the population mean (using $z_{0.025} = 1.96$), the margin of error (half-width of the interval) is: **[1 mark]**

(A) 1.96

(B) 2.00

(C) 3.92

(D) 0.245

Q8. For a bivariate data set the sums of squares and cross-products about the means are $S_{xx} = 25$, $S_{yy} = 100$ and $S_{xy} = 40$, and the positively sloped scatter is shown. The Pearson correlation coefficient r between x and y is: **[1 mark]**

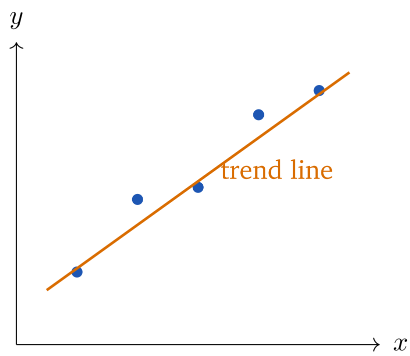

(A) 0.80

(B) 1.60

(C) 0.40

(D) 0.64

Q9. A fair coin is tossed 5 times independently. The probability of obtaining *exactly* 2 heads is: **[1 mark]**

(A) 0.1563

(B) 0.3125

(C) 0.5000

(D) 0.2500

Q10. A continuous random variable X is uniformly distributed on the interval $[0, 10]$. The probability $P(X > 7)$ is: **[1 mark]**

(A) 0.7

(B) 0.5

(C) 0.3

(D) 0.1

Linear Algebra, Calculus & Optimization

Q11. The trace of the matrix $A = \begin{bmatrix} 2 & 1 & 3 \\ 0 & 4 & 1 \\ 5 & 2 & 6 \end{bmatrix}$ is: **[1 mark]**

(A) 8

(B) 6

(C) 10

(D) 12

Q12. The determinant of the upper-triangular matrix $B = \begin{bmatrix} 1 & 2 & 0 \\ 0 & 3 & 4 \\ 0 & 0 & 5 \end{bmatrix}$ is: **[1 mark]**

(A) 8

(B) 15

(C) 20

(D) 9

Q13. The *smaller* eigenvalue of the matrix $C = \begin{bmatrix} 2 & 2 \\ 1 & 3 \end{bmatrix}$ is: **[1 mark]**

(A) 4

(B) 1

(C) 5

(D) 3

Q14. A square matrix A has eigenvalues 2, 3 and 5. The determinant of A (the product of its eigenvalues) is: **[1 mark]**

(A) 30

(B) 10

(C) 25

(D) 15

Q15. For the vector $v = (6,\ 8)$, the Euclidean (ℓ_2) norm $\|v\|_2$ is: **[1 mark]**

(A) 14

(B) 48

(C) 8

(D) 10

Q16. The function $f(x) = x^2 - 6x + 11$ is minimized over the real line, as suggested by its convex graph below. The value of x at which the minimum occurs is: **[1 mark]**

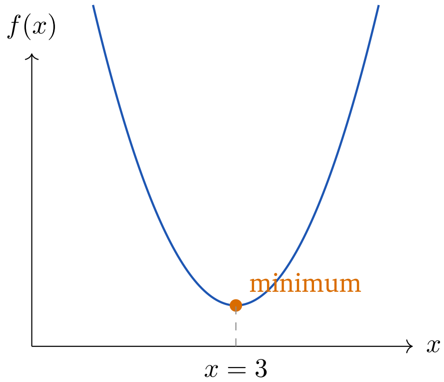

(A) 3

(B) 2

(C) 6

(D) 0

Q17. The angle θ between the vectors $\vec{a} = (1,\ 0)$ and $\vec{b} = (1,\ 1)$, shown below, is: **[1 mark]**

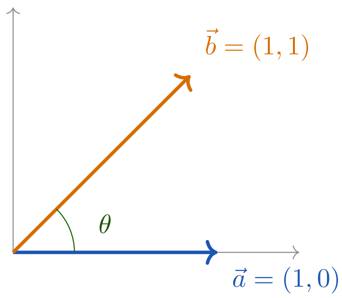

(A) $90°$

(B) $30°$

(C) $60°$

(D) $45°$

Q18. Gradient descent is used to minimize $f(x) = x^2$ with a fixed learning rate $\eta = 0.25$. Starting from $x_0 = 2$, the value of x_1 after one update $x_1 = x_0 - \eta f'(x_0)$ is: **[1 mark]**

(A) 2

(B) 1

(C) 1.5

(D) 0

Q19. Consider the function $f(x, y) = x^2 - y^2$ on $\mathbb{R}^2$. Based on its Hessian, the function is: **[1 mark]**

(A) convex everywhere

(B) neither convex nor concave

(C) concave everywhere

(D) linear (affine)

Programming, Data Structures & Algorithms

Q20. What is the output of the following Python snippet?

```
a = [1, 2, 2, 3, 3, 3]
print(len(set(a)))
```

[1 mark]

(A) 3

(B) 6

(C) 1

(D) 2

Q21. What is the output of the following Python snippet?

```
print(list(range(2, 10, 2)))
```

[2 marks]

(A) `[2, 4, 6, 8]`

(B) `[2, 4, 6, 8, 10]`

(C) `[0, 2, 4, 6, 8]`

(D) `[2, 3, 4, 5, 6, 7, 8, 9]`

Q22. The following operations are performed on an initially empty *queue* (first-in, first-out), in order: `enqueue(5)`, `enqueue(10)`, `dequeue()`, `enqueue(15)`, `dequeue()`. The value returned by the *last* `dequeue()` is: [2 marks]

(A) 5

(B) 10

(C) 15

(D) 20

Q23. Starting from an empty binary search tree, the keys 40, 20, 60, 10, 30 are inserted in the given order, producing the tree shown. The *height* of the resulting tree (measured as the number of edges on the longest root-to-leaf path) is: [2 marks]

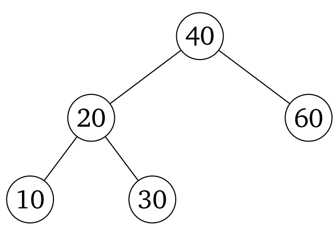

(A) 1

(B) 2

(C) 3

(D) 4

Q24. A hash table of size 7 uses the hash function $h(k) = k \bmod 7$ with *linear probing* for collision resolution. The keys 15, 22, 8 are inserted in this order into an otherwise empty table. The final index occupied by key 8 is: **[2 marks]**

(A) 1

(B) 2

(C) 3

(D) 4

Q25. Consider the loop below, where `n = 32`:

```
i = 1
while i <= n:
i = i * 2
```

The number of times the loop body executes is: **[2 marks]**

(A) 32

(B) 5

(C) 6

(D) 16

Q26. The worst-case time complexity of the Merge Sort algorithm on an array of n elements is: **[2 marks]**

(A) $O(n^2)$

(B) $O(n)$

(C) $O(\log n)$

(D) $O(n \log n)$

Q27. A depth-first search (DFS) is started from vertex A in the undirected graph shown, always visiting adjacent vertices in increasing alphabetical order and backtracking when no unvisited neighbour remains. The order in which the vertices are first visited is: **[2 marks]**

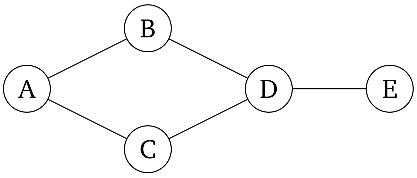

(A) A, B, D, C, E

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

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

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

Q28. The running time of an algorithm satisfies the recurrence $T(n) = T(n/2) + c$ for a constant c, with $T(1) = c$. By the Master theorem, $T(n)$ is: **[2 marks]**

(A) $\Theta(n)$

(B) $\Theta(n \log n)$

(C) $\Theta(\log n)$

(D) $\Theta(n^2)$

Database Management & Warehousing

Q29. Two union-compatible relations R and S contain 10 and 6 tuples respectively, and they share exactly 4 tuples in common. The number of tuples in $R \cup S$ is: **[2 marks]**

(A) 16

(B) 60

(C) 10

(D) 12

Q30. A table `EMP` has a `dept` column with the values {`A, A, B, B, B, C, NULL`} across its 7 rows. What does the query `SELECT COUNT(DISTINCT dept) FROM EMP;` return? **[2 marks]**

(A) 7

(B) 6

(C) 3

(D) 4

Q31. A relation $R(A, B, C)$ has *only* the functional dependencies $A \to B$ and $A \to C$, with A as the sole candidate key. The highest normal form that R satisfies is: **[2 marks]**

(A) BCNF

(B) 2NF

(C) 3NF

(D) 1NF

Q32. Relation $R(A, B)$ contains the tuples $\{(1, p), (2, p), (3, q)\}$ and relation $S(B, C)$ contains $\{(p, x), (p, y), (q, z)\}$. The number of tuples in the natural join $R \bowtie S$ (on the common attribute B) is: **[2 marks]**

(A) 3

(B) 9

(C) 5

(D) 6

Q33. In OLAP, the operation that aggregates a data cube from a finer level to a coarser (more summarized) level, such as from monthly totals to yearly totals, is called: **[2 marks]**

(A) drill-down

(B) roll-up

(C) slice

(D) pivot

Q34. The figure shows a data-warehouse schema in which a single central table stores the numeric measures and references four surrounding tables through foreign keys. In this star schema, the central table is called the: **[2 marks]**

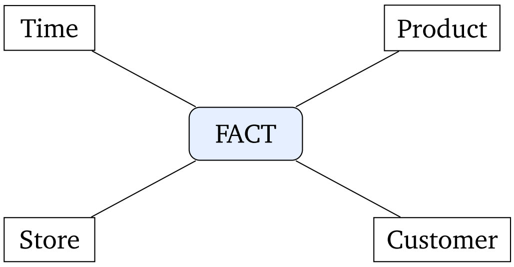

(A) dimension table

(B) fact table

(C) index table

(D) materialized view

Q35. A relation $R(A, B, C, D)$ has the functional dependencies $AB \to C$, $C \to D$ and $D \to A$. The attribute closure $(AB)^+$ is: **[2 marks]**

(A) $\{A, B\}$

(B) $\{A, B, C\}$

(C) $\{A, B, D\}$

(D) $\{A, B, C, D\}$

Machine Learning & Artificial Intelligence

Q36. A logistic regression model has weight vector $w = (2, -2)$ and bias $b = 0$, and uses the sigmoid $\sigma(z) = 1/(1 + e^{-z})$. For the input $x = (1, 1)$, the predicted probability $\sigma(w^\top x + b)$ is: **[2 marks]**

(A) 0.50

(B) 0.88

(C) 0.12

(D) 0.73

Q37. A very simple linear model performs poorly on *both* the training set and the test set, with roughly equal (high) error on each. In terms of the bias–variance decomposition, this behaviour is best described as: **[2 marks]**

(A) low bias, high variance

(B) low bias, low variance

(C) high bias, high variance

(D) high bias, low variance

Q38. A 1-nearest-neighbour classifier (Euclidean distance) is trained on the labelled points $A(1,1)$: Class 0, $B(5,5)$: Class 1, $C(4,3)$: Class 1, $D(0,4)$: Class 0. The query point $Q(4,4)$ is classified as: **[2 marks]**

(A) Class 0

(B) a tie (undecided)

(C) Class 1

(D) cannot be determined

Q39. A training set for a binary classification task contains 10 examples, of which 5 are positive and 5 are negative. The entropy of this set (in bits) is: **[2 marks]**

(A) 0.811

(B) 0.722

(C) 1.000

(D) 0.500

Q40. A node in a decision tree contains 10 examples, of which 6 are positive and 4 are negative, as shown. The Gini impurity of this node is: **[2 marks]**

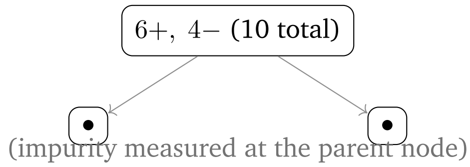

(A) 0.48

(B) 0.50

(C) 0.24

(D) 0.60

Q41. For the linear support vector machine shown, the separating hyperplane $w^{\mathsf{T}}x + b = 0$ lies midway between its two margin boundaries. If $\|w\| = 2$, the geometric margin (the distance from the hyperplane to a support vector, equal to $1/\|w\|$) is: **[2 marks]**

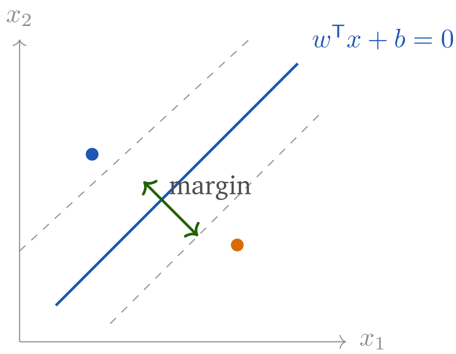

(A) 4

(B) 2

(C) 1

(D) 0.5

Q42. In one iteration of the k-means algorithm, a cluster is assigned the three points $(0, 0)$, $(2, 4)$ and $(4, 2)$, shown below. The updated centroid of this cluster is: **[2 marks]**

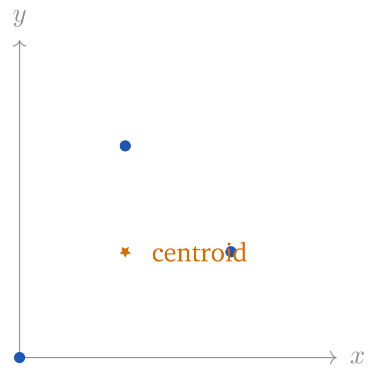

(A) $(3, 3)$

(B) $(2, 3)$

(C) $(2, 2)$

(D) $(6, 6)$

Q43. Principal component analysis is applied to a data set whose covariance matrix has eigenvalues $5, 3, 2$. The fraction of the total variance explained by the *first* principal component alone is: **[2 marks]**

(A) 30%

(B) 20%

(C) 80%

(D) 50%

Q44. A single neuron with a ReLU activation computes ReLU$(w_1x_1 + w_2x_2 + b)$, as shown, where ReLU$(z) = \max(0, z)$. With inputs $x_1 = 2$, $x_2 = 1$, weights $w_1 = 1$, $w_2 = 2$ and bias $b = -1$, the output of the neuron is: **[2 marks]**

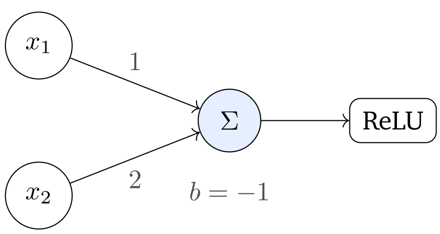

(A) 0

(B) 3

(C) -3

(D) 1

Q45. A* search maintains three frontier nodes A, B and C with g (path cost so far) and h (heuristic) values shown in the tree below, where $f = g + h$. The node A* selects for expansion next is: **[2 marks]**

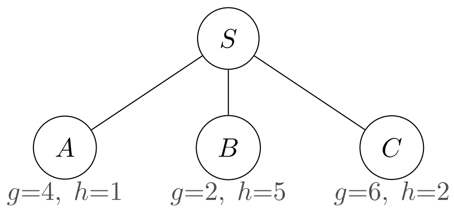

(A) A

(B) B

(C) C

(D) all three at once

Q46. In the two-ply game tree shown, the root is a MAX node and the second level are MIN nodes; the leaf values are given. The minimax value backed up to the root is: **[2 marks]**

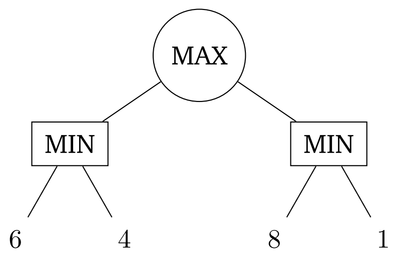

(A) 4

(B) 1

(C) 6

(D) 8

Detailed Solutions

Q1.

Solution

Concept – probability without replacement: Multiply the probability of the first red draw by the probability of a second red draw given the first.

Step 1 – total balls: $4 + 6 = 10$ balls.

Step 2 – probability the first ball is red: $P(\text{1st red}) = \frac{4}{10}$

Step 3 – probability the second ball is red (one red removed): $P(\text{2nd red} \mid \text{1st red}) = \frac{3}{9}$

Step 4 – multiply the two: $P(\text{both red}) = \frac{4}{10} \times \frac{3}{9}$

Step 5 – simplify: $= \frac{12}{90} = \frac{2}{15}$

Final Answer: $\frac{2}{15} \Rightarrow \boxed{\text{C}}$

Answer: (C) Go Back to Q1

Q2.

Solution

Concept – union of two events: $P(A \cup B) = P(A) + P(B) - P(A \cap B)$, or simply add the three non-overlapping regions.

Step 1 – read the regions from the Venn diagram: A only $= 0.30$, overlap $A \cap B = 0.20$, B only $= 0.20$.

Step 2 – add the parts inside the union: $P(A \cup B) = 0.30 + 0.20 + 0.20$

Step 3 – total: $P(A \cup B) = 0.70$

Step 4 – sanity check with the outside region: $0.70+0.30 = 1.00$, so the diagram is consistent.

Final Answer: $P(A \cup B) = 0.70 \Rightarrow \boxed{\text{D}}$

Answer: (D) Go Back to Q2

Q3.

Solution

Concept – Bayes' theorem (reverse the tree): $P(C \mid +) = \dfrac{P(C)\,P(+ \mid C)}{P(C)\,P(+ \mid C) + P(\bar{C})\,P(+ \mid \bar{C})}$.

Step 1 – probability along the "has condition" branch: $P(C)\,P(+ \mid C) = 0.01 \times 0.90 = 0.0090$

Step 2 – probability along the "no condition" branch: $P(\bar{C})\,P(+ \mid \bar{C}) = 0.99 \times 0.05 = 0.0495$

Step 3 – total probability of a positive test: $P(+) = 0.0090 + 0.0495 = 0.0585$

Step 4 – apply Bayes' theorem: $P(C \mid +) = \dfrac{0.0090}{0.0585}$

Step 5 – divide: $P(C \mid +) = 0.1538 \approx 0.154$

Final Answer: $P(C \mid +) \approx 0.154 \Rightarrow \boxed{\text{A}}$

Answer: (A) Go Back to Q3

Q4.

Solution

Concept – standardize, then use the normal table: Convert X to the standard normal $Z = \dfrac{X - \mu}{\sigma}$.

Step 1 – standardize the cut-off $X = 90$: $Z = \dfrac{90 - 70}{10}$

$Z = \dfrac{20}{10} = 2$

Step 2 – write the required probability in Z: $P(X > 90) = P(Z > 2)$

Step 3 – use the standard normal value: $P(Z \leq 2) = 0.9772$

Step 4 – take the right tail: $P(Z > 2) = 1 - 0.9772 = 0.0228$

Final Answer: $P(X > 90) = 0.0228 \Rightarrow \boxed{\text{B}}$

Answer: (B) Go Back to Q4

Q5.

Solution

Concept – Poisson probability: $P(X = k) = \dfrac{e^{-\lambda}\,\lambda^k}{k!}$, with mean rate λ.

Step 1 – list the parameters: $\lambda = 2$ per hour, $k = 3$.

Step 2 – substitute into the formula: $P(X=3) = \dfrac{e^{-2}\,2^3}{3!}$

Step 3 – evaluate the powers and factorial: $2^3 = 8$ and $3! = 6$

$P(X=3) = \dfrac{e^{-2} \times 8}{6}$

Step 4 – substitute $e^{-2} = 0.135$**:** $P(X=3) = \dfrac{0.135 \times 8}{6} = \dfrac{1.08}{6}$

Step 5 – divide: $P(X=3) = 0.180$

Final Answer: $P(X=3) \approx 0.180 \Rightarrow \boxed{\text{D}}$

Answer: (D) Go Back to Q5

Q6.

Solution

Concept – variance from the definition: $\mathrm{Var}(X) = E[X^2] - \big(E[X]\big)^2$.

Step 1 – compute the mean $E[X]$ **(each value has probability** $\frac{1}{4}$**):** $E[X] = \dfrac{1+2+3+4}{4}$

$= \dfrac{10}{4}$

$= 2.5$

Step 2 – compute $E[X^2]$**:** $E[X^2] = \dfrac{1^2+2^2+3^2+4^2}{4}$

$= \dfrac{1+4+9+16}{4}$

$= \dfrac{30}{4} = 7.5$

Step 3 – square the mean: $\big(E[X]\big)^2 = (2.5)^2 = 6.25$

Step 4 – subtract: $\mathrm{Var}(X) = 7.5 - 6.25 = 1.25$

Final Answer: $\mathrm{Var}(X) = 1.25 \Rightarrow \boxed{\text{B}}$

Answer: (B) Go Back to Q6

Q7.

Solution

Concept – margin of error of a confidence interval: For a known-σ mean, the margin of error is $E = z\,\dfrac{\sigma}{\sqrt{n}}$.

Step 1 – compute $\sqrt{n}$**:** $\sqrt{64} = 8$

Step 2 – compute the standard error: $\dfrac{\sigma}{\sqrt{n}} = \dfrac{16}{8} = 2$

Step 3 – multiply by the critical value: $E = 1.96 \times 2$

$E = 3.92$

Step 4 – interpret: The 95% interval is $\bar{x} \pm 3.92$; its half-width is 3.92.

Final Answer: margin of error $= 3.92 \Rightarrow \boxed{\text{C}}$

Answer: (C) Go Back to Q7

Q8.

Solution

Concept – Pearson correlation coefficient: $r = \dfrac{S_{xy}}{\sqrt{S_{xx}\,S_{yy}}}$.

Step 1 – list the sums of squares and cross-products: $S_{xx} = 25$, $S_{yy} = 100$, $S_{xy} = 40$.

Step 2 – compute the product under the root: $S_{xx}\,S_{yy} = 25 \times 100 = 2500$

Step 3 – take the square root: $\sqrt{2500} = 50$

Step 4 – form the ratio: $r = \dfrac{40}{50}$

Step 5 – simplify: $r = 0.80$

Final Answer: $r = 0.80 \Rightarrow \boxed{\text{A}}$

Answer: (A) Go Back to Q8

Q9.

Solution

Concept – binomial probability: $P(X = k) = \binom{n}{k} p^k (1-p)^{n-k}$.

Step 1 – list the parameters: $n = 5$, $k = 2$, $p = 0.5$, $1 - p = 0.5$.

Step 2 – evaluate the binomial coefficient: $\binom{5}{2} = \dfrac{5 \times 4}{2 \times 1} = 10$

Step 3 – evaluate the powers: $p^2(1-p)^3 = (0.5)^2(0.5)^3 = (0.5)^5$

$= 0.03125$

Step 4 – multiply: $P(X = 2) = 10 \times 0.03125$

$= 0.3125$

Final Answer: $P(X = 2) = 0.3125 \Rightarrow \boxed{\text{B}}$

Answer: (B) Go Back to Q9

Q10.

Solution

Concept – uniform distribution probability: For X uniform on $[a, b]$, $P(X > c) = \frac{b-c}{b-a}$ for $a \le c \le b$.

Step 1 – identify the interval: $a = 0$, $b = 10$, so the total length is 10.

Step 2 – length of the favourable region $X > 7$**:** $b - c = 10 - 7 = 3$

Step 3 – form the ratio: $P(X > 7) = \frac{3}{10}$

Step 4 – evaluate: $P(X > 7) = 0.3$

Final Answer: $P(X > 7) = 0.3 \Rightarrow \boxed{\text{C}}$

Answer: (C) Go Back to Q10

Q11.

Solution

Concept – trace of a matrix: The trace is the sum of the diagonal entries $a_{11} + a_{22} + a_{33}$.

Step 1 – read the diagonal entries: $a_{11} = 2$, $a_{22} = 4$, $a_{33} = 6$.

Step 2 – add them: $\text{tr}(A) = 2 + 4 + 6$

Step 3 – evaluate: $= 12$

Final Answer: $\text{tr}(A) = 12 \Rightarrow \boxed{\text{D}}$

Answer: (D) Go Back to Q11

Q12.

Solution

Concept – determinant of a triangular matrix: For an upper- (or lower-) triangular matrix, the determinant is the product of the diagonal entries.

Step 1 – confirm the matrix is upper triangular: All entries below the main diagonal are 0.

Step 2 – read the diagonal entries: $1,\ 3,\ 5$.

Step 3 – multiply them: $\det B = 1 \times 3 \times 5$

Step 4 – evaluate: $= 15$

Final Answer: $\det B = 15 \Rightarrow \boxed{\text{B}}$

Answer: (B) Go Back to Q12

Q13.

Solution

Concept – eigenvalues from the characteristic equation: Solve $\det(C - \lambda I) = 0$; use trace and determinant as a shortcut.

Step 1 – trace and determinant of C: $\mathrm{tr}(C) = 2 + 3 = 5$ and $\det(C) = (2)(3) - (2)(1) = 6 - 2 = 4$.

Step 2 – write the characteristic equation: $\lambda^2 - \mathrm{tr}(C)\,\lambda + \det(C) = 0$

$\lambda^2 - 5\lambda + 4 = 0$

Step 3 – factor: $(\lambda - 4)(\lambda - 1) = 0$

Step 4 – read the roots: $\lambda = 4$ or $\lambda = 1$.

Step 5 – pick the smaller: The smaller eigenvalue is 1.

Final Answer: $\lambda_{\min} = 1 \Rightarrow \boxed{\text{B}}$

Answer: (B) Go Back to Q13

Q14.

Solution

Concept – determinant equals the product of eigenvalues: $\det(A) = \lambda_1\lambda_2\cdots\lambda_n$.

Step 1 – list the eigenvalues: $\lambda_1 = 2$, $\lambda_2 = 3$, $\lambda_3 = 5$.

Step 2 – multiply them: $\det(A) = 2 \times 3 \times 5$

Step 3 – evaluate step by step: $2 \times 3 = 6$

$6 \times 5 = 30$

Final Answer: $\det(A) = 30 \Rightarrow \boxed{\text{A}}$

Answer: (A) Go Back to Q14

Q15.

Solution

Concept – the ℓ_2 (Euclidean) norm: $\|v\|_2 = \sqrt{\sum_i v_i^2}$.

Step 1 – square each component: $6^2 = 36$ and $8^2 = 64$.

Step 2 – add the squares: $36 + 64 = 100$

Step 3 – take the square root: $\|v\|_2 = \sqrt{100} = 10$

Step 4 – contrast with the other norms: $\|v\|_1 = 6 + 8 = 14$ and $\|v\|_\infty = \max(6, 8) = 8$; the Euclidean value is 10.

Final Answer: $\|v\|_2 = 10 \Rightarrow \boxed{\text{D}}$

Answer: (D) Go Back to Q15

Q16.

Solution

Concept – minimize by setting the derivative to zero: For a convex parabola, the stationary point is the global minimum.

Step 1 – differentiate: $f'(x) = 2x - 6$

Step 2 – set the derivative to zero: $2x - 6 = 0$

Step 3 – solve for x: $2x = 6$

$x = 3$

Step 4 – confirm it is a minimum: $f''(x) = 2 > 0$, so the stationary point is a minimum.

Final Answer: minimizing $x = 3 \Rightarrow \boxed{\text{A}}$

Answer: (A) Go Back to Q16

Q17.

Solution

Concept – angle between vectors from the dot product: $\cos\theta = \dfrac{\vec{a} \cdot \vec{b}}{\|\vec{a}\|\,\|\vec{b}\|}$.

Step 1 – compute the dot product: $\vec{a} \cdot \vec{b} = (1)(1) + (0)(1) = 1$

Step 2 – compute the magnitudes: $\|\vec{a}\| = \sqrt{1^2 + 0^2} = 1$

$\|\vec{b}\| = \sqrt{1^2 + 1^2} = \sqrt{2}$

Step 3 – form the cosine: $\cos\theta = \dfrac{1}{1 \times \sqrt{2}} = \dfrac{1}{\sqrt{2}}$

Step 4 – take the inverse cosine: $\theta = \cos^{-1}\left(\dfrac{1}{\sqrt{2}}\right) = 45°$

Final Answer: $\theta = 45° \Rightarrow \boxed{\text{D}}$

Answer: (D) Go Back to Q17

Q18.

Solution

Concept – one gradient-descent step: $x_1 = x_0 - \eta\, f'(x_0)$.

Step 1 – differentiate the objective: $f(x) = x^2 \Rightarrow f'(x) = 2x$

Step 2 – evaluate the gradient at $x_0 = 2$: $f'(2) = 2 \times 2 = 4$

Step 3 – apply the update with $\eta = 0.25$: $x_1 = 2 - 0.25 \times 4$

Step 4 – finish the arithmetic: $x_1 = 2 - 1 = 1$

Final Answer: $x_1 = 1 \Rightarrow \boxed{\text{B}}$

Answer: (B) Go Back to Q18

Q19.

Solution

Concept – convexity via the Hessian: A twice-differentiable function is convex iff its Hessian is positive semi-definite everywhere, and concave iff it is negative semi-definite; if the Hessian is indefinite it is neither.

Step 1 – compute the second partials: $f_{xx} = 2$, $f_{yy} = -2$, $f_{xy} = f_{yx} = 0$.

Step 2 – write the Hessian: $H = \begin{bmatrix} 2 & 0 \\ 0 & -2 \end{bmatrix}$

Step 3 – read the eigenvalues: The Hessian is diagonal, so its eigenvalues are 2 and -2.

Step 4 – classify: One eigenvalue is positive and one is negative, so H is indefinite (a saddle shape); f is neither convex nor concave.

Final Answer: neither convex nor concave $\Rightarrow \boxed{\text{B}}$

Answer: (B) Go Back to Q19

Q20.

Solution

Concept – a Python set stores only distinct values: Converting a list to a set removes duplicates; `len` then counts the distinct elements.

Step 1 – list the elements: `a = [1, 2, 2, 3, 3, 3]`.

Step 2 – form the set: `set(a)` keeps one copy of each value: $\{1, 2, 3\}$.

Step 3 – count the distinct values: There are 3 distinct values.

Final Answer: $3 \Rightarrow \boxed{\text{A}}$

Answer: (A) Go Back to Q20

Q21.

Solution

Concept – `range(start, stop, step)`: It yields values beginning at `start`, incrementing by `step`, and stopping *before* `stop`.

Step 1 – start value: Begin at 2.

Step 2 – add the step repeatedly: $2,\ 2+2=4,\ 4+2=6,\ 6+2=8$.

Step 3 – apply the stop rule: The next value would be 10, but `stop = 10` is excluded, so it is not generated.

Step 4 – collect into a list: `[2, 4, 6, 8]`.

Final Answer: `[2, 4, 6, 8]` $\Rightarrow \boxed{\text{A}}$

Answer: (A) Go Back to Q21

Q22.

Solution

Concept – a queue is first-in, first-out (FIFO): Elements leave in the same order they entered; track the front of the queue.

Step 1 – `enqueue(5)`, `enqueue(10)`: Queue (front→rear): $[5,\ 10]$.

Step 2 – `dequeue()`: Removes the front 5; queue becomes $[10]$.

Step 3 – `enqueue(15)`: Queue becomes $[10,\ 15]$.

Step 4 – `dequeue()` (the last one): Removes the front 10; this is the value returned.

Final Answer: last `dequeue()` returns $10 \Rightarrow \boxed{\text{B}}$

Answer: (B) Go Back to Q22

Q23.

Solution

Concept – height counts edges on the longest root-to-leaf path: Build the BST and follow its deepest branch.

Step 1 – place the root: 40 is inserted first and becomes the root (level 0).

Step 2 – place the next level: $20 < 40$ goes left; $60 > 40$ goes right (level 1).

Step 3 – place the remaining keys: $10 < 20$ (left of 20, level 2); 30 is between 20 and 40 (right of 20, level 2).

Step 4 – find the longest path: The deepest leaves are 10 and 30, reached by $40 \to 20 \to 10$ (or $\to 30$), which is 2 edges.

Step 5 – state the height: Height $= 2$ edges.

Final Answer: height $= 2 \Rightarrow \boxed{\text{B}}$

Answer: (B) Go Back to Q23

Q24.

Solution

Concept – linear probing scans the next free slot: On a collision at index i, try $i + 1, i + 2, \ldots$ (modulo table size).

Step 1 – insert 15: $h(15) = 15 \bmod 7 = 1$; slot 1 is free, so 15 goes to index 1.

Step 2 – insert 22: $h(22) = 22 \bmod 7 = 1$; index 1 is occupied, so probe index 2; it is free, so 22 goes to index 2.

Step 3 – insert 8: $h(8) = 8 \bmod 7 = 1$; index 1 occupied $\to$ probe index 2 (occupied) $\to$ probe index 3 (free).

Step 4 – place 8: 8 is stored at index 3.

Final Answer: index $3 \Rightarrow \boxed{\text{C}}$

Answer: (C) Go Back to Q24

Q25.

Solution

Concept – a doubling loop runs about $\log_2 n$ times: Track i as it doubles until the condition $i \le n$ fails.

Step 1 – list the value of i that enters each iteration ($n = 32$): $i = 1$ (run), then i becomes 2.

Step 2 – continue doubling while $i \le 32$: $i = 2$ (run) $\to 4$; $i = 4$ (run) $\to 8$; $i = 8$ (run) $\to 16$.

Step 3 – last executions: $i = 16$ (run) $\to 32$; $i = 32$ (run) $\to 64$.

Step 4 – next test fails: $i = 64 > 32$, so the loop stops.

Step 5 – count the runs: The body ran for $i = 1, 2, 4, 8, 16, 32$, which is 6 times.

Final Answer: 6 executions $\Rightarrow$ $\boxed{\text{C}}$

Answer: (C) Go Back to Q25

Q26.

Solution

Concept – Merge Sort divides and merges: It splits the array in half recursively and merges the sorted halves in linear time.

Step 1 – write the recurrence: $T(n) = 2\,T(n/2) + O(n)$

Step 2 – count the levels of recursion: Halving n down to 1 takes $\log_2 n$ levels.

Step 3 – work per level: Merging at each level costs $O(n)$ in total.

Step 4 – multiply levels by per-level work: $T(n) = O(n) \times \log_2 n = O(n \log n)$

Step 5 – note the consistency across cases: Merge Sort is $O(n \log n)$ in the best, average, and worst cases alike.

Final Answer: worst case $= O(n \log n) \Rightarrow \boxed{\text{D}}$

Answer: (D) Go Back to Q26

Q27.

Solution

Concept – DFS goes deep before backtracking: Visit a neighbour, then *immediately* recurse from it, exploring alphabetically and backtracking only when stuck.

Step 1 – start at A: Visit A. Its neighbours are B and C; go to B (alphabetically first).

Step 2 – recurse from B**:** Visit B. Its unvisited neighbour is D; go to D.

Step 3 – recurse from D**:** Visit D. Unvisited neighbours are C and E; go to C (alphabetically first).

Step 4 – recurse from C**:** Visit C. Its neighbours A and D are already visited, so backtrack to D.

Step 5 – finish from D**:** The remaining unvisited neighbour of D is E; visit E.

Step 6 – assemble the order: $A,\ B,\ D,\ C,\ E$.

Final Answer: $A,\ B,\ D,\ C,\ E \Rightarrow \boxed{\text{A}}$

Answer: (A) Go Back to Q27

Q28.

Solution

Concept – Master theorem: For $T(n) = a\,T(n/b) + f(n)$, compare $f(n)$ with $n^{\log_b a}$.

Step 1 – identify the parameters: $a = 1$, $b = 2$, $f(n) = c = \Theta(1)$.

Step 2 – compute the critical exponent: $n^{\log_b a} = n^{\log_2 1} = n^0 = 1$

Step 3 – compare $f(n)$ **with** $n^{\log_b a}$**:** $f(n) = \Theta(1) = \Theta\big(n^{\log_b a}\big)$, which is Master-theorem *Case 2*.

Step 4 – apply Case 2: $T(n) = \Theta\big(n^{\log_b a} \log n\big) = \Theta(\log n)$

Step 5 – sanity check: This is the familiar cost of binary search, which halves the input doing constant work per level.

Final Answer: $T(n) = \Theta(\log n) \Rightarrow \boxed{\text{C}}$

Answer: (C) Go Back to Q28

Q29.

Solution

Concept – size of a set union (inclusion–exclusion): $|R \cup S| = |R| + |S| - |R \cap S|$; the relational UNION also removes duplicate tuples.

Step 1 – read the counts: $|R| = 10$, $|S| = 6$, $|R \cap S| = 4$.

Step 2 – add the two relation sizes: $10 + 6 = 16$

Step 3 – subtract the shared tuples once: $16 - 4 = 12$

Step 4 – interpret: The 4 common tuples were counted twice, so removing them once gives the distinct total.

Final Answer: 12 tuples $\Rightarrow \boxed{\text{D}}$

Answer: (D) Go Back to Q29

Q30.

Solution

Concept – `COUNT(DISTINCT col)` counts distinct non-null values: NULLs are ignored, and repeated values are counted only once.

Step 1 – list the column values: {`A`, `A`, `B`, `B`, `B`, `C`, `NULL`}.

Step 2 – drop the NULL: NULL is not counted, leaving {`A`, `A`, `B`, `B`, `B`, `C`}.

Step 3 – keep only distinct values: Distinct values are `A`, `B`, `C`.

Step 4 – count them: There are 3 distinct values.

Final Answer: `COUNT(DISTINCT dept)` returns $3 \Rightarrow \boxed{\text{C}}$

Answer: (C) Go Back to Q30

Q31.

Solution

Concept – BCNF requires every determinant to be a superkey: Check each FD's left-hand side; if all are candidate keys, the relation is in BCNF.

Step 1 – identify the candidate key: A is the sole candidate key, so A is prime and B, C are non-prime.

Step 2 – check 2NF (partial dependency): The key $\{A\}$ is a single attribute, so no proper subset exists and no partial dependency is possible. R is in 2NF.

Step 3 – check 3NF (transitive dependency): The only FDs are $A \to B$ and $A \to C$; there is no chain like $A \to X \to Y$, so no transitive dependency exists. R is in 3NF.

Step 4 – check BCNF (every determinant a superkey): Both FDs have the determinant A, which is the candidate key. Every determinant is a superkey, so R is in BCNF.

Final Answer: highest form $=$ BCNF $\Rightarrow \boxed{\text{A}}$

Answer: (A) Go Back to Q31

Q32.

Solution

Concept – natural join matches tuples on the common attribute: Combine a tuple of R with a tuple of S whenever their B values are equal.

Step 1 – match R's tuple $(1, p)$: S has (p, x) and (p, y) with $B = p$, giving $(1, p, x)$ and $(1, p, y)$ – 2 tuples.

Step 2 – match R's tuple $(2, p)$: S has (p, x) and (p, y) with $B = p$, giving $(2, p, x)$ and $(2, p, y)$ – 2 tuples.

Step 3 – match R's tuple $(3, q)$: S has (q, z) with $B = q$, giving $(3, q, z)$ – 1 tuple.

Step 4 – total the matches: $2 + 2 + 1 = 5$

Final Answer: 5 tuples $\Rightarrow$ $\boxed{\text{C}}$

Answer: (C) Go Back to Q32

Q33.

Solution

Concept – OLAP navigation operations: Roll-up aggregates to a coarser level; drill-down does the opposite, moving to finer detail.

Step 1 – describe the direction of movement: Going from monthly totals to yearly totals *removes* detail (more aggregation).

Step 2 – match the operation: Moving to a more summarized level is **roll-up**.

Why the other options are wrong: Drill-down is the reverse (more detail); slice fixes one dimension to a single value; pivot rotates the cube's axes.

Final Answer: roll-up $\Rightarrow$ $\boxed{\text{B}}$

Answer: (B) Go Back to Q33

Q34.

Solution

Concept – the star schema: A central *fact table* holds the numeric measures and foreign keys, surrounded by *dimension tables* that describe the context.

Step 1 – identify the central table's contents: It stores measures and references the surrounding tables through foreign keys.

Step 2 – name that table: A table of measures at the centre of a star schema is the **fact table**.

Why the other options are wrong: Dimension tables are the surrounding descriptive tables; an index table and a materialized view are not the schema's central measure table.

Final Answer: fact table $\Rightarrow \boxed{\text{B}}$

Answer: (B) Go Back to Q34

Q35.

Solution

Concept – attribute closure: $(AB)^+$ is the set of all attributes functionally determined by AB, found by repeatedly applying the FDs.

Step 1 – start with the attributes themselves: $(AB)^+ = \{A, B\}$

Step 2 – apply $AB \to C$: AB is present, so C is added: $(AB)^+ = \{A, B, C\}$.

Step 3 – apply $C \to D$: C is now in the closure, so D is added: $(AB)^+ = \{A, B, C, D\}$.

Step 4 – apply $D \to A$: A is already present, so nothing new is added.

Step 5 – check for more: All four attributes are present; no further FD adds anything.

Final Answer: $(AB)^+ = \{A, B, C, D\} \Rightarrow \boxed{\text{D}}$

Answer: (D) Go Back to Q35

Q36.

Solution

Concept – logistic regression prediction: Compute the linear score $z = w^\top x + b$, then apply the sigmoid.

Step 1 – compute the linear score: $z = (2)(1) + (-2)(1) + 0$

$= 2 - 2 + 0$

$= 0$

Step 2 – write the sigmoid at $z = 0$: $\sigma(0) = \dfrac{1}{1 + e^0}$

Step 3 – evaluate e^0: $e^0 = 1$

Step 4 – finish: $\sigma(0) = \dfrac{1}{1+1} = \dfrac{1}{2} = 0.50$

Final Answer: $\sigma(0) = 0.50 \Rightarrow \boxed{\text{A}}$

Answer: (A) Go Back to Q36

Q37.

Solution

Concept – underfitting in bias–variance terms: A model too simple to capture the pattern does poorly everywhere: high bias, low variance.

Step 1 – interpret the high training error: Failing to fit even the training data means the model is too rigid, i.e. *high bias*.

Step 2 – interpret the similar test error: Nearly equal train and test error means the model is stable across samples, i.e. *low variance*.

Step 3 – combine: This is the classic signature of underfitting: high bias, low variance.

Final Answer: high bias, low variance $\Rightarrow \boxed{\text{D}}$

Answer: (D) Go Back to Q37

Q38.

Solution

Concept – 1-NN takes the label of the single nearest point: Compute the distance from Q to every training point and pick the smallest.

Step 1 – distance $Q(4,4)$ to $A(1,1)$: $\sqrt{(4-1)^2+(4-1)^2} = \sqrt{9+9} = \sqrt{18} = 4.24$

Step 2 – distance to $B(5,5)$: $\sqrt{(4-5)^2+(4-5)^2} = \sqrt{1+1} = \sqrt{2} = 1.41$

Step 3 – distance to $C(4,3)$: $\sqrt{(4-4)^2+(4-3)^2} = \sqrt{0+1} = 1.00$

Step 4 – distance to $D(0,4)$: $\sqrt{(4-0)^2+(4-4)^2} = \sqrt{16+0} = 4.00$

Step 5 – pick the nearest point: The smallest distance is 1.00 to C, which is Class 1.

Step 6 – assign the label: With $k=1$, Q takes C's label, Class 1.

Final Answer: Q is Class $1 \Rightarrow \boxed{\text{C}}$

Answer: (C) Go Back to Q38

Q39.

Solution

Concept – binary entropy: $H = -p\log_2 p - (1-p)\log_2(1-p)$, with p the fraction of positives.

Step 1 – compute the class proportions: $p = \dfrac{5}{10} = 0.5$ positive, $1-p = \dfrac{5}{10} = 0.5$ negative.

Step 2 – take the base-2 logs: $\log_2 0.5 = -1$ (for both terms).

Step 3 – form each term: $-0.5 \times (-1) = 0.5$

$-0.5 \times (-1) = 0.5$

Step 4 – add the terms: $H = 0.5 + 0.5 = 1.000$ bit

Step 5 – interpret: A perfectly balanced two-class set has maximum entropy, exactly 1 bit.

Final Answer: $H = 1.000$ bit $\Rightarrow \boxed{\text{C}}$

Answer: (C) Go Back to Q39

Q40.

Solution

Concept – Gini impurity: Gini $= 1 - \sum_i p_i^2$, where p_i is the fraction of class i.

Step 1 – compute the class fractions: $p_+ = \frac{6}{10} = 0.6$, $p_- = \frac{4}{10} = 0.4$.

Step 2 – square each fraction: $p_+^2 = 0.6^2 = 0.36$

$p_-^2 = 0.4^2 = 0.16$

Step 3 – sum the squares: $0.36 + 0.16 = 0.52$

Step 4 – subtract from 1: Gini $= 1 - 0.52 = 0.48$

Final Answer: Gini $= 0.48 \Rightarrow \boxed{\text{A}}$

Answer: (A) Go Back to Q40

Q41.

Solution

Concept – SVM geometric margin: For the hyperplane $w^\top x + b = 0$, the distance from the hyperplane to a support vector is $\frac{1}{\|w\|}$ (the full margin between the two boundaries is $\frac{2}{\|w\|}$).

Step 1 – write the geometric-margin formula: $\gamma = \frac{1}{\|w\|}$

Step 2 – substitute $\|w\| = 2$: $\gamma = \frac{1}{2}$

Step 3 – evaluate: $\gamma = 0.5$

Final Answer: geometric margin $= 0.5 \Rightarrow \boxed{\text{D}}$

Answer: (D) Go Back to Q41

Q42.

Solution

Concept – k-means centroid update: The new centroid is the coordinate-wise mean of the points assigned to the cluster.

Step 1 – average the x-coordinates: $\bar{x} = \dfrac{0+2+4}{3} = \dfrac{6}{3} = 2$

Step 2 – average the y-coordinates: $\bar{y} = \dfrac{0+4+2}{3} = \dfrac{6}{3} = 2$

Step 3 – assemble the centroid: $(\bar{x}, \bar{y}) = (2,\ 2)$

Final Answer: centroid $= (2, 2) \Rightarrow \boxed{\text{C}}$

Answer: (C) Go Back to Q42

Q43.

Solution

Concept – variance explained in PCA: Each principal component explains variance proportional to its eigenvalue; the fraction is the component's eigenvalue over the total.

Step 1 – total variance (sum of all eigenvalues): $5 + 3 + 2 = 10$

Step 2 – variance in the first component: The largest eigenvalue is 5.

Step 3 – form the fraction: $\dfrac{5}{10} = 0.50$

Step 4 – express as a percentage: $0.50 = 50\%$

Final Answer: 50% of the variance $\Rightarrow \boxed{\text{D}}$

Answer: (D) Go Back to Q43

Q44.

Solution

Concept – neuron forward pass with ReLU: Compute the weighted sum plus bias, then apply $\text{ReLU}(z) = \max(0, z)$.

Step 1 – weighted sum of inputs: $w_1x_1 + w_2x_2 = (1)(2) + (2)(1)$

$= 2 + 2$

$= 4$

Step 2 – add the bias: $z = 4 + (-1) = 3$

Step 3 – apply ReLU: $\text{ReLU}(3) = \max(0,\ 3) = 3$

Final Answer: output $= 3 \Rightarrow \boxed{\text{B}}$

Answer: (B) Go Back to Q44

Q45.

Solution

Concept – A* expands the node of least $f = g + h$: Compute f for each frontier node and pick the minimum.

Step 1 – f for node A: $f(A) = g + h = 4 + 1 = 5$

Step 2 – f for node B: $f(B) = 2 + 5 = 7$

Step 3 – f for node C: $f(C) = 6 + 2 = 8$

Step 4 – choose the smallest f: $\min(5, 7, 8) = 5$, which belongs to node A.

Final Answer: A* expands $A \Rightarrow \boxed{\text{A}}$

Answer: (A) Go Back to Q45

Q46.

Solution

Concept – minimax back-up: MIN nodes take the minimum of their children; the MAX root takes the maximum of the MIN results.

Step 1 – evaluate the left MIN node: $\min(6, 4) = 4$

Step 2 – evaluate the right MIN node: $\min(8, 1) = 1$

Step 3 – evaluate the MAX root: $\max(4, 1) = 4$

Final Answer: minimax value $= 4 \Rightarrow \boxed{\text{A}}$

Answer: (A) Go Back to Q46

Answer Key

Q	Ans	Q	Ans	Q	Ans	Q	Ans	Q	Ans
1	C	2	D	3	A	4	B	5	D
6	B	7	C	8	A	9	B	10	C
11	D	12	B	13	B	14	A	15	D
16	A	17	D	18	B	19	B	20	A
21	A	22	B	23	B	24	C	25	C
26	D	27	A	28	C	29	D	30	C
31	A	32	C	33	B	34	B	35	D
36	A	37	D	38	C	39	C	40	A
41	D	42	C	43	D	44	B	45	A
46	A								