IAE ACADEMY INSTITUTE FOR ENGINEERS

GATE DA · 2026

IAE Academy free study material · Updated 27 October 2026

GATE DA 2026 Solved Paper

INSIDE

- All 65 questions, in the order of the paper
- The official answer to every question
- Step-by-step worked solutions
- The answer key on one page

Question paper and official answer key: IIT Guwahati

Online, with the questions linked to their topic notes: **iaeacademy.in/study/papers/da-2026**

FREE COUNSELLING

Preparing for GATE DA? Talk to IAE Academy.

Or call **+91 84204 36524**

Every day, 10:30 am to 7:30 pm

Chat on WhatsApp

Free. Not for sale. Pass it on.

About this paper

65 questions, 100 marks. Question types: 33 MCQ · 14 MSQ · 18 NAT.

Official answer key: IIT Guwahati · https://gate2026.iitg.ac.in/doc/download/2026/Keys/DA_Keys.pdf

We placed each question under a subject and topic ourselves; the exam paper does not label them.

This paper online, with the questions linked to their topic notes:
iaeacademy.in/study/papers/da-2026

Answer key

Q	TYPE	MARKS	ANSWER
Q.1	MCQ	1	**B**
Q.2	MCQ	1	**B**
Q.3	MCQ	1	**A**
Q.4	MCQ	1	**C**
Q.5	MCQ	1	**D**
Q.6	MCQ	2	**A**
Q.7	MCQ	2	**B**
Q.8	MCQ	2	**D**
Q.9	MCQ	2	**C**
Q.10	MCQ	2	**C**
Q.11	MCQ	1	**B**
Q.12	MCQ	1	**C**
Q.13	MCQ	1	**C**
Q.14	MCQ	1	**C**
Q.15	MCQ	1	**D**
Q.16	MCQ	1	**B**
Q.17	MCQ	1	**A**
Q.18	MCQ	1	**D**
Q.19	MCQ	1	**C**
Q.20	MCQ	1	**A**
Q.21	MCQ	1	**B**
Q.22	MCQ	1	**A**
Q.23	MSQ	1	**B**
Q.24	MSQ	1	**A, B, C**
Q.25	MSQ	1	**A, B**
Q.26	MSQ	1	**A, B**
Q.27	MSQ	1	**B, D**
Q.28	MSQ	1	**A, C**
Q.29	NAT	1	**8.90 to 9.10**
Q.30	NAT	1	**2**
Q.31	NAT	1	**10**
Q.32	NAT	1	**195**
Q.33	NAT	1	**10**
Q.34	NAT	1	**0.34 to 0.36**
Q.35	NAT	1	**1**
Q.36	MCQ	2	**D**
Q.37	MCQ	2	**D**
Q.38	MCQ	2	**B**
Q.39	MCQ	2	**C**
Q.40	MCQ	2	**D**
Q.41	MCQ	2	**A**
Q.42	MCQ	2	**D**
Q.43	MCQ	2	**A**
Q.44	MCQ	2	**A**
Q.45	MCQ	2	**A**
Q.46	MCQ	2	**C**
Q.47	MSQ	2	**A, B**
Q.48	MSQ	2	**B, C, D**
Q.49	MSQ	2	**B, C, D**
Q.50	MSQ	2	**B, D**
Q.51	MSQ	2	**A**
Q.52	MSQ	2	**A, D**
Q.53	MSQ	2	**A, B, C**
Q.54	MSQ	2	**B, C**
Q.55	NAT	2	**6.90 to 7.10**
Q.56	NAT	2	**135**
Q.57	NAT	2	**0.76 to 0.78**
Q.58	NAT	2	**8**
Q.59	NAT	2	**2 to 3**
Q.60	NAT	2	**3**
Q.61	NAT	2	**4**
Q.62	NAT	2	**10**
Q.63	NAT	2	**0**
Q.64	NAT	2	**0.08 to 0.12**
Q.65	NAT	2	**1**

General Aptitude Q.1 to Q.10

Q.1 MCQ · 1 MARK

General Aptitude › Reasoning

Verbosity : Brevity :: Insolence : _____

Choose the word that best fills the blank.

(A) Innocence

(B) Respect

(C) Solace

(D) Wealth

Official answer: **B**

SOLUTION

1. Read the first pair. Verbosity means using too many words, and brevity means using few words. The two words are opposites, so this is an antonym analogy.
2. Apply the same relation to the second pair. The blank must be the antonym of insolence.
3. Insolence means rude, disrespectful behaviour. Its opposite is respect, which means polite regard for others.
4. Check the other options. Innocence (freedom from guilt) is not the opposite of rudeness. Solace (comfort) and wealth (riches) have no antonym link with insolence.
5. So the analogy reads Verbosity : Brevity :: Insolence : Respect, which is option (B).

KEY IDEA

First name the relation between the given pair. Here the two words are antonyms. Then pick the option that has the same relation with the third word.

WATCH OUT FOR

Choosing Innocence because it sounds like a good-behaviour word. Innocence is about guilt, not about courtesy, so it is not the direct opposite of insolence.

Q.2 MCQ · 1 MARK

General Aptitude › Numerical

The product of the digits of a three-digit number is 70. The sum of the digits of this three-digit number is _____

(A) 12

(B) 14

(C) 16

(D) 18

Official answer: **B**

SOLUTION

1. Factorise the product: $70 = 2 \times 5 \times 7$, so the prime factors are 2, 5 and 7.
2. Each digit is at most 9, so every digit is a single-digit divisor of 70: 1, 2, 5 or 7 (the other divisors 10, 14, 35 and 70 are too large to be digits). Digit 0 is impossible because the product is not 0.
3. Suppose one digit is 1. The other two digits must then have product 70. The factor pairs of 70 are $(1, 70)$, $(2, 35)$, $(5, 14)$ and $(7, 10)$, and none has both factors at most 9. So 1 cannot be a digit.
4. With no digit equal to 1, each of the three digits is at least 2. The three digits have product $2 \times 5 \times 7$, which has exactly three prime factors counted with multiplicity. Three digits each at least 2 use at least one prime factor each, so each digit takes exactly one prime factor. The digits are therefore 2, 5 and 7 in some order.

5. The sum of the digits is $2 + 5 + 7 = 14$, whatever the order. The answer is option (B).

KEY IDEA
Factorise the product into primes and use the fact that a digit is at most 9. Since 70 has exactly three prime factors, the three digits must be 2, 5 and 7, so the sum does not depend on their order.

WATCH OUT FOR
Using a digit 1 with a pair such as 7 and 10, or treating 14 or 35 as a digit. Also, forgetting that the sum is the same for every arrangement of 2, 5 and 7.

Q.3 MCQ · 1 MARK

General Aptitude › Spatial

Q.3 The four pieces of a puzzle are shown in the figure below.



Which one of the figures labelled as P, Q, R, and S can be constructed by using each of the four pieces only once without overlaps?

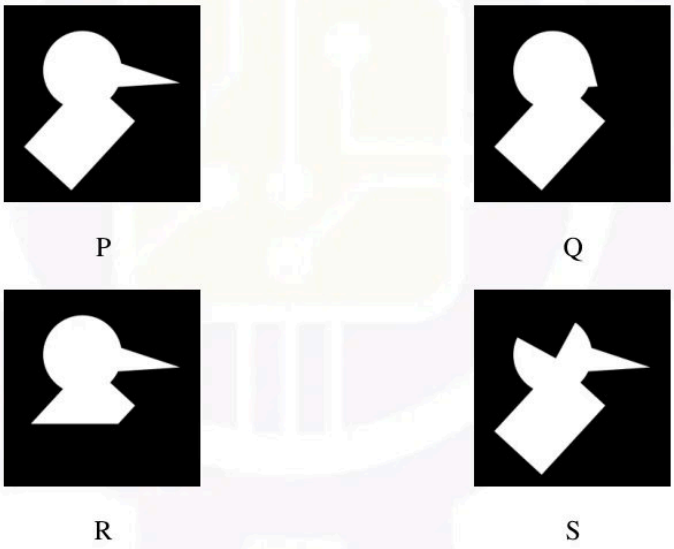

(A)	P
(B)	Q
(C)	R
(D)	S

Official answer: **A**

Q.4 MCQ · 1 MARK

General Aptitude › Numerical

Consider two distinct positive real numbers m, n, with m > n. Let x = n^{log10(m)} and y = m^{log10(n)}. The relation between x and y is _____.

(A) x > y

(B) x < y

(C) x = y

(D) x = log10(y)

Official answer: **C**

SOLUTION

1. Both x and y are positive, so they can be compared by taking $\log_{10}$ of each.
2. Power rule: $\log_{10}(a^b) = b \log_{10} a$.
3. $\log_{10} x = \log_{10} \left(n^{\log_{10} m}\right) = \log_{10} m \cdot \log_{10} n$.
4. $\log_{10} y = \log_{10} \left(m^{\log_{10} n}\right) = \log_{10} n \cdot \log_{10} m$.
5. Multiplication is commutative, so $\log_{10} x = \log_{10} y$.
6. The logarithm is one-to-one on positive numbers, so $x = y$. The condition $m > n$ plays no role.
7. Quick check: $m = 100$, $n = 10$ gives $x = 10^2 = 100$ and $y = 100^1 = 100$, so $x = y$.

KEY IDEA

Take $\log_{10}$ of both quantities: each becomes the product $\log_{10} m \cdot \log_{10} n$, so they are equal. In general $a^{\log b} = b^{\log a}$.

WATCH OUT FOR

Assuming $m > n$ forces an inequality between x and y, or trying to compare the two powers by their sizes of base and exponent instead of taking logarithms.

Q.5 MCQ · 1 MARK

General Aptitude › Reasoning

'If his latest movie had been a commercial success, the actor would have made enough money to sponsor his next movie.'

Based only on the above sentence, which one of the following statements is true?

(A) The actor will certainly sponsor his next movie.

(B) His latest movie was a commercial success.

(C) The actor made enough money from his latest movie.

(D) His latest movie was not commercially successful.

Official answer: **D**

SOLUTION

1. The sentence is a past (third) conditional: 'If A had been, B would have'. It describes an unreal past situation, so it implies that A did not happen.
2. Here A = 'his latest movie was a commercial success' and B = 'the actor made enough money to sponsor his next movie'. The sentence implies that A was false, and so B did not happen either.
3. (B) says the movie was a success. This contradicts the implication, so it is false.
4. (C) says he made enough money. B did not happen, so this is false.
5. (A) is about the future and goes beyond the sentence. The sentence says nothing certain about what he will do, and 'certainly' is unsupported, so it is false.

6. (D) says the movie was not commercially successful. This is exactly what the counterfactual 'had been' implies, so it is true.

KEY IDEA
A past conditional ('If X had been..., Y would have...') states that X did not happen, and so Y did not happen either. The false condition is the one thing the sentence lets us conclude.

WATCH OUT FOR
Reading the sentence as a plain 'if' statement and choosing (B) or (C), which assert the very things the sentence denies. Another slip is choosing (A), which looks ahead to the future and so goes beyond what the sentence says.

Q.6 MCQ · 2 MARKS

General Aptitude › Grammar & vocabulary

'My friend and I parted _____ the door _____ the cabin that I had rented _____ the night.' Choose the option with the correct sequence of words to fill the blanks.

(A) at; of; for
(B) for; at; of
(C) of; for; in
(D) in; of; for

Official answer: **A**

SOLUTION

1. Blank 1: in 'parted _____ the door', the door is the point where the two people separated. A point or spot takes the preposition 'at', so 'parted at the door'.
2. Blank 2: in 'the door _____ the cabin', the door is a part of the cabin. Belonging of a part to a whole takes 'of', so 'the door of the cabin'.
3. Blank 3: in 'rented _____ the night', the phrase gives the length of time for which the cabin was rented. 'for' marks a period of time, so 'rented for the night'.
4. The sequence is at; of; for, and the sentence reads: 'My friend and I parted at the door of the cabin that I had rented for the night.'
5. Check the others. B starts with 'parted for the door', which makes no sense. C starts with 'parted of the door', which is not idiomatic, and ends with 'rented in the night', which does not give the rental period. D has 'parted in the door', which is not the way to name a meeting or parting point, and ends with 'the door for the cabin', which does not show that the door is part of the cabin. Only A fits all three blanks.

KEY IDEA
Fill each blank by the relation it expresses. 'at' names a point or spot, 'of' links a part to its whole, and 'for' gives the length of time.

WATCH OUT FOR
Picking 'in' for the third blank because 'in the night' sounds natural. 'in the night' names a time of day. 'rented _____ the night' states how long the cabin was hired, which needs 'for'.

Q.7 MCQ · 2 MARKS

General Aptitude › Statistics & probability

Five integers are picked from 0 to 20, with possible repetitions, such that their mean is 12, median is 18, and they have a single mode of 20.

Ignoring permutations, the number of ways to pick these five integers is _____

(A) 0
(B) 1
(C) 2

(D) 3

Official answer: **B**

SOLUTION

1. Order the integers $a \le b \le c \le d \le e$. A mean of 12 gives a total of $5 \times 12 = 60$.
2. A median of 18 means the middle value is $c = 18$, so $d, e \ge 18$, and $a, b \le 18$.
3. Suppose 20 appeared only once. For 20 to be the single mode, every other value would have to occur fewer than once, which is impossible. So 20 must appear at least twice.
4. Since $a \le b \le c = 18$, only d and e can equal 20. So $d = e = 20$. A third 20 is impossible.
5. Sum condition: $a + b + 18 + 20 + 20 = 60$, so $a + b = 2$.
6. With $0 \le a \le b$ integers, the pairs are $(0, 2)$ and $(1, 1)$.
7. Check $(1, 1)$: the set $\{1, 1, 18, 20, 20\}$ has two values each appearing twice (1 and 20), so 20 is not a single mode. Reject.
8. Check $(0, 2)$: the set $\{0, 2, 18, 20, 20\}$ has mean $60/5 = 12$, median 18, and only 20 repeats, so 20 is the single mode. Valid.
9. Exactly one set works, so the number of ways is 1, option (B).

KEY IDEA
Sort the numbers. The median fixes the middle value, which forces the repeated mode into the top two places. The mean then fixes the sum of the two smallest values, leaving a short list to check.

WATCH OUT FOR
Accepting $\{1, 1, 18, 20, 20\}$ because it has the right sum, without noticing that 1 and 20 both appear twice, so there is no single mode.

Q.8 MCQ · 2 MARKS

General Aptitude › Reasoning

Rishi and Swathi are students of Class 5. Pavan and Tanvi are students of Class 4. Rishi and Pavan are boys. Swathi and Tanvi are girls. The four students played a total of three games of chess. The games were played one after another. A player who lost a game did not participate in any more games. It was observed that:
(i) the first game was the only game where two students of the same class played against each other,
(ii) the students of Class 5 won more games than the students of Class 4, and
(iii) the boys won two games and the girls won one game.
The student who did not lose any game is _____.

(A) Pavan

(B) Rishi

(C) Swathi

(D) Tanvi

Official answer: **D**

SOLUTION

1. Every game has exactly one loser, and a loser takes no further part. So the three games remove three different players, and the one student who never loses is the winner of game 3.
2. By (i), game 1 is the only game between two students of the same class. So game 1 is either Rishi vs Swathi (both Class 5) or Pavan vs Tanvi (both Class 4), and games 2 and 3 must each be between students of different classes.
3. Suppose game 1 is Pavan vs Tanvi. Its winner is a Class 4 student, and the loser is out. Game 2 must then be that winner against a Class 5 student. If the Class 5 student won, the two players left for game 3 would be Rishi and Swathi, both Class 5, which (i) forbids. So the Class 4 student wins game 2 and

has 2 wins. Class 5 can then win at most the 1 remaining game, which contradicts (ii). So this case is impossible.

4. Hence game 1 is Rishi vs Swathi. Its winner W is a Class 5 student, and the loser is out. Game 2 must be W against a Class 4 student X. If X won, game 3 would be X against the other Class 4 student, which (i) forbids. So W wins game 2 as well.
5. Game 3 is W against the other Class 4 student Y. If W won, W would have 3 wins, but (iii) allows only 2 wins for the boys and 1 for the girls, so the wins could not be split that way. Hence Y beats W. Class 5 has 2 wins and Class 4 has 1 win, which agrees with (ii).
6. W has two wins and Y has one. By (iii), the two wins belong to the boys and the one win to the girls. Since both of the boys' wins are W's, W is a boy, so W = Rishi. Then Y, whose single win is the girls' win, is a girl, so Y = Tanvi.
7. Check: Swathi lost game 1, Pavan lost game 2, Rishi lost game 3, and Tanvi played only game 3 and won it. So the student who did not lose any game is Tanvi.

KEY IDEA

Each game eliminates exactly one player, so the only student who never loses is the winner of the last game. Condition (i) fixes which pairs can meet in games 2 and 3, and (ii) and (iii) then fix the winners.

WATCH OUT FOR

Starting with Pavan vs Tanvi in game 1, or letting a Class 4 student win game 2. Both lead to two students of the same class meeting again in game 3, or to Class 4 outscoring Class 5. Another slip is to forget that the loser of each game is out, which is what makes game 3's pair forced.

Q.9 MCQ · 2 MARKS

General Aptitude › Numerical

P, Q, R, S, X, and Y are distinct single-digit whole numbers taking values from 0 to 9.

PQ is a two-digit number with Q being in the units place and P in the tens place. Similarly, RS is a two-digit number.

It is known that PQ and RS are consecutive numbers and $(PQ)^2 + (RS)^2 = XYP$, with XYP being a three-digit number.

The value of Y is _____

(A) 4

(B) 5

(C) 6

(D) 7

Official answer: **C**

SOLUTION

1. Let the smaller of the two consecutive numbers be n, so the pair is n and $n+1$. Then $n^2 + (n+1)^2 = 2n^2 + 2n + 1$.
2. Both numbers are two-digit, so $n \geq 10$. The sum must be a three-digit number, so $2n^2 + 2n + 1 \leq 999$. For $n = 21$ it is $882 + 42 + 1 = 925$, and for $n = 22$ it is $968 + 44 + 1 = 1013 > 999$. Hence $10 \leq n \leq 21$.
3. The letters are distinct, so $P \neq R$: the two numbers cannot share a tens digit. For n from 10 to 21, the numbers n and $n+1$ have different tens digits only when n ends in 9, so $n = 19$ and the pair is $19, 20$.
4. The sum is $19^2 + 20^2 = 361 + 400 = 761$, so $XYP = 761$, giving $X = 7$, $Y = 6$, and the last digit $P = 1$.

5. The last digit of the sum must equal P, the tens digit of PQ. If $PQ = 19$ then $P = 1$, which matches the units digit 1 of 761. If $PQ = 20$ then $P = 2$, but the sum is still 761 and ends in 1, so this order fails. Hence $PQ = 19$ and $RS = 20$.
6. So $P = 1, Q = 9, R = 2, S = 0, X = 7, Y = 6$. These six digits are all distinct and satisfy every condition.
7. Therefore $Y = 6$.

KEY IDEA

Write the consecutive numbers as n and $n + 1$, bound n using the three-digit sum, then use distinctness of the digits (different tens digits) and the condition that the sum ends in P to find a unique pair.

WATCH OUT FOR

Forgetting that all six letters must be distinct, so the pair cannot share a tens digit, or ignoring that the sum XYP must end in the tens digit P of the first number, which fixes the order as 19 then 20.

Q.10 MCQ · 2 MARKS

General Aptitude › Geometry

Q.10	In the given figure, P, Q, and R are three points on a circle of radius 10 cm with O as its center, $\overline{PQ} = \overline{RQ}$, and $\angle PQR = 45°$. The figure is representative. The area of the shaded region $PQRO$ is _____ cm^2. 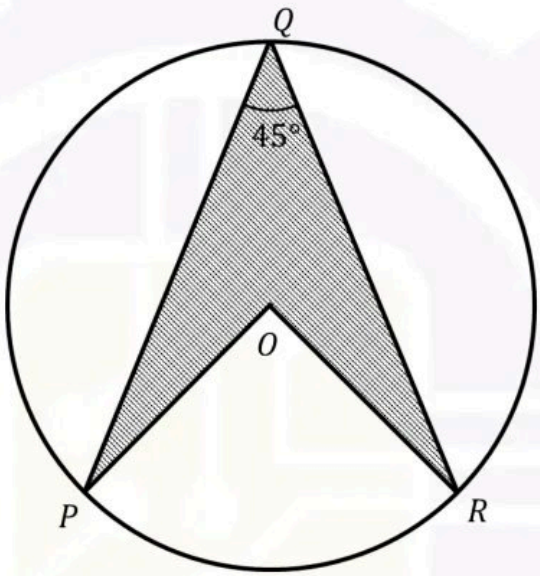
(A)	50
(B)	$25\sqrt{2}$
(C)	$50\sqrt{2}$
(D)	100

Official answer: **C**

SOLUTION

1. The inscribed angle $\angle PQR = 45°$ stands on arc PR (the arc not containing Q), so the central angle is double: $\angle POR = 90°$.
2. The rest of the circle is $360° - 90° = 270°$. Since $PQ = RQ$, the chords cut off equal arcs, so arcs PQ and QR are each $135°$.
3. A central angle equals its arc, so $\angle POQ = \angle QOR = 135°$.
4. The shaded dart $PQRO$ is made of triangles OQP and OQR sharing the side OQ.

5. Each triangle has two sides equal to the radius, $OQ = OP = 10$, with included angle 135°: $\frac{1}{2} \cdot 10 \cdot 10 \cdot \sin 135^\circ = 50 \cdot \frac{\sqrt{2}}{2} = 25\sqrt{2}$.
6. The two triangles are congruent, so the shaded area is $2 \times 25\sqrt{2} = 50\sqrt{2}\ \text{cm}^2$.
7. Check by another route: $PQ^2 = 2 \cdot 10^2(1 - \cos 135^\circ) = 200 + 100\sqrt{2}$, so $\text{area}(PQR) = \frac{1}{2}PQ^2 \sin 45^\circ = 50 + 50\sqrt{2}$. Triangle POR has area $\frac{1}{2} \cdot 10 \cdot 10 = 50$. The dart is triangle PQR minus triangle POR, which is $50\sqrt{2}$. The answer is option (C).

KEY IDEA

An inscribed angle is half the central angle on the same arc. Use this to find the arcs, then split the dart into two triangles at O with the radius as two sides and use $\frac{1}{2}r^2 \sin\theta$ for each.

WATCH OUT FOR

Treating the 45 degree angle as the angle at O, or taking the area as the sector or as triangle PQR alone. The dart is triangle PQR minus triangle POR, or equivalently the two 135 degree triangles at O.

Data Science and Artificial Intelligence Q.11 to Q.65

Q.11 MCQ · 1 MARK

Machine Learning › PCA

For a classification problem, Principal Component Analysis (PCA) has been used to reduce the dimensionality of a feature space from 100 to 10.

Which of the following options is true about the angle θ between the first and the tenth principal components?

(A) $\theta = 0^\circ$

(B) $\theta = 90^\circ$

(C) $90^\circ < \theta \leq 180^\circ$

(D) $0 < \theta < 90^\circ$

Official answer: **B**

SOLUTION

1. The principal components of the 100 features are the unit eigenvectors of the covariance matrix Σ (real and symmetric), ordered by decreasing eigenvalue: $\mathbf{v}_1$ has the largest eigenvalue and $\mathbf{v}_{10}$ the tenth largest.
2. A real symmetric matrix has an orthonormal set of eigenvectors (spectral theorem). Directly: if $\Sigma\mathbf{v}_1 = \lambda_1\mathbf{v}_1$ and $\Sigma\mathbf{v}_{10} = \lambda_{10}\mathbf{v}_{10}$ with $\lambda_1 \neq \lambda_{10}$, then
$$\lambda_1\, \mathbf{v}_1^\top \mathbf{v}_{10} = (\Sigma\mathbf{v}_1)^\top \mathbf{v}_{10} = \mathbf{v}_1^\top \Sigma^\top \mathbf{v}_{10} = \mathbf{v}_1^\top \Sigma \mathbf{v}_{10} = \lambda_{10}\, \mathbf{v}_1^\top \mathbf{v}_{10},$$
so $(\lambda_1 - \lambda_{10})\, \mathbf{v}_1^\top \mathbf{v}_{10} = 0$, which gives $\mathbf{v}_1^\top \mathbf{v}_{10} = 0$.
3. If some eigenvalues are repeated, the eigenvectors within a repeated eigenspace can still be chosen orthonormal, and PCA does choose them so. Hence the principal components are always mutually orthogonal.
4. Therefore $\cos\theta = \dfrac{\mathbf{v}_1^\top \mathbf{v}_{10}}{\|\mathbf{v}_1\|\,\|\mathbf{v}_{10}\|} = 0$, so $\theta = 90^\circ$.
5. Reducing the dimension from 100 to 10 only keeps the top 10 components and discards the other 90. It does not change the directions of the kept components, so they remain orthogonal and $\theta = 90^\circ$. This is option (B); (A), (C) and (D) all require a non-zero dot product or parallel vectors.

KEY IDEA

Principal components are eigenvectors of the symmetric covariance matrix, so they form an orthonormal set. Any two distinct components are perpendicular, whichever two you pick and however many you keep.

WATCH OUT FOR

Thinking that the first and the last retained components are somehow "far apart" or correlated, or that dimensionality reduction tilts the retained axes. PCA axes are always mutually orthogonal, and discarding components does not change that.

Q.12 MCQ · 1 MARK

Machine Learning › Bias–variance & CV

Consider that you are training a classifier for a 10-class classification problem. Each input is represented as a 512-dimensional vector. There are 1000 samples, out of which first 100 will be used for testing. Let Leave-One-Out-Cross-Validation (LOOCV) be used for selection of the classifier model before testing.

Which of the following options is the correct number of validation splits that will be generated?

(A) 10

(B) 512

(C) 900

(D) 1000

Official answer: **C**

SOLUTION

1. The 100 test samples are held out and must not take part in model selection, so the pool available for LOOCV is $1000 - 100 = 900$ samples.
2. Leave-one-out cross-validation is k-fold cross-validation with $k = n$, where n is the size of the pool: each sample is held out once as the validation set and the model is trained on the other $n - 1$ samples.
3. So the number of validation splits equals the pool size: $n = 900$.
4. The number of classes (10) and the input dimension (512) do not change how many splits LOOCV makes, so options (A) and (B) are distractors.
5. The count 1000 would include the 100 test samples, which would leak test data into model selection, so (D) is wrong.
6. Hence 900 splits, each with 899 training samples and 1 validation sample.

KEY IDEA

LOOCV makes one split per sample in the pool used for model selection. The held-out test set is excluded from that pool, so the number of splits is $1000 - 100 = 900$.

WATCH OUT FOR

Answering 1000 by counting every sample, which pulls the test data into model selection. Another slip is to link the split count to the number of classes (10) or the feature dimension (512), which have no bearing on it.

Q.13 MCQ · 1 MARK

Artificial Intelligence › Search

Which of the following algorithms is NOT an example of uninformed search?

(A) Breadth First Search

(B) Depth First Search

(C) A* Search

(D) Depth-limited Search

Official answer: **C**

SOLUTION

1. Uninformed (blind) search uses only the problem definition: the start state, the successor function and the goal test. It has no estimate of how far a state is from the goal.
2. Breadth First Search expands nodes level by level using only the graph structure, so it is uninformed.
3. Depth First Search expands the deepest node first, again with no goal-distance information, so it is uninformed.
4. Depth-limited Search is Depth First Search with a fixed depth cutoff, so it is also uninformed.
5. A* Search orders the frontier by $f(n) = g(n) + h(n)$. The heuristic $h(n)$ estimates the remaining cost to the goal, so A* is an informed (heuristic) search. It is therefore the one algorithm that is NOT uninformed, which is option (C).

KEY IDEA

Uninformed search uses no goal-distance estimate. Only a heuristic $h(n)$, as in $f(n) = g(n) + h(n)$ for A*, makes a search informed.

WATCH OUT FOR

Misreading the question: it asks which algorithm is NOT uninformed, so the answer is the informed one. Choosing a blind search such as DFS answers the opposite question.

Q.14 MCQ · 1 MARK

Artificial Intelligence › Logic

Which of the following statements is NOT true?

(The names of the predicates are intuitive.)

(A) $\forall x \forall y\ Classmate(x, y) \Rightarrow Classmate(y, x)$

(B) $\forall x\ Likes(x, Icecream) \Rightarrow \neg\exists x\ \neg Likes(x, Icecream)$

(C) "Each king is a person" is equivalent to $\forall x\ IsKing(x) \wedge IsPerson(x)$

(D) "All humans are mortal" is equivalent to $\forall x\ IsHuman(x) \Rightarrow IsMortal(x)$

Official answer: **C**

SOLUTION

1. We need the statement that is NOT true, so check each option in turn.
2. (A) $\forall x \forall y\ Classmate(x, y) \Rightarrow Classmate(y, x)$ says that being classmates is symmetric. If x is a classmate of y, then y is a classmate of x. This is true.
3. (B) By quantifier duality, $\neg\exists x\ \neg P(x) \equiv \forall x\ P(x)$. With $P(x) = Likes(x, Icecream)$, the right side is equivalent to $\forall x\ Likes(x, Icecream)$. The statement therefore has the form $Q \Rightarrow Q$, which is always true.
4. (D) "All humans are mortal" is the standard universal-implication form $\forall x\ (IsHuman(x) \Rightarrow IsMortal(x))$. This is true.
5. (C) "Each king is a person" translates to $\forall x\ (IsKing(x) \Rightarrow IsPerson(x))$.
6. The option instead gives $\forall x\ (IsKing(x) \wedge IsPerson(x))$. This says every object in the domain is both a king and a person. It fails in any domain that contains a non-king, for instance a person who is not a king. In such a domain the intended sentence is true and the option's sentence is false, so they are not equivalent. The statement is NOT true.
7. Hence the answer is (C).

KEY IDEA

A universal quantifier over a restricted class ("each king", "all humans") pairs with an implication $\Rightarrow$, never a conjunction $\wedge$. An existential statement ("some king is a person") pairs with $\wedge$.

WATCH OUT FOR

Writing $\forall x\ (P(x) \wedge Q(x))$ for "every P is a Q". That claims every object in the domain is a P and a Q, which is far stronger than the English sentence.

Q.15 MCQ · 1 MARK

Data Structures and Algorithms › Design techniques

Consider that the quick sort algorithm is used to sort an array of n distinct randomly ordered elements. In every call, the pivot is chosen as the first element of the current subarray.

Let $T(n)$ denote the expected time to sort the array. Assume that the time to partition is linear in the size of the current subarray.

Which of the following recurrence relations correctly represents $T(n)$ in this scenario?

(A) $T(n) = T(1) + T(n-1) + O(n)$

(B) $T(n) = T\left(\frac{n}{4}\right) + T\left(\frac{3n}{4}\right) + O(n)$

(C) $T(n) = 2T\left(\frac{n}{2}\right) + O(n)$

(D) $T(n) = \frac{1}{n}\sum_{k=0}^{n-1}[T(k) + T(n-k-1)] + O(n)$

Official answer: **D**

SOLUTION

1. Because the input is randomly ordered and the pivot is the first element of the current subarray, the pivot is equally likely to have any rank $1, 2, \ldots, n$ in that subarray. This holds again in every recursive call, because each subarray of a random permutation is itself randomly ordered.
2. If the pivot has rank $k+1$, partitioning leaves k elements on the left and $n-k-1$ on the right, for $k = 0, 1, \ldots, n-1$. Each of these n splits has probability $\frac{1}{n}$.
3. Partitioning scans the current subarray once, so its cost is $O(n)$ (given in the question as linear in the subarray size).
4. The expected time is the partition cost plus the expected cost of the two subproblems, averaged over all splits:

$$T(n) = \frac{1}{n}\sum_{k=0}^{n-1}[T(k) + T(n-k-1)] + O(n).$$

5. Option (A) is a single fixed split of sizes 1 and $n-1$, the extreme case of always picking a minimum or maximum pivot. It is not an expected value. Options (B) and (C) fix the split at $\frac{1}{4} : \frac{3}{4}$ or $\frac{1}{2} : \frac{1}{2}$, which assumes the pivot always lands at one chosen position. Only option (D) averages over every possible split with equal probability, so it is the recurrence for the expected time. The answer is (D).

KEY IDEA

An expected-time recurrence for quick sort averages over all n equally likely pivot ranks. Each split $(k, n-k-1)$ has probability $\frac{1}{n}$, and the linear partition cost is added on top.

WATCH OUT FOR

Choosing a fixed-split recurrence. Option (A) is the worst case, while (B) and (C) assume a constant-ratio split, such as the balanced $\frac{1}{2} : \frac{1}{2}$ split. Expected time needs the average over every split, not one assumed split.

Q.16 MCQ · 1 MARK

Data Structures and Algorithms › Python

Consider the given Python program.

```
def append_to_lst(val, lst=[]):
    lst.append(val)
    return lst
print(append_to_lst(1))
print(append_to_lst(2))
print(append_to_lst(3, []))
```

Which of the following is the correct output of this program?

(A) [1]
[2]
[3]

(B) [1]
[1, 2]
[3]

(C) [1]
[2]
[1, 2, 3]

(D) [1]
[1, 2]
[1, 3]

Official answer: **B**

SOLUTION

1. The default value `lst=[]` is evaluated once, when the `def` statement runs. The function object keeps that single list, and every call that omits the second argument uses it. Changes made to it persist between calls.
2. Call `append_to_lst(1)` : no list is passed, so the shared default list (initially empty) is used. After `lst.append(1)` it is `[1]` , and the function returns it. The program prints `[1]` .
3. Call `append_to_lst(2)` : no list is passed, so the same shared default list is used. It already holds `1` . After `lst.append(2)` it is `[1, 2]` . The program prints `[1, 2]` .
4. Call `append_to_lst(3, [])` : a new empty list is created at this call site and passed explicitly. The default is not used, and the shared list is untouched. After `lst.append(3)` this new list is `[3]` . The program prints `[3]` .
5. The three lines printed are `[1]` , `[1, 2]` and `[3]` . This is option (B).

KEY IDEA
A default argument such as `lst=[]` is evaluated once, when the function is defined, so every call that omits it shares one mutable list. A list passed explicitly is a different object and does not touch the shared default.

WATCH OUT FOR
Assuming `lst=[]` creates a fresh empty list on every call, which gives [1], [2], [3] (option A). The second call reuses the list that already holds 1, so it prints [1, 2].

Q.17 MCQ · 1 MARK

Database Management and Warehousing › Normal forms

Let R(A,B,C,D,E) be a relational schema with functional dependency set F={A→BC, CD→E, E→A}.

Which of the following statements is correct?

(A) AD, ED and CD are the only candidate keys of R.

(B) AD and ED are the only candidate keys of R.

(C) A, E and CD are the only candidate keys of R.

(D) A and CD are the only candidate keys of R.

Official answer: **A**

SOLUTION

1. Attribute D appears on no right-hand side of F, so no other attribute can derive it and D belongs to every candidate key. Attribute B appears on no left-hand side, so it is never needed to derive anything and belongs to no candidate key.
2. Every candidate key is therefore D together with a subset of $\{A, C, E\}$. Since $D^+ = \{D\}$, D alone is not a key.
3. $\{A, D\}^+$: $A \to BC$ gives $\{A, B, C, D\}$. Then $CD \to E$ gives $\{A, B, C, D, E\}$. So AD is a key.
4. $\{E, D\}^+$: $E \to A$ gives $\{A, D, E\}$. Then $A \to BC$ gives $\{A, B, C, D, E\}$. So ED is a key.
5. $\{C, D\}^+$: $CD \to E$ gives $\{C, D, E\}$. Then $E \to A$ gives $\{A, C, D, E\}$. Then $A \to BC$ gives $\{A, B, C, D, E\}$. So CD is a key.
6. Each of AD, ED, CD is minimal: removing an attribute leaves a single attribute, and A^+, E^+, C^+ do not contain D (so none is a key), while $D^+ = \{D\}$. Any other candidate key must contain D plus at least one of A, C, E (and not B), so it contains one of AD, CD, ED and would not be minimal.

7. Hence the candidate keys are exactly AD, ED and CD, which is option (A). Options (B), (C) and (D) are false: (B) omits CD, and (C) and (D) name A, E or C alone as keys, but $A^+ = \{A, B, C\}$ and $E^+ = \{E, A, B, C\}$ never contain D.

KEY IDEA
First classify attributes: one that never appears on a right-hand side must be in every key, and one that never appears on a left-hand side is in no key. Then take attribute closures of the remaining combinations and keep the minimal ones.

WATCH OUT FOR
Treating A, E or CD as keys because $A \to BC$ and $E \to A$ look like they reach everything. These never derive D, which appears on no right-hand side, so A and E alone cannot be keys. Another slip is stopping after finding AD and CD and missing ED.

Q.18 MCQ · 1 MARK

Database Management and Warehousing › Data warehousing

Consider that the visualization of a 3-dimensional data cube is showing *Sales Quantity* for each combination of the attributes *Product Type*, *Month* and *Country*.

From this, if we want to further visualize the *Sales Quantity* for each combination of *Product Type*, *Month* and *State*, which of the following OLAP operations should be performed?

(A) Slicing

(B) Dicing

(C) Roll-up

(D) Drill-down

Official answer: **D**

SOLUTION

1. The cube keeps the same three dimensions: Product Type, Month and a location attribute. Only the level of the location dimension changes, from Country to State.
2. In the location hierarchy $\text{Country} \to \text{State} \to \text{City}$, State is a finer level than Country, because each country contains many states.
3. Moving from a coarser level to a finer level of detail along a dimension hierarchy is drill-down. The reverse, moving to a coarser level, is roll-up.
4. Slicing fixes one dimension to a single value and so produces a cube with one dimension fewer. Dicing selects a sub-cube by restricting values on two or more dimensions. Neither changes the level of granularity of a dimension.
5. Going from Country to State gives more detailed data, so the operation is drill-down: option (D).

KEY IDEA
Changing the level of a dimension's hierarchy is drill-down (to a finer level) or roll-up (to a coarser level). Slicing and dicing only select part of the cube and keep the level of detail.

WATCH OUT FOR
Picking dicing or slicing because the question says 'for each combination of' attributes. These operations only restrict values. They do not change the level of detail. Another slip is reversing the direction and choosing roll-up, but State is finer than Country, so the move is down the hierarchy.

Q.19 MCQ · 1 MARK

Probability and Statistics › Counting & axioms

Let M be a randomly chosen non-empty subset of $S = \{1, 2, 3, \ldots, 2026\}$. Which of the following is the probability that the product of all the elements of M is even?

(A) $2^{1013}(2^{1013} - 1)/2^{2026}$

(B) $2^{1013}/2^{2026}$

(C) $2^{1013}(2^{1013}-1)/(2^{2026}-1)$

(D) $1/(2^{2026}-1)$

Official answer: **C**

SOLUTION

1. M is chosen uniformly from the non-empty subsets of $S=\{1,2,\ldots,2026\}$. S has 2^{2026} subsets and one of them is empty, so the sample space has $2^{2026}-1$ equally likely outcomes.
2. A product of integers is even exactly when at least one factor is even. So the product is odd exactly when every element of M is odd. Count the favourable outcomes through this complement.
3. Among $1,2,\ldots,2026$ there are 1013 odd numbers and 1013 even numbers.
4. The non-empty subsets of S made only of odd numbers are the non-empty subsets of a 1013-element set. There are $2^{1013}-1$ of them.
5. The non-empty subsets with an even product number
$$(2^{2026}-1)-(2^{1013}-1)=2^{2026}-2^{1013}=2^{1013}\left(2^{1013}-1\right).$$
6. Divide by the size of the sample space:
$$P(\text{product is even})=\frac{2^{1013}\left(2^{1013}-1\right)}{2^{2026}-1}.$$
7. This is option (C). As a check, $2^{1013}\cdot 2^{1013}=2^{2026}$, so the ratio is slightly less than 1, as it should be. Option (A) wrongly divides by 2^{2026}, which includes the empty set.

KEY IDEA

Use the complement. A product is odd only if every element is odd. Count non-empty subsets of the 1013 odd numbers, and divide by the number of non-empty subsets of S.

WATCH OUT FOR

Dividing by 2^{2026} (counting the empty set as an outcome) gives option (A). The question says M is non-empty, so the denominator is $2^{2026}-1$. A second slip is to forget to remove the empty set from the all-odd count, which should be $2^{1013}-1$.

Q.20 MCQ · 1 MARK

Probability and Statistics › Counting & axioms

Suppose that a computer program provides a non-negative and integer-valued random solution to the equation $n_1+n_2+n_3+n_4=20$. Which of the following is the probability that all of n_1,n_2,n_3,n_4 in the provided solution are positive?

(A) $\binom{19}{3}/\binom{23}{3}$

(B) $\binom{20}{4}/\binom{24}{4}$

(C) $\binom{20}{3}/\binom{23}{3}$

(D) $\binom{19}{4}/\binom{24}{4}$

Official answer: **A**

SOLUTION

1. The question treats the program's output as a random solution chosen uniformly from all non-negative integer solutions of $n_1+n_2+n_3+n_4=20$. This is the usual reading, so the probability is $\dfrac{\text{favourable solutions}}{\text{total solutions}}$.
2. Total: by stars and bars, the number of non-negative integer solutions of $n_1+n_2+n_3+n_4=20$ is $\binom{20+4-1}{4-1}=\binom{23}{3}=1771$.
3. Favourable: put $m_i=n_i-1$ for each i. Then $n_i\geq 1$ exactly when $m_i\geq 0$, and the equation becomes $m_1+m_2+m_3+m_4=20-4=16$. Its number of non-negative solutions is $\binom{16+3}{3}=\binom{19}{3}=969$.

4. Probability $= \dfrac{\binom{19}{3}}{\binom{23}{3}} = \dfrac{969}{1771} \approx 0.547$, which is option (A).
5. Check on the other options. (B) and (D) have denominator $\binom{24}{4}$. That would count non-negative solutions of a sum with five variables, or of $n_1 + \cdots + n_4 = 20$ with $n_i \leq$ something, and neither applies here. (C) has numerator $\binom{20}{3}$, which is the number of solutions with all $n_i \geq 1$ of a sum equal to 21, not 20.

KEY IDEA
Count with stars and bars. All non-negative solutions: $\binom{n+k-1}{k-1}$. For positive solutions, subtract 1 from each variable first, so the sum falls from n to $n-k$.

WATCH OUT FOR
Using $\binom{20}{3}$ or $\binom{19}{4}$ for the count, which comes from mixing up "$n+k-1$ choose $k-1$" with "n choose k". Another slip is forgetting to subtract 4 from 20 for the positive case.

Q.21 MCQ · 1 MARK

Linear Algebra › Special matrices

Let $M = \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix}$ be a 2×2 matrix, where $\theta = \frac{2\pi}{5}$, and $I_2 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$. Which of the following options is equal to M^{2026} ?

(A) M^2

(B) M

(C) M^{-1}

(D) I_2

Official answer: **B**

SOLUTION

1. M is the rotation matrix by angle θ. By the angle-addition formulas, $\begin{bmatrix} \cos a & -\sin a \\ \sin a & \cos a \end{bmatrix} \begin{bmatrix} \cos b & -\sin b \\ \sin b & \cos b \end{bmatrix} = \begin{bmatrix} \cos(a+b) & -\sin(a+b) \\ \sin(a+b) & \cos(a+b) \end{bmatrix}$, so $M^n = \begin{bmatrix} \cos n\theta & -\sin n\theta \\ \sin n\theta & \cos n\theta \end{bmatrix}$.
2. Here $\theta = \frac{2\pi}{5}$, so M^5 is a rotation by $5\theta = 2\pi$, giving $M^5 = I_2$. Hence M^n depends only on $n \bmod 5$.
3. Reduce the exponent: $2026 = 5 \times 405 + 1$, so $2026 \equiv 1 \pmod 5$.
4. Therefore $M^{2026} = (M^5)^{405}\, M = I_2^{405}\, M = M$.
5. Option (B) is correct. The other options are different rotations: M^2 is a rotation by $\frac{4\pi}{5}$, $M^{-1} = M^4$ is a rotation by $\frac{8\pi}{5}$ (equivalently $-\frac{2\pi}{5}$), and I_2 is a rotation by 0. None equals the rotation by $\frac{2\pi}{5}$.

KEY IDEA
A rotation matrix by $\theta = \frac{2\pi}{5}$ has order 5, since $M^5 = I_2$. Reduce the exponent modulo 5: $2026 \equiv 1$.

WATCH OUT FOR
Reducing the exponent incorrectly, for example reading $2026 \equiv 2$ or $2026 \equiv 0 \pmod 5$. That gives M^2 or I_2. The remainder of 2026 on division by 5 is 1.

Q.22 MCQ · 1 MARK

Linear Algebra › Vector spaces

Consider a set $S_1 = \{x = (x_1, x_2, x_3)^T \in \mathbb{R}^3 \mid x^T x \leq 16\}$. Let S_2 be another set which is a subspace of $\mathbb{R}^3$ with dimension two.

Which of the following gives the area of $S_1 \cap S_2$?

(A) 16π

(B) 4π

(C) $4\pi^2$

(D) $16\pi^2$

Official answer: **A**

SOLUTION

1. Write $x^Tx = x_1^2 + x_2^2 + x_3^2$. The condition $x^Tx \le 16$ describes the solid ball in $\mathbb{R}^3$ of radius $\sqrt{16} = 4$ centred at the origin.
2. S_2 is a subspace of $\mathbb{R}^3$ of dimension two, so it is a plane. Every subspace contains the zero vector, so this plane passes through the origin, which is the centre of the ball.
3. Take any point x of the plane. It lies in S_1 exactly when $\|x\| \le 4$. Inside the plane, with an orthonormal basis u_1, u_2 of S_2 and $x = au_1 + bu_2$, we have $\|x\|^2 = a^2 + b^2$. So $S_1 \cap S_2 = \{(a, b) : a^2 + b^2 \le 16\}$, a full disc of radius 4. This holds for every two-dimensional subspace, so the answer does not depend on which S_2 is chosen.
4. Area of the disc:

$$\pi r^2 = \pi \cdot 4^2 = 16\pi.$$

This is option (A).

KEY IDEA

A two-dimensional subspace of $\mathbb{R}^3$ is a plane through the origin. The origin is the centre of the ball $x^Tx \le 16$, so every such plane cuts the ball in a disc of the full radius, 4.

WATCH OUT FOR

Treating S_2 as a general plane that might miss the centre, and concluding the area depends on S_2. A subspace must contain the origin, so the cross-section is always the largest disc. Another slip is taking the radius as 16 instead of $\sqrt{16} = 4$, which gives 256π. A third is confusing the disc area πr^2 with the surface area or volume of the ball.

Q.23 MSQ · 1 MARK

Machine Learning › LDA, SVM & trees

In the following table, the `Task` column lists a few tasks related to machine learning. The `Algorithm` column lists a few algorithms.

Each entry "t" from the `Task` column is to be matched with an appropriate entry "a" from the `Algorithm` column such that the task "t" can be solved using the algorithm "a". Denote such a match as `t:a`.

TASK	ALGORITHM
T1 – Clustering	A1 – Markov Chain Monte Carlo
T2 – Classification	A2 – K-Medoid
T3 – Sampling	A3 – Linear Discriminant Analysis
T4 – Feature Extraction	A4 – Naive Bayes

Which of the following options is/are the correct matching(s)?

(A) T1:A4, T2:A3, T3:A1, T4:A2

(B) T1:A2, T2:A4, T3:A1, T4:A3

(C) T1:A3, T2:A4, T3:A1, T4:A2

(D) T1:A4, T2:A2, T3:A1, T4:A3

Official answer: **B**

SOLUTION

1. Each task must be paired with an algorithm that can actually solve it. Fix T1 first, because it is the task that separates the four options.
2. T1, Clustering: K-Medoid partitions unlabelled points into clusters around representative points (medoids). So T1:A2. Naive Bayes and Linear Discriminant Analysis need class labels, so they are supervised and are not clustering methods. Markov Chain Monte Carlo is a sampling method, not a clustering method.
3. T2, Classification: Naive Bayes is a generative classifier. It applies Bayes' theorem with a conditional-independence assumption on the features, so T2:A4.
4. T3, Sampling: Markov Chain Monte Carlo builds a Markov chain whose stationary distribution is the target distribution, then draws samples from it. So T3:A1.
5. T4, Feature Extraction: Linear Discriminant Analysis projects the data onto directions that maximise between-class scatter relative to within-class scatter. The projections form a reduced set of new features, so T4:A3.
6. The full matching is T1:A2, T2:A4, T3:A1, T4:A3.
7. Compare with the options. (A) has T1:A4 and (D) has T1:A4, but Naive Bayes is a supervised classifier and cannot cluster. (C) has T1:A3, but Linear Discriminant Analysis is also supervised. Option (D) also has T2:A2, but K-Medoid is a clustering algorithm. Only (B) gives the full matching above.

KEY IDEA

Match each task to the algorithm built for it. Clustering goes with K-Medoid, classification with Naive Bayes, sampling with Markov Chain Monte Carlo, and feature extraction with Linear Discriminant Analysis, which projects onto class-separating directions. Clustering is the quickest way to eliminate the wrong options.

WATCH OUT FOR

Pairing clustering with a supervised method such as Naive Bayes or Linear Discriminant Analysis, or pairing Linear Discriminant Analysis with classification only and forgetting its use as a feature-extraction (dimensionality-reduction) step. It can also be used to classify, but that does not make T1:A3 or T1:A4 valid.

Q.24 MSQ · 1 MARK

Artificial Intelligence › Logic

Sentence X is said to entail Sentence Y if whenever X is TRUE, Y also must hold TRUE.

Which of the following statements is/are correct if X entails Y?

(A) $X \Rightarrow Y$

(B) $X \wedge \neg Y$ is FALSE

(C) if X then Y

(D) if Y then X

Official answer: **A, B, C**

SOLUTION

1. Entailment means every model that makes X true also makes Y true. So no model has X true and Y false.
2. (A) $X \Rightarrow Y$ is equivalent to $\neg X \vee Y$. It is false only when X is true and Y is false, and entailment rules that case out. So the implication is true in every model. Correct.
3. (B) $X \wedge \neg Y$ is true exactly when X is true and Y is false. Entailment says this never happens, so $X \wedge \neg Y$ is FALSE in every model. Correct. This is the same condition as (A), since $\neg(X \Rightarrow Y) \equiv X \wedge \neg Y$.
4. (C) 'If X then Y' is the plain-English reading of $X \Rightarrow Y$, so it is correct for the same reason as (A).

5. (D) 'If Y then X' is the converse, $Y \Rightarrow X$. Entailment is not symmetric. Take $X = P \wedge Q$ and $Y = P$. Then X entails Y. But with P true and Q false, Y is true and X is false, so $Y \Rightarrow X$ fails. So (D) does not follow. Not correct.
6. Hence A, B and C are correct.

KEY IDEA
X entails Y exactly when $X \Rightarrow Y$ is true in every model. That is the same as saying $X \wedge \neg Y$ is never true. Options A, B and C are three ways of writing this one condition.

WATCH OUT FOR
Treating 'if Y then X' as also following. That is the converse, and entailment goes in one direction only.

Q.25 MSQ · 1 MARK

Data Structures and Algorithms › Trees & heaps

You are given the following Pre-order and In-order traversals of a Binary Tree T with nodes E, F, G, P, Q, R, S.

Pre-order: P Q S E R F G
In-order: S Q E P F R G

Which of the following statements is/are true about the Binary Tree T?

(A) Node P is the root of T
(B) The Post-order traversal of T is: S E Q F G R P
(C) Node Q has only one child
(D) The left subtree of node R contains the node G

Official answer: **A, B**

SOLUTION

1. The first node in the pre-order is the root, so the root of T is P. Statement (A) is true.
2. In the in-order, the nodes to the left of P are S, Q, E (left subtree, 3 nodes) and the nodes to the right are F, R, G (right subtree, 3 nodes). So the pre-order splits as P, then Q, S, E (left subtree), then R, F, G (right subtree).
3. Left subtree: its pre-order is Q, S, E, so its root is Q. In its in-order S, Q, E, S lies left of Q and E lies right of Q. Hence Q has left child S and right child E. Node Q has two children, so (C) is false.
4. Right subtree: its pre-order is R, F, G, so its root is R. In its in-order F, R, G, F lies left of R and G lies right of R. Hence R has left child F and right child G.
5. The left subtree of R is just $\{F\}$. Node G is in the right subtree of R, so (D) is false.
6. Post-order is left subtree, right subtree, root. Left subtree $(Q; S, E)$ gives S, E, Q. Right subtree $(R; F, G)$ gives F, G, R. Then the root P. The post-order is S, E, Q, F, G, R, P, which matches (B), so (B) is true.
7. Only (A) and (B) are true.

KEY IDEA
The first node of the pre-order is the root; its position in the in-order splits the remaining nodes into the left and right subtrees. Repeat on each part, using the subtree sizes to cut the pre-order.

WATCH OUT FOR
Splitting the pre-order at the wrong place, or treating a node as having only one child without checking both sides of it in the in-order. Here Q has S on its left and E on its right, so (C) is false. Another slip is thinking G is in R's left subtree, when it appears after R in the in-order.

Q.26 MSQ · 1 MARK

Database Management and Warehousing › Normal forms

Consider two relations r and s defined on the relational schemas R(A,B) and S(E,C), respectively. A is the primary key of R and E is a foreign key of S referencing A in R.

Which of the following operations will NEVER violate the foreign key constraint?

(A) Inserting records into relation r

(B) Deleting records from relation s

(C) Deleting records from relation r

(D) Inserting records into relation s

Official answer: **A, B**

SOLUTION

1. S is the referencing (child) relation and R is the referenced (parent) relation. The foreign key requires that every non-null value of S.E appears as some value of R.A .
2. (A) Inserting into r (parent): a new parent row only adds a possible key value. No existing child row loses its parent and no child row changes, so the foreign key cannot be violated. (Only the primary key of R could be violated, for example by a duplicate A , and that is a different constraint.) Correct.
3. (B) Deleting from s (child): removing referencing rows cannot leave any remaining S.E value without a parent. The foreign key cannot be violated. Correct.
4. (C) Deleting from r (parent): if a deleted A value is still referenced by some S.E , that child row becomes dangling. This can violate the constraint, unless a cascade or set-null action is defined. Not correct.
5. (D) Inserting into s (child): a new row whose E is non-null and absent from R.A violates the constraint. Not correct.
6. Rule: inserting into the parent and deleting from the child are always safe. Inserting into the child and deleting from the parent can violate the constraint. So the operations that never violate it are (A) and (B).

KEY IDEA

A foreign key only constrains the child's values. Adding a parent row or removing a child row cannot create a dangling reference. Adding a child row or removing a parent row can.

WATCH OUT FOR

Treating deletion from the parent as safe because the key is "still unique", or forgetting that inserting into the child needs a matching parent value. Both of these can violate the constraint.

Q.27 MSQ · 1 MARK

Calculus and Optimization › Maxima & minima

Let $f(x) = x^3 - 3x^2 + 2$ be a function defined on $(-1, 3]$.

Which of the following statements is/are correct?

(A) $f(x)$ has exactly two roots in $[-0.9, 0]$.

(B) $f(x)$ has a minimum at 2 only.

(C) $f(x)$ has maximum at 0 only.

(D) $f(x)$ has a root at 1.

Official answer: **B, D**

SOLUTION

1. Check (D) first. $f(1) = 1 - 3 + 2 = 0$, so $x = 1$ is a root and (D) is true. Dividing out $(x - 1)$ gives $f(x) = (x - 1)(x^2 - 2x - 2)$. Expanding confirms it: $x^3 - 2x^2 - 2x - x^2 + 2x + 2 = x^3 - 3x^2 + 2$.

2. The quadratic factor $x^2 - 2x - 2$ has roots $x = 1 \pm \sqrt{3}$, which are about 2.732 and -0.732. So the roots of f are $-0.732,\ 1,\ 2.732$.
3. In $[-0.9, 0]$ only $x \approx -0.732$ lies in the interval. So f has exactly one root there, not two, and (A) is false.
4. Critical points: $f'(x) = 3x^2 - 6x = 3x(x-2)$, so $x = 0$ or $x = 2$. Then $f'(x) > 0$ for $x < 0$, $f'(x) < 0$ for $0 < x < 2$ and $f'(x) > 0$ for $x > 2$. Hence f increases on $(-1, 0)$, decreases on $(0, 2)$ and increases on $(2, 3]$.
5. Values: $f(0) = 2$, $f(2) = 8 - 12 + 2 = -2$, $f(3) = 27 - 27 + 2 = 2$. At the open end, $f(-1) = -1 - 3 + 2 = -2$, but -1 is excluded.
6. Minimum: on $(-1, 0)$ the function is increasing and tends to -2 as $x \to -1^+$, so every value there is greater than -2. On $(0, 2)$ the values fall from 2 to -2, and on $(2, 3]$ they rise from -2 to 2. So the least value -2 is attained only at $x = 2$, and (B) is true.
7. Maximum: the largest value is 2, attained at both $x = 0$ and $x = 3$ (the point 3 is included). So 'maximum at 0 only' is false, and (C) is false.
8. Correct statements: (B) and (D).

KEY IDEA

On a restricted domain, find the global extrema by comparing the critical-point values with the values at the endpoints. An endpoint can tie with a critical value, and an excluded endpoint is never attained.

WATCH OUT FOR

Taking $x = 0$ as the only maximum and ignoring that $f(3) = 2$ equals $f(0)$ because 3 is included. A second slip is to treat $f(-1) = -2$ as a second minimum, although -1 is excluded from the domain. A third is to count both roots $1 \pm \sqrt{3}$ as lying in $[-0.9, 0]$, when only $1 - \sqrt{3} \approx -0.732$ does.

Q.28 MSQ · 1 MARK

Probability and Statistics › Continuous distributions

Suppose a random variable Z follows $Normal(\mu = 0, \sigma^2 = 1)$ distribution with probability density function $g(z)$ and cumulative distribution function $G(z)$. Another random variable Y follows t_1 distribution with probability density function $h(y)$ and cumulative distribution function $H(y)$. Let c be the positive real number for which $g(c) = h(c)$.

Which of the following statements is/are correct?

(A) $G(0) = H(0)$

(B) $G(c) < H(c)$

(C) $G(-c) < H(-c)$

(D) $g(0) = h(0)$

Official answer: **A, C**

SOLUTION

1. A t_1 distribution is the standard Cauchy distribution: $h(y) = \dfrac{1}{\pi(1+y^2)}$ and $H(y) = \dfrac{1}{2} + \dfrac{\arctan y}{\pi}$. The normal density is $g(z) = \dfrac{1}{\sqrt{2\pi}} e^{-z^2/2}$.
2. (A) Both densities are symmetric about 0, so each CDF equals $\frac{1}{2}$ at 0: $G(0) = H(0) = \frac{1}{2}$. True.
3. (D) $g(0) = \dfrac{1}{\sqrt{2\pi}} \approx 0.399$ while $h(0) = \dfrac{1}{\pi} \approx 0.318$. They differ, so (D) is false.
4. Locate c. The condition $g(c) = h(c)$ is equivalent to $(1 + c^2)e^{-c^2/2} = \sqrt{2\pi}/\pi \approx 0.798$. The left side equals 1 at $c = 0$, rises to about 1.21 at $c = 1$, then decreases to 0, so it meets 0.798 exactly once for $c > 0$. At $c = 1$ we have $g = 0.242 > h = 0.159$ and at $c = 2$ we have $g = 0.054 < h = 0.064$, so

$1 < c < 2$. Solving numerically gives $c \approx 1.85$, where $g \approx 0.0721$ and $h \approx 0.0720$. For $z > c$ the Cauchy density stays larger, because its tail decays polynomially while the normal tail decays like $e^{-z^2/2}$.

5. (B) $G(c) \approx 0.968$, while $H(c) = \frac{1}{2} + \frac{\arctan 1.85}{\pi} \approx 0.5 + 0.343 = 0.843$. So $G(c) > H(c)$ and (B) is false. Equivalently, the Cauchy has more mass in the tail (c, ∞), so less mass below c.
6. (C) By symmetry $G(-c) = 1 - G(c) \approx 0.032$ and $H(-c) = 1 - H(c) \approx 0.157$. So $G(-c) < H(-c)$ and (C) is true. This is the same tail comparison, seen on the left.
7. Correct statements: (A) and (C).

KEY IDEA

A t_1 distribution is the standard Cauchy distribution, whose tails are heavier than the normal's. The two densities cross once for $c > 0$, at $c \approx 1.85$. Beyond that point the Cauchy has more tail mass, so $G(-c) < H(-c)$ and $G(c) > H(c)$. At 0, both CDFs equal $\frac{1}{2}$ by symmetry, but the peak densities differ.

WATCH OUT FOR

Assuming that because $g(c) = h(c)$ the CDFs must also be ordered the same way at c and $-c$. Both curves are symmetric, so the lower-tail comparison is the opposite of the comparison at $+c$. Another slip is to think $g(0) = h(0)$ because both bells are centred at 0.

Q.29 NAT · 1 MARK

Machine Learning › Regression

Consider that for a supervised learning task, the objective function being minimized is $f_w(x) = wx$, where $x \in \mathbb{R}$ is the input and $w \in \mathbb{R}$ is the parameter. Stochastic Gradient Descent with learning rate of 0.10 is used for parameter updates.

Suppose that at the end of iteration i, the value of w becomes 10.00.

Let $x = 10.00$ be the input for iteration $(i + 1)$.

The value of w at the end of iteration $(i + 1)$ is _____. (Rounded off to two decimal places)

Official answer: **8.90 to 9.10**

SOLUTION

1. The objective is $f_w(x) = wx$, which is linear in the parameter w. Stochastic gradient descent updates w using the gradient of this objective with respect to w for the current sample.
2. Differentiate with respect to w, treating x as a constant: $\frac{\partial f_w(x)}{\partial w} = x$.
3. At iteration $(i + 1)$ the input is $x = 10.00$, so the gradient is 10.00. The value $w = 10.00$ from the end of iteration i does not enter the gradient, because the objective is linear in w.
4. SGD update rule: $w_{\text{new}} = w_{\text{old}} - \eta \frac{\partial f}{\partial w}$ with $\eta = 0.10$ and $w_{\text{old}} = 10.00$.
5. Substitute: $w_{\text{new}} = 10.00 - 0.10 \times 10.00 = 10.00 - 1.00 = 9.00$.

KEY IDEA

For a linear model $f_w(x) = wx$ the gradient with respect to w is just the input x. One SGD step is therefore $w \leftarrow w - \eta x$.

WATCH OUT FOR

Differentiating with respect to x instead of w gives w, i.e. 10, and a wrong update of $10 - 0.1 \times 10$ for the wrong reason. Another slip is to multiply the learning rate by $wx = 100$ and get 0.

Q.30 NAT · 1 MARK

Artificial Intelligence › Search

Q.30 Consider the game tree for a two-player turn-taking minimax game as shown in the figure. The value of a terminal node represents the utility of the game state if the game ends there. The numbers written next to the edges denote the strategies.

There are two players MAX and MIN. At any particular state of the game, MAX prefers to move to a state of maximum value. On the other hand, MIN prefers to move to a state of minimum value.

Suppose MAX starts the game at the root and has three strategies: 1, 2 and 3. Next, MIN plays and also has three strategies: 1, 2 and 3. The game ends there. Both players always take optimal strategies throughout the game.

At the root, the best strategy for MAX is _____ . (Answer in **integer**)

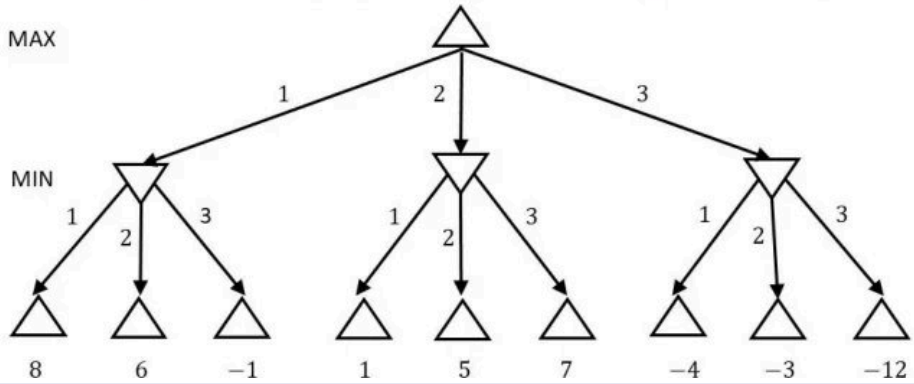

Official answer: **2**

SOLUTION

1. MIN moves at the second level and picks the smallest value among its children, so first back up the value of each MIN node.
2. Left MIN node (reached by root strategy 1): $\min(8, 6, -1) = -1$.
3. Middle MIN node (reached by root strategy 2): $\min(1, 5, 7) = 1$.
4. Right MIN node (reached by root strategy 3): $\min(-4, -3, -12) = -12$.
5. MAX at the root picks the largest backed-up value: $\max(-1, 1, -12) = 1$.
6. The value 1 comes from the middle MIN node, which is the child reached by the edge labelled 2.
7. So the best strategy for MAX is strategy 2, and the game value is 1.

KEY IDEA

In minimax, back up values from the leaves. MIN nodes take the minimum of their children and MAX nodes take the maximum. The question asks for the edge label (the strategy), not the game value.

WATCH OUT FOR

Giving the game value (1) instead of the strategy number (2). The other slip is taking the maximum at the MIN level, which would pick strategy 1 (leaf 8 beats 7 and -3).

Q.31 NAT · 1 MARK

Data Structures and Algorithms › Sorting & hashing

Let A be a sorted array containing 1000 distinct integers. You perform a recursive binary search on A to find an element y. Suppose each comparison checks whether the middle element computed during the current recursive step is equal to, less than, or greater than y.

The maximum number of comparisons that may have to be performed if y is not an element of A is _____ . (Answer in integer)

Official answer: **10**

SOLUTION

1. Each recursive step makes one three-way comparison of y with the middle element. It then recurses on the left or right part. For a subarray of m elements, these parts have $\lfloor (m-1)/2 \rfloor$ and $\lceil (m-1)/2 \rceil$ elements. Since $\lceil (m-1)/2 \rceil = \lfloor m/2 \rfloor$, the larger part has $\lfloor m/2 \rfloor$ elements.
2. If y is absent, the search can never stop early on equality. It keeps going until the subarray is empty, so the worst case follows the larger part every time.
3. Worst-case subarray sizes: $1000 \to 500 \to 250 \to 125 \to 62 \to 31 \to 15 \to 7 \to 3 \to 1 \to 0$.
4. One comparison is made at every non-empty subarray, with sizes $1000, 500, 250, 125, 62, 31, 15, 7, 3, 1$. That is 10 comparisons. The search then reaches the empty subarray and ends without any further comparison.
5. Check with the closed form: the worst case is $\lfloor \log_2 n \rfloor + 1 = \lfloor \log_2 1000 \rfloor + 1 = 9 + 1 = 10$, because $2^9 = 512 \le 1000 < 1024 = 2^{10}$.

KEY IDEA

Binary search on n elements halves the range at each comparison. An unsuccessful search runs until the range is empty, so the worst case is $\lfloor \log_2 n \rfloor + 1$ comparisons when each step uses one three-way comparison.

WATCH OUT FOR

Answering $\lceil \log_2 1000 \rceil = 10$ by coincidence, or counting the final empty-range check as an extra comparison and answering 11. Another slip is using $\log_2 1000 \approx 9.97$ and writing 9 or 10 without checking the halving sizes.

Q.32 NAT · 1 MARK

Database Management and Warehousing › Indexing

In a relational database, a B+ Tree Index is to be constructed for a relation on a key field. In a B+ Tree, a Node Pointer points to a sub-tree and a Data Record Pointer points to a block of database records.

Let, Node size = 4096 bytes, Node Pointer size = 10 bytes, Search Key Field size = 11 bytes and Data Record Pointer size = 12 bytes.

The maximum number of Node Pointers that can be present in a non-leaf node of the B+ Tree is _____ . (Answer in integer)

Official answer: **195**

SOLUTION

1. In a B+ tree, a non-leaf (internal) node stores only node pointers and search keys. Data record pointers appear only in leaf nodes, so the 12-byte data record pointer plays no part in this question.
2. A non-leaf node with n node pointers holds $n-1$ search keys. The node must fit in one block, so $10n + 11(n-1) \le 4096$.
3. Simplify: $21n - 11 \le 4096$, so $21n \le 4107$ and $n \le \dfrac{4107}{21} = 195.57\ldots$
4. n must be a whole number, so the largest value is $n_{\max} = 195$.
5. Check: for $n = 195$ the node uses $195 \times 10 + 194 \times 11 = 1950 + 2134 = 4084 \le 4096$ bytes, so it fits. For $n = 196$ it uses $1960 + 2145 = 4105 > 4096$ bytes, so it does not fit.

KEY IDEA

An internal node with n node pointers has $n-1$ search keys and no data record pointers. Solve $n \cdot p + (n-1) \cdot k \leq \text{block size}$ and take the floor.

WATCH OUT FOR

Counting n keys instead of $n-1$ gives $21n \leq 4096$, so $n = 195$ by coincidence here but a wrong method. Adding the 12-byte data record pointer to a non-leaf node is also wrong, because those pointers are held only in leaf nodes. Forgetting to take the floor of 195.57 is a third slip.

Q.33 NAT · 1 MARK

Probability and Statistics › Counting & axioms

The number of bijections $f(\cdot)$ from the set $S = \{1, 2, 3, 4\}$ to itself such that $f(f(n)) = n$, for all $n \in S$, is _____. (Answer in integer)

Official answer: **10**

SOLUTION

1. The condition $f(f(n)) = n$ for every $n \in S$ says $f \circ f = \mathrm{id}$, so f is its own inverse. A bijection with this property is an involution, and every cycle of its cycle decomposition has length 1 (a fixed point) or length 2 (a swapped pair).
2. Count by the number k of swapped pairs. The elements not in any pair are fixed points.
3. $k = 0$: every element is fixed, so f is the identity. This gives 1 bijection.
4. $k = 1$: choose the two elements that swap in $\binom{4}{2} = 6$ ways. The other two are fixed. This gives 6 bijections.
5. $k = 2$: split the 4 elements into two unordered pairs. Choose the partner of element 1 in 3 ways; the remaining two elements form the other pair. This equals $\frac{1}{2}\binom{4}{2} = 3$, so there are 3 bijections.
6. A 3-cycle or a 4-cycle is excluded, because applying f twice would not return every element to itself. Each of those has order 3 or 4 rather than dividing 2.
7. Total: $1 + 6 + 3 = 10$.
8. Check with the involution recurrence $a_n = a_{n-1} + (n-1)a_{n-2}$. Element n is either fixed (a_{n-1} ways) or paired with one of the other $n-1$ elements ($(n-1)a_{n-2}$ ways). With $a_1 = 1$ and $a_2 = 2$ this gives $a_3 = 2 + 2 \cdot 1 = 4$ and $a_4 = 4 + 3 \cdot 2 = 10$.

KEY IDEA

$f(f(n)) = n$ means f is an involution, so its cycles have length 1 or 2. Count permutations by the number of 2-cycles: $1 + 6 + 3 = 10$.

WATCH OUT FOR

Counting only the permutations made of swaps and forgetting the identity and the single-swap cases. Another slip is counting the two-pair case as $\binom{4}{2} = 6$ instead of 3, which counts each unordered pair of pairs twice.

Q.34 NAT · 1 MARK

Probability and Statistics › Continuous distributions

Let X be an exponentially distributed random variable with mean $\lambda (> 0)$. If $P(X > 5) = 0.35$, then the conditional probability $P(X > 10 \mid X > 5)$ is _____. (Rounded off to two decimal places.)

Official answer: **0.34 to 0.36**

SOLUTION

1. For an exponential random variable with mean λ, the survival function is $P(X > x) = e^{-x/\lambda}$ for $x \geq 0$.
2. Apply the definition of conditional probability. Since $\{X > 10\} \subset \{X > 5\}$, the intersection is $\{X > 10\}$:

$$P(X > 10 \mid X > 5) = \frac{P(X > 10)}{P(X > 5)} = \frac{e^{-10/\lambda}}{e^{-5/\lambda}} = e^{-5/\lambda}.$$

3. $e^{-5/\lambda} = P(X > 5)$. This is the memoryless property: $P(X > s + t \mid X > s) = P(X > t)$, here with $s = t = 5$.
4. The problem gives $P(X > 5) = 0.35$, so $P(X > 10 \mid X > 5) = 0.35$. The value of λ is not needed (it would be $\lambda = 5/\ln(1/0.35) \approx 4.76$). To two decimal places the answer is 0.35.

KEY IDEA

The exponential distribution is memoryless, so $P(X > 10 \mid X > 5) = P(X > 5)$. The given 0.35 is the answer directly, and λ never has to be found.

WATCH OUT FOR

Solving for λ first and then computing $P(X > 10)$ alone, which gives $0.35^2 \approx 0.12$. That forgets to divide by $P(X > 5)$, so it is not a conditional probability.

Q.35 NAT · 1 MARK

Calculus and Optimization › Taylor series

The value of $\sum_{i=0}^{\infty} \sum_{j=1}^{\infty} 2^{-i}\, 3^{-j}$ is _____. (Answer in integer)

Official answer: **1**

SOLUTION

1. Every term is positive and the summand $2^{-i}3^{-j}$ is a function of i times a function of j, with independent index ranges. So the double series splits into a product of two single series:
$$\sum_{i=0}^{\infty}\sum_{j=1}^{\infty} 2^{-i}3^{-j} = \left(\sum_{i=0}^{\infty} 2^{-i}\right)\left(\sum_{j=1}^{\infty} 3^{-j}\right).$$
2. The first factor is a geometric series with first term 1 and ratio $\frac{1}{2}$. Since $|r| < 1$ it converges:
$$\sum_{i=0}^{\infty} \left(\tfrac{1}{2}\right)^{i} = \frac{1}{1 - \frac{1}{2}} = 2.$$
3. The second factor is a geometric series that starts at $j = 1$, with first term $\frac{1}{3}$ and ratio $\frac{1}{3}$:
$$\sum_{j=1}^{\infty} \left(\tfrac{1}{3}\right)^{j} = \frac{\frac{1}{3}}{1 - \frac{1}{3}} = \frac{1/3}{2/3} = \frac{1}{2}.$$
4. Multiply the two sums: $2 \times \frac{1}{2} = 1$.

KEY IDEA

A double sum of positive terms whose summand factors into a function of i times a function of j equals the product of the two single sums. Each factor here is a geometric series, $\sum_{k \geq k_0} r^k = r^{k_0}/(1 - r)$ for $|r| < 1$.

WATCH OUT FOR

Using the formula $1/(1 - r)$ for the j series, which starts at $j = 1$, not $j = 0$. That gives $3/2$ in place of $1/2$ and a final answer of 3. The first term of the j series is $\frac{1}{3}$, not 1.

Q.36 MCQ · 2 MARKS

Machine Learning › Clustering

Let four points in three-dimensional space be: P1: [2, 3, −1], P2: [3, 1, 1], P3: [5, −2, 3] and P4: [3, 3, 3].

Hierarchical Agglomerative Clustering is used to cluster the above points. If Manhattan Distance is used as the distance metric during clustering, which of the following options indicates the two points that will be merged first?

(A) P1, P2

(B) P2, P3

(C) P3, P4

(D) P2, P4

Official answer: **D**

SOLUTION

1. Agglomerative clustering starts with every point as its own cluster, so the first merge joins the pair with the smallest pairwise distance. At this stage every cluster is a single point, so the linkage rule (single, complete or average) makes no difference.
2. Manhattan distance is $d(a,b) = \sum_i |a_i - b_i|$. There are $\binom{4}{2} = 6$ pairs, so compute all six.
3. $d(P1,P2) = |2-3| + |3-1| + |-1-1| = 1+2+2 = 5$.
4. $d(P1,P3) = |2-5| + |3-(-2)| + |-1-3| = 3+5+4 = 12$ and $d(P1,P4) = |2-3| + |3-3| + |-1-3| = 1+0+4 = 5$.
5. $d(P2,P3) = |3-5| + |1-(-2)| + |1-3| = 2+3+2 = 7$ and $d(P3,P4) = |5-3| + |-2-3| + |3-3| = 2+5+0 = 7$.
6. $d(P2,P4) = |3-3| + |1-3| + |1-3| = 0+2+2 = 4$.
7. The distances are $5, 12, 5, 7, 7, 4$. The smallest is 4, for $P2$ and $P4$, so they merge first. This is option (D).

KEY IDEA
The first merge in agglomerative clustering takes the closest pair of individual points. The linkage rule only matters from the second merge onward, so just compute all pairwise distances in the stated metric and take the minimum.

WATCH OUT FOR
Using Euclidean distance instead of Manhattan, or dropping a minus sign when subtracting coordinates such as -1 and -2. Another slip is to stop at the first small value, such as $d(P1,P2) = 5$, without checking all six pairs. $P1, P4$ also has distance 5, and the true minimum is 4.

Q.37 MCQ · 2 MARKS

Machine Learning › Regression

Which of the following statements is true for Ridge Regression?

(A) The regularizer in the objective function of Ridge Regression is used to guard against scenarios where the model works well for the test data, but poorly for the training data.

(B) The regularizer of Ridge Regression uses L_1 norm.

(C) Ridge Regression aims to reduce the number of parameters that have negative values.

(D) The regularizer of Ridge Regression may increase the bias of the model, but it helps in reducing the variance in predictions.

Official answer: **D**

SOLUTION

1. Ridge regression minimises $\|y - X\mathbf{w}\|^2 + \lambda\|\mathbf{w}\|_2^2$, so its penalty is the squared L_2 norm of the weights. This rules out (B), which names the L_1 norm (that is Lasso).
2. The penalty shrinks every weight towards zero, whatever its sign. It does not count weights or single out negative ones, so (C) is false.
3. Overfitting is a model that does well on training data but poorly on test data, and regularisation guards against that case. (A) states the reverse (good on test data, poor on training data), so it is false.
4. Setting the gradient to zero gives the closed form $\hat{\mathbf{w}} = (X^\top X + \lambda I)^{-1} X^\top y$. For $\lambda > 0$ this differs from the unbiased least-squares solution $(X^\top X)^{-1} X^\top y$, and the gap grows with λ. So the penalty can increase the bias.
5. Adding λI also makes $X^\top X + \lambda I$ better conditioned, so the estimate changes less from one training set to another. Variance falls. This is the bias-variance trade-off, so (D) is true.

KEY IDEA

Ridge adds a squared L_2 penalty $\lambda\|\mathbf{w}\|_2^2$. It shrinks the weights, which adds a little bias and lowers variance (the bias-variance trade-off). It is a guard against overfitting.

WATCH OUT FOR

Reading (A) carelessly. Overfitting means good on training data and poor on test data, and (A) says the opposite. Another slip is mixing up Ridge (L_2) with Lasso (L_1).

Q.38 MCQ · 2 MARKS

Artificial Intelligence › Logic

Assume that a Creative (C) person will Succeed (S) if the person is also Disciplined (D), but will not succeed otherwise. Now, consider the following statements:

(i) $C \wedge S \Leftrightarrow D$
(ii) $C \Rightarrow (S \Leftrightarrow D)$
(iii) $C \Leftrightarrow ((D \Rightarrow S) \vee \neg S)$

Which of the following options is correct?

(A) Both (i) and (ii) are TRUE

(B) Only (ii) is TRUE

(C) Both (ii) and (iii) are TRUE

(D) Only (iii) is TRUE

Official answer: **B**

SOLUTION

1. The assumption says: a creative person succeeds if disciplined and does not succeed otherwise. So whenever C holds, S and D have the same truth value: $C \Rightarrow (S \Leftrightarrow D)$. This is exactly statement (ii), so (ii) is TRUE.
2. Statement (i), with $\wedge$ binding tighter than $\Leftrightarrow$, reads $(C \wedge S) \Leftrightarrow D$. The assumption says nothing about a person who is not creative, so D can be true while C is false.
3. Counter-case for (i): take $C = \mathrm{F}$, $D = \mathrm{T}$, and any S. The assumption still holds, because $C \Rightarrow (S \Leftrightarrow D)$ is vacuously true when C is false. But $C \wedge S$ is F and D is T, so $(C \wedge S) \Leftrightarrow D$ is F. Statement (i) is therefore not guaranteed, so it is FALSE.
4. Simplify the right side of (iii): $(D \Rightarrow S) \vee \neg S = \neg D \vee S \vee \neg S$, which is a tautology (T) because $S \vee \neg S$ is always true.
5. So (iii) becomes $C \Leftrightarrow \mathrm{T}$, which is just C. The assumption does not force C to be true: with $C = \mathrm{F}$ the assumption holds but (iii) is FALSE. So (iii) is not guaranteed, so it is FALSE.
6. Only (ii) is TRUE, which is option (B).

KEY IDEA

A statement of the form "if X then (S iff D)" is a conditional. It says nothing when the condition C is false, so any statement that must hold for non-creative people as well is not guaranteed.

WATCH OUT FOR

Reading the assumption as a full biconditional between D and success for everyone, and so accepting (i). The assumption applies only to creative people, so $C = \mathrm{F}$, $D = \mathrm{T}$ breaks (i). Another slip is missing that the right side of (iii) is a tautology, which turns (iii) into plain C.

Q.39 MCQ · 2 MARKS

Data Structures and Algorithms › Python

A recursive function in Python is given.

```
def mystery(n):
  if n <= 0:
    return 1
  else:
    return mystery(n-1) + mystery(n-2)
```

Now, consider the following function call:

```
mystery(4)
```

Assume that a typical runtime stack is used to manage function calls. Each function call is pushed onto the stack and removed only after it finishes execution.

Which of the following options denotes the total number of function calls (i.e., the total number of stack activations), including the initial call, to compute `mystery(4)`?

(A) 5

(B) 9

(C) 15

(D) 17

Official answer: **C**

SOLUTION

1. Let $C(n)$ be the number of activations (including the first one) when `mystery(n)` is called.
2. Base case: if `n <= 0` the function returns `1` at once and makes no further calls, so $C(n) = 1$ for every $n \leq 0$. In particular $C(0) = C(-1) = 1$.
3. Recursive case: for `n > 0` the call counts once and then makes `mystery(n-1)` and `mystery(n-2)`, so $C(n) = 1 + C(n-1) + C(n-2)$.
4. $C(1) = 1 + C(0) + C(-1) = 1 + 1 + 1 = 3$.
5. $C(2) = 1 + C(1) + C(0) = 1 + 3 + 1 = 5$.
6. $C(3) = 1 + C(2) + C(1) = 1 + 5 + 3 = 9$.
7. $C(4) = 1 + C(3) + C(2) = 1 + 9 + 5 = 15$.
8. Check with the call tree: every call with $n > 0$ has exactly two children, and every call with $n \leq 0$ is a leaf. Let $L(n)$ be the number of leaves. Then $L(n) = 1$ for $n \leq 0$ and $L(n) = L(n-1) + L(n-2)$ otherwise, giving $L(1) = 2$, $L(2) = 3$, $L(3) = 5$, $L(4) = 8$. A full binary tree with 8 leaves has $8 - 1 = 7$ internal calls, so the total is $8 + 7 = 15$.

KEY IDEA

Count activations, not the returned value. The recurrence is $C(n) = 1 + C(n-1) + C(n-2)$ with $C(n) = 1$ for $n \leq 0$. Note that `n <= 0` makes both $n = 0$ and negative n base cases.

WATCH OUT FOR

Treating only `n == 0` as a leaf, or forgetting that `mystery(n-2)` can reach $n = -1$, which is also a leaf. Another slip is answering 5 (the depth plus one) or 9 (the count for `mystery(3)`), or counting only the leaves (8) or only the internal calls (7).

Q.40 MCQ · 2 MARKS

Data Structures and Algorithms › Graph algorithms

Consider a directed graph $G = (V, E)$, where V is the finite set of vertices and E is the set of directed edges between the vertices. G may contain cycles but there is no self-loop. Further, G may not be strongly connected.

Let G^R be the graph obtained by reversing the directions of all the edges in G without changing the set of vertices.

Assume that Breadth First Search (BFS) or Depth First Search (DFS) from any given vertex v of a graph visits only the reachable vertices from v in that graph.

Which of the following statements must always be true, regardless of the structure of G?

(A) If u is a reachable vertex in the BFS of G^R from v, then u is also a reachable vertex in the DFS of G from v.

(B) In G^R, the BFS traversal from v will visit exactly the same set of vertices as the DFS from v in G.

(C) The order of vertices visited in the BFS of G^R from v is the reverse of the order of vertices visited in the DFS of G from v.

(D) If u is a reachable vertex in the DFS of G from v, then v is also a reachable vertex in the BFS of G^R from u.

Official answer: **D**

SOLUTION

1. BFS and DFS from a vertex both visit exactly the vertices reachable from it, so each statement reduces to a claim about paths. Reversing every edge turns a path $a \rightsquigarrow b$ in G into a path $b \rightsquigarrow a$ in G^R, and conversely.
2. (A) u reached by BFS of G^R from v means there is a path $v \rightsquigarrow u$ in G^R, which is a path $u \rightsquigarrow v$ in G. DFS of G from v reaches u only if there is a path $v \rightsquigarrow u$ in G. Counterexample: G has the single edge $u \to v$. Then G^R has the edge $v \to u$, so BFS of G^R from v reaches u, but DFS of G from v visits only v. False.
3. (B) BFS in G^R from v visits the vertices that can reach v in G, while DFS in G from v visits the vertices reachable from v. Take G with only the edge $v \to w$. DFS in G visits $\{v, w\}$, but G^R has only the edge $w \to v$, so BFS in G^R from v visits $\{v\}$. The sets differ. False.
4. (C) Take G with edges $v \to a$ and $v \to b$. In G^R the edges are $a \to v$ and $b \to v$, so BFS from v visits only v, while DFS of G from v visits v, a, b. The orders are not reverses of each other. False.
5. (D) If DFS of G from v reaches u, there is a path $v \rightsquigarrow u$ in G. Reversing every edge gives a path $u \rightsquigarrow v$ in G^R. BFS of G^R from u visits every vertex reachable from u, so it visits v. This uses only the path correspondence, so it holds for every G, with or without cycles and whether or not G is strongly connected. True.

KEY IDEA

Reversing all edges reverses every path, so $v \rightsquigarrow u$ in G if and only if $u \rightsquigarrow v$ in G^R. BFS and DFS both visit exactly the reachable set, so only statements phrased through this path reversal always hold.

WATCH OUT FOR

Treating reachability from v in G^R as the same as reachability from v in G (options A and B), or assuming BFS and DFS orders are related by reversal (option C). Only the start and end of the path swap, not the starting vertex.

Q.41 MCQ · 2 MARKS

Database Management and Warehousing › Indexing

Q.41 Consider a B+ Tree where the maximum number of key values in each leaf node is 2 and the maximum number of pointers in each non-leaf node is 3. Let the content of the B+ Tree be as shown in the figure.

Which of the following options denotes the key value(s) stored in the root node after inserting a key value 3 in the given B+ Tree?

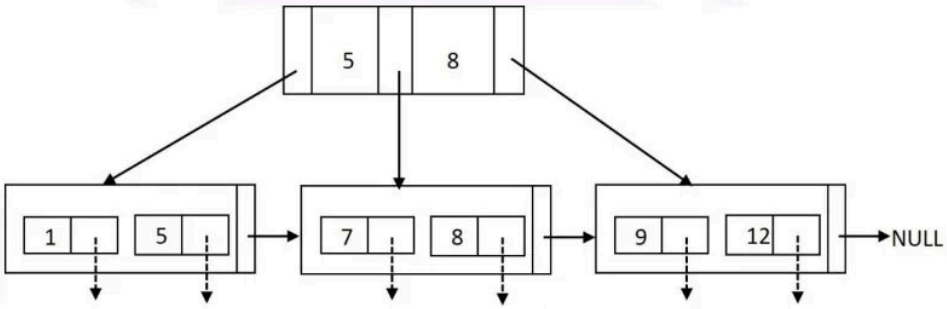

(A) 5

(B) 8

(C) 3 and 5

(D) 3, 5 and 8

Official answer: **A**

SOLUTION

1. Each non-leaf node holds at most 3 pointers, so at most 2 keys. Each leaf holds at most 2 keys. In the figure the separator 5 sits above a leaf that contains 5 ($[1, 5]$), and the separator 8 sits above a leaf that contains 8 ($[7, 8]$). So keys $\leq$ a separator lie in the subtree to its left.
2. The key 3 is ≤ 5, so it goes into the leftmost leaf $[1, 5]$. The leaf would become $[1, 3, 5]$ with 3 keys, which exceeds the limit of 2, so the leaf splits.

3. Split the leaf into $[1, 3]$ and $[5]$. In a B+ tree the separator is copied up and stays in the leaf. With the keys $\leq$ separator convention the separator is 3, the largest key of the left leaf.
4. The root would now hold keys $3, 5, 8$ with 4 child pointers. This is more than the 3 pointers allowed, so the root overflows and splits.
5. Splitting a non-leaf node moves the middle key up instead of copying it. The 4 pointers divide 2 and 2. The middle key 5 moves up. The left internal node holds key 3 with children $[1, 3]$ and $[5]$. The right internal node holds key 8 with children $[7, 8]$ and $[9, 12]$.
6. The new root holds the single key 5. If the leaf were split as $[1]$ and $[3, 5]$ instead, the copied-up separator would be 1, the root would briefly hold $1, 5, 8$, and the middle key would still be 5. So the root is $[5]$ under either split, and the answer is (A).

KEY IDEA

A leaf split copies the separator up and keeps it in the leaf. A non-leaf split moves the middle key up. Here the inserted 3 overflows the leaf, which makes the root overflow, and the middle key 5 becomes the new root.

WATCH OUT FOR

Stopping after the leaf split and concluding the root is $3, 5, 8$ or $3, 5$ (options D or C). With at most 3 pointers per non-leaf node, 3 keys and 4 pointers cannot stay in one node, so the root must split too. Another mistake is to copy the middle key up from the internal node instead of moving it, which would leave 5 duplicated in the internal level.

Q.42 MCQ · 2 MARKS

Database Management and Warehousing › Relational algebra

Q.42 Consider the given relations X, Y and Z. The relation X has three columns P, Q and R. The relation Y has three columns P, Q and S. The relation Z has two columns P and T.

X

P	Q	R
P1	Q1	R1
P2	Q2	R2
P3	Q3	R2

Y

P	Q	S
P1	Q1	2
P1	Q2	5
P2	Q1	6
P3	Q3	1

Z

P	T
P1	T1
P3	T2
P4	T3
P4	NULL

Consider the relational algebra expression

$$\prod_{PRS}\left[\left(\sigma_{(Q = Q3 \vee R = R2)}[X \bowtie Y]\right) \bowtie \left(\sigma_{(S > 1)}[Y \bowtie Z]\right)\right]$$

where $\bowtie$ denotes natural join operation.

Which of the following options is the correct output for the given expression?

(A) Two rows (P1, R1, 2) and (P1, R1, 5)

(B) Three rows (P1, R1, 2), (P1, R1, 5) and (P2, R2, 6)

(C) One row (P1, R1, 2)

(D) Zero rows

Official answer: **D**

SOLUTION

1. Left part: $X \bowtie Y$ is a natural join on the common attributes P, Q. The pairs (P, Q) in X are $(P1, Q1)$, $(P2, Q2)$, $(P3, Q3)$. The pairs in Y are $(P1, Q1)$, $(P1, Q2)$, $(P2, Q1)$, $(P3, Q3)$. Only $(P1, Q1)$ and $(P3, Q3)$ occur in both, so the join has the rows $(P1, Q1, R1, 2)$ and $(P3, Q3, R2, 1)$ over (P, Q, R, S).
2. Apply $\sigma_{Q=Q3 \vee R=R2}$. The row $(P1, Q1, R1, 2)$ has $Q = Q1$ and $R = R1$, so both conditions fail and it is removed. The row $(P3, Q3, R2, 1)$ satisfies both and stays. Left result: $L = \{(P3, Q3, R2, 1)\}$.

3. Right part: $Y \bowtie Z$ is a natural join on P. Z contains $P1, P3, P4$. Y rows with $P1$ give $(P1, Q1, 2, T1)$ and $(P1, Q2, 5, T1)$. The Y row with $P3$ gives $(P3, Q3, 1, T2)$. The Y row with $P2$ has no partner in Z and is dropped. $P4$ appears in no Y row.
4. Apply $\sigma_{S>1}$. The row with $S = 1$ is removed, leaving $R = \{(P1, Q1, 2, T1),\ (P1, Q2, 5, T1)\}$ over (P, Q, S, T).
5. Final natural join $L \bowtie R$ is on the common attributes P, Q, S. Every tuple of L has $P = P3$, and every tuple of R has $P = P1$, so no pair of tuples agrees on P. The join is empty.
6. The projection π_{PRS} of an empty relation is empty, so the output has zero rows. This is option (D).

KEY IDEA

A natural join keeps only tuples that agree on every common attribute. Evaluate each side of the final join separately and compare them on the shared attributes P, Q, S.

WATCH OUT FOR

Reading the final join as being on P alone, or forgetting that the selection with $Q = Q3 \vee R = R2$ removes the $P1$ row from the left side. Another slip is to treat $Y \bowtie Z$ as the only source of rows and read off $(P1, R1, 2)$ and $(P1, R1, 5)$, ignoring that the left side contains only $P3$.

Q.43 MCQ · 2 MARKS

Q.43	Consider the concept hierarchies as shown in the figure. Which of the following options denotes the total number of possible data cuboids from these concept hierarchies? 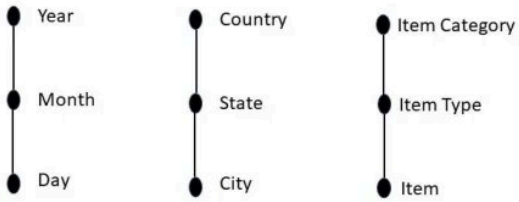
(A)	4^3
(B)	2^3
(C)	2!
(D)	4!

Official answer: **A**

SOLUTION

1. The figure shows three dimensions, each with a concept hierarchy of 3 levels: time (Year, Month, Day), location (Country, State, City) and item (Item Category, Item Type, Item).
2. A cuboid is obtained by choosing, for every dimension, one abstraction level to group by. Besides the three named levels, each dimension also has the virtual top level 'all', where that dimension is aggregated away completely.

3. So each dimension offers $L_i + 1 = 3 + 1 = 4$ choices of level, and the choices for different dimensions are independent.
4. By the multiplication rule, the total number of cuboids is $\prod_{i=1}^{n}(L_i + 1)$.
5. Here $n = 3$ and every $L_i = 3$, so the total is $4 \times 4 \times 4 = 4^3 = 64$.
6. The value $2^3 = 8$ would be the count if no dimension had a hierarchy, so each dimension offered only 'present' or 'aggregated away'. That does not apply, because every dimension has several levels.
7. $2!$ and $4!$ equal 2 and 24, and neither equals 64. Hence the answer is option (A), 4^3.

KEY IDEA

With concept hierarchies, a cuboid picks one level per dimension, including the top level 'all'. The count is $\prod(L_i + 1)$, where L_i is the number of levels of dimension i excluding 'all'.

WATCH OUT FOR

Using 2^n, the count for dimensions without hierarchies, and choosing option (B). Another slip is counting only the 3 named levels per dimension and forgetting the virtual level 'all', which gives 3^3 and is not an option.

Q.44 MCQ · 2 MARKS

Probability and Statistics › Random variables

Let X and Y be two independent random variables. X follows $Bernoulli(p = 0.3)$ distribution and Y follows $Normal(\mu = 0, \sigma^2 = 100)$ distribution.

Which of the following options is the variance of $(2X - 1)Y$?

(A) 100

(B) 90

(C) 49

(D) 21

Official answer: **A**

SOLUTION

1. Let $W = 2X - 1$. Since $X \in \{0, 1\}$, $W = -1$ when $X = 0$ and $W = 1$ when $X = 1$. So $W^2 = 1$ with certainty, and $E[W^2] = 1$ whatever the value of p.
2. W is a function of X alone and X is independent of Y, so W and Y are independent.
3. Because W and Y are independent and $E[Y] = 0$: $E[WY] = E[W]\,E[Y] = 0$.
4. W^2 and Y^2 are also independent, so $E[W^2Y^2] = E[W^2]\,E[Y^2]$.
5. Hence $\mathrm{Var}(WY) = E[W^2Y^2] - (E[WY])^2 = E[W^2]\,E[Y^2] - 0$.
6. $E[Y^2] = \mathrm{Var}(Y) + (E[Y])^2 = 100 + 0 = 100$, so $\mathrm{Var}\left((2X - 1)Y\right) = 1 \times 100 = 100$.
7. Check with the general formula for independent W and Y: $\mathrm{Var}(WY) = E[W^2]\sigma_Y^2 + \mathrm{Var}(W)\mu_Y^2$. Here $E[W] = 2(0.3) - 1 = -0.4$ and $\mathrm{Var}(W) = 4 \cdot 0.3 \cdot 0.7 = 0.84$, so the value is $1 \cdot 100 + 0.84 \cdot 0 = 100$. The answer is option (A).

KEY IDEA

For independent W and Y with $E[Y] = 0$, $\mathrm{Var}(WY) = E[W^2]\,E[Y^2]$. Here $2X - 1$ takes only the values ± 1, so $E[W^2] = 1$ and the variance equals $\mathrm{Var}(Y)$, whatever p is.

WATCH OUT FOR

Using $\mathrm{Var}(2X - 1)\,\mathrm{Var}(Y) = 0.84 \times 100 = 84$, which wrongly treats the variance of a product as the product of the variances. Another slip is using $4p(1 - p) \cdot 100$ and ignoring that $E[(2X - 1)^2] = 1$, not 0.84.

Q.45 MCQ · 2 MARKS

Probability and Statistics › CLT & tests

Let

$$L = \lim_{n\to\infty} \sum_{k=0}^{n} \frac{e^{-n} n^k}{k!}$$

Which of the following is the value of L?

(A) 0.5

(B) 1.0

(C) 0

(D) e^{-1}

Official answer: **A**

SOLUTION

1. The term $\frac{e^{-n} n^k}{k!}$ is the probability mass function of a $\text{Poisson}(n)$ random variable at k. So the sum equals $P(X_n \le n)$, where $X_n \sim \text{Poisson}(n)$.
2. A $\text{Poisson}(n)$ variable has the same distribution as the sum of n independent $\text{Poisson}(1)$ variables $Y_1, \ldots, Y_n$, because Poisson variables add by adding their means. Each Y_i has mean 1 and variance 1, so $E[X_n] = n$ and $\text{Var}(X_n) = n$.
3. Standardise: $P(X_n \le n) = P\left(\frac{X_n - n}{\sqrt{n}} \le 0\right)$.
4. By the central limit theorem, $\frac{X_n - n}{\sqrt{n}} \to Z \sim N(0,1)$ in distribution as $n \to \infty$.
5. Convergence in distribution gives convergence of the cdf at every continuity point of the limit cdf. The normal cdf is continuous everywhere, in particular at 0, so $L = \Phi(0) = \frac{1}{2}$.
6. Hence $L = 0.5$, which is option (A).

KEY IDEA

The sum is the Poisson(n) cdf evaluated at its own mean n. A Poisson(n) variable is a sum of n independent Poisson(1) variables, so the central limit theorem makes it approximately normal, and the normal cdf at its mean is $\frac{1}{2}$.

WATCH OUT FOR

Reading the sum as the whole Poisson series and answering 1 (option B). That would need the sum to run to infinity. It stops at $k = n$, which is the mean, so only about half the probability mass is included. Another trap is answering e^{-1} by looking only at the e^{-n} factor.

Q.46 MCQ · 2 MARKS

Linear Algebra › Eigenvalues

Let $\gamma_1, \gamma_2, \gamma_3$ be the eigenvalues of the matrix $\begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos t & \sin t \\ 0 & -\sin t & \cos t \end{bmatrix}$, where $t \in [-\pi, \pi]$ is in radians.

Which one of the following options lists all the possible values of t satisfying $\gamma_1 + \gamma_2 + \gamma_3 = 1 + \sqrt{2}$?

(A) $\{\frac{\pi}{3}, -\frac{\pi}{4}\}$

(B) $\{\frac{\pi}{4}, -\frac{\pi}{3}\}$

(C) $\{\frac{\pi}{4}, -\frac{\pi}{4}\}$

(D) $\{\frac{\pi}{3}, -\frac{\pi}{3}\}$

Official answer: **C**

SOLUTION

1. The sum of the eigenvalues of a square matrix equals its trace, whether the eigenvalues are real or complex: $\gamma_1 + \gamma_2 + \gamma_3 = \text{tr}(A)$.
2. The diagonal entries are 1, $\cos t$ and $\cos t$, so

$$\mathrm{tr}(A) = 1 + 2\cos t.$$

3. Set the trace equal to the given sum: $1 + 2\cos t = 1 + \sqrt{2}$. This gives $\cos t = \dfrac{\sqrt{2}}{2}$.
4. On $[-\pi, \pi]$, $\cos t = \dfrac{\sqrt{2}}{2}$ holds exactly at $t = \dfrac{\pi}{4}$ and $t = -\dfrac{\pi}{4}$. Cosine is even, so the two values are symmetric. No other angle in this interval has this cosine: $\cos t$ is positive only for $|t| < \dfrac{\pi}{2}$, and on $\left[0, \dfrac{\pi}{2}\right)$ it is strictly decreasing, so it takes the value $\dfrac{\sqrt{2}}{2}$ once.
5. The full set of values is $\left\{\dfrac{\pi}{4}, -\dfrac{\pi}{4}\right\}$, which is option (C).

KEY IDEA

The sum of the eigenvalues is the trace, so there is no need to find the eigenvalues. Here the trace is $1 + 2\cos t$. Solve $\cos t = \dfrac{\sqrt{2}}{2}$ on $[-\pi, \pi]$, where both t and $-t$ qualify because cosine is even.

WATCH OUT FOR

Computing the characteristic polynomial and the complex eigenvalues $\cos t \pm i \sin t$ takes longer and invites sign errors. Another slip is to give only the positive angle, or to pair $\dfrac{\pi}{4}$ with $-\dfrac{\pi}{3}$ because $\sin t$ also appears in the matrix. The trace involves only $\cos t$, so t and $-t$ always give the same sum.

Q.47 MSQ · 2 MARKS

Machine Learning › Bias–variance & CV

Consider that 20 stories of Author X and 10 stories of Author Y were kept together without mentioning the names of the authors. A classifier was then asked to predict the author (X or Y) of each of these stories. Let, out of X's stories, 6 were classified as that of Y. On the other hand, out of Y's stories, 2 were classified as that of X.

Considering X and Y as two classes, which of the following statements is/are true?

(A) Accuracy of the classifier is 11/15.

(B) Precision of Class X is higher than the Precision of Class Y.

(C) Recall of Class X is higher than the Recall of Class Y.

(D) Accuracy of the classifier is 14/15.

Official answer: **A, B**

SOLUTION

1. Build the confusion matrix. Of the 20 stories by X, 6 were classified as Y, so 14 were classified as X. Of the 10 stories by Y, 2 were classified as X, so 8 were classified as Y.

	PREDICTED X	PREDICTED Y
Actual X	14	6
Actual Y	2	8

2. Accuracy is the number of correct predictions over the total: $\dfrac{14 + 8}{30} = \dfrac{22}{30} = \dfrac{11}{15}$. So (A) is true and (D), which claims $\dfrac{14}{15}$, is false.
3. Precision of a class is the fraction of stories predicted as that class that truly belong to it. For X, 14 of the $14 + 2 = 16$ stories predicted as X are correct: $P_X = \dfrac{14}{16} = 0.875$.
4. For Y, 8 of the $8 + 6 = 14$ stories predicted as Y are correct: $P_Y = \dfrac{8}{14} \approx 0.571$. Since $0.875 > 0.571$, (B) is true.

5. Recall of a class is the fraction of its actual stories that were found. $R_X = \frac{14}{20} = 0.7$ and $R_Y = \frac{8}{10} = 0.8$. Recall of X is lower than recall of Y, so (C) is false.
6. The true statements are (A) and (B).

KEY IDEA

Build the 2x2 confusion matrix from the counts. Precision divides by the column total (everything predicted as that class); recall divides by the row total (everything that truly is that class); accuracy divides the diagonal by the grand total.

WATCH OUT FOR

Mixing up precision and recall by dividing by the wrong total. Another slip is computing accuracy as 14/15 or from one class only. Here the class with the higher recall (Y) has the lower precision, so swapping the two definitions gives the wrong verdict on (B) and (C).

Q.48 MSQ · 2 MARKS

Artificial Intelligence › Logic

Let $P(x)$ be a predicate.

Which of the following statements is/are NOT valid in first-order logic?

(A) $\forall x\, P(x) \Rightarrow \exists x\, P(x)$

(B) $\exists x\, P(x) \Rightarrow \forall x\, P(x)$

(C) $\exists x\, P(x) \Leftrightarrow \forall x\, P(x)$

(D) $\forall x\, P(x) \Rightarrow \exists x\, \neg P(x)$

Official answer: **B, C, D**

SOLUTION

1. A formula is valid if it is true in every interpretation. In standard first-order logic every domain is non-empty. We check each option for a counter-model.
2. (A) $\forall x\, P(x) \Rightarrow \exists x\, P(x)$: if P holds for every element and the domain has at least one element, then P holds for some element. No counter-model exists, so (A) is valid.
3. (B) $\exists x\, P(x) \Rightarrow \forall x\, P(x)$: take domain $\{1, 2\}$ with $P(1)$ true and $P(2)$ false. Then $\exists x\, P(x)$ is true and $\forall x\, P(x)$ is false, so the implication is false in this interpretation. (B) is NOT valid.
4. (C) $\exists x\, P(x) \Leftrightarrow \forall x\, P(x)$: a biconditional needs both directions. The direction $\forall \rightarrow \exists$ holds, but $\exists \rightarrow \forall$ fails in the same counter-model as in (B). So (C) is false there and NOT valid.
5. (D) $\forall x\, P(x) \Rightarrow \exists x\, \neg P(x)$: take any non-empty domain in which P is true everywhere. Then $\forall x\, P(x)$ is true, while $\exists x\, \neg P(x)$ is false, so the implication is false. (D) is NOT valid. (It is still satisfiable: it is true in any interpretation where P fails for some element, because the antecedent is then false. Satisfiable is not the same as valid.)
6. Hence the statements that are NOT valid are (B), (C) and (D).

KEY IDEA

A formula is valid only if no interpretation falsifies it. Because first-order domains are non-empty, $\forall \rightarrow \exists$ is valid, but $\exists \rightarrow \forall$ and anything claiming $\forall x\, P(x)$ forces $\exists x\, \neg P(x)$ each have a simple counter-model.

WATCH OUT FOR

Reading "NOT valid" as "never true", or confusing satisfiable with valid. A statement can be true in some models (as (B) and (D) are) and still not be valid. Also, forgetting that the question asks for the NOT valid ones and picking (A).

Q.49 MSQ · 2 MARKS

Data Structures and Algorithms › Sorting & hashing

Consider the problem of sorting the given array in ascending order:

P = [1, 2, 3, 5, 4]

Consider two sorting algorithms Bubble Sort (BS) and Insertion Sort (IS).

Let N1 be the total number of comparisons done by BS on the elements of P and N2 be the total number of comparisons done by IS on the elements of P.

Which of the following options is/are correct?

(A) N1 = 10, N2 = 4

(B) N1 > N2

(C) IS on P will perform only one swap

(D) Both BS and IS on P will make at least one unnecessary comparison (i.e., comparing elements that are already in correct order)

Official answer: **B, C, D**

SOLUTION

1. Insertion sort (IS) on P = [1, 2, 3, 5, 4] . Inserting 2, 3 and 5 each costs exactly one comparison. The key is already larger than its left neighbour, so the inner loop stops at once. That is 3 comparisons and no shifts.
2. Inserting 4: compare with 5. Since 4 < 5, shift 5 one place right (1 swap/shift). Then compare with 3. Since 4 > 3, stop. That is 2 comparisons and 1 shift.
3. So $N_2 = 3 + 2 = 5$, and IS performs exactly one swap/shift. Option (A) is false because $N_2 = 5$, not 4. Option (C) is true.
4. Bubble sort (BS), pass 1: compare (1,2), (2,3), (3,5), (5,4). Only the last pair swaps, giving [1, 2, 3, 4, 5] . That is 4 comparisons.
5. Pass 2 makes 3 comparisons, (1,2), (2,3), (3,4), with no swap. With the early-exit (swapped flag) version BS stops here, so $N_1 = 4 + 3 = 7$. Without early exit BS makes $\frac{n(n-1)}{2} = 10$ comparisons. Either way $N_1 \geq 7 > 5 = N_2$, so (B) is true. Note also that (A) fails on N_2 regardless of which BS version is used.
6. Unnecessary comparisons: BS compares (1,2), (2,3) and (3,5) in pass 1, and these are already in order. IS compares 2 with 1, 3 with 2 and 5 with 3, which are also already in order. Only the comparison of 4 with 5 involves a misplaced pair. So both algorithms make at least one such comparison and (D) is true.
7. Correct options: B, C, D.

KEY IDEA

Count comparisons directly. Insertion sort on an almost-sorted array costs one comparison per already-placed element plus the extra comparisons for the one displaced element. Bubble sort needs at least two passes here, because the early-exit pass that confirms sortedness is still a full pass of comparisons.

WATCH OUT FOR

Counting only the comparisons that lead to a swap, which gives N2 = 1 or 4. Also assuming IS needs only n - 1 = 4 comparisons. Inserting 4 needs two comparisons, so N2 = 5. Another slip is taking N1 = 4 by forgetting the second pass that confirms the array is sorted.

Q.50 MSQ · 2 MARKS

Data Structures and Algorithms › Python

Consider the given Python program.

```
def outer():
  x = []
  def inner(val):
    x.append(val)
    return x
  return inner
f1 = outer()
f2 = outer()
print(f1(10))  # Line P
print(f1(20))  # Line Q
print(f2(30))  # Line R
print(f1(40))  # Line S
```

Which of the following options is/are correct?

(A) f1 and f2 share the same list x

(B) Output of Line Q is [10, 20]

(C) Output of Line R is [10, 20, 30]

(D) Output of Line S is [10, 20, 40]

Official answer: **B, D**

SOLUTION

1. Each call to outer() runs its body afresh, so it creates a new list x and a new function object inner that captures that list.
2. f1 = outer() makes f1 a closure over a list, call it x_1. f2 = outer() makes f2 a closure over a different list x_2. Both lists start empty.
3. Line P: f1(10) appends 10 to x_1 and returns x_1 itself, so it prints [10] .
4. Line Q: f1(20) appends 20 to x_1, which now holds 10 and 20, so it prints [10, 20] . Option (B) is true.
5. Line R: f2(30) appends 30 to x_2, which was still empty, so it prints [30] , not [10, 20, 30] . Option (C) is false.
6. Line S: f1(40) appends 40 to x_1, which held [10, 20] , so it prints [10, 20, 40] . Option (D) is true.
7. Option (A) is false: f1 and f2 capture separate lists, as the output of Line R shows. The correct options are B and D.

KEY IDEA

A closure captures the variable of the specific call to outer() that created it. Every call to outer() makes a fresh list, so each returned function keeps its own list. append changes that list in place, so the earlier values stay in it.

WATCH OUT FOR

Assuming f1 and f2 share one list because they come from the same outer definition, which leads to choosing (A) and (C). The list is created inside outer , so it is new on every call and is not shared.

Q.51 MSQ · 2 MARKS

Database Management and Warehousing › SQL

Q.51

Consider a table `Employee(EmpID, TeamID)`, where the column `EmpID` (ID of an employee) is the primary key. The column `TeamID` denotes the team ID of the team of which the employee is a member. `TeamID` is a NOT NULL column.

We want to display the size of the team (denoted as `TeamSize`) in which each employee is a member by using SQL. As an example, the desired output for the given `Employee` table is also shown in tabular form.

Which of the following is/are correct?

Employee

EmpID	TeamID
1	8
2	8
3	8
4	7
5	7
6	9

Output

EmpID	TeamSize
1	3
2	3
3	3
4	2
5	2
6	1

(A)

```
SELECT E.EmpID, B.TeamSize
FROM Employee AS E, (SELECT TeamID, COUNT(TeamID) AS
TeamSize FROM Employee GROUP BY TeamID) AS B
WHERE E.TeamID = B.TeamID
```

(B)

```
SELECT A.EmpID, COUNT(B.TeamID) AS TeamSize
FROM Employee AS A, Employee AS B
WHERE A.TeamID = B.TeamID AND A.EmpID = B.EmpID
GROUP BY A.EmpID
```

(C)

```
SELECT B.EmpID, B.TeamSize
FROM (SELECT EmpID, COUNT(TeamID) AS TeamSize
FROM Employee GROUP BY EmpID) AS B
```

(D)

```
SELECT A.EmpID, B.TeamSize
FROM Employee AS A, (SELECT COUNT(TeamID) AS TeamSize FROM
Employee GROUP BY TeamID) AS B
WHERE A.TeamID = B.TeamID
```

Official answer: **A**

SOLUTION

1. Target output: one row per employee, with the size of that employee's own team. From the Employee table, team 8 has 3 members, team 7 has 2 and team 9 has 1, so the required rows are $(1,3),(2,3),(3,3),(4,2),(5,2),(6,1)$.
2. Option (A): the derived table B is built with GROUP BY TeamID and holds $(\text{TeamID}, \text{TeamSize}) = (8,3),(7,2),(9,1)$. The condition E.TeamID = B.TeamID matches each employee to the row for their

own team. Row by row, employee 1, 2, 3 (team 8) get 3, employees 4, 5 (team 7) get 2, and employee 6 (team 9) gets 1. This is exactly the required output. Correct.

3. Option (B): the condition `A.EmpID = B.EmpID` pairs each row with itself, because `EmpID` is the primary key and so unique. The extra condition `A.TeamID = B.TeamID` is then always true. Each `GROUP BY A.EmpID` group therefore holds exactly one row, so `COUNT(B.TeamID)` is 1 for every employee. The result is $(1,1),(2,1),(3,1),(4,1),(5,1),(6,1)$, which differs from the required output. Wrong.
4. Option (C): `GROUP BY EmpID` makes one group per employee, since `EmpID` is unique. `COUNT(TeamID)` is 1 in each group, giving $(1,1),\ldots,(6,1)$. It never counts the other members of the team. Wrong.
5. Option (D): the derived table `B` selects only `COUNT(TeamID) AS TeamSize` with `GROUP BY TeamID`, so its only column is `TeamSize`. It has no `TeamID` column, so `B.TeamID` in the `WHERE` clause refers to a column that does not exist and the query is rejected. Even if the join column were available, a bare count would not carry the team it belongs to. Wrong.
6. Only (A) produces the required output.

KEY IDEA

To attach a group size to every row, aggregate with GROUP BY on the grouping column (`TeamID`) in a derived table, keep that column in the derived table's SELECT list, and join it back to the base table on that column.

WATCH OUT FOR

Grouping by the primary key (`EmpID`), as in (B) and (C), gives one row per group and a count of 1. Dropping `TeamID` from the derived table, as in (D), leaves nothing to join on, so the query is invalid.

Q.52 MSQ · 2 MARKS

Linear Algebra › Special matrices

Let $M = \left(I_n - \frac{1}{n}\mathbf{1}\mathbf{1}^T\right)$ be a matrix, where $\mathbf{1} = (1,1,1,\ldots,1)^T \in \mathbb{R}^n$ and I_n is the identity matrix of order n. Which of the following options is/are correct?

(A) $M^T = M$

(B) $M^2 = I_n$

(C) $\text{Trace}(M) = n$

(D) M is a projection matrix

Official answer: **A, D**

SOLUTION

1. Since $\mathbf{1}^T\mathbf{1} = n$, the matrix $P = \frac{1}{n}\mathbf{1}\mathbf{1}^T$ projects onto the line spanned by $\mathbf{1}$, and $M = I_n - P$.
2. (A) $P^T = \frac{1}{n}(\mathbf{1}\mathbf{1}^T)^T = \frac{1}{n}\mathbf{1}\mathbf{1}^T = P$, so $M^T = I_n - P^T = I_n - P = M$. True.
3. (B) $P^2 = \frac{1}{n^2}\mathbf{1}(\mathbf{1}^T\mathbf{1})\mathbf{1}^T = \frac{n}{n^2}\mathbf{1}\mathbf{1}^T = P$, so $M^2 = I_n - 2P + P^2 = I_n - P = M$. Since $P \neq 0$, $M \neq I_n$, so $M^2 = M \neq I_n$. False.
4. (C) $\text{tr}(P) = \frac{1}{n}\sum_{i=1}^{n} 1 = 1$, so $\text{tr}(M) = n - 1 \neq n$. False.
5. (D) M is symmetric ($M^T = M$) and idempotent ($M^2 = M$), which is exactly an orthogonal projection matrix. It projects onto the subspace orthogonal to $\mathbf{1}$, and its rank $n-1$ equals its trace. True.
6. Correct options: A and D.

KEY IDEA

$\frac{1}{n}\mathbf{1}\mathbf{1}^T$ is the orthogonal projection onto $\mathbf{1}$, so I_n minus it is the orthogonal projection onto the complement: symmetric, idempotent, with trace equal to its rank $n-1$.

WATCH OUT FOR

Forgetting that a projection matrix satisfies $M^2 = M$ rather than $M^2 = I_n$, or taking $\text{tr}(M) = n$ by ignoring the subtracted rank-one part, whose trace is 1.

Q.53 MSQ · 2 MARKS

Probability and Statistics › Continuous distributions

Let $X_1, X_2, \ldots, X_n$ be n independent random variables. Each of the random variables follows $\text{Normal}(\mu = 0, \sigma^2 = 1)$ distribution. Define $\bar{X} = \frac{1}{n}\sum_{i=1}^{n} X_i$.

Which of the following statements is/are correct?

(A) $\sum_{i=1}^{n} X_i^2$ follows Chi-square distribution with n degrees of freedom.

(B) $\sum_{i=1}^{n}(X_i - \bar{X})^2$ follows Chi-square distribution with $(n-1)$ degrees of freedom.

(C) $X_1^2 + X_n^2$ follows exponential distribution with mean 2.

(D) $(\sqrt{n}\,\bar{X})^2$ follows Chi-square distribution with 2 degrees of freedom.

Official answer: **A, B, C**

SOLUTION

1. Each $X_i \sim N(0,1)$ independently, so $X_i^2 \sim \chi_1^2$. A sum of independent chi-square variables is chi-square with the degrees of freedom added.
2. (A) $\sum_{i=1}^{n} X_i^2$ is a sum of n independent χ_1^2 variables, so it follows χ_n^2. This is the definition of the chi-square distribution with n degrees of freedom. (A) is correct.
3. (B) Because $\sigma^2 = 1$, the sample-variance theorem gives $\sum_{i=1}^{n}(X_i - \bar{X})^2 = (n-1)S^2 \sim \chi_{n-1}^2$. One degree of freedom is lost because the deviations $X_i - \bar{X}$ always sum to zero. (B) is correct.
4. (C) For $n \geq 2$, X_1 and X_n are different, independent variables, so $X_1^2 + X_n^2 \sim \chi_2^2$. Its density is $f(x) = \frac{1}{2}e^{-x/2}$ for $x > 0$. This is the exponential distribution with rate $\frac{1}{2}$, whose mean is 2. (C) is correct.
5. (D) $\bar{X} \sim N(0, 1/n)$, so $\sqrt{n}\,\bar{X} \sim N(0,1)$. Hence $(\sqrt{n}\,\bar{X})^2 \sim \chi_1^2$, with 1 degree of freedom, not 2. (D) is incorrect.
6. Correct options: A, B and C.

KEY IDEA

A sum of squares of independent standard normals is chi-square with one degree of freedom per independent term. Replacing the true mean with $\bar{X}$ costs one degree of freedom. χ_2^2 is exponential with mean 2.

WATCH OUT FOR

Judging (D) from the factor $\sqrt{n}$ or from the sum, and giving it 2 or n degrees of freedom. $\sqrt{n}\,\bar{X}$ is a single standard normal, so its square is χ_1^2. Another slip is to treat (B) as χ_n^2 and forget the lost degree of freedom.

Q.54 MSQ · 2 MARKS

Probability and Statistics › Random variables

Let X be a discrete valued random variable with cumulative distribution function $F(x)$.

Which of the following statements is/are correct?

(A) $F(x)$ is always a positive function.

(B) $F(x)$ is a non-decreasing function.

(C) $F(x)$ has jump discontinuity.

(D) $F(x)$ is a left continuous function.

Official answer: **B, C**

SOLUTION

1. Recall that the CDF is $F(x) = P(X \leq x)$. It is non-decreasing, right-continuous, with $F(-\infty) = 0$ and $F(+\infty) = 1$.
2. (A) For any x smaller than every value that X can take, $\{X \leq x\}$ is empty, so $F(x) = 0$. Zero is not a positive number, so F is not always positive. (A) is false.

3. (B) If $a < b$, then $\{X \le a\} \subseteq \{X \le b\}$, so $F(a) = P(X \le a) \le P(X \le b) = F(b)$. Hence F is non-decreasing. (B) is true.
4. (C) A discrete random variable takes some value x_0 with $P(X = x_0) > 0$. At that point $F(x_0) - F(x_0^-) = P(X \le x_0) - P(X < x_0) = P(X = x_0) > 0$, so F has a jump discontinuity of size $P(X = x_0)$. (C) is true.
5. (D) Because $F(x_0) = P(X \le x_0)$ includes the point mass at x_0, we get $F(x_0) = F(x_0^+)$, so F is right continuous. At a jump point, $F(x_0^-) = P(X < x_0) = F(x_0) - P(X = x_0) \ne F(x_0)$, so F is not left continuous. (D) is false.
6. The correct statements are (B) and (C).

KEY IDEA

Any CDF $F(x) = P(X \le x)$ is non-decreasing and right continuous, and its jump at a point equals the point mass there. A discrete variable has such jumps, so F has jump discontinuities and is not left continuous.

WATCH OUT FOR

Mixing up right and left continuity. Because F uses $\le$, it is right continuous, not left continuous. Another slip is calling F "positive" when it is only non-negative, since it equals 0 to the left of the smallest value of X.

Q.55 NAT · 2 MARKS

Machine Learning › Regression

Consider that Linear Ridge Regression is being used to learn a prediction function $y_{pred} = w^T x$, where $w, x \in \mathbb{R}^2$ and Mean Absolute Error (MAE) is used to measure the prediction error. A weight of 0.20 is associated with the regularizer.

At an intermediate step of the training process, assume that the parameter $w = [-3.00, 4.00]^T$. In the next step, for the input $x = [1.00, 2.00]^T$, the predicted value of y is noted. Let the relation between $x = [x_1, x_2]^T$ and the true value of y be $y_{true} = x_1 + x_2$.

The value of the overall regularized loss function for this instance is _____ . (Rounded off to two decimal places)

Official answer: **6.90 to 7.10**

SOLUTION

1. Ridge regression adds an L_2 penalty on the weights to the data-fit term. Here the data-fit term is the absolute error (MAE for one instance), so
$$L = |y_{pred} - y_{true}| + \lambda \|w\|_2^2, \qquad \lambda = 0.20.$$
2. Prediction: $y_{pred} = w^T x = (-3)(1) + (4)(2) = -3 + 8 = 5$.
3. True value: $y_{true} = x_1 + x_2 = 1 + 2 = 3$.
4. Error for this single instance: $|y_{pred} - y_{true}| = |5 - 3| = 2$.
5. Regulariser: $\|w\|_2^2 = (-3)^2 + 4^2 = 9 + 16 = 25$, so the penalty is $0.20 \times 25 = 5$.
6. Total:
$$L = 2 + 0.20 \times 25 = 2 + 5 = 7.00$$

KEY IDEA

Regularized loss = data-fit term (here absolute error, because MAE is specified) + $\lambda\|w\|_2^2$. Compute the prediction first, then the error, then the penalty on the given weight vector.

WATCH OUT FOR

Using squared error instead of absolute error (giving $4 + 5 = 9$), or penalising with $\lambda\|w\|_2$ (not squared) or $\frac{\lambda}{2}\|w\|^2$, or forgetting the penalty term altogether (giving 2).

Q.56 NAT · 2 MARKS

Machine Learning › Neural networks

Consider a fully-connected feed-forward multi-layer perceptron. It has 30 neurons in the input layer, followed by two hidden layers and an output layer. The first hidden layer has 4 neurons and the second hidden layer has 3 neurons. The output layer has only one neuron. Assume that no bias parameters are used.

The number of learnable parameters in the multi-layer perceptron is _____. (Answer in integer)

Official answer: **135**

SOLUTION

1. With no bias parameters, the learnable parameters are only the weights on the connections between consecutive layers.
2. In a fully-connected layer pair, every neuron of one layer connects to every neuron of the next, so a pair with m and n neurons has $m \times n$ weights.
3. Input to hidden layer 1: $30 \times 4 = 120$.
4. Hidden layer 1 to hidden layer 2: $4 \times 3 = 12$.
5. Hidden layer 2 to output: $3 \times 1 = 3$.
6. Total:

$$120 + 12 + 3 = 135.$$

KEY IDEA
With no biases, count one weight per connection between adjacent layers: the sum of the products of consecutive layer sizes, $30 \cdot 4 + 4 \cdot 3 + 3 \cdot 1$.

WATCH OUT FOR
Adding a bias for each neuron out of habit (which would give $135 + 4 + 3 + 1 = 143$), or adding the layer sizes instead of multiplying consecutive pairs.

Q.57 NAT · 2 MARKS

Probability and Statistics › Conditional & Bayes

A clinic specializes in testing for a disease D. The result of the test can be either positive or negative.

A study revealed that if a person suffers from the disease D, the test result in that clinic comes out positive 80% of the time, and negative 20% of the time. If a person is not suffering from the disease D, the test comes out positive 10% of the time and negative 90% of the time. It is also known that among the general population, the disease D occurs in 30% of the individuals.

If a person tests positive for D in that clinic, the probability that he/she actually suffers from the disease D is _____. (Rounded off to two decimal places)

Official answer: **0.76 to 0.78**

SOLUTION

1. Let D be the event that the person has the disease and $+$ the event that the test is positive. The population rate gives $P(D) = 0.30$, so $P(\bar{D}) = 0.70$.
2. The test results given are $P(+ \mid D) = 0.80$ (and $P(- \mid D) = 0.20$), and $P(+ \mid \bar{D}) = 0.10$ (and $P(- \mid \bar{D}) = 0.90$).
3. By the law of total probability,
$$P(+) = P(+ \mid D)P(D) + P(+ \mid \bar{D})P(\bar{D}) = 0.8 \times 0.3 + 0.1 \times 0.7 = 0.24 + 0.07 = 0.31.$$
4. By Bayes' theorem,
$$P(D \mid +) = \frac{P(+ \mid D)P(D)}{P(+)} = \frac{0.24}{0.31} = \frac{24}{31} \approx 0.7742.$$
5. Rounded to two decimal places, $P(D \mid +) = 0.77$.

KEY IDEA

The question asks for $P(D \mid +)$ while the data give $P(+ \mid D)$, so reverse the conditioning with Bayes' theorem. Find the denominator $P(+)$ first, using total probability over D and $\bar{D}$.

WATCH OUT FOR

Reporting $P(+ \mid D) = 0.80$ as the answer, which confuses the two conditional probabilities. Another slip is dividing 0.24 by $0.8 + 0.1$ and ignoring the prior weights 0.3 and 0.7 in the denominator.

Q.58 NAT · 2 MARKS

Data Structures and Algorithms › Python

Consider the given Python program.

```
def fun(L, i=0):
    if i >= len(L)-1:
        return 0
    if L[i] > L[i+1]:
        L[i+1], L[i] = L[i], L[i+1]
        return 1+fun(L, i+1)
    else:
        return fun(L, i+1)
data = [5, 3, 4, 1, 2]
count = 0
for _ in range(len(data)):
    count += fun(data)
print(count)
```

The output of the program is _____ . (Answer in integer)

Official answer: **8**

SOLUTION

1. Each call `fun(data)` starts with `i=0` and walks `i` up to `len(L)-2 = 3`. At each `i`, if `L[i] > L[i+1]` it swaps the pair and adds 1 to the result. The swap line `L[i+1], L[i] = L[i], L[i+1]` evaluates the right side first, so it is a true swap. One call is therefore one bubble-sort pass, and it returns the number of swaps made. `data` is the same list object every time, so each call sees the result of the previous one.
2. The loop runs `len(data) = 5` times, so `count` is the sum of the swap counts of 5 passes.
3. Pass 1 on `[5, 3, 4, 1, 2]`: 5 > 3 swaps to `[3, 5, 4, 1, 2]`; 5 > 4 swaps to `[3, 4, 5, 1, 2]`; 5 > 1 swaps to `[3, 4, 1, 5, 2]`; 5 > 2 swaps to `[3, 4, 1, 2, 5]`; at `i = 4` the base case returns 0. The pass returns 4.
4. Pass 2 on `[3, 4, 1, 2, 5]`: 3 < 4 no swap; 4 > 1 swaps to `[3, 1, 4, 2, 5]`; 4 > 2 swaps to `[3, 1, 2, 4, 5]`; 4 < 5 no swap. The pass returns 2.
5. Pass 3 on `[3, 1, 2, 4, 5]`: 3 > 1 swaps to `[1, 3, 2, 4, 5]`; 3 > 2 swaps to `[1, 2, 3, 4, 5]`; the remaining pairs are in order. The pass returns 2.
6. Passes 4 and 5 act on the sorted list `[1, 2, 3, 4, 5]`, so no swap happens and each returns 0.
7. Total: $4 + 2 + 2 + 0 + 0 = 8$, so the program prints 8.
8. Check: every bubble-sort swap removes exactly one inversion, and 5 passes are enough to sort 5 elements, so the total number of swaps equals the number of inversions of the original list. In $[5, 3, 4, 1, 2]$, 5 is above 4 later elements, 3 is above 2 (1 and 2), 4 is above 2 (1 and 2), and 1 is above none. That gives $4 + 2 + 2 + 0 = 8$.

KEY IDEA

`fun` is one bubble-sort pass that returns its swap count, and the loop makes 5 passes on the same mutated list. The total is the number of inversions of the original list: $4 + 2 + 2 = 8$.

WATCH OUT FOR

Treating each call as acting on a fresh copy of `[5, 3, 4, 1, 2]`, which would give $5 \times 4 = 20$. The list is mutated in place and the loop passes the same object every time. Another slip is stopping after the first pass (4) and forgetting that later passes still swap.

Q.59 NAT · 2 MARKS

Database Management and Warehousing › Relational algebra

Q.59 Let there be two relations X and Y as shown. X has three columns P, Q and R. Y has two columns P and S.

X

P	Q	R
P1	Q1	R1
P2	Q2	R2
P3	Q3	R2

Y

P	S
P1	10
P1	15
P2	20
P3	1

Consider that the following tuple relational calculus expression is evaluated.

$$\{t | t \in X \wedge \exists z \in X(t[P] = z[P]) \wedge \exists m \in Y\ (m[P] = t[P] \wedge m[S] > 1)\}$$

The number of tuples that will be returned is _____. (Answer in **integer**)

Official answer: **2 to 3**

SOLUTION

1. The expression is $\{t \mid t \in X \wedge \exists z \in X\ (t[P] = z[P]) \wedge \exists m \in Y\ (m[P] = t[P] \wedge m[S] > 1)\}$. Every returned tuple t comes from X, and it must satisfy both existential conditions.
2. The condition $\exists z \in X\ (t[P] = z[P])$ is always true. Take $z = t$, which is in X because $t \in X$. So it removes no tuple.
3. The real filter is $\exists m \in Y\ (m[P] = t[P] \wedge m[S] > 1)$. Y must contain a row with the same P value as t and with $S > 1$.
4. For $t = (P1, Q1, R1)$: Y has $(P1, 10)$ and $(P1, 15)$. Both have $S > 1$, so t is returned.
5. For $t = (P2, Q2, R2)$: Y has $(P2, 20)$ and $20 > 1$, so t is returned.
6. For $t = (P3, Q3, R2)$: the only row of Y with $P = P3$ is $(P3, 1)$, and $1 > 1$ is false. So t is not returned.
7. Tuple relational calculus returns a set of tuples t taken from X. The two matching rows of Y for $P1$ do not duplicate $(P1, Q1, R1)$, because it is listed once as a single tuple t (unlike a join).
8. Exactly 2 tuples are returned: $(P1, Q1, R1)$ and $(P2, Q2, R2)$.

KEY IDEA

The first existential is trivially true because t can serve as z. The second is a semi-join-like test on X against Y with the filter $S > 1$. The result is a set of tuples from X, so each qualifying tuple of X appears once.

WATCH OUT FOR

Counting 3 tuples by treating $(P3, 1)$ as a match even though $1 > 1$ is false. Counting 3 by treating the expression as a join, so that $P1$ appears twice. A join would repeat P1, but this expression returns tuples t of X only.

Q.60 NAT · 2 MARKS

Database Management and Warehousing › SQL

Q.60 Let Account be a relation as shown.

Account

AccNo	Balance
A1	5000
A2	5000
A3	10000
A4	15000
A5	18000

Consider the given SQL query.

```
SELECT AccNo FROM Account AS A
WHERE (SELECT COUNT(*) FROM Account AS B
WHERE A.Balance < B.Balance) >= (SELECT COUNT(*)
FROM Account AS C WHERE A.Balance > C.Balance)
```

The number of rows returned by the SQL query is _____ . (Answer in **integer**)

Official answer: **3**

SOLUTION

1. The query is correlated. For each row of A , the first subquery counts the rows B with B.Balance strictly greater than A.Balance . The second subquery counts the rows C with C.Balance strictly less than A.Balance . The row is returned when the first count is at least the second.
2. The balances are $5000, 5000, 10000, 15000, 18000$. Write g for the count strictly greater and l for the count strictly less than the outer row's balance.
3. A1 (5000): $g = 3$ (10000, 15000, 18000) and $l = 0$. Since $3 \geq 0$, A1 is returned.
4. A2 (5000): the balance equals A1's, so $g = 3$ and $l = 0$. A2 is returned. The tie does not matter, because both comparisons are strict and the other 5000 row counts on neither side.
5. A3 (10000): $g = 2$ (15000, 18000) and $l = 2$ (5000, 5000). Since $2 \geq 2$, A3 is returned.
6. A4 (15000): $g = 1$ (18000) and $l = 3$ (5000, 5000, 10000). Since $1 \geq 3$ is false, A4 is not returned.
7. A5 (18000): $g = 0$ and $l = 4$. Since $0 \geq 4$ is false, A5 is not returned.
8. The rows returned are A1, A2 and A3, so the query returns 3 rows.

KEY IDEA
For each outer row, a correlated subquery counts the rows with a strictly larger balance and the rows with a strictly smaller balance. The row qualifies when the larger-count is at least the smaller-count. Rows with equal balance fall in neither count.

WATCH OUT FOR
Counting tied balances (the two 5000 rows) as greater or smaller. Both comparisons are strict, so a tied row counts on neither side. Another slip is reversing the two counts, which would keep A3, A4 and A5 and still give 3 rows by coincidence of a different set. Evaluate each row explicitly.

Q.61 NAT · 2 MARKS

Database Management and Warehousing › ER model

Consider an ER model with the entities E1(A11, A12, A13) and E2(A21, A22, A23), where A11, A12, A13 are the attributes of E1, and A21, A22, A23 are the attributes of E2. Let A22 be a multi-valued attribute. A11 and A21 are the primary keys of E1 and E2, respectively.

Let R12 be a many-to-many relationship between E1 and E2. Participation of both E1 and E2 in R12 is total.

The minimum number of relations required to convert the ER model into relational model (assuming there is no other functional dependency) where each relation is in third normal form (3NF) is _____. (Answer in integer)

Official answer: **4**

SOLUTION

1. Entity $E1$ has the key A_{11} and only single-valued attributes, so it becomes one relation $E1(A_{11}, A_{12}, A_{13})$.
2. Entity $E2$ has the key A_{21} and one multi-valued attribute A_{22}. A multi-valued attribute cannot be stored as a column of $E2$, because a cell would hold a set and break first normal form. So $E2$ becomes $E2(A_{21}, A_{23})$.
3. The multi-valued attribute A_{22} gets its own relation $E2_A22(A_{21}, A_{22})$. Its key is the pair (A_{21}, A_{22}), and A_{21} is a foreign key to $E2$.
4. The many-to-many relationship $R12$ needs its own relation $R12(A_{11}, A_{21})$. Its key is the pair (A_{11}, A_{21}), and both parts are foreign keys. A many-to-many relationship cannot be folded into $E1$ or $E2$, because that would again store a set of values in one cell.
5. Total participation of $E1$ and $E2$ in $R12$ only makes a foreign key non-null when the relationship is merged into an entity relation, as in the 1:1 or 1:N case. It does not remove the separate relation for a many-to-many relationship, so it changes nothing here.
6. The question says there is no other functional dependency. So in each relation the only determinant is its key, and every relation is in 3NF (indeed BCNF).
7. Count: $E1$, $E2$, $E2_A22$ and $R12$ give 4 relations. None can be dropped without losing the multi-valued attribute, the relationship, or 3NF.

KEY IDEA
In ER-to-relational mapping, each entity gets a relation, each multi-valued attribute gets its own relation, and each many-to-many relationship gets its own relation. Total participation only affects foreign-key nullability when a relationship is merged, so it does not change the count here.

WATCH OUT FOR
Counting only 3 relations by folding the multi-valued attribute A_{22} into $E2$, or by merging $R12$ into an entity relation because participation is total. Both break first normal form. Another slip is to count only the two entities plus the relationship and miss the separate relation for A_{22}.

Q.62 NAT · 2 MARKS

Probability and Statistics › Descriptive statistics

For a given data set $\{x_1, x_2, \ldots, x_n\}$, where $n = 100$, it is known that

$$\frac{1}{2000}\sum_{i=1}^{n}\sum_{j=1}^{n}(x_i - x_j)^2 = 99.$$

Let us denote $\bar{x} = \frac{1}{n}\sum_{i=1}^{n} x_i$. The value of $\frac{1}{99}\sum_{i=1}^{n}(x_i - \bar{x})^2$ is _____. (Answer in integer)

Official answer: **10**

SOLUTION

1. Let $S_{xx} = \sum_{i=1}^{n}(x_i - \bar{x})^2$. The double sum is a pairwise squared-difference sum, so rewrite it through the mean.
2. Insert $\bar{x}$: $(x_i - x_j)^2 = \big((x_i - \bar{x}) - (x_j - \bar{x})\big)^2 = (x_i - \bar{x})^2 - 2(x_i - \bar{x})(x_j - \bar{x}) + (x_j - \bar{x})^2$.
3. Sum over i and j. The first term gives $n\,S_{xx}$, because it does not depend on j and the sum over j contributes a factor n. The last term gives $n\,S_{xx}$ in the same way. The cross term gives $-2\left(\sum_i (x_i - \bar{x})\right)\left(\sum_j (x_j - \bar{x})\right) = 0$, since deviations from the mean sum to zero.
4. Hence
$$\sum_{i=1}^{n}\sum_{j=1}^{n}(x_i - x_j)^2 = 2n\,S_{xx}.$$
5. The given condition says this sum equals $2000 \times 99 = 198000$. With $n = 100$: $200\,S_{xx} = 198000$, so $S_{xx} = 990$.
6. Required value: $\frac{1}{99}S_{xx} = \frac{990}{99} = 10$.

KEY IDEA

The double sum of squared pairwise differences equals $2n$ times the sum of squared deviations from the mean, $\sum_i\sum_j (x_i - x_j)^2 = 2n\sum_i (x_i - \bar{x})^2$. This turns the given condition into the sample variance (divisor $n - 1 = 99$).

WATCH OUT FOR

Forgetting the factor 2 (writing $n\,S_{xx}$ instead of $2n\,S_{xx}$), which gives 20. Another slip is to divide the double sum by n^2 and treat it as a variance with divisor n instead of $n - 1$.

Q.63 NAT · 2 MARKS

Probability and Statistics › Descriptive statistics

Let X be a random variable that follows $\text{Uniform}(-1, 1)$ distribution. The conditional distribution of the random variable Y given $X = x$ is the $\text{Uniform}(x^2 - 0.1,\ x^2 + 0.1)$ distribution.

The value of $\text{correlation}(X, Y)$ is _____. (Answer in integer)

Official answer: **0**

SOLUTION

1. Given $X = x$, Y is uniform on $(x^2 - 0.1,\ x^2 + 0.1)$, an interval centred at x^2. So $E[Y \mid X = x] = x^2$, that is $E[Y \mid X] = X^2$.
2. The correlation is $\rho = \dfrac{\text{Cov}(X, Y)}{\sigma_X \sigma_Y}$. Here $\sigma_X = \sqrt{1/3} > 0$. Also $\sigma_Y > 0$, because Y is not constant. So $\rho = 0$ exactly when $\text{Cov}(X, Y) = 0$.
3. $X \sim \text{Uniform}(-1, 1)$ has density $\frac{1}{2}$ on $(-1, 1)$, which is symmetric about 0. Hence $E[X] = \int_{-1}^{1} \frac{1}{2}x\,dx = 0$ and $E[X^3] = \int_{-1}^{1} \frac{1}{2}x^3\,dx = 0$.
4. By the tower property, $E[XY] = E\big[X\,E[Y \mid X]\big] = E[X \cdot X^2] = E[X^3] = 0$.
5. Then $\text{Cov}(X, Y) = E[XY] - E[X]E[Y] = 0 - 0 \cdot E[Y] = 0$.
6. Therefore $\rho = 0/(\sigma_X \sigma_Y) = 0$. This holds even though Y depends strongly on X through X^2: the dependence is not linear, and correlation measures only linear association.

KEY IDEA
Find $E[Y \mid X] = X^2$, then use the tower property to get $E[XY] = E[X^3] = 0$. X is symmetric about 0, so the covariance and hence the correlation are zero, although Y depends on X.

WATCH OUT FOR
Assuming that a strong dependence of Y on X (here through X^2) means a non-zero correlation. Correlation detects only linear association. Zero correlation does not imply independence.

Q.64 NAT · 2 MARKS

Probability and Statistics › Discrete distributions

Let $A_{5\times5}$ be a matrix such that each of its elements follows $Bernoulli(p = 0.50)$ distribution independently.

The probability that the row-sum of the second row and the column-sum of the third column are both equal to 3 is _____ . (Rounded off to **two** decimal places)

Official answer: **0.08 to 0.12**

SOLUTION

1. The second row and the third column of the 5×5 matrix share exactly one entry, a_{23}. The other 4 entries of the row and the other 4 entries of the column are 8 distinct entries, and all of them are independent $\text{Bernoulli}(\frac{1}{2})$.
2. Condition on $a_{23} = x$. Given x, the row needs its other 4 entries to sum to $3 - x$, and the column needs its other 4 entries to sum to $3 - x$. These two events involve disjoint sets of entries, so they are independent given x.
3. Let $S \sim \text{Binomial}(4, \frac{1}{2})$ be the sum of 4 such entries. Then $P(S = k) = \binom{4}{k}/16$.
4. Case $x = 1$, which has probability $\frac{1}{2}$: both the row and the column need $S = 2$. $P(S = 2) = \dfrac{6}{16}$, so the contribution is $\dfrac{1}{2}\left(\dfrac{6}{16}\right)^2 = \dfrac{18}{256}$.
5. Case $x = 0$, which has probability $\frac{1}{2}$: both the row and the column need $S = 3$. $P(S = 3) = \dfrac{4}{16}$, so the contribution is $\dfrac{1}{2}\left(\dfrac{4}{16}\right)^2 = \dfrac{8}{256}$.
6. By total probability, $P = \dfrac{18 + 8}{256} = \dfrac{26}{256} \approx 0.1016$.
7. Rounded to two decimal places, the answer is 0.10. (Treating the two events as independent would give $\left(\frac{10}{32}\right)^2 \approx 0.098 \approx 0.10$ here by coincidence, but that method is not valid because the events share a_{23}.)

KEY IDEA
The row and the column overlap in one entry, a_{23}. Condition on that entry. Given it, the other 8 entries are independent, so the row event and the column event are independent and each is a Binomial$(4, \frac{1}{2})$ probability.

WATCH OUT FOR
Multiplying $P(\text{row sum} = 3) \cdot P(\text{column sum} = 3) = (10/32)^2$ as if the events were independent. They are not, because they share a_{23}. The result happens to round to the same 0.10 here, but the method is wrong.

Q.65 NAT · 2 MARKS

Linear Algebra › Special matrices

Let $A = \left(I_n - \frac{1}{n}\mathbf{1}\mathbf{1}^T\right)$ be a matrix, where $\mathbf{1} = (1, 1, 1, \ldots, 1)^T \in \mathbb{R}^n$ and I_n is the identity matrix of order n.

The value of $\max_S x^T Ax$, where $S = \{x \in \mathbb{R}^n \mid x^T x = 1\}$, is _____. (Answer in integer)

Official answer: **1**

SOLUTION

1. Write $P = \frac{1}{n}\mathbf{1}\mathbf{1}^T$, so $A = I_n - P$. Since $\mathbf{1}^T\mathbf{1} = n$, we get $P^2 = \frac{1}{n^2}\mathbf{1}(\mathbf{1}^T\mathbf{1})\mathbf{1}^T = P$, and $P^T = P$. So P is an orthogonal projection of rank 1.
2. Then $A^T = A$ and $A^2 = I - 2P + P^2 = I - P = A$. So A is symmetric and idempotent, i.e. the orthogonal projection onto the subspace perpendicular to $\mathbf{1}$.
3. An idempotent matrix has eigenvalues only 0 or 1. Here $A\mathbf{1} = \mathbf{1} - \frac{1}{n}\mathbf{1}(\mathbf{1}^T\mathbf{1}) = \mathbf{1} - \mathbf{1} = 0$, so 0 is an eigenvalue. The other $n-1$ eigenvalues are 1, because $\mathrm{rank}(A) = n - 1$ (for $n \geq 2$).
4. For a real symmetric matrix, $\max_{x^T x = 1} x^T Ax = \lambda_{\max}(A)$ (Rayleigh quotient), so the maximum is 1.
5. Direct check: for any unit x, $x^T Ax = x^T x - \frac{1}{n}(\mathbf{1}^T x)^2 = 1 - \frac{(\mathbf{1}^T x)^2}{n} \leq 1$, with equality exactly when $\mathbf{1}^T x = 0$. Such unit vectors exist for $n \geq 2$, for example $x = \frac{1}{\sqrt{2}}(1, -1, 0, \ldots, 0)^T$.
6. Hence $\max_S x^T Ax = 1$.

KEY IDEA

$I_n - \frac{1}{n}\mathbf{1}\mathbf{1}^T$ is the centering matrix, a symmetric projection with eigenvalues 0 (once) and 1 ($n-1$ times). The Rayleigh quotient maximum is the largest eigenvalue, 1.

WATCH OUT FOR

Taking the single eigenvalue 0 (from the direction $\mathbf{1}$) as the answer, or forgetting that the maximum is the largest eigenvalue rather than the smallest. Another slip is to think the value depends on n.

FROM IAE ACADEMY

Preparing for GATE DA?

IAE Academy teaches GATE DA in Ballygunge, Kolkata, and online. Counselling is free: ask us about batches, fees or your preparation.

FREE COUNSELLING

Chat on WhatsApp

Call +91 84204 36524

Open every day, 10:30 am to 7:30 pm.

Scan to chat on WhatsApp

THE GATE DA COURSE

Classroom batches in Ballygunge and online batches. Fees and details are on the course page.

Scan: course page

See the course

FREE ON IAEACADEMY.IN

More GATE DA papers, question by question with the official answers, and free notes by subject.

Scan: more papers

More solved papers

Found this useful? Send it to a friend preparing for GATE.

Send to a friend on WhatsApp

IAE Academy free study material. The questions are those of the GATE DA 2026 question paper, organised by IIT Guwahati; the answers are from its official answer key. Prepared and published by IAE Academy, Kolkata. Free to read and free to share, unaltered and with attribution. Not affiliated with, endorsed by or produced by IIT Guwahati or the GATE committee.

IAE ACADEMY INSTITUTE FOR ENGINEERS

GATE coaching in Kolkata since 2012 · EC, EE, IN, CS and IT, DA, BT and CE · in the classroom and online

227A Rash Behari Ave, Ballygunge, Kolkata 700019

+91 84204 36524 · iaeacademy.in · Updated 27 October 2026

Get directions