IAE ACADEMY INSTITUTE FOR ENGINEERS

GATE DA · 2024

IAE Academy free study material · Updated 27 October 2026

GATE DA 2024 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: IISc Bengaluru

Online, with the questions linked to their topic notes: **iaeacademy.in/study/papers/da-2024**

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: 37 MCQ · 12 MSQ · 16 NAT.

Official answer key: IISc Bengaluru · https://gate2024.iisc.ac.in/wp-content/uploads/2024/DAFinalAnswerKey.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-2024

Answer key

Q	TYPE	MARKS	ANSWER
Q.1	MCQ	1	**D**
Q.2	MCQ	1	**A**
Q.3	MCQ	1	**B**
Q.4	MCQ	1	**B**
Q.5	MCQ	1	**B**
Q.6	MCQ	2	**D**
Q.7	MCQ	2	**A**
Q.8	MCQ	2	**B**
Q.9	MCQ	2	**A**
Q.10	MCQ	2	**A**
Q.11	MCQ	1	**A**
Q.12	MCQ	1	**A**
Q.13	MCQ	1	**B**
Q.14	MCQ	1	**A**
Q.15	MCQ	1	**A**
Q.16	MCQ	1	**A**
Q.17	MCQ	1	**D**
Q.18	MCQ	1	**C**
Q.19	MCQ	1	**D**
Q.20	MCQ	1	**B**
Q.21	MCQ	1	**B**
Q.22	MCQ	1	**A**
Q.23	MCQ	1	**D**
Q.24	MCQ	1	**C**
Q.25	MCQ	1	**C**
Q.26	MCQ	1	**C**
Q.27	MCQ	1	**B**
Q.28	MSQ	1	**A, B, C**
Q.29	MSQ	1	**B, C, D**
Q.30	NAT	1	**0**
Q.31	NAT	1	**3**
Q.32	NAT	1	**17**
Q.33	NAT	1	**0.24**
Q.34	NAT	1	**42**
Q.35	NAT	1	**0**
Q.36	MCQ	2	**C**
Q.37	MCQ	2	**A**
Q.38	MCQ	2	**D**
Q.39	MCQ	2	**C**
Q.40	MCQ	2	**A**
Q.41	MCQ	2	**C**
Q.42	MCQ	2	**A**
Q.43	MCQ	2	**A**
Q.44	MCQ	2	**C**
Q.45	MCQ	2	**B**
Q.46	MSQ	2	**A, B, D**
Q.47	MSQ	2	**A, C**
Q.48	MSQ	2	**B, D**
Q.49	MSQ	2	**B, C, D**
Q.50	MSQ	2	**A, B**
Q.51	MSQ	2	**B, D**
Q.52	MSQ	2	**A, B, C**
Q.53	MSQ	2	**A, D**
Q.54	MSQ	2	**A, C**
Q.55	MSQ	2	**A, B**
Q.56	NAT	2	**0.125**
Q.57	NAT	2	**0.2**
Q.58	NAT	2	**0.25**
Q.59	NAT	2	**Marks to all**
Q.60	NAT	2	**0.5**
Q.61	NAT	2	**55**
Q.62	NAT	2	**0.12 to 0.13**
Q.63	NAT	2	**5**
Q.64	NAT	2	**0.125**
Q.65	NAT	2	**0.062 to 0.063**

General Aptitude Q.1 to Q.10

Q.1 MCQ · 1 MARK

General Aptitude › Grammar & vocabulary

If '→' denotes increasing order of intensity, then the meaning of the words [sick → infirm → moribund] is analogous to [silly → _____ → daft]. Which one of the given options is appropriate to fill the blank?

(A) frown

(B) fawn

(C) vein

(D) vain

Official answer: **D**

SOLUTION

1. The words in the first group rise in intensity: sick (unwell) → infirm (weak, frail) → moribund (close to death). All three are adjectives, and each is stronger than the one before.
2. The second group must follow the same pattern: silly (mildly foolish) → blank → daft (very foolish). The blank therefore needs an adjective of foolishness, stronger than silly and weaker than daft.
3. Option (A) frown is a verb or noun (a facial expression) and does not describe foolishness. Option (B) fawn is a verb (to flatter) or a noun (a young deer) and does not describe foolishness either.
4. Option (C) vein is a noun (a blood vessel or a streak). It is the wrong part of speech and has no link to silly or daft.
5. Option (D) vain is an adjective. In the sense of a foolish, futile or empty attitude it fits between silly and daft, and it is the only option that is an adjective with a suitable meaning.
6. Hence the blank is filled by vain, option (D).

KEY IDEA

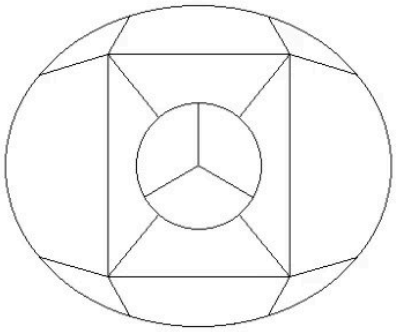
In an intensity analogy, the missing word must be the same part of speech as its neighbours and sit between them in degree. Here that means an adjective between silly and daft.

WATCH OUT FOR

Choosing a word by its sound or spelling, such as fawn or vein, without checking that it is an adjective of the right meaning. The other three options are not adjectives of foolishness at all.

Q.2 MCQ · 1 MARK

General Aptitude › Spatial

Q.2	The 15 parts of the given figure are to be painted such that no two adjacent parts with shared boundaries (excluding corners) have the same color. The minimum number of colors required is 
(A)	4
(B)	3
(C)	5
(D)	6

GATE 2024

Official answer: **A**

SOLUTION

1. The innermost circle is split by three radii into 3 sectors. Each pair of sectors shares a radius, so all three are mutually adjacent and need 3 different colors. Call the left, bottom and right sectors S_1, S_2, S_3 and give them colors 1, 2, 3.
2. The 4 ring trapezoids each touch the circle's arc, and neighbouring trapezoids share a diagonal line. The top trapezoid touches S_1 and S_3. The left one touches S_1 and S_2. The right one touches S_2 and S_3. The bottom one touches only S_2.

3. Try 3 colors. The top trapezoid must differ from colors 1 and 3, so it gets 2. The left one must differ from 1 and 2, so it gets 3. The right one must differ from 2 and 3, so it gets 1. The bottom trapezoid touches S_2 (color 2), the left trapezoid (3) and the right trapezoid (1), so it needs a fourth color. Hence 3 colors are impossible, and at least 4 are needed.
4. Show 4 colors suffice. Give the bottom trapezoid color 4. Each of the 4 outer side regions touches only the trapezoid across the square's edge, plus the two corner regions beside it. Color the top outer region 1, the right outer region 2, the bottom outer region 1 and the left outer region 1. Each differs from the trapezoid it touches (colors 2, 1, 4, 3).
5. Each corner region touches only the two side regions next to it, and these two are not adjacent to each other. The top-right and bottom-right corner regions touch colors 1 and 2, so give them 3. The top-left and bottom-left corner regions touch only color 1, so give them 2. Every pair of adjacent parts now has different colors.
6. Minimum number of colors $= 4$, which is option (A).

KEY IDEA

Find a group of mutually adjacent parts to set a lower bound, then show the lower bound is reachable by an explicit coloring. Here the three circle sectors force 3 colors, and the ring of trapezoids around them forces a fourth.

WATCH OUT FOR

Stopping at 3 colors because the three circle sectors need only 3, or assuming the outer regions need extra colors. The fourth color is forced by the bottom trapezoid, which touches the bottom sector and both of its neighbouring trapezoids.

Q.3 MCQ · 1 MARK

General Aptitude › Numerical

How many 4-digit positive integers divisible by 3 can be formed using only the digits {1, 3, 4, 6, 7}, such that no digit appears more than once in a number?

(A) 24

(B) 48

(C) 72

(D) 12

Official answer: **B**

SOLUTION

1. A number is divisible by 3 exactly when its digit sum is divisible by 3. The order of the digits does not change the digit sum.
2. A 4-digit number with no repeated digit, made from the 5 given digits, uses four of them. Exactly one digit d is left out.
3. The five digits add up to $1 + 3 + 4 + 6 + 7 = 21$, which is divisible by 3.
4. The four chosen digits therefore add up to $21 - d$. This is divisible by 3 exactly when d is divisible by 3.
5. The digits in the set that are divisible by 3 are 3 and 6. Leaving out 3 gives the digits $\{1, 4, 6, 7\}$ with sum 18. Leaving out 6 gives $\{1, 3, 4, 7\}$ with sum 15.
6. None of the digits is 0, so every arrangement of four distinct digits is a valid 4-digit number. Each set gives $4! = 24$ numbers.
7. Total $= 24 + 24 = 48$, which is option (B).

KEY IDEA
Use the digit-sum test for 3. The five digits add up to 21, so a four-digit selection is divisible by 3 only when the digit left out is itself a multiple of 3. Count the arrangements of each valid selection.

WATCH OUT FOR
Counting all $5 \cdot 4 \cdot 3 \cdot 2 = 120$ arrangements and dividing by 3 to get 40, or leaving out any digit without checking the digit sum. Divisibility by 3 depends on which digits are chosen, so the numbers do not split evenly into thirds.

Q.4 MCQ · 1 MARK

General Aptitude › Numerical

The sum of the following infinite series is

$$2 + \frac{1}{2} + \frac{1}{3} + \frac{1}{4} + \frac{1}{8} + \frac{1}{9} + \frac{1}{16} + \frac{1}{27} + \cdots$$

(A) 11/3

(B) 7/2

(C) 13/4

(D) 9/2

Official answer: **B**

SOLUTION

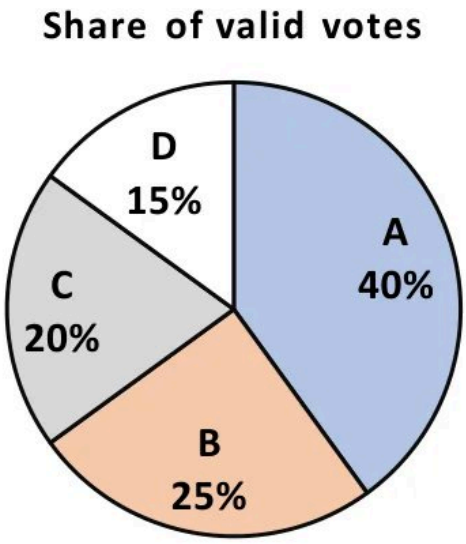
1. Read the terms after the leading 2: $\frac{1}{2}, \frac{1}{3}, \frac{1}{4}, \frac{1}{8}, \frac{1}{9}, \frac{1}{16}, \frac{1}{27}, \ldots$. Every denominator is a power of 2 or a power of 3 (2, 4, 8, 16 and 3, 9, 27), listed in increasing order of size. So the series is the leading 2 plus all terms $\frac{1}{2^k}$ ($k \geq 1$) plus all terms $\frac{1}{3^k}$ ($k \geq 1$).
2. Both groups are geometric series with ratio of magnitude less than 1, so each converges absolutely. Absolute convergence means the terms can be regrouped in any order without changing the sum.
3. Sum of an infinite geometric series with first term a and ratio r, $|r| < 1$: $S = \frac{a}{1-r}$.
4. Powers of 2: $a = \frac{1}{2}$, $r = \frac{1}{2}$, so $S_1 = \frac{1/2}{1-1/2} = 1$.
5. Powers of 3: $a = \frac{1}{3}$, $r = \frac{1}{3}$, so $S_2 = \frac{1/3}{1-1/3} = \frac{1}{2}$.
6. Total: $2 + 1 + \frac{1}{2} = \frac{7}{2}$, which is option (B).

KEY IDEA
Split the series into two geometric series, one for the powers of 2 and one for the powers of 3, and sum each with $S = \frac{a}{1-r}$. Absolute convergence makes the regrouping valid.

WATCH OUT FOR
Leaving out the leading 2, or applying $\frac{a}{1-r}$ to the powers of 2 and 3 as if they were a single geometric series. Another slip is taking the first term as 1 instead of $\frac{1}{2}$ or $\frac{1}{3}$, which gives 2 and $\frac{3}{2}$.

Q.5 MCQ · 1 MARK

General Aptitude › Data interpretation

Q.5	In an election, the share of valid votes received by the four candidates A, B, C, and D is represented by the pie chart shown. The total number of votes cast in the election were 1,15,000, out of which 5,000 were invalid.  Based on the data provided, the total number of valid votes received by the candidates B and C is
(A)	45,000
(B)	49,500
(C)	51,750
(D)	54,000

Official answer: **B**

SOLUTION

1. Valid votes = total votes cast minus invalid votes: $115000 - 5000 = 110000$. The pie chart shows shares of valid votes only, so 110000 is the base for the percentages.
2. Read the shares from the chart: $B = 25\%$ and $C = 20\%$.
3. Combined share of B and C: $25\% + 20\% = 45\%$.
4. Votes for B and C: $0.45 \times 110000 = 49500$.

5. Check: the shares add up to $40 + 25 + 20 + 15 = 100\%$, so 110000 is the right base. The answer is option (B).

KEY IDEA
The pie chart shows shares of valid votes only. Subtract the invalid votes from the total cast to get the base, then apply the combined percentage of B and C to that base.

WATCH OUT FOR
Applying 45% to the total votes cast, 115000, gives 51750 (option C). That is wrong because the chart's percentages are shares of the valid votes only, and 5000 votes were invalid.

Q.6 MCQ · 2 MARKS

General Aptitude › Reading

Thousands of years ago, some people began dairy farming. This coincided with a number of mutations in a particular gene that resulted in these people developing the ability to digest dairy milk.

Based on the given passage, which of the following can be inferred?

(A) All human beings can digest dairy milk.

(B) No human being can digest dairy milk.

(C) Digestion of dairy milk is essential for human beings.

(D) In human beings, digestion of dairy milk resulted from a mutated gene.

Official answer: **D**

SOLUTION

1. An inference question asks what must follow from the passage alone, not what is plausible or true in general.
2. The passage says that mutations in a particular gene resulted in some people developing the ability to digest dairy milk.
3. (A) says all human beings can digest dairy milk. The passage speaks only of some people, so this over-generalises.
4. (B) says no human being can digest dairy milk. This contradicts the passage, which says some people can.
5. (C) says digestion of dairy milk is essential for human beings. The passage says nothing about necessity.
6. (D) says that in human beings, digestion of dairy milk resulted from a mutated gene. This restates the passage's causal claim (mutations in a gene gave the ability), so it is supported.
7. Only (D) follows from the passage.

KEY IDEA
An inference must follow from the passage alone. Pick the option that restates what the passage says, and reject options that go beyond it or contradict it.

WATCH OUT FOR
Choosing (A) by reading "some people" as "all people", or choosing (C) because dairy digestion sounds important. The passage supports neither.

Q.7 MCQ · 2 MARKS

General Aptitude › Statistics & probability

The probability of a boy or a girl being born is 1/2. For a family having only three children, what is the probability of having two girls and one boy?

(A) 3/8

(B) 1/8

(C) 1/4

(D) 1/2

Official answer: **A**

SOLUTION

1. Each child is independently a girl or a boy with probability $\frac{1}{2}$, so the three births form a binomial experiment with $n = 3$ and $p = \frac{1}{2}$.
2. The sample space has $2^3 = 8$ equally likely outcomes.
3. The favourable outcomes are those with exactly two girls and one boy. The boy can be the first, second or third child, so there are $\binom{3}{2} = 3$ outcomes: GGB, GBG, BGG.
4. Each outcome has probability $\left(\frac{1}{2}\right)^3 = \frac{1}{8}$.
5. $$P(2 \text{ girls}, 1 \text{ boy}) = \binom{3}{2}\left(\frac{1}{2}\right)^2\left(\frac{1}{2}\right)^1 = 3 \times \frac{1}{8} = \frac{3}{8}.$$
6. This matches option (A).

KEY IDEA

Count the arrangements. Exactly 2 girls in 3 independent births with p = 1/2 is a binomial case, so the probability is C(3,2) times (1/2)^3, which is 3/8.

WATCH OUT FOR

Using only one arrangement (GGB) and answering 1/8, or forgetting that the boy can be born in any of the three positions.

Q.8 MCQ · 2 MARKS

General Aptitude › Numerical

Q.8

Person 1 and Person 2 invest in three mutual funds A, B, and C. The amounts they invest in each of these mutual funds are given in the table.

	Mutual fund A	Mutual fund B	Mutual fund C
Person 1	₹10,000	₹20,000	₹20,000
Person 2	₹20,000	₹15,000	₹15,000

At the end of one year, the total amount that Person 1 gets is ₹500 more than Person 2. The annual rate of return for the mutual funds B and C is 15% each. What is the annual rate of return for the mutual fund A?

(A) 7.5%

(B) 10%

(C) 15%

(D) 20%

GATE 2024

Official answer: **B**

SOLUTION

1. Let the annual rate of return of fund A be r. Rs. 1 in fund A becomes $(1+r)$ after one year, and Rs. 1 in fund B or C becomes 1.15.
2. Person 1 receives $10000(1+r)+20000(1.15)+20000(1.15)=10000(1+r)+46000$.
3. Person 2 receives $20000(1+r)+15000(1.15)+15000(1.15)=20000(1+r)+34500$.
4. Person 1 gets Rs. 500 more than Person 2, so

$$10000(1+r)+46000-20000(1+r)-34500=500.$$

5. Simplify: $11500-10000(1+r)=500$, so $10000(1+r)=11000$.
6. Hence $1+r=1.1$, so $r=0.10=10\%$. This is option (B).

KEY IDEA

Write each person's amount after one year as the sum of principal times growth factor for each fund, then use the stated difference of Rs. 500 to solve for the one unknown growth factor.

WATCH OUT FOR

Using the rate r directly instead of the growth factor $(1+r)$. Another slip is comparing only the returns (interest earned) while using total amounts, or mixing up which person invested more in fund A. Person 2 put more into fund A, so that fund's term carries the negative sign in the difference.

Q.9 MCQ · 2 MARKS

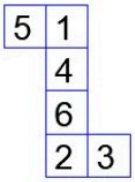
General Aptitude › Spatial

Q.9 Three different views of a dice are shown in the figure below.

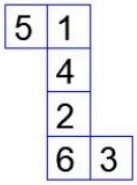


The piece of paper that can be folded to make this dice is

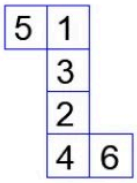
(A)



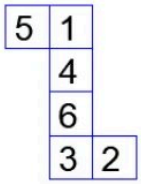
(B)



(C)



(D)



Official answer: **A**

SOLUTION

1. Faces that meet at one corner of the dice are adjacent, never opposite. View 1 shows 5, 4, 1 together; view 2 shows 4, 6, 3; view 3 shows 2, 6, 5.
2. Face 5 is adjacent to 4 and 1 (view 1) and to 2 and 6 (view 3). The only remaining face is 3, so 5 is opposite 3.

GATE 2024

3. Face 4 is adjacent to 5 and 1 (view 1) and to 6 and 3 (view 2). The only remaining face is 2, so 4 is opposite 2.
4. Face 6 is adjacent to 4 and 3 (view 2) and to 2 and 5 (view 3). The only remaining face is 1, so 6 is opposite 1.
5. Opposite pairs: $5 \leftrightarrow 3$, $4 \leftrightarrow 2$, $6 \leftrightarrow 1$.
6. Each option is a column of four squares with one extra square on the side of the first and one on the side of the last, on opposite sides of the column. When folded, the four column squares form a ring around the dice, so squares 1 and 3 of the column are opposite, and squares 2 and 4 are opposite. The two side squares become the two end caps of the ring, so they are opposite each other.
7. Option (A): the column is 1, 4, 6, 2. The pairs (1, 6) and (4, 2) are opposite, which agrees with the pairs found above. The side squares are 5 and 3, which are also opposite. All three pairs match.
8. Option (B): the column is 1, 4, 2, 6, which puts 1 and 2 opposite. This is wrong. Option (C): the column is 1, 3, 2, 4, which again puts 1 and 2 opposite. This is wrong. Option (D): the column is 1, 4, 6, 3, which puts 4 and 3 opposite. This is wrong.
9. Only option (A) respects all the opposite-face pairs, so the answer is (A).

KEY IDEA

Faces seen together at a corner are adjacent. From three views, find each face's opposite as the one face never seen beside it, then test each net: in a four-square column, squares one place apart in the sequence (1st and 3rd, 2nd and 4th) are opposite, and the two side squares on opposite sides are opposite.

WATCH OUT FOR

Checking only the faces that touch the first square, or forgetting that squares two places apart in the column fold to opposite faces. Options (B), (C) and (D) each look plausible but place an adjacent pair, or a pair that must be opposite, in the wrong positions.

Q.10 MCQ · 2 MARKS

General Aptitude › Spatial

Visualize two identical right circular cones such that one is inverted over the other and they share a common circular base. If a cutting plane passes through the vertices of the assembled cones, what shape does the outer boundary of the resulting cross-section make?

(A) A rhombus

(B) A triangle

(C) An ellipse

(D) A hexagon

Official answer: **A**

SOLUTION

1. Picture the solid: two identical right circular cones joined base to base. Let the apexes be P (top) and Q (bottom), both on the common axis, and let O be the centre of the shared base.
2. The cutting plane passes through both P and Q, so it contains the whole line PQ, which is the axis of the solid. The axis is perpendicular to the base at O.
3. A plane containing the axis meets the common circular base in a diameter AB with O as its midpoint, and $AB \perp PQ$ because the axis is perpendicular to the base.
4. In the upper cone the section is triangle PAB. The cone is right circular, so $PA = PB$.
5. The lower cone is identical to the upper one, so its section is triangle QAB with $QA = QB$. Since $PO = QO$ and $AO = BO$ with both axes perpendicular at O, the four right triangles POA, POB, QOA, QOB are congruent, so $PA = PB = QA = QB$.
6. The two triangles share the side AB, so the outer boundary of the section is the quadrilateral $PAQB$, with all four sides equal.

7. A quadrilateral with four equal sides is a rhombus. Its diagonals PQ and AB bisect each other at right angles at O. It becomes a square only if $PQ = AB$, which is not required here.
8. So the outer boundary of the cross-section is a rhombus, option (A).

KEY IDEA

A plane through both apexes of a double cone must contain the axis. It then cuts each cone in an isosceles triangle on the shared base diameter, and the two congruent triangles together form a rhombus.

WATCH OUT FOR

Thinking of a plane cutting across the cone, which gives a conic such as an ellipse. Here the plane contains the axis, so the section is bounded by straight lines. Another slip is to take one triangle (option B) and forget that the section of the second cone is joined to it.

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

Q.11 MCQ · 1 MARK

Probability and Statistics › Continuous distributions

Consider the following statements:
(i) The mean and variance of a Poisson random variable are equal.
(ii) For a standard normal random variable, the mean is zero and the variance is one.

Which ONE of the following options is correct?

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

(B) (i) is true and (ii) is false

(C) (ii) is true and (i) is false

(D) Both (i) and (ii) are false

Official answer: **A**

SOLUTION

1. Statement (i): let $X \sim \text{Poisson}(\lambda)$ with $P(X=k) = e^{-\lambda}\lambda^k/k!$ for $k = 0, 1, 2, \ldots$. Then $E[X] = \sum_{k\geq 1} k\, e^{-\lambda}\lambda^k/k! = \lambda e^{-\lambda}\sum_{k\geq 1}\lambda^{k-1}/(k-1)! = \lambda$.
2. The factorial moment is $E[X(X-1)] = \sum_{k\geq 2} e^{-\lambda}\lambda^k/(k-2)! = \lambda^2 e^{-\lambda}\sum_{j\geq 0}\lambda^j/j! = \lambda^2$.
3. So $E[X^2] = E[X(X-1)] + E[X] = \lambda^2 + \lambda$.
4. Then $\text{Var}(X) = E[X^2] - (E[X])^2 = \lambda^2 + \lambda - \lambda^2 = \lambda$, which equals the mean. Statement (i) is true.
5. Statement (ii): the normal distribution $N(\mu, \sigma^2)$ has mean μ and variance σ^2. The standard normal is the case $\mu = 0$, $\sigma^2 = 1$, so its mean is zero and its variance is one. Statement (ii) is true.
6. Both statements are true, so the answer is option (A).

KEY IDEA
A Poisson(λ) variable has mean λ and variance λ. The standard normal $N(0,1)$ is the normal with mean 0 and variance 1. Both statements hold.

WATCH OUT FOR
Confusing the second parameter of $N(\mu, \sigma^2)$ with the standard deviation, or doubting that a Poisson variable's variance equals its mean. Both statements are standard facts, so there is nothing to trap you.

Q.12 MCQ · 1 MARK

Probability and Statistics › Counting & axioms

Three fair coins are tossed independently. T is the event that two or more tosses result in heads. S is the event that two or more tosses result in tails.

What is the probability of the event $T \cap S$?

(A) 0

(B) 0.5

(C) 0.25

(D) 1

Official answer: **A**

SOLUTION

1. Three tosses are made and each toss is a head or a tail, so $(\text{number of heads}) + (\text{number of tails}) = 3$.
2. T means at least 2 heads: $T = \{HHH, HHT, HTH, THH\}$, so $P(T) = \frac{4}{8} = 0.5$.
3. S means at least 2 tails: $S = \{TTT, TTH, THT, HTT\}$, so $P(S) = \frac{4}{8} = 0.5$.
4. $T \cap S$ needs at least 2 heads and at least 2 tails together, so at least $2 + 2 = 4$ tosses. Only 3 tosses are made, so no outcome can satisfy both. The two outcome sets above also share no element.

5. Hence $T \cap S = \varnothing$ and
$$P(T \cap S) = 0.$$
6. Check on independence: $P(T)\,P(S) = 0.5 \times 0.5 = 0.25 \neq 0 = P(T \cap S)$. So T and S are mutually exclusive but not independent. The coins are independent, but the events T and S built from them are not.

KEY IDEA
With 3 tosses, heads plus tails equals 3. Two events that each need at least 2 of the 3 tosses cannot both occur, so they are mutually exclusive and the intersection is empty.

WATCH OUT FOR
Multiplying $P(T)P(S) = 0.25$ as if T and S were independent, and picking (C). Independent coins do not make T and S independent. In fact S is determined by the same tosses and cannot occur when T does.

Q.13 MCQ · 1 MARK

Linear Algebra › Eigenvalues

Consider the matrix $\mathbf{M} = \begin{bmatrix} 2 & -1 \\ 3 & 1 \end{bmatrix}$.

Which ONE of the following statements is TRUE?

(A) The eigenvalues of $\mathbf{M}$ are non-negative and real.

(B) The eigenvalues of $\mathbf{M}$ are complex conjugate pairs.

(C) One eigenvalue of $\mathbf{M}$ is positive and real, and another eigenvalue of $\mathbf{M}$ is zero.

(D) One eigenvalue of $\mathbf{M}$ is non-negative and real, and another eigenvalue of $\mathbf{M}$ is negative and real.

Official answer: **B**

SOLUTION

1. For a 2×2 matrix the characteristic equation is $\lambda^2 - (\operatorname{tr}\mathbf{M})\lambda + \det\mathbf{M} = 0$.
2. Trace: $\operatorname{tr}\mathbf{M} = 2 + 1 = 3$.
3. Determinant: $\det\mathbf{M} = (2)(1) - (-1)(3) = 2 + 3 = 5$.
4. So the characteristic equation is $\lambda^2 - 3\lambda + 5 = 0$.
5. Discriminant: $D = 3^2 - 4(5) = 9 - 20 = -11 < 0$, so there are no real roots.
6. $\lambda = \dfrac{3 \pm i\sqrt{11}}{2}$, a complex conjugate pair. Check: the sum is $3 = \operatorname{tr}\mathbf{M}$ and the product is $\dfrac{9 + 11}{4} = 5 = \det\mathbf{M}$.
7. Options (A), (C) and (D) each require at least one real eigenvalue, so they are false. In (C), $\det\mathbf{M} = 5 \neq 0$ also rules out a zero eigenvalue. Only (B) is true.

KEY IDEA
For a 2×2 real matrix, find the sign of the discriminant of $\lambda^2 - \operatorname{tr}(\mathbf{M})\lambda + \det(\mathbf{M})$. A negative discriminant means a complex conjugate pair, with no real eigenvalue.

WATCH OUT FOR
Computing the determinant as $2 - 3 = -1$ by dropping the sign of the -1 entry. That gives $\lambda^2 - 3\lambda - 1 = 0$, which has real roots of opposite sign and wrongly points to (D). The correct determinant is $(2)(1) - (-1)(3) = 5$.

Q.14 MCQ · 1 MARK

Data Structures and Algorithms › Graph algorithms

Consider performing depth-first search (DFS) on an undirected and unweighted graph G starting at vertex s. For any vertex u in G, $d[u]$ is the length of the shortest path from s to u. Let (u, v) be an edge in G such that $d[u] < d[v]$. If the edge (u, v) is explored first in the direction from u to v during the above DFS, then (u, v) becomes a _____ edge.

(A) tree

(B) cross

(C) back

(D) gray

Official answer: **A**

Q.15 MCQ · 1 MARK

Calculus and Optimization › Maxima & minima

For any twice differentiable function $f : \mathbb{R} \to \mathbb{R}$, if at some $x^* \in \mathbb{R}$, $f'(x^*) = 0$ and $f''(x^*) > 0$, then the function f necessarily has a _____ at $x = x^*$.

Note: $\mathbb{R}$ denotes the set of real numbers.

(A) local minimum

(B) global minimum

(C) local maximum

(D) global maximum

Official answer: **A**

SOLUTION

1. Since $f'(x^*) = 0$, the point x^* is a stationary (critical) point of f.
2. Because $f''(x^*) = \lim_{x \to x^*} \dfrac{f'(x) - f'(x^*)}{x - x^*} = \lim_{x \to x^*} \dfrac{f'(x)}{x - x^*} > 0$, the quotient is positive for x close to x^*. So $f'(x) < 0$ just to the left of x^* and $f'(x) > 0$ just to the right. Hence f decreases and then increases, and x^* is a strict local minimum (the second-derivative test).
3. It cannot be a local maximum, since f would have to stop increasing after x^*, but $f' > 0$ immediately to the right.
4. The conclusion is local only, because $f''(x^*) > 0$ gives no information about f far from x^*. To show that a global minimum is not guaranteed, take $f(x) = x^3 - 3x$ and $x^* = 1$. Then $f'(x) = 3x^2 - 3$, so $f'(1) = 0$, and $f''(x) = 6x$, so $f''(1) = 6 > 0$. Yet $f(1) = -2$ while $f(-3) = -27 + 9 = -18 < -2$ (indeed $f(x) \to -\infty$ as $x \to -\infty$). So $x = 1$ is a local minimum but not a global one. This also rules out a global maximum, since a global maximum would be a local maximum, which was excluded in step 3.
5. So f necessarily has a local minimum at x^*, which is option A.

KEY IDEA

At a stationary point, $f''(x^*) > 0$ (second-derivative test) guarantees a local minimum only. It says nothing about global behaviour.

WATCH OUT FOR

Choosing "global minimum" because the second derivative is positive. $f''(x^*) > 0$ is a local condition. Global minimality needs convexity everywhere, i.e. $f'' \geq 0$ on all of $\mathbb{R}$.

Q.16 MCQ · 1 MARK

Data Structures and Algorithms › Linear structures

Q.16 Match the items in **Column 1** with the items in **Column 2** in the following table:

Column 1	Column 2
(p) First In First Out	(i) Stacks
(q) Lookup Operation	(ii) Queues
(r) Last In First Out	(iii) Hash Tables

(A)	(p) – (ii), (q) – (iii), (r) – (i)
(B)	(p) – (ii), (q) – (i), (r) – (iii)
(C)	(p) – (i), (q) – (ii), (r) – (iii)
(D)	(p) – (i), (q) – (iii), (r) – (ii)

Official answer: **A**

SOLUTION

1. First In First Out (FIFO) means the earliest inserted item leaves first. That is the queue discipline, so (p) matches (ii) Queues.
2. Last In First Out (LIFO) means the most recently inserted item leaves first. That is the stack discipline, so (r) matches (i) Stacks.
3. Lookup (search by key) is the defining operation of a hash table, which gives expected constant-time lookup. So (q) matches (iii) Hash Tables.
4. The matching is (p)-(ii), (q)-(iii), (r)-(i), which is option A.
5. Checking the other options against this matching: B pairs lookup with stacks and LIFO with hash tables, so (q) and (r) are wrong. C pairs FIFO with stacks, lookup with queues and LIFO with hash tables, so all three pairs are wrong. D gets lookup right but swaps FIFO and LIFO, pairing FIFO with stacks and LIFO with queues.

KEY IDEA

Match each item to the structure that defines it: FIFO is the queue discipline, LIFO is the stack discipline, and key-based lookup is what a hash table is built for.

WATCH OUT FOR

Swapping FIFO and LIFO, so that stacks are paired with first-in-first-out (option D). A queue serves the oldest item and a stack serves the newest.

Q.17 MCQ · 1 MARK

Machine Learning › LDA, SVM & trees

Consider the dataset with six datapoints: $\{(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \ldots, (\mathbf{x}_6, y_6)\}$, where $\mathbf{x}_1 = \begin{bmatrix}1\\0\end{bmatrix}$, $\mathbf{x}_2 = \begin{bmatrix}0\\1\end{bmatrix}$, $\mathbf{x}_3 = \begin{bmatrix}0\\-1\end{bmatrix}$, $\mathbf{x}_4 = \begin{bmatrix}-1\\0\end{bmatrix}$, $\mathbf{x}_5 = \begin{bmatrix}2\\2\end{bmatrix}$, $\mathbf{x}_6 = \begin{bmatrix}-2\\-2\end{bmatrix}$ and the labels are given by $y_1 = y_2 = y_5 = 1$, and $y_3 = y_4 = y_6 = -1$. A hard margin linear support vector machine is trained on the above dataset. Which ONE of the following sets is a possible set of support vectors?

(A) $\{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_5\}$

(B) $\{\mathbf{x}_3, \mathbf{x}_4, \mathbf{x}_5\}$

(C) $\{\mathbf{x}_4, \mathbf{x}_5\}$

(D) $\{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4\}$

Official answer: **D**

SOLUTION

1. Positives: $\mathbf{x}_1 = (1, 0)$, $\mathbf{x}_2 = (0, 1)$, $\mathbf{x}_5 = (2, 2)$. Negatives: $\mathbf{x}_3 = (0, -1)$, $\mathbf{x}_4 = (-1, 0)$, $\mathbf{x}_6 = (-2, -2)$. The data are linearly separable, since $x_1 + x_2 \geq 1$ for every positive and ≤ -1 for every negative.
2. A hard-margin SVM finds the separating line with the largest margin. Equivalently, the boundary is the perpendicular bisector of the shortest segment joining the convex hull of the positives to the convex hull of the negatives.
3. Compute $x_1 + x_2$ for each point. Positives: $1, 1, 4$ for $\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_5$. Negatives: $-1, -1, -4$ for $\mathbf{x}_3, \mathbf{x}_4, \mathbf{x}_6$.
4. Take $\mathbf{w} = (1, 1)$ and $b = 0$. Then $y_i(\mathbf{w}^\top \mathbf{x}_i + b) = 1$ for $i = 1, 2, 3, 4$, and it equals 4 for $i = 5, 6$. All constraints $y_i(\mathbf{w}^\top \mathbf{x}_i + b) \geq 1$ hold, and the points with value exactly 1 are $\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4$.
5. Check that this is the maximum margin. The point $(0.5, 0.5)$ lies on the segment $\mathbf{x}_1\mathbf{x}_2$ (in the positive hull) and $(-0.5, -0.5)$ lies on the segment $\mathbf{x}_3\mathbf{x}_4$ (in the negative hull). Their distance is $\sqrt{2}$, so no separating margin can be wider than $\sqrt{2}$. The solution above has margin width $2/\|\mathbf{w}\| = 2/\sqrt{2} = \sqrt{2}$, so it is optimal. The hard-margin SVM solution is unique, so this is the trained classifier.
6. Support vectors are the training points lying exactly on the margin. Here these are $\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4$. Indeed $\mathbf{w} = \sum_i \alpha_i y_i \mathbf{x}_i = (1, 1)$ is attainable, for example with $\alpha_1 = \alpha_2 = \alpha_3 = \alpha_4 = \frac{1}{2}$, which also satisfies $\sum_i \alpha_i y_i = 0$.
7. Options A, B and C each contain $\mathbf{x}_5$, but $y_5(\mathbf{w}^\top \mathbf{x}_5 + b) = 4 > 1$, so $\mathbf{x}_5$ lies strictly outside the margin and cannot be a support vector. The answer is D.

KEY IDEA

Support vectors are the points exactly on the margin. Find the maximum-margin hyperplane (the bisector of the closest pair of points in the two convex hulls), then pick the points with $y_i(\mathbf{w}^\top \mathbf{x}_i + b) = 1$.

WATCH OUT FOR

Assuming the far points $(2, 2)$ and $(-2, -2)$ must be support vectors because they look extreme. Support vectors are the points closest to the boundary, not the farthest ones; points well outside the margin have zero multiplier.

Q.18 MCQ · 1 MARK

Machine Learning › Logistic, k-NN, naive Bayes

Q.18 Match the items in **Column 1** with the items in **Column 2** in the following table:

Column 1	Column 2
(p) Principal Component Analysis	(i) Discriminative Model
(q) Naïve Bayes Classification	(ii) Dimensionality Reduction
(r) Logistic Regression	(iii) Generative Model

(A) (p) – (iii), (q) – (i), (r) – (ii)

(B) (p) – (ii), (q) – (i), (r) – (iii)

(C) (p) – (ii), (q) – (iii), (r) – (i)

(D) (p) – (iii), (q) – (ii), (r) – (i)

Official answer: **C**

SOLUTION

1. Principal Component Analysis projects the data onto the directions of largest variance, which are the top eigenvectors of the covariance matrix. It reduces the number of features, so it is a dimensionality-reduction method: (p) - (ii).
2. Naive Bayes models how the data arise. It estimates the class prior $P(y)$ and the class-conditional likelihood $P(\mathbf{x} \mid y)$, assuming the features are conditionally independent given the class. It then applies Bayes' theorem to get $P(y \mid \mathbf{x}) \propto P(y)\, P(\mathbf{x} \mid y)$. Because it models the joint distribution $P(\mathbf{x}, y)$, it is a generative model: (q) - (iii).
3. Logistic regression models the posterior directly, $P(y = 1 \mid \mathbf{x}) = \dfrac{1}{1 + e^{-\mathbf{w}^T \mathbf{x}}}$, and never models $P(\mathbf{x})$. Modelling the conditional distribution directly makes it a discriminative model: (r) - (i).
4. The matching is (p)-(ii), (q)-(iii), (r)-(i). This is option (C).

KEY IDEA

A generative model learns the joint distribution $P(\mathbf{x}, y)$, usually as $P(y)P(\mathbf{x} \mid y)$ (Naive Bayes). A discriminative model learns $P(y \mid \mathbf{x})$ directly (logistic regression). PCA is unsupervised and has no class model, so it is dimensionality reduction.

WATCH OUT FOR

Swapping Naive Bayes and logistic regression because both are called classifiers that output class probabilities. What separates them is whether $P(\mathbf{x} \mid y)$ is modelled (generative) or only $P(y \mid \mathbf{x})$ (discriminative). Another slip is labelling PCA a generative model.

Q.19 MCQ · 1 MARK

Machine Learning › Clustering

Euclidean distance based k-means clustering algorithm was run on a dataset of 100 points with $k = 3$. If the points $\begin{bmatrix}1\\1\end{bmatrix}$ and $\begin{bmatrix}-1\\1\end{bmatrix}$ are both part of cluster 3, then which ONE of the following points is necessarily also part of cluster 3?

(A) $\begin{bmatrix}0\\0\end{bmatrix}$

(B) $\begin{bmatrix}0\\2\end{bmatrix}$

(C) $\begin{bmatrix}2\\0\end{bmatrix}$

(D) $\begin{bmatrix}0\\1\end{bmatrix}$

Official answer: **D**

SOLUTION

1. In k-means with Euclidean distance, a point belongs to cluster 3 when its nearest centroid is μ_3. So cluster 3 is the set $\{x : \|x - \mu_3\| \leq \|x - \mu_j\| \text{ for } j = 1, 2\}$.
2. Expand $\|x - \mu_3\|^2 \leq \|x - \mu_j\|^2$. The $\|x\|^2$ terms cancel, leaving $2(\mu_j - \mu_3)^\top x \leq \|\mu_j\|^2 - \|\mu_3\|^2$. This is a half-plane, so cluster 3's region is an intersection of half-planes (a Voronoi cell) and is convex.
3. Both $\begin{bmatrix}1\\1\end{bmatrix}$ and $\begin{bmatrix}-1\\1\end{bmatrix}$ lie in this convex region, so every point of the segment joining them also lies in it. The segment is $\{(t, 1) : -1 \leq t \leq 1\}$.
4. The midpoint is $\left(\frac{1+(-1)}{2}, \frac{1+1}{2}\right) = (0, 1)$, which is option (D). So (D) is necessarily in cluster 3, apart from the degenerate case of an exact tie between centroids.
5. The other options are not guaranteed. Take $\mu_3 = (0, 1)$ and put another centroid exactly at the option point p. Then $\|(\pm 1, 1) - \mu_3\| = 1$, while the distance from $(\pm 1, 1)$ to p is at least $\sqrt{2}$ for $p = (0, 0)$ and $p = (0, 2)$, and at least $\sqrt{2}$ for $p = (2, 0)$ (distances $\sqrt{2}$ and $\sqrt{10}$). So both given points stay in cluster 3. But p is at distance 0 from its own centroid, so p is in a different cluster. Hence (A), (B) and (C) are not necessarily in cluster 3.

KEY IDEA

Each k-means cluster is a Voronoi cell, which is convex. Any point on the segment between two members of a cluster is therefore also in that cluster.

WATCH OUT FOR

Picking $(0, 0)$ because it is the origin or the "centre" of the data. The origin is not on the segment between the two points, so a centroid can sit near it and claim it for another cluster.

Q.20 MCQ · 1 MARK

Machine Learning › Logistic, k-NN, naive Bayes

Given a dataset with K binary-valued attributes (where $K > 2$) for a two-class classification task, the number of parameters to be estimated for learning a naïve Bayes classifier is

(A) $2^K + 1$

(B) $2K + 1$

(C) $2^{K+1} + 1$

(D) $K^2 + 1$

Official answer: **B**

SOLUTION

1. Naïve Bayes assumes the attributes are conditionally independent given the class, so the model is $P(Y)\prod_{i=1}^{K} P(X_i \mid Y)$. Only these factors need to be estimated.
2. Class prior: Y takes 2 values and $P(Y=0)+P(Y=1)=1$, so $P(Y)$ has $2-1=1$ free parameter.
3. Each attribute X_i is binary, so for each class value y the distribution $P(X_i \mid Y=y)$ is a Bernoulli distribution with one free parameter, $P(X_i=1 \mid Y=y)$. The value $P(X_i=0 \mid Y=y)$ follows from it.
4. Conditional parameters: K attributes $\times$ 2 classes $= 2K$.
5. Total number of parameters: $2K+1$, which is option (B).
6. Without the independence assumption, the joint distribution of K binary attributes would need $2^K - 1$ free parameters per class, so $2(2^K-1)+1$ in all. The naïve assumption therefore cuts the growth from exponential to linear in K.

KEY IDEA
Conditional independence given the class turns the joint over K attributes into K separate one-parameter Bernoulli distributions per class. Count the free parameters: 1 for the prior, plus $K \times 2$ for the likelihoods.

WATCH OUT FOR
Counting the full joint table 2^K per class, which ignores the independence assumption. Another slip is counting both $P(X_i=0 \mid y)$ and $P(X_i=1 \mid y)$, or both class priors, instead of only the free parameters.

Q.21 MCQ · 1 MARK

Data Structures and Algorithms › Sorting & hashing

Consider performing uniform hashing on an open address hash table with load factor $\alpha = \frac{n}{m} < 1$, where n elements are stored in the table with m slots. The expected number of probes in an unsuccessful search is at most $\frac{1}{1-\alpha}$.

Inserting an element in this hash table requires at most _____ probes, on average.

(A) $\ln\left(\frac{1}{1-\alpha}\right)$
(B) $\frac{1}{1-\alpha}$
(C) $1+\frac{\alpha}{2}$
(D) $\frac{1}{1+\alpha}$

Official answer: **B**

SOLUTION

1. In open addressing, to insert a key you follow its probe sequence, examining one slot after another, until you reach an empty slot, and then place the key there.
2. This is exactly the work of an unsuccessful search: that search also follows the probe sequence over occupied slots and stops at the first empty one. So the number of probes for an insertion has the same distribution as for an unsuccessful search; the only extra work is the single write into the empty slot, which is not a probe.
3. The question states that the expected number of probes in an unsuccessful search is at most $\frac{1}{1-\alpha}$. Hence inserting an element needs at most $\frac{1}{1-\alpha}$ probes on average.
4. Check by derivation. Let X be the number of probes. Under uniform hashing, the probability that the first i probes all land on occupied slots is $\frac{n}{m}\cdot\frac{n-1}{m-1}\cdots\frac{n-i+1}{m-i+1} \le \alpha^i$, because $\frac{n-k}{m-k} \le \frac{n}{m}$ when $n \le m$. Therefore $E[X] = \sum_{i\ge 0} P(X > i) \le \sum_{i\ge 0}\alpha^i = \frac{1}{1-\alpha}$.
5. Check the other options. (D) $\frac{1}{1+\alpha} < 1$ for every $\alpha > 0$, but at least one probe is always needed, so it cannot be an average. (A) $\ln\frac{1}{1-\alpha}$ equals 0 at $\alpha = 0$ and is below 1 for $\alpha < 1 - \frac{1}{e}$, again less than the

one probe an insertion always needs. (C) $1+\frac{\alpha}{2}$ stays below 2, yet the true expected number of probes, $\frac{m+1}{m-n+1}$, grows without limit as $\alpha \to 1$ (for large m), so $1+\frac{\alpha}{2}$ cannot be an upper bound.

6. Hence the answer is option (B), $\frac{1}{1-\alpha}$.

KEY IDEA

In open addressing, an insertion probes exactly as an unsuccessful search does, until the first empty slot, so it has the same bound $\frac{1}{1-\alpha}$.

WATCH OUT FOR

Picking $\ln \frac{1}{1-\alpha}$ (option A) by mixing up the insertion cost with the successful-search cost, which is about $\frac{1}{\alpha}\ln\frac{1}{1-\alpha}$. An insertion always needs at least one probe and its cost blows up as $\alpha \to 1$.

Q.22 MCQ · 1 MARK

Machine Learning › LDA, SVM & trees

For any binary classification dataset, let $S_B \in \mathbb{R}^{d\times d}$ and $S_W \in \mathbb{R}^{d\times d}$ be the between-class and within-class scatter (covariance) matrices, respectively. The Fisher linear discriminant is defined by $u^* \in \mathbb{R}^d$, that maximizes

$$J(u) = \frac{u^T S_B u}{u^T S_W u}$$

If $\lambda = J(u^*)$, S_W is non-singular and $S_B \neq 0$, then (u^*, λ) must satisfy which ONE of the following equations?

Note: $\mathbb{R}$ denotes the set of real numbers.

(A) $S_W^{-1} S_B u^* = \lambda u^*$

(B) $S_W u^* = \lambda S_B u^*$

(C) $S_B S_W u^* = \lambda u^*$

(D) $u^{*T} u^* = \lambda^2$

Official answer: **A**

SOLUTION

1. $J(u)$ is a ratio of quadratic forms (a Rayleigh quotient) and is unchanged when u is scaled. At the maximiser $u^* \neq 0$ the gradient must vanish: $\nabla J(u^*) = 0$.
2. Differentiate with the quotient rule, using that S_B and S_W are symmetric:
$$\nabla J = \frac{2S_B u\,(u^T S_W u) - 2S_W u\,(u^T S_B u)}{(u^T S_W u)^2}.$$
3. Set this to zero at u^* and divide by $u^{*T} S_W u^*$, which is positive because S_W is a non-singular scatter matrix (positive definite) and $u^* \neq 0$:
$$S_B u^* = \frac{u^{*T} S_B u^*}{u^{*T} S_W u^*}\, S_W u^* = \lambda\, S_W u^*.$$
4. This is the generalised eigenvalue problem $S_B u^* = \lambda S_W u^*$, where $\lambda = J(u^*)$ is the maximum value of the quotient. Since $S_B \neq 0$ is positive semi-definite, some u gives $u^T S_B u > 0$, so $\lambda > 0$.
5. Since S_W is non-singular, multiply on the left by S_W^{-1}:
$$S_W^{-1} S_B u^* = \lambda u^*.$$
So u^* is the eigenvector of $S_W^{-1} S_B$ belonging to its largest eigenvalue λ. This is option (A).
6. Check the others. (B) has S_W and S_B swapped: $S_W u^* = \lambda S_B u^*$ would mean $S_B u^* = \lambda^{-1} S_W u^*$, which contradicts step 3 unless $\lambda^2 = 1$. (C) has the matrices in the wrong order and no inverse, so it does not follow. (D) is not implied, because u^* can be scaled freely without changing J, so $u^{*T} u^*$ is not fixed by λ. Only (A) must hold.

KEY IDEA

Maximising the ratio $\frac{u^T S_B u}{u^T S_W u}$ gives the generalised eigenproblem $S_B u = \lambda S_W u$. With S_W invertible this becomes the ordinary eigenproblem for $S_W^{-1} S_B$, and the maximum J is its largest eigenvalue.

WATCH OUT FOR

Swapping the matrices and writing $S_W u = \lambda S_B u$, or multiplying in the wrong order ($S_B S_W$). The inverse must multiply S_B from the left, so that it comes from $S_B u = \lambda S_W u$.

Q.23 MCQ · 1 MARK

Artificial Intelligence › Search

Let h_1 and h_2 be two admissible heuristics used in A^* search.

Which ONE of the following expressions is always an admissible heuristic?

(A) $h_1 + h_2$

(B) $h_1 \times h_2$

(C) $h_1/h_2,\ (h_2 \neq 0)$

(D) $|h_1 - h_2|$

Official answer: **D**

SOLUTION

1. A heuristic h is admissible if $0 \leq h(n) \leq h^*(n)$ for every node n, where $h^*(n)$ is the true cheapest cost from n to a goal. We need the expression that satisfies this bound for every pair of admissible h_1, h_2 at every node.
2. (A) $h_1 + h_2$: take $h_1 = h_2 = h^*$, which is admissible because h^* never overestimates itself. The sum is $2h^*(n) > h^*(n)$ at any node with $h^*(n) > 0$. It overestimates, so it is not always admissible.
3. (B) $h_1 \times h_2$: take $h_1 = h_2 = h^*$ at a node where $h^*(n) = 5$. The product is $25 > 5$. Not admissible.
4. (C) h_1/h_2: take $h_1 = h^* = 5$ and $h_2 = 0.1$, which is admissible since $0 \leq 0.1 \leq 5$. The ratio is $50 > 5$. Not admissible.
5. (D) $|h_1 - h_2|$: both $h_1(n)$ and $h_2(n)$ lie in $[0, h^*(n)]$. The absolute difference of two non-negative numbers is at most the larger of them, so $0 \leq |h_1 - h_2| \leq \max(h_1, h_2) \leq h^*(n)$. It never overestimates, so it is admissible for every pair of admissible heuristics.
6. Therefore only $|h_1 - h_2|$ is always admissible, which is option (D).

KEY IDEA

Admissible means $0 \leq h \leq h^*$ at every node. The absolute difference of two values in $[0, h^*]$ stays in $[0, h^*]$, but a sum, product or ratio can exceed h^*.

WATCH OUT FOR

Choosing $h_1 + h_2$ because combining two good heuristics seems to give a stronger one. Adding two admissible estimates can double-count the remaining cost and overestimate. Only the maximum (or minimum) of admissible heuristics is guaranteed admissible.

Q.24 MCQ · 1 MARK

Artificial Intelligence › Uncertainty

Consider five random variables U, V, W, X, and Y whose joint distribution satisfies:

$$P(U, V, W, X, Y) = P(U)\, P(V)\, P(W \mid U, V)\, P(X \mid W)\, P(Y \mid W)$$

Which ONE of the following statements is FALSE?

(A) Y is conditionally independent of V given W

(B) X is conditionally independent of U given W

(C) U and V are conditionally independent given W

(D) Y and X are conditionally independent given W

Official answer: **C**

SOLUTION

1. Read the factorisation as a Bayesian network. Each factor $P(\text{child} \mid \text{parents})$ gives edges from parent to child, so the edges are $U \to W$, $V \to W$, $W \to X$ and $W \to Y$. The statement that is FALSE is the one that d-separation (or direct algebra) does not support.
2. (A) The only path from V to Y is $V \to W \to Y$. Here W is a chain (pass-through) node, so conditioning on W blocks the path. Hence $Y \perp V \mid W$ holds, and (A) is true.
3. (B) The only path from U to X is $U \to W \to X$. Again W is a chain node, so conditioning on W blocks it. Hence $X \perp U \mid W$ holds, and (B) is true.
4. (D) The only path from X to Y is $X \leftarrow W \to Y$, where W is a common cause (fork). Conditioning on W blocks it. Algebraically, $P(X, Y \mid W) = \dfrac{\sum_{U,V} P(U)P(V)P(W \mid U, V)\, P(X \mid W)P(Y \mid W)}{P(W)} = P(X \mid W)\, P(Y \mid W)$, because the factors $P(X \mid W)$ and $P(Y \mid W)$ do not involve U or V. So (D) is true.
5. (C) The path $U \to W \leftarrow V$ has W as a collider. Marginally $P(U, V) = P(U)P(V)$, so U and V are independent. Conditioning on a collider opens the path, so U and V are in general dependent given W.
6. Check with a counterexample. Let U, V be independent fair bits and let $W = U \oplus V$, which is allowed because $P(W \mid U, V)$ is an arbitrary conditional distribution. Given $W = 0$, we have $U = V$, so $P(U = 1, V = 1 \mid W = 0) = \frac{1}{2}$. But $P(U = 1 \mid W = 0)\, P(V = 1 \mid W = 0) = \frac{1}{2} \cdot \frac{1}{2} = \frac{1}{4}$. These differ, so $U \not\perp V \mid W$.
7. In general $P(U, V \mid W) = \dfrac{P(U)P(V)P(W \mid U, V)}{P(W)}$, which factors into a function of U times a function of V only if $P(W \mid U, V)$ itself factors. The network does not guarantee this, so statement (C) is the false one.

KEY IDEA

Read the factorisation as a Bayesian network and apply d-separation. Chains and forks are blocked by conditioning on the middle node, but a collider such as $U \to W \leftarrow V$ is opened by conditioning on W (or on a descendant of it). Marginally independent parents therefore become dependent given the child.

WATCH OUT FOR

Assuming that because U and V are independent a priori (the factorisation has $P(U)P(V)$), they stay independent once W is observed. This ignores the explaining-away effect at a collider.

Q.25 MCQ · 1 MARK

Artificial Intelligence › Search

Consider the following statement:

In adversarial search, α-β pruning can be applied to game trees of any depth where α is the (m) value choice we have formed so far at any choice point along the path for the MAX player and β is the (n) value choice we have formed so far at any choice point along the path for the MIN player.

Which ONE of the following choices of (m) and (n) makes the above statement valid?

(A) (m) = highest, (n) = highest

(B) (m) = lowest, (n) = highest

(C) (m) = highest, (n) = lowest

(D) (m) = lowest, (n) = lowest

Official answer: **C**

SOLUTION

1. In alpha-beta pruning, α is the value of the best choice found so far along the path for MAX. MAX wants the largest value, so the best choice is the **highest** value. Hence α can only increase as the search proceeds, and (m) = highest.
2. β is the value of the best choice found so far along the path for MIN. MIN wants the smallest value, so the best choice is the **lowest** value. Hence β can only decrease as the search proceeds, and (n) = lowest.
3. Check against the pruning rule. At a MIN node, once a child value is $\leq \alpha$, MAX already has a better option elsewhere (a value of at least α) and will never choose this branch, so the remaining children are pruned. At a MAX node, once a child value is $\geq \beta$, MIN already has an option no larger than β and will never allow this branch, so the remaining children are pruned. In both cases pruning happens exactly when $\alpha \geq \beta$. This only makes sense if α is the highest value MAX can guarantee and β is the lowest value MIN can force.
4. Option (A) makes β a highest value, which is not what MIN seeks. Option (B) swaps the two roles. Option (D) makes α a lowest value, which is not what MAX seeks. Only (m) = highest, (n) = lowest describes each player's best guarantee, so the answer is (C).

KEY IDEA

α is the highest value MAX can already guarantee along the path, and β is the lowest value MIN can already force. A branch is pruned as soon as $\alpha \geq \beta$.

WATCH OUT FOR

Reading "best choice for MIN" as the highest value. Best for MIN is the lowest value, so β tracks the lowest value and only decreases. Mixing up the two bounds also flips which comparison ($\leq \alpha$ at MIN nodes, $\geq \beta$ at MAX nodes) triggers pruning.

Q.26 MCQ · 1 MARK

Database Management and Warehousing › Relational algebra

Consider a database that includes the following relations:

Defender(name, rating, side, goals)
Forward(name, rating, assists, goals)
Team(name, club, price)

Which ONE of the following relational algebra expressions checks that every name occurring in Team appears in either Defender or Forward, where ϕ denotes the empty set?

(A) $\pi_{\text{name}}(\text{Team}) \setminus (\pi_{\text{name}}(\text{Defender}) \cap \pi_{\text{name}}(\text{Forward})) = \phi$

(B) $(\pi_{\text{name}}(\text{Defender}) \cap \pi_{\text{name}}(\text{Forward})) \setminus \pi_{\text{name}}(\text{Team}) = \phi$

(C) $\pi_{\text{name}}(\text{Team}) \setminus (\pi_{\text{name}}(\text{Defender}) \cup \pi_{\text{name}}(\text{Forward})) = \phi$

(D) $(\pi_{\text{name}}(\text{Defender}) \cup \pi_{\text{name}}(\text{Forward})) \setminus \pi_{\text{name}}(\text{Team}) = \phi$

Official answer: **C**

SOLUTION

1. The condition is a subset test: every name in $\pi_{\text{name}}(\text{Team})$ must belong to $\pi_{\text{name}}(\text{Defender}) \cup \pi_{\text{name}}(\text{Forward})$.
2. In relational algebra, $R \subseteq S$ holds exactly when $R \setminus S = \phi$, because the difference keeps only the tuples of R that are absent from S.
3. "Appears in either Defender or Forward" means membership in at least one of the two relations. That is membership in the union, so $S = \pi_{\text{name}}(\text{Defender}) \cup \pi_{\text{name}}(\text{Forward})$.

4. The test is therefore $\pi_{\text{name}}(\text{Team}) \setminus (\pi_{\text{name}}(\text{Defender}) \cup \pi_{\text{name}}(\text{Forward})) = \phi$, which is option (C).
5. Option (A) subtracts the intersection. It is empty only when every Team name is in both Defender and Forward, which is stricter than required. For example, a Team name present only in Defender leaves a non-empty result, although the condition holds.
6. Options (B) and (D) put Team on the right of the difference. They test whether every name in Defender and Forward (the intersection in (B), the union in (D)) appears in Team. That is the reverse direction and says nothing about the Team names.

KEY IDEA
$R \subseteq S$ is checked as $R \setminus S = \phi$, with the relation that must be contained on the left. "Either ... or" means the union.

WATCH OUT FOR
Using the intersection for "either or" (option A), which demands membership in both relations. Or reversing the difference (options B and D), which tests whether Defender/Forward names are in Team.

Q.27 MCQ · 1 MARK

Database Management and Warehousing › Data transformation

Let the minimum, maximum, mean and standard deviation values for the attribute income of data scientists be 46000, 170000, 96000, and 21000, respectively. The z-score normalized income value of 106000 is closest to which ONE of the following options?

(A) 0.217
(B) 0.476
(C) 0.623
(D) 2.304

Official answer: **B**

SOLUTION

1. Z-score normalization rescales a value using the attribute's mean and standard deviation: $z = \dfrac{x - \mu}{\sigma}$.
2. The minimum (46000) and maximum (170000) are not needed. They belong to min-max normalization, so here they are distractors.
3. Substitute $x = 106000$, $\mu = 96000$, $\sigma = 21000$:
$$z = \frac{106000 - 96000}{21000} = \frac{10000}{21000} = \frac{10}{21}.$$
4. Compute: $\frac{10}{21} \approx 0.47619$, which is closest to 0.476, option (B).
5. Check against the distractor: min-max normalization would give $\frac{106000-46000}{170000-46000} = \frac{60000}{124000} \approx 0.484$. That is a different quantity and matches no option exactly, so the z-score value 0.476 is the answer. The other options do not fit: 0.217, 0.623 and 2.304 are far from $\frac{10}{21}$.

KEY IDEA
Z-score normalization uses only the mean and standard deviation: z = (x − mean) / standard deviation. Minimum and maximum are used by min-max normalization, not here.

WATCH OUT FOR
Using the minimum and maximum (min-max scaling) gives about 0.484, which is near 0.476 and can tempt you to pick the wrong method. Check that the question asks for the z-score.

Q.28 MSQ · 1 MARK

Data Structures and Algorithms › Trees & heaps

Consider the following tree traversals on a full binary tree:

(i) Preorder
(ii) Inorder
(iii) Postorder

Which of the following traversal options is/are sufficient to uniquely reconstruct the full binary tree?

(A) (i) and (ii)

(B) (ii) and (iii)

(C) (i) and (iii)

(D) (ii) only

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

SOLUTION

1. Assume, as is standard for reconstruction questions, that every node carries a distinct label. A full binary tree is one in which every node has either 0 or 2 children, so no node has exactly one child.
2. **Preorder + inorder (i, ii):** the first preorder element is the root. Its position in the inorder sequence splits the remaining nodes into the left and right subtrees, and the subtree sizes tell us which preorder elements belong to each. Recursing on both parts rebuilds the tree uniquely. This works for any binary tree, so (A) is correct.
3. **Inorder + postorder (ii, iii):** the last postorder element is the root. Its position in the inorder sequence splits the nodes into the left and right subtrees, and the left size tells us which postorder elements belong to each. Recursing rebuilds the tree uniquely, for any binary tree. So (B) is correct.
4. **Preorder + postorder (i, iii):** for a general binary tree this pair is ambiguous, because a node with exactly one child cannot be classed as a left or a right child. For example, a root a with only a left child b and a root a with only a right child b both give preorder a, b and postorder b, a. In a full binary tree no node has exactly one child, so this ambiguity cannot arise.
5. Reconstruction from this pair: the root is the first preorder element and also the last postorder element. If the tree has more than one node, the root has two children, and the second preorder element is the root of the left subtree. In the postorder sequence, the left subtree occupies the positions up to and including that element, because a subtree's root comes last in its own postorder block. The remaining positions, apart from the final root, form the right subtree. Recurse on each part. Every step is forced, so the tree is unique. So (C) is correct.
6. **Inorder only (ii):** it is not sufficient. Take the 5-node full trees. Tree 1: root 4, left child 2 with leaves 1 and 3, right child a leaf 5. Tree 2: root 2, left child a leaf 1, right child 4 with leaves 3 and 5. Both are full binary trees, and both have inorder sequence $1, 2, 3, 4, 5$, yet they are different trees. So the inorder sequence alone does not determine a full binary tree, and (D) is wrong.
7. Hence (A), (B) and (C) are each sufficient and (D) is not, so the answer is A, B, C.

KEY IDEA
Inorder together with either preorder or postorder always rebuilds a binary tree. Preorder with postorder fails only because of nodes with a single child, and a full binary tree has none, so it works there too.

WATCH OUT FOR
Remembering the rule that preorder plus postorder cannot rebuild a tree and so rejecting (C). That rule applies to general binary trees, and the question specifies a full binary tree.

Q.29 MSQ · 1 MARK

Artificial Intelligence › Logic

Let x and y be two propositions. Which of the following statements is a tautology/are tautologies?

(A) $(\neg x \wedge y) \Rightarrow (y \Rightarrow x)$

(B) $(x \wedge \neg y) \Rightarrow (\neg x \Rightarrow y)$

(C) $(\neg x \wedge y) \Rightarrow (\neg x \Rightarrow y)$

(D) $(x \wedge \neg y) \Rightarrow (y \Rightarrow x)$

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

SOLUTION

1. An implication $P \Rightarrow Q$ is a tautology exactly when no assignment makes P true and Q false. So for each option, assume the antecedent is true and check whether the consequent can be false.
2. (A) The antecedent $\neg x \wedge y$ is true only for $x = F,\ y = T$. Then $y \Rightarrow x = T \Rightarrow F = F$. The antecedent is true and the consequent false, so (A) is not a tautology.
3. (B) The antecedent $x \wedge \neg y$ is true only for $x = T,\ y = F$. Then $\neg x = F$, so $\neg x \Rightarrow y$ has a false premise and is true. No falsifying assignment exists, so (B) is a tautology.
4. (C) The antecedent $\neg x \wedge y$ is true only for $x = F,\ y = T$. Then $\neg x \Rightarrow y = T \Rightarrow T = T$. No falsifying assignment exists, so (C) is a tautology. Equivalently, the consequent is $x \vee y$, which follows from y.
5. (D) The antecedent $x \wedge \neg y$ is true only for $x = T,\ y = F$. Then $y \Rightarrow x = F \Rightarrow T = T$. In fact $y \Rightarrow x$ is true whenever x is true. No falsifying assignment exists, so (D) is a tautology.
6. Only (A) has a falsifying assignment, so the tautologies are (B), (C) and (D).

KEY IDEA
An implication fails only when its antecedent is true and its consequent false. So make the antecedent true, which here fixes both x and y, and test the consequent at that single assignment.

WATCH OUT FOR
Writing $y \Rightarrow x$ when the converse $x \Rightarrow y$ is meant, or the reverse. In (A) the antecedent forces $y = T$ and $x = F$, which makes $y \Rightarrow x$ false. In (D) the antecedent forces $x = T$, which makes $y \Rightarrow x$ true.

Q.30 NAT · 1 MARK

Data Structures and Algorithms › Design techniques

Q.30 Consider sorting the following array of integers in ascending order using an in-place Quicksort algorithm that uses the last element as the pivot.

60	70	80	90	100

The minimum number of swaps performed during this Quicksort is _____.

Official answer: **0**

SOLUTION

1. The array $[60, 70, 80, 90, 100]$ is already in ascending order, and the pivot is always the last element of the current subarray.
2. A swap exchanges two elements at different positions. Partitioning only has to put every element $\leq$ pivot to the left of the pivot and every element $>$ pivot to its right, and then place the pivot at its boundary.
3. First partition, pivot 100: it is the maximum, so all four other elements are ≤ 100 and already lie to its left. No element is on the wrong side, so no two distinct positions need to be exchanged. The pivot's boundary position is index 4, which is where it already sits, so even the final pivot placement moves nothing.
4. The problem splits into $[60, 70, 80, 90]$ and an empty right part. In $[60, 70, 80, 90]$ the pivot 90 is again the maximum, so the same argument applies: no swap. Likewise pivot 80 in $[60, 70, 80]$ and pivot 70 in

$[60, 70]$.

5. The last subarray $[60]$ has one element and needs no work.
6. Every partition leaves its pivot where it already is and moves nothing, so no two distinct elements are ever exchanged. The minimum number of swaps is 0.
7. Note: a textbook Lomuto loop that blindly executes a swap even when the two indices coincide would perform swaps of an element with itself. Those are not exchanges of two elements, and the question asks for the minimum, which is attained when swaps are made only between distinct positions.

KEY IDEA

With the last element as pivot, an already sorted array makes every pivot the maximum of its subarray, so every element is already on the correct side of it. Nothing needs to be exchanged, so the number of real swaps is zero (the sorted input is the best case for swaps, though the worst case for the recursion depth and comparisons).

WATCH OUT FOR

Counting the self-swaps of a blind Lomuto loop (4+3+2+1 element-with-itself exchanges, plus the final pivot placement) and answering 10 or 14. Those do not exchange two distinct elements, and the question asks for the minimum. Another slip is to treat sorted input as the worst case and expect many swaps; it is the worst case for comparisons and depth, not for real swaps.

Q.31 NAT · 1 MARK

Database Management and Warehousing › SQL

Q.31 Consider the following two tables named Raider and Team in a relational database maintained by a Kabaddi league. The attribute *ID* in table Team references the primary key of the Raider table, *ID*.

Raider

ID	*Name*	*Raids*	*RaidPoints*
1	Arjun	200	250
2	Ankush	190	219
3	Sunil	150	200
4	Reza	150	190
5	Pratham	175	220
6	Gopal	193	215

Team

City	*ID*	*BidPoints*
Jaipur	2	200
Patna	3	195
Hyderabad	5	175
Jaipur	1	250
Patna	4	200
Jaipur	6	200

The SQL query described below is executed on this database:

```
SELECT *
FROM Raider, Team
WHERE Raider.ID=Team.ID AND City="Jaipur" AND
RaidPoints > 200;
```

The number of rows returned by this query is _____.

Official answer: **3**

SOLUTION

1. The query lists FROM Raider, Team with Raider.ID=Team.ID , so it is an equi-join. Every Team.ID (2, 3, 5, 1, 4, 6) matches exactly one Raider.ID , because Team.ID references the primary key of Raider . The join therefore has 6 rows, one per Team row.
2. Apply City="Jaipur" . The Team rows in Jaipur are those with ID 2, 1 and 6. The Patna rows (ID 3 and 4) and the Hyderabad row (ID 5) are removed.
3. Apply RaidPoints > 200 to the matching Raider rows. ID 1 (Arjun) has 250, ID 2 (Ankush) has 219 and ID 6 (Gopal) has 215. All three are greater than 200, so all three stay.
4. The result has 3 rows. The strict comparison > would matter only for a row with exactly 200 raid points. That is Sunil (ID 3), who is in the Patna team and is already removed by the city condition.

KEY IDEA
A comma join with an equality condition on the key is an equi-join. Apply each filter to the joined rows, and count the rows that satisfy all three conditions together.

WATCH OUT FOR
Counting the three Jaipur rows without checking `RaidPoints`, or treating `RaidPoints > 200` as `>= 200` and then wrongly including Sunil. Sunil belongs to Patna, so he is excluded either way.

Q.32 NAT · 1 MARK

Data Structures and Algorithms › Linear structures

The fundamental operations in a double-ended queue D are:

insertFirst(e) – Insert a new element e at the beginning of D.
insertLast(e) – Insert a new element e at the end of D.
removeFirst() – Remove and return the first element of D.
removeLast() – Remove and return the last element of D.

In an empty double-ended queue, the following operations are performed:

```
insertFirst(10)
insertLast(32)
a ← removeFirst()
insertLast(28)
insertLast(17)
a ← removeFirst()
a ← removeLast()
```

The value of a is _____.

Official answer: **17**

SOLUTION

1. Start with the empty deque: D = [] .
2. insertFirst(10) puts 10 at the front, so D = [10] .
3. insertLast(32) appends at the back, so D = [10, 32] .
4. a ← removeFirst() removes and returns the front element, so a = 10 and D = [32] .
5. insertLast(28) gives D = [32, 28] , and insertLast(17) gives D = [32, 28, 17] .
6. a ← removeFirst() removes the front element, so a = 32 and D = [28, 17] .
7. a ← removeLast() removes the back element, so a = 17 and D = [28] .
8. Each assignment overwrites a , so the final value of a is the one returned by the last operation, which is 17.

KEY IDEA
Trace the deque operation by operation. `a` is overwritten at every assignment, so only the last removal (`removeLast()`) decides the final value.

WATCH OUT FOR
Reporting the value from the first or second removal (10 or 32), or treating `insertFirst` as an append at the back. Another slip is to forget that `a` is reassigned, so only the last removal counts.

Q.33 NAT · 1 MARK

Calculus and Optimization › Limits & continuity

Let $f : \mathbb{R} \to \mathbb{R}$ be the function $f(x) = \dfrac{1}{1+e^{-x}}$. The value of the derivative of f at x where $f(x) = 0.4$ is

_____ (rounded off to two decimal places). Note: $\mathbb{R}$ denotes the set of real numbers.

Official answer: **0.24**

SOLUTION

1. Recognise $f(x) = \dfrac{1}{1+e^{-x}}$ as the logistic sigmoid, defined for every real x and taking values in $(0, 1)$, so the value $f(x) = 0.4$ is attained at exactly one x.
2. Differentiate with the chain rule. Write $f(x) = (1+e^{-x})^{-1}$, so $f'(x) = -(1+e^{-x})^{-2} \cdot (-e^{-x}) = \dfrac{e^{-x}}{(1+e^{-x})^2}$.
3. Split the fraction: $f'(x) = \dfrac{1}{1+e^{-x}} \cdot \dfrac{e^{-x}}{1+e^{-x}}$.
4. Since $\dfrac{e^{-x}}{1+e^{-x}} = \dfrac{(1+e^{-x})-1}{1+e^{-x}} = 1 - \dfrac{1}{1+e^{-x}}$, this gives $f'(x) = f(x)\big(1-f(x)\big)$.
5. The derivative depends only on the value $f(x)$, so there is no need to solve for x (it is $x = \ln(2/3)$, which is not required).
6. At $f(x) = 0.4$: $f'(x) = 0.4 \times 0.6 = 0.24$.

KEY IDEA

For the logistic sigmoid, $f'(x) = f(x)\big(1-f(x)\big)$, so the derivative can be read straight from the function value without finding x.

WATCH OUT FOR

Solving $\dfrac{1}{1+e^{-x}} = 0.4$ for x and then substituting into the derivative is slower and invites rounding errors. Another slip is dropping the minus sign from the chain rule on e^{-x}, or forgetting the square in the denominator.

Q.34 NAT · 1 MARK

Probability and Statistics › Descriptive statistics

The sample average of 50 data points is 40. The updated sample average after including a new data point taking the value of 142 is _____.

Official answer: **42**

SOLUTION

1. The sample average of 50 points is 40, so the sum of the 50 points is $50 \times 40 = 2000$.
2. Adding the new point 142 gives a new sum of $2000 + 142 = 2142$ over 51 points.
3. The updated mean is $\bar{x}_{\text{new}} = \dfrac{2142}{51}$.
4. Since $51 \times 42 = 2142$, the updated mean is 42.
5. Check with the update formula $\bar{x}_{\text{new}} = \bar{x} + \dfrac{x_{\text{new}} - \bar{x}}{n+1} = 40 + \dfrac{142-40}{51} = 40 + \dfrac{102}{51} = 42$.

KEY IDEA

A mean is the total divided by the count, so recover the total ($n\bar{x}$), add the new value, and divide by $n+1$. Equivalently, the mean moves by $\dfrac{x_{\text{new}} - \bar{x}}{n+1}$.

WATCH OUT FOR

Dividing the new sum by 50 instead of 51, or averaging the old mean with the new value as $(40+142)/2 = 91$. Both ignore that the old mean carries the weight of 50 points.

Q.35 NAT · 1 MARK

Linear Algebra › Eigenvalues

Consider the 3×3 matrix $\mathbf{M} = \begin{bmatrix} 1 & 2 & 3 \\ 3 & 1 & 3 \\ 4 & 3 & 6 \end{bmatrix}$. The determinant of $(\mathbf{M}^2 + 12\mathbf{M})$ is _____.

Official answer: **0**

SOLUTION

1. Factor the matrix polynomial: $\mathbf{M}^2 + 12\mathbf{M} = \mathbf{M}(\mathbf{M} + 12\mathbf{I})$.
2. The determinant is multiplicative, so $\det(\mathbf{M}^2 + 12\mathbf{M}) = \det(\mathbf{M})\ \det(\mathbf{M} + 12\mathbf{I})$.
3. Expand $\det(\mathbf{M})$ along the first row: $1(1 \cdot 6 - 3 \cdot 3) - 2(3 \cdot 6 - 3 \cdot 4) + 3(3 \cdot 3 - 1 \cdot 4)$.
4. This equals $1(-3) - 2(6) + 3(5) = -3 - 12 + 15 = 0$.
5. So $\mathbf{M}$ is singular, and 0 is an eigenvalue of $\mathbf{M}$. Then $\mathbf{M}^2 + 12\mathbf{M}$ has the eigenvalue $\lambda^2 + 12\lambda = 0$ for $\lambda = 0$, which agrees with the factor $\det(\mathbf{M}) = 0$ in step 2.
6. Hence $\det(\mathbf{M}^2 + 12\mathbf{M}) = 0 \cdot \det(\mathbf{M} + 12\mathbf{I}) = 0$. The second factor is a finite number, so it does not need to be computed.

KEY IDEA

Factor $\mathbf{M}^2 + 12\mathbf{M} = \mathbf{M}(\mathbf{M} + 12\mathbf{I})$. If $\det(\mathbf{M}) = 0$, the whole product has determinant 0, so the second factor never has to be computed.

WATCH OUT FOR

Computing $\mathbf{M}^2 + 12\mathbf{M}$ entry by entry and then taking its determinant. This is long and error-prone. Another slip is to check only whether the trace or an individual entry is zero, instead of checking $\det(\mathbf{M})$.

Q.36 MCQ · 2 MARKS

Probability and Statistics › Random variables

A fair six-sided die (with faces numbered 1, 2, 3, 4, 5, 6) is repeatedly thrown independently.

What is the expected number of times the die is thrown until **two** consecutive throws of even numbers are seen?

(A) 2

(B) 4

(C) 6

(D) 8

Official answer: **C**

SOLUTION

1. Each throw is even (2, 4 or 6) with probability $p = \frac{1}{2}$ and odd with probability $\frac{1}{2}$, independently of the other throws. We wait for two even throws in a row.
2. Define two states. E_0 is the expected number of further throws when the last throw was not even (or at the start). E_1 is the expected number of further throws when the last throw was even.
3. From state 0 one throw is made. With probability $\frac{1}{2}$ it is even and we move to state 1. With probability $\frac{1}{2}$ it is odd and we stay in state 0:
$$E_0 = 1 + \tfrac{1}{2}E_1 + \tfrac{1}{2}E_0.$$
4. From state 1 one throw is made. With probability $\frac{1}{2}$ it is even and we stop. With probability $\frac{1}{2}$ it is odd and we fall back to state 0:
$$E_1 = 1 + \tfrac{1}{2} \cdot 0 + \tfrac{1}{2}E_0.$$
5. Substitute E_1 into the first equation:
$$E_0 = 1 + \tfrac{1}{2}\left(1 + \tfrac{1}{2}E_0\right) + \tfrac{1}{2}E_0 = \tfrac{3}{2} + \tfrac{3}{4}E_0.$$

6. So $\frac{1}{4}E_0 = \frac{3}{2}$, which gives $E_0 = 6$ (and $E_1 = 4$). The die starts in state 0, so the expected number of throws is 6.
7. Check with the standard result for r consecutive successes with success probability p:

$$E = \frac{1-p^r}{p^r(1-p)} = \frac{1-\frac{1}{4}}{\frac{1}{4}\cdot\frac{1}{2}} = \frac{3/4}{1/8} = 6.$$

This agrees, so the answer is option (C).

KEY IDEA
Model the wait as a two-state Markov chain (last throw even or not) and solve the two first-step equations for the expected remaining throws.

WATCH OUT FOR
Answering $\frac{1}{p^2} = 4$ by treating pairs of throws as independent trials. An odd throw after one even throw wipes out the progress, and overlapping runs are allowed, so the true expectation is larger, at 6.

Q.37 MCQ · 2 MARKS

Calculus and Optimization › Limits & continuity

Let $f : \mathbb{R} \to \mathbb{R}$ be a function. Note: $\mathbb{R}$ denotes the set of real numbers.

$$f(x) = \begin{cases} -x, & \text{if } x < -2 \\ ax^2 + bx + c, & \text{if } x \in [-2, 2] \\ x, & \text{if } x > 2 \end{cases}$$

Which ONE of the following choices gives the values of a, b, c that make the function f continuous and differentiable?

(A) $a = \frac{1}{4},\ b = 0,\ c = 1$
(B) $a = \frac{1}{2},\ b = 0,\ c = 0$
(C) $a = 0,\ b = 0,\ c = 0$
(D) $a = 1,\ b = 1,\ c = -4$

Official answer: **A**

SOLUTION

1. The outer pieces are $-x$ for $x < -2$ and x for $x > 2$. Both are differentiable, so only the joins at $x = \pm 2$ need checking. The middle piece is $q(x) = ax^2 + bx + c$, with $q'(x) = 2ax + b$.
2. Continuity at $x = -2$: the left limit is $-(-2) = 2$, so $q(-2) = 4a - 2b + c = 2$.
3. Continuity at $x = 2$: the right limit is 2, so $q(2) = 4a + 2b + c = 2$.
4. Subtracting the first equation from the second gives $4b = 0$, so $b = 0$. Then $4a + c = 2$.
5. Differentiability at $x = -2$: the left derivative is $\frac{d}{dx}(-x) = -1$ and the right derivative is $q'(-2) = -4a + b$. Setting them equal gives $-4a + b = -1$. With $b = 0$ this gives $a = \frac{1}{4}$.
6. Differentiability at $x = 2$: the right derivative is $\frac{d}{dx}(x) = 1$ and the left derivative is $q'(2) = 4a + b$. Setting them equal gives $4a + b = 1$. With $b = 0$ this again gives $a = \frac{1}{4}$, so the two conditions agree.
7. From $4a + c = 2$ with $a = \frac{1}{4}$: $c = 2 - 1 = 1$.
8. Hence $a = \frac{1}{4},\ b = 0,\ c = 1$, which is option (A). Check: $q(x) = \frac{1}{4}x^2 + 1$ gives $q(\pm 2) = 2$, matching the outer pieces. Its slope $q'(x) = \frac{x}{2}$ is -1 at $x = -2$ and $+1$ at $x = 2$, matching the slopes of $-x$ and x.
9. Options (B), (C) and (D) fail. (B) has $a = \frac{1}{2}$, which violates $a = \frac{1}{4}$. (C) gives $q \equiv 0$, so $q(2) = 0 \neq 2$. (D) gives $q(-2) = 4 - 2 - 4 = -2 \neq 2$.

KEY IDEA

Differentiable implies continuous. At each junction, match both the function value and the derivative of the middle quadratic with those of the outer line. This gives four linear equations in a, b, c (two of them consistent).

WATCH OUT FOR

Imposing only continuity. Continuity alone gives $b = 0$ and $4a + c = 2$, which leaves a whole family of solutions, and options (B) and (C) are tempting for that reason. The derivative conditions fix $a = \frac{1}{4}$.

Q.38 MCQ · 2 MARKS

Data Structures and Algorithms › Python

Consider the following Python code:

```
def count(child_dict, i):
  if i not in child_dict.keys():
     return 1
  ans = 1
  for j in child_dict[i]:
     ans += count(child_dict, j)
  return ans

child_dict = dict()
child_dict[0] = [1,2]
child_dict[1] = [3,4,5]
child_dict[2] = [6,7,8]
print(count(child_dict,0))
```

Which ONE of the following is the output of this code?

(A) 6

(B) 1

(C) 8

(D) 9

Official answer: **D**

SOLUTION

1. The dictionary stores a tree as child lists: key `0` has children `1` and `2`, key `1` has children `3, 4, 5`, and key `2` has children `6, 7, 8`. Keys `3` to `8` are not in the dictionary, so they are leaves.
2. `count(child_dict, i)` returns `1` at once when `i` is not a key. Otherwise it starts with `ans = 1` (the node itself) and adds `count(child_dict, j)` for every child `j`. So it returns the number of nodes in the subtree rooted at `i`.
3. For each leaf `3, 4, 5, 6, 7, 8`, the test `i not in child_dict.keys()` is true, so `count` returns `1`.
4. `count(child_dict, 1)`: `ans = 1`, then `ans += 1` three times (children `3, 4, 5`), so it returns $1 + 3 = 4$.
5. `count(child_dict, 2)`: `ans = 1`, then `ans += 1` three times (children `6, 7, 8`), so it returns $1 + 3 = 4$.
6. `count(child_dict, 0)`: `ans = 1`, then `ans += count(child_dict, 1)` gives $1 + 4 = 5$, then `ans += count(child_dict, 2)` gives $5 + 4 = 9$. The function returns `9`, and `print` shows `9`. This equals the number of nodes, `0` to `8`, in the tree. The answer is option (D).

KEY IDEA

The recursion counts the nodes of a tree stored as a dictionary of child lists. Each call returns 1 for itself plus the counts of its children, and a key missing from the dictionary is a leaf that returns 1. The output is 1 + (1 + 3) + (1 + 3) = 9.

WATCH OUT FOR

Counting only the keys of the dictionary (3) or only the leaves (6), or forgetting that a missing key still returns 1 and so counts as a node. Another slip is returning 0 for leaves, which would give 3 instead of 9.

Q.39 MCQ · 2 MARKS

Data Structures and Algorithms › Python

Consider the function `computeS(X)` whose pseudocode is given below:

```
computeS(X)
S[1] ← 1
for i ← 2 to length(X)
  S[i] ← 1
  if X[i-1] ≤ X[i]
      S[i] ← S[i] + S[i-1]
  end if
end for
return S
```

Which ONE of the following values is returned by the function `computeS(X)` for X = [6, 3, 5, 4, 10]?

(A) [1, 1, 2, 3, 4]

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

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

(D) [1, 1, 2, 1, 5]

Official answer: **C**

SOLUTION

1. The pseudo-code is:

```
computeS(X)
S[1] ← 1
for i ← 2 to length(X)
  S[i] ← 1
  if X[i-1] ≤ X[i]
      S[i] ← S[i] + S[i-1]
  end if
end for
return S
```

   For each position $i \geq 2$, `S[i]` is first set to 1. It grows by `S[i-1]` only when `X[i-1] <= X[i]`. So `S[i]` is the length of the run of non-decreasing consecutive elements that ends at position i. It restarts at 1 whenever `X[i-1] > X[i]`.

2. Input: $X = [6, 3, 5, 4, 10]$ (1-indexed), so `length(X) = 5`. Start with $S[1] = 1$.
3. $i = 2$: `S[2] = 1`. The test `X[1] <= X[2]` is $6 \leq 3$, which is false. So $S[2] = 1$.
4. $i = 3$: `S[3] = 1`. The test `X[2] <= X[3]` is $3 \leq 5$, which is true. So $S[3] = 1 + S[2] = 1 + 1 = 2$.
5. $i = 4$: `S[4] = 1`. The test `X[3] <= X[4]` is $5 \leq 4$, which is false. So $S[4] = 1$.
6. $i = 5$: `S[5] = 1`. The test `X[4] <= X[5]` is $4 \leq 10$, which is true. So $S[5] = 1 + S[4] = 1 + 1 = 2$.
7. The function returns $S = [1, 1, 2, 1, 2]$, which is option (C).

KEY IDEA

`S[i]` is the length of the non-decreasing run ending at position `i`. The comparison is `X[i-1] <= X[i]`, and `S[i]` is reset to 1 whenever it fails. The value added is `S[i-1]`, which is the previous position's already-computed run length.

WATCH OUT FOR

Using `S[i-1] + 1` unconditionally, or counting a run only when the sequence strictly increases, gives the wrong array. Another slip is forgetting that `S[i]` restarts at 1 after every descent (6 to 3, and 5 to 4). That slip leads to the cumulative-looking distractors in options (A), (B) and (D).

Q.40 MCQ · 2 MARKS

Data Structures and Algorithms › Sorting & hashing

Let $F(n)$ denote the maximum number of comparisons made while searching for an entry in a sorted array of size n using binary search.

Which ONE of the following options is TRUE?

(A) $F(n) = F(\lfloor n/2 \rfloor) + 1$

(B) $F(n) = F(\lfloor n/2 \rfloor) + F(\lceil n/2 \rceil)$

(C) $F(n) = F(\lfloor n/2 \rfloor)$

(D) $F(n) = F(n-1) + 1$

Official answer: **A**

SOLUTION

1. Binary search compares the key with the middle element once. If it is not a match, the search continues in exactly one of the two remaining parts, never both.
2. With the middle element removed from an array of size n, the two parts have sizes $\lfloor (n-1)/2 \rfloor$ and $\lfloor n/2 \rfloor$. The larger part therefore has exactly $\lfloor n/2 \rfloor$ elements. The worst case is when the key lies in the larger part (or is absent).
3. So the worst case is one comparison plus the worst case on $\lfloor n/2 \rfloor$ elements: $F(n) = F(\lfloor n/2 \rfloor) + 1$ with $F(1) = 1$. This solves to $F(n) = \lfloor \log_2 n \rfloor + 1$. Values: $F(1) = 1$, $F(2) = 2$, $F(3) = 2$, $F(4) = 3$, $F(5) = 3$.
4. Check (B): $F(n) = F(\lfloor n/2 \rfloor) + F(\lceil n/2 \rceil)$ would add the costs of both halves, as an algorithm does when it processes both halves. It gives $F(2) = F(1) + F(1) = 2$, but $F(4) = F(2) + F(2) = 4$, whereas the true value is $F(4) = 3$. So (B) is false.
5. Check (C): $F(n) = F(\lfloor n/2 \rfloor)$ ignores the comparison at the middle. At $n = 2$ it says $F(2) = F(1) = 1$, but binary search on two elements can need 2 comparisons. So (C) is false.
6. Check (D): $F(n) = F(n-1) + 1$ would give $F(3) = F(2) + 1 = 3$, but the true value is $F(3) = 2$. It describes linear search, not binary search. So (D) is false.
7. Only (A) is true.

KEY IDEA

Binary search makes one comparison at the middle, then continues in only one part, and the larger part has $\lfloor n/2 \rfloor$ elements. The worst case therefore satisfies $F(n) = F(\lfloor n/2 \rfloor) + 1$, which gives $\lfloor \log_2 n \rfloor + 1$ comparisons.

WATCH OUT FOR

Choosing (B) by copying the merge sort recurrence, which adds the costs of both halves, or choosing (C) by forgetting that the middle comparison costs one step each time.

Q.41 MCQ · 2 MARKS

Data Structures and Algorithms › Python

Consider the following Python function:

```
def fun(D, s1, s2):
    if s1 < s2:
        D[s1], D[s2] = D[s2], D[s1]
        fun(D, s1+1, s2-1)
```

What does this Python function `fun()` do? Select the ONE appropriate option below.

(A) It finds the smallest element in D from index s1 to s2, both inclusive.

(B) It performs a merge sort in-place on this list D between indices s1 and s2, both inclusive.

(C) It reverses the list D between indices s1 and s2, both inclusive.

(D) It swaps the elements in D at indices s1 and s2, and leaves the remaining elements unchanged.

Official answer: **C**

SOLUTION

1. The function is:

```
def fun(D, s1, s2):
    if s1 < s2:
        D[s1], D[s2] = D[s2], D[s1]
        fun(D, s1+1, s2-1)
```

   Both the swap and the recursive call sit inside the `if`, so when `s1 < s2` is false the function does nothing and returns. This is the base case, reached when the indices meet (`s1 == s2`) or cross (`s1 > s2`).

2. When `s1 < s2`, the tuple assignment `D[s1], D[s2] = D[s2], D[s1]` evaluates the right-hand side first and then assigns, so it exchanges the two end elements of the window with no temporary variable needed.
3. The call `fun(D, s1+1, s2-1)` moves both indices one step inward. It shrinks the window by two and swaps the next pair counted from the outside. The list `D` is mutable and passed by reference, so every swap changes the caller's list.
4. Trace with `D = [a, b, c, d, e]`, `s1 = 0`, `s2 = 4`. The pair (0, 4) gives `[e, b, c, d, a]`. The pair (1, 3) gives `[e, d, c, b, a]`. At (2, 2) the test `2 < 2` fails, so the call stops. The result is the reverse of the list.
5. Even length, `D = [a, b, c, d]` with `s1 = 0`, `s2 = 3`. The pair (0, 3) gives `[d, b, c, a]`. The pair (1, 2) gives `[d, c, b, a]`. At (2, 1) the test `2 < 1` fails and the call stops. This is again the full reverse.
6. General argument by induction on the window length $n = s2 - s1 + 1$. For $n \leq 1$ nothing happens, and a window of length 0 or 1 is its own reverse. For $n \geq 2$ the first swap puts the old last element at index `s1` and the old first element at index `s2`. The recursion then reverses the inner window `s1+1 .. s2-1`, which has length $n - 2$, by the induction hypothesis. Together this gives the reverse of the whole window `s1 .. s2`. Elements outside the window are never touched.
7. Nothing in the code compares elements, so it cannot find a minimum (A) or sort (B). It performs about $(n)/2$ swaps rather than a single swap, so (D) is wrong, since (D) would leave the interior unchanged. Only option (C) matches. The answer is C.

KEY IDEA

A recursive two-pointer swap: swap the ends of the window, then recurse on the window shrunk by one at each end, until the pointers meet or cross. That is in-place reversal of the segment.

WATCH OUT FOR

Reading the single swap `D[s1], D[s2] = D[s2], D[s1]` and picking (D) without noticing the recursive call repeats the swap on the inner pairs. Another error is thinking swaps imply sorting. No comparison between elements ever happens, so the result depends only on positions, never on values.

Q.42 MCQ · 2 MARKS

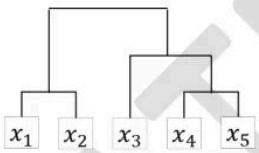
Machine Learning › Clustering

Q.42 Consider the table below, where the $(i,j)^{th}$ element of the table is the distance between points x_i and x_j. Single linkage clustering is performed on data points, x_1, x_2, x_3, x_4, x_5.

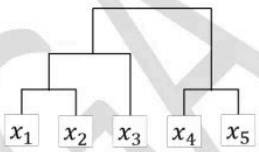
	x_1	x_2	x_3	x_4	x_5
x_1	0	1	4	3	6
x_2	1	0	3	5	3
x_3	4	3	0	2	5
x_4	3	5	2	0	1
x_5	6	3	5	1	0

Which **ONE** of the following is the correct representation of the clusters produced?

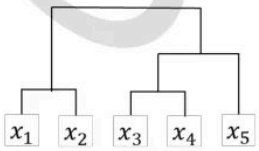
(A)



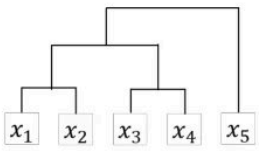
(B)



(C)



(D)



Official answer: **A**

SOLUTION

1. Single linkage merges the two clusters with the smallest minimum pairwise distance. The smallest entries in the table are $d(x_1, x_2) = 1$ and $d(x_4, x_5) = 1$.
2. Merge $\{x_1, x_2\}$ at height 1 and $\{x_4, x_5\}$ at height 1.
3. Recompute the cluster distances with the minimum over pairs: $d(\{x_1, x_2\}, x_3) = \min(4, 3) = 3$, $d(\{x_1, x_2\}, \{x_4, x_5\}) = \min(3, 6, 5, 3) = 3$, $d(x_3, \{x_4, x_5\}) = \min(2, 5) = 2$.

4. The smallest is 2, so x_3 joins $\{x_4, x_5\}$ at height 2, giving $\{x_3, x_4, x_5\}$.
5. Finally $d(\{x_1, x_2\}, \{x_3, x_4, x_5\}) = \min(4, 3, 3, 5, 6, 3)$ over all cross pairs $= 3$, so the two groups merge at height 3.
6. The tree is $((x_1, x_2), (x_3, (x_4, x_5)))$: a pair x_1, x_2, a pair x_4, x_5 with x_3 joining it, and the two groups joined at the top. This is option A.
7. Option B has x_3 joining $\{x_1, x_2\}$ while x_4, x_5 stay a separate pair, but x_3 is closer to $\{x_4, x_5\}$ (distance 2) than to $\{x_1, x_2\}$ (distance 3). Option C pairs x_3 with x_4 and then adds x_5, but x_4 and x_5 (distance 1) merge before x_3 can join either. Option D merges x_3, x_4 as a pair and leaves x_5 until last, but x_5 is closest to x_4 (distance 1).

KEY IDEA

In single linkage, the distance between two clusters is the smallest pairwise distance across them. Merge the closest pair repeatedly, recomputing cluster distances with a minimum.

WATCH OUT FOR

Merging x_3 with x_4 as a pair, because $d(x_3, x_4) = 2$ looks small. That ignores that x_4 and x_5 already merged at distance 1, so x_3 joins the cluster $\{x_4, x_5\}$ rather than x_4 alone.

Q.43 MCQ · 2 MARKS

Machine Learning › Neural networks

Q.43 Consider the two neural networks (NNs) shown in Figures 1 and 2, with $ReLU$ activation ($ReLU(z) = \max\{0, z\}, \forall z \in \mathbb{R}$). $\mathbb{R}$ denotes the set of real numbers. The connections and their corresponding weights are shown in the Figures. The biases at every neuron are set to 0. For what values of p, q, r in Figure 2 are the two NNs equivalent, when x_1, x_2, x_3 are positive?

Figure 1

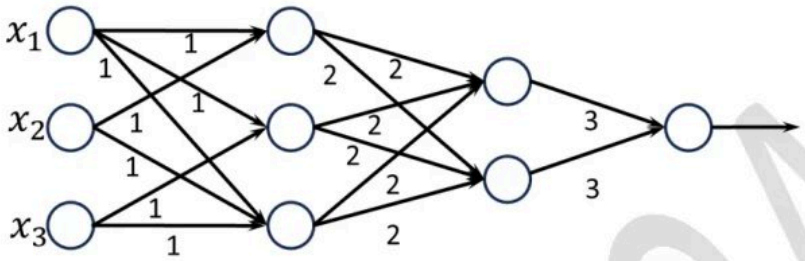
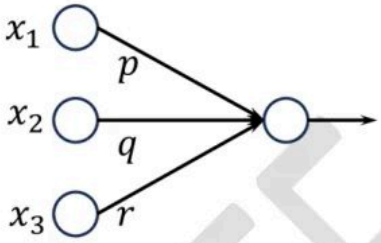


Figure 2



(A)	$p = 36, q = 24, r = 24$
(B)	$p = 24, q = 24, r = 36$
(C)	$p = 18, q = 36, r = 24$
(D)	$p = 36, q = 36, r = 36$

Official answer: **A**

SOLUTION

1. Read the first layer of Figure 1 edge by edge, since it is not fully connected. x_1 feeds all three first-layer neurons. x_2 feeds only the top and bottom neurons. x_3 feeds only the middle and bottom neurons. Every one of these edges has weight 1, and all biases are 0.
2. Pre-activations of the first layer: top $= x_1 + x_2$, middle $= x_1 + x_3$, bottom $= x_1 + x_2 + x_3$. All three are positive because $x_1, x_2, x_3 > 0$, so ReLU passes each one unchanged: $h_1 = x_1 + x_2$, $h_2 = x_1 + x_3$, $h_3 = x_1 + x_2 + x_3$.

3. Add them up: $h_1 + h_2 + h_3 = 3x_1 + 2x_2 + 2x_3$.
4. The second layer is fully connected with weight 2 on every edge, so each of its two neurons has pre-activation $2(h_1 + h_2 + h_3) = 6x_1 + 4x_2 + 4x_3 > 0$. ReLU leaves this unchanged: $g_1 = g_2 = 6x_1 + 4x_2 + 4x_3$.
5. The output neuron has weight 3 from each second-layer neuron:
$$3g_1 + 3g_2 = 6(6x_1 + 4x_2 + 4x_3) = 36x_1 + 24x_2 + 24x_3,$$
which is positive, so a ReLU at the output would change nothing.
6. Figure 2 computes $px_1 + qx_2 + rx_3$. For the two networks to agree for every positive x_1, x_2, x_3, the coefficients must match: $p = 36$, $q = 24$, $r = 24$. This is option (A).

KEY IDEA

With positive inputs, zero biases and positive weights, every pre-activation is positive. ReLU then acts as the identity, so Figure 1 collapses to a linear map. The coefficient of each input equals the sum, over all paths from that input to the output, of the product of the weights along the path. x_1 reaches 3 first-layer neurons and each first-layer neuron gives $2 \cdot 2 \cdot 3 = 12$ downstream, so $p = 3 \times 12 = 36$. x_2 and x_3 each reach only 2 first-layer neurons, so $q = r = 2 \times 12 = 24$.

WATCH OUT FOR

Assuming the first layer is fully connected (9 edges) because the crossing lines look dense. That gives $36(x_1 + x_2 + x_3)$ and option (D). Count the edges leaving each input: x_1 has 3, but x_2 and x_3 have only 2 each.

Q.44 MCQ · 2 MARKS

Artificial Intelligence › Search

Consider a state space where the start state is number 1. The successor function for the state numbered n returns two states numbered $n + 1$ and $n + 2$. Assume that the states in the unexpanded state list are expanded in the ascending order of numbers and the previously expanded states are not added to the unexpanded state list.

Which ONE of the following statements about breadth-first search (BFS) and depth-first search (DFS) is true, when reaching the goal state number 6?

(A) BFS expands more states than DFS.

(B) DFS expands more states than BFS.

(C) Both BFS and DFS expand equal number of states.

(D) Both BFS and DFS do not reach the goal state number 6.

Official answer: **C**

SOLUTION

1. The question fixes the order in which the unexpanded list is processed: always the smallest state number first. BFS normally takes the oldest entry (a queue) and DFS the newest (a stack), but here that choice is overridden by the ascending-number rule, so both searches pick the same state at every step.
2. Start: expand 1. Successors 2 and 3, so the unexpanded list is $\{2, 3\}$.
3. Expand 2 (smallest). Successors 3 and 4. State 3 is already waiting in the list, so the list is $\{3, 4\}$.
4. Expand 3. Successors 4 and 5. State 4 is already waiting, so the list is $\{4, 5\}$.
5. Expand 4. Successors 5 and 6, so the list is $\{5, 6\}$.
6. Expand 5. Successors 6 and 7, so the list is $\{6, 7\}$. The smallest state is now 6, the goal, and it is selected next.

7. The expansion sequence is $1, 2, 3, 4, 5$, then the goal 6 is reached. This sequence was produced using only the ascending-order rule, so it is exactly the same for BFS and for DFS. Even if a state were listed twice, expanding the smallest number first and never re-expanding an expanded state gives the same sequence for both.
8. Hence BFS and DFS expand the same states (1 to 5, five expansions) before reaching 6. Options (A) and (B) are false because neither expands more than the other, and (D) is false because both reach 6. The correct option is (C).

KEY IDEA
When the question dictates the expansion order (smallest number first), that rule overrides the usual queue or stack behaviour of BFS and DFS. With the same order, both searches expand the same states and reach the goal after the same number of expansions.

WATCH OUT FOR
Applying the textbook behaviour of BFS (level by level) and DFS (deepest first) and concluding they differ, without noticing that the question forces the ascending-number order on both. Another slip is to count the goal itself as an expansion for one search and not the other.

Q.45 MCQ · 2 MARKS

Data Structures and Algorithms › Sorting & hashing

Consider the following sorting algorithms:

(i) Bubble sort
(ii) Insertion sort
(iii) Selection sort

Which ONE among the following choices of sorting algorithms sorts the numbers in the array [4, 3, 2, 1, 5] in increasing order after exactly two passes over the array?

(A) (i) only
(B) (iii) only
(C) (i) and (iii) only
(D) (ii) and (iii) only

Official answer: **B**

SOLUTION

1. Trace each algorithm on $[4, 3, 2, 1, 5]$ and look at the array after exactly two passes. A pass of bubble sort is one left-to-right sweep of adjacent compares. A pass of insertion sort inserts the next element into the sorted prefix. A pass of selection sort puts the minimum of the unsorted part into its final place.
2. Bubble sort, pass 1: compare $(4, 3)$ and swap, giving $[3, 4, 2, 1, 5]$. Compare $(4, 2)$ and swap, giving $[3, 2, 4, 1, 5]$. Compare $(4, 1)$ and swap, giving $[3, 2, 1, 4, 5]$. Compare $(4, 5)$ and leave it. The array is $[3, 2, 1, 4, 5]$.
3. Bubble sort, pass 2: compare $(3, 2)$ and swap, giving $[2, 3, 1, 4, 5]$. Compare $(3, 1)$ and swap, giving $[2, 1, 3, 4, 5]$. Compare $(3, 4)$ and $(4, 5)$ and leave both. The array is $[2, 1, 3, 4, 5]$. The first two entries are still out of order, so bubble sort is not finished.
4. Insertion sort, pass 1: insert 3 into the sorted prefix $[4]$, giving $[3, 4, 2, 1, 5]$.
5. Insertion sort, pass 2: insert 2 into $[3, 4]$, giving $[2, 3, 4, 1, 5]$. The element 1 has not yet been inserted, so the array is not sorted.
6. Selection sort, pass 1: the minimum of the whole array is 1. Swap it with the first position, giving $[1, 3, 2, 4, 5]$.
7. Selection sort, pass 2: the minimum of the remaining part $[3, 2, 4, 5]$ is 2. Swap it with the second position, giving $[1, 2, 3, 4, 5]$. The array is sorted.

8. Only selection sort has the array fully sorted after exactly two passes, so the answer is (iii) only, which is option (B).

KEY IDEA
Simulate each algorithm for exactly two passes and compare with the sorted array. Selection sort fixes one element per pass in its final place. Here the two smallest elements, 1 and 2, are fixed in two passes, and the rest, $[3, 4, 5]$, already happens to be in order.

WATCH OUT FOR
Assuming that bubble sort or insertion sort also finishes in two passes. Bubble sort carries only the largest element to the end each pass, so the small elements 1 and 2 near the right move left by just one place per pass. Insertion sort after two passes has only sorted the first three elements and has not yet looked at 1.

Q.46 MSQ · 2 MARKS

Database Management and Warehousing › Normal forms

Given the relational schema $R = (U, V, W, X, Y, Z)$ and the set of functional dependencies:

$\{U \to V,\ U \to W,\ WX \to Y,\ WX \to Z,\ V \to X\}$

Which of the following functional dependencies can be derived from the above set?

(A) $VW \to YZ$
(B) $WX \to YZ$
(C) $VW \to U$
(D) $VW \to Y$

Official answer: **A, B, D**

SOLUTION

1. A dependency $P \to Q$ can be derived from a set F exactly when $Q \subseteq P^+$, where P^+ is the closure of P under F. So compute the closures needed.
2. Closure of VW. Start with $\{V, W\}$.
3. Apply $V \to X$: the left side V is in the set, so add X. The set becomes $\{V, W, X\}$.
4. Apply $WX \to Y$ and $WX \to Z$: W and X are both in the set, so add Y and Z. The set becomes $\{V, W, X, Y, Z\}$.
5. U appears on the right side of no dependency, so it can never be added. The other dependencies ($U \to V, U \to W$) need U on the left and add nothing new. Hence $(VW)^+ = \{V, W, X, Y, Z\}$.
6. (A) $VW \to YZ$: $\{Y, Z\} \subseteq (VW)^+$, so it can be derived.
7. (C) $VW \to U$: $U \notin (VW)^+$, so it cannot be derived.
8. (D) $VW \to Y$: $Y \in (VW)^+$, so it can be derived.
9. (B) $WX \to YZ$: $WX \to Y$ and $WX \to Z$ are given. By the union rule (if $P \to Q$ and $P \to R$ then $P \to QR$), $WX \to YZ$ holds. Equivalently, $(WX)^+ \supseteq \{W, X, Y, Z\}$.
10. The derivable dependencies are (A), (B) and (D).

KEY IDEA
$P \to Q$ is derivable exactly when $Q \subseteq P^+$. Compute the closure of the left side by repeatedly firing every dependency whose left side is already inside the set. Here $V \to X$ is what lets VW reach WX and so pull in Y and Z.

WATCH OUT FOR
Stopping the closure of VW at $\{V, W, X\}$ and forgetting to fire $WX \to Y$ and $WX \to Z$ afterwards. Another slip is to think $VW \to U$ holds because U determines V and W. The converse does not follow, since nothing determines U.

Q.47 MSQ · 2 MARKS

Linear Algebra › Vector spaces

Select all choices that are subspaces of $\mathbb{R}^3$.

Note: $\mathbb{R}$ denotes the set of real numbers.

(A) $\left\{ \mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \in \mathbb{R}^3 : \mathbf{x} = \alpha \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} + \beta \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \alpha, \beta \in \mathbb{R} \right\}$

(B) $\left\{ \mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \in \mathbb{R}^3 : \mathbf{x} = \alpha^2 \begin{bmatrix} 1 \\ 2 \\ 0 \end{bmatrix} + \beta^2 \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}, \alpha, \beta \in \mathbb{R} \right\}$

(C) $\left\{ \mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \in \mathbb{R}^3 : 5x_1 + 2x_3 = 0,\ 4x_1 - 2x_2 + 3x_3 = 0 \right\}$

(D) $\left\{ \mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \in \mathbb{R}^3 : 5x_1 + 2x_3 + 4 = 0 \right\}$

Official answer: **A, C**

SOLUTION

1. A subset $S \subseteq \mathbb{R}^3$ is a subspace if it contains $\mathbf{0}$ and is closed under addition and under scalar multiplication. Test each option.
2. (A) The set is $\text{span}\left\{ \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \right\}$, because α, β range over all of $\mathbb{R}$. A span of vectors is always a subspace, so (A) is a subspace (it is the plane $x_3 = 0$).
3. (B) Here $x_1 = \alpha^2 + \beta^2$, $x_2 = 2\alpha^2$, $x_3 = \beta^2$, so every vector in the set has $x_1 \geq 0$, $x_2 \geq 0$ and $x_3 \geq 0$. Take $\mathbf{u} = \begin{bmatrix} 1 \\ 2 \\ 0 \end{bmatrix}$ (with $\alpha = 1, \beta = 0$), which is in the set. Then $-\mathbf{u} = \begin{bmatrix} -1 \\ -2 \\ 0 \end{bmatrix}$ has $x_1 < 0$, so it is not in the set. Closure under scalar multiplication fails, so (B) is not a subspace. The squares allow only non-negative multipliers, giving a cone rather than a subspace.
4. (C) The set is the solution set of the homogeneous system $A\mathbf{x} = \mathbf{0}$ with $A = \begin{bmatrix} 5 & 0 & 2 \\ 4 & -2 & 3 \end{bmatrix}$, that is, the null space of A. If $A\mathbf{x} = \mathbf{0}$ and $A\mathbf{y} = \mathbf{0}$, then $A(a\mathbf{x} + b\mathbf{y}) = \mathbf{0}$, so a null space is always a subspace. Hence (C) is a subspace (here a line: the two rows are independent, so the rank is 2 and the nullity is $3 - 2 = 1$).
5. (D) Put $\mathbf{x} = \mathbf{0}$: $5(0) + 2(0) + 4 = 4 \neq 0$. The zero vector is not in the set, so (D) is not a subspace. It is a plane that does not pass through the origin.
6. Hence the subspaces are (A) and (C).

KEY IDEA

A subspace must contain the zero vector and be closed under addition and all real scalar multiples. A span of vectors and the solution set of a homogeneous linear system $A\mathbf{x} = \mathbf{0}$ are always subspaces.

WATCH OUT FOR

Treating (B) as a span because it is built from two vectors. The squared coefficients α^2, β^2 can only be non-negative, so negative multiples are missing. For (D), the mistake is forgetting that a plane $5x_1 + 2x_3 = -4$ misses the origin.

Q.48 MSQ · 2 MARKS

Linear Algebra › Linear systems

Which of the following statements is/are TRUE?

Note: $\mathbb{R}$ denotes the set of real numbers.

(A) There exist $\mathbf{M} \in \mathbb{R}^{3\times3}$, $\mathbf{p} \in \mathbb{R}^3$, and $\mathbf{q} \in \mathbb{R}^3$ such that $\mathbf{Mx} = \mathbf{p}$ has a unique solution and $\mathbf{Mx} = \mathbf{q}$ has infinite solutions.

(B) There exist $\mathbf{M} \in \mathbb{R}^{3\times3}$, $\mathbf{p} \in \mathbb{R}^3$, and $\mathbf{q} \in \mathbb{R}^3$ such that $\mathbf{Mx} = \mathbf{p}$ has no solutions and $\mathbf{Mx} = \mathbf{q}$ has infinite solutions.

(C) There exist $\mathbf{M} \in \mathbb{R}^{2\times3}$, $\mathbf{p} \in \mathbb{R}^2$, and $\mathbf{q} \in \mathbb{R}^2$ such that $\mathbf{Mx} = \mathbf{p}$ has a unique solution and $\mathbf{Mx} = \mathbf{q}$ has infinite solutions.

(D) There exist $\mathbf{M} \in \mathbb{R}^{3\times2}$, $\mathbf{p} \in \mathbb{R}^3$, and $\mathbf{q} \in \mathbb{R}^3$ such that $\mathbf{Mx} = \mathbf{p}$ has a unique solution and $\mathbf{Mx} = \mathbf{q}$ has no solutions.

Official answer: **B, D**

SOLUTION

1. Recall: if $\mathbf{Mx} = \mathbf{b}$ is consistent, its solution is unique exactly when the columns of $\mathbf{M}$ are independent (nullity 0). If it is consistent and the nullity is positive, there are infinitely many solutions. There is no solution when $\mathbf{b}$ lies outside the column space of $\mathbf{M}$.
2. (A) A unique solution of a square 3×3 system means nullity 0, so $\operatorname{rank}\mathbf{M} = 3$ and $\mathbf{M}$ is invertible. Then $\mathbf{Mx} = \mathbf{q}$ has exactly one solution for every $\mathbf{q}$, never infinitely many. **False.**
3. (B) Take $\mathbf{M} = \begin{bmatrix}1&0&0\\0&0&0\\0&0&0\end{bmatrix}$, $\mathbf{p} = \begin{bmatrix}0\\1\\0\end{bmatrix}$, $\mathbf{q} = \begin{bmatrix}0\\0\\0\end{bmatrix}$. The second row of $\mathbf{Mx} = \mathbf{p}$ reads $0 = 1$, so there is no solution. For $\mathbf{q} = \mathbf{0}$ the system is consistent and the nullity is $3 - 1 = 2 > 0$, so there are infinitely many solutions. **True.**
4. (C) For $\mathbf{M} \in \mathbb{R}^{2\times3}$, $\operatorname{rank}\mathbf{M} \le 2$, so by rank-nullity the nullity is $3 - \operatorname{rank}\mathbf{M} \ge 1$. If $\mathbf{Mx} = \mathbf{p}$ has a solution, adding any non-zero null-space vector gives another, so the solution is never unique. **False.**
5. (D) Take $\mathbf{M} = \begin{bmatrix}1&0\\0&1\\0&0\end{bmatrix}$, whose columns are independent (nullity 0). With $\mathbf{p} = \begin{bmatrix}1\\1\\0\end{bmatrix}$, $\mathbf{Mx} = \mathbf{p}$ gives $x_1 = 1$, $x_2 = 1$, and the third row $0 = 0$ holds, so the solution is unique. With $\mathbf{q} = \begin{bmatrix}0\\0\\1\end{bmatrix}$ the third row reads $0 = 1$, so there is no solution. **True.**
6. Hence the true statements are (B) and (D).

KEY IDEA

Uniqueness of a solution depends only on the nullity of $\mathbf{M}$ (independent columns), while existence depends on whether the right-hand side lies in the column space. So the same $\mathbf{M}$ can give different outcomes for different right-hand sides only when the nullity is positive (infinite versus none) or when the system is tall (unique versus none).

WATCH OUT FOR

Thinking that one matrix cannot give different outcomes for different right-hand sides. For a square invertible matrix every right-hand side behaves alike, but a singular square matrix (B) or a tall matrix with independent columns (D) can have a solution for one vector and none for another. In (C), forgetting that a wide matrix always has a non-trivial null space.

Q.49 MSQ · 2 MARKS

Linear Algebra › Special matrices

Let $\mathbb{R}$ be the set of real numbers, U be a subspace of $\mathbb{R}^3$ and $\mathbf{M} \in \mathbb{R}^{3\times3}$ be the matrix corresponding to the projection on to the subspace U.

Which of the following statements is/are TRUE?

(A) If U is a 1-dimensional subspace of $\mathbb{R}^3$, then the null space of $\mathbf{M}$ is a 1-dimensional subspace.

(B) If U is a 2-dimensional subspace of $\mathbb{R}^3$, then the null space of $\mathbf{M}$ is a 1-dimensional subspace.

(C) $\mathbf{M}^2 = \mathbf{M}$

(D) $\mathbf{M}^3 = \mathbf{M}$

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

SOLUTION

1. A projection onto U sends every vector $\mathbf{x}$ to a vector $\mathbf{Mx} \in U$, and leaves every vector of U unchanged. So the column space of $\mathbf{M}$ is exactly U, and $\text{rank}(\mathbf{M}) = \dim U$.
2. By rank-nullity for the 3×3 matrix $\mathbf{M}$: $\dim N(\mathbf{M}) = 3 - \text{rank}(\mathbf{M}) = 3 - \dim U$.
3. (A) If $\dim U = 1$, then $\dim N(\mathbf{M}) = 3 - 1 = 2$, not 1. So (A) is false.
4. (B) If $\dim U = 2$, then $\dim N(\mathbf{M}) = 3 - 2 = 1$. So (B) is true.
5. (C) For any $\mathbf{x}$, the vector $\mathbf{Mx}$ lies in U, and $\mathbf{M}$ fixes every vector of U. Hence $\mathbf{M}(\mathbf{Mx}) = \mathbf{Mx}$ for all $\mathbf{x}$, that is $\mathbf{M}^2 = \mathbf{M}$. So (C) is true.
6. (D) Using (C): $\mathbf{M}^3 = \mathbf{M}\,\mathbf{M}^2 = \mathbf{M}\,\mathbf{M} = \mathbf{M}^2 = \mathbf{M}$. So (D) is true.
7. The true statements are B, C and D.

KEY IDEA

A projection matrix is idempotent ($\mathbf{M}^2 = \mathbf{M}$), its rank equals $\dim U$, and rank-nullity gives $\dim N(\mathbf{M}) = 3 - \dim U$.

WATCH OUT FOR

Confusing the dimension of the subspace being projected onto with the dimension of the null space. For a 1-dimensional U the null space has dimension 2, not 1.

Q.50 MSQ · 2 MARKS

Calculus and Optimization › Maxima & minima

Consider the function $f : \mathbb{R} \to \mathbb{R}$ where $\mathbb{R}$ is the set of all real numbers.

$$f(x) = \frac{x^4}{4} - \frac{2x^3}{3} - \frac{3x^2}{2} + 1$$

Which of the following statements is/are TRUE?

(A) $x = 0$ is a local maximum of f

(B) $x = 3$ is a local minimum of f

(C) $x = -1$ is a local maximum of f

(D) $x = 0$ is a local minimum of f

Official answer: **A, B**

SOLUTION

1. Differentiate: $f'(x) = x^3 - 2x^2 - 3x = x(x^2 - 2x - 3) = x(x-3)(x+1)$.
2. Stationary points (where $f'(x) = 0$): $x = -1,\ 0,\ 3$.
3. Second derivative: $f''(x) = 3x^2 - 4x - 3$.
4. At $x = 0$: $f''(0) = -3 < 0$, so $x = 0$ is a local maximum. Option (A) is true and option (D) is false.
5. At $x = 3$: $f''(3) = 27 - 12 - 3 = 12 > 0$, so $x = 3$ is a local minimum. Option (B) is true.
6. At $x = -1$: $f''(-1) = 3 + 4 - 3 = 4 > 0$, so $x = -1$ is a local minimum, not a local maximum. Option (C) is false.
7. Cross-check with the sign of $f'(x) = x(x-3)(x+1)$: it is negative for $x < -1$ (for example $f'(-2) = -10$), positive on $(-1, 0)$ (for example $f'(-0.5) = 0.875$), negative on $(0, 3)$ (for example $f'(1) = -4$) and positive for $x > 3$. So f falls then rises at -1 (minimum), rises then falls at 0 (maximum), and falls then rises at 3 (minimum). This agrees with the second-derivative test.

KEY IDEA

Find the stationary points from $f'(x) = 0$, then classify each one with the second-derivative test (or the sign change of f'): $f'' > 0$ gives a local minimum and $f'' < 0$ gives a local maximum.

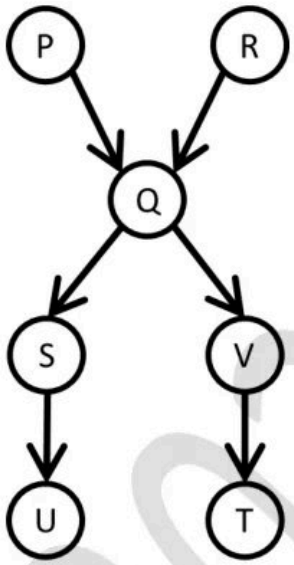
WATCH OUT FOR

Mixing up the sign of f''. Here $f''(-1) = 4 > 0$, so $x = -1$ is a local minimum, and option (C) is a trap. Another slip is to treat $x = 0$ as a minimum, when $f''(0) = -3 < 0$ makes it a maximum.

Q.51 MSQ · 2 MARKS

Data Structures and Algorithms › Graph algorithms

Q.51 Consider the directed acyclic graph (DAG) below:



Which of the following is/are valid vertex orderings that can be obtained from a topological sort of the DAG?

(A)	P Q R S T U V
(B)	P R Q V S U T
(C)	P Q R S V U T
(D)	P R Q S V T U

Official answer: **B, D**

SOLUTION

1. Read the edges from the figure: $P \to Q$, $R \to Q$, $Q \to S$, $Q \to V$, $S \to U$, $V \to T$.
2. A topological order must place the tail of every edge before its head. So: P and R both come before Q; Q comes before S and V; S comes before U; V comes before T.
3. (A) P Q R S T U V : Q is at position 2 and R at position 3, so $R \to Q$ is violated. Invalid. (It also puts T before V, violating $V \to T$.)

4. (B) P R Q V S U T : P (1) and R (2) come before Q (3); Q comes before V (4) and S (5); S (5) comes before U (6); V (4) comes before T (7). Every edge is respected. Valid.
5. (C) P Q R S V U T : Q is at position 2 and R at position 3, so $R \to Q$ is violated. Invalid.
6. (D) P R Q S V T U : P (1) and R (2) come before Q (3); Q comes before S (4) and V (5); S (4) comes before U (7); V (5) comes before T (6). Every edge is respected. Valid.
7. So the valid topological orderings are (B) and (D).

KEY IDEA
An ordering is a valid topological sort exactly when, for every directed edge $u \to v$, u appears before v. Check each edge against each option's positions. The only freedom is the relative order of independent branches, here P versus R, and the $S \to U$ chain versus the $V \to T$ chain.

WATCH OUT FOR
Checking only that P comes first and Q comes early, and missing that R must also precede Q. Options (A) and (C) put Q right after P, before R. Another slip is forgetting that $S \to U$ and $V \to T$ are separate chains, so their order relative to each other is free.

Q.52 MSQ · 2 MARKS

Data Structures and Algorithms › Trees & heaps

Let H, I, L, and N represent height, number of internal nodes, number of leaf nodes, and the total number of nodes respectively in a rooted binary tree.

Which of the following statements is/are always **TRUE**?

(A) $L \leq I + 1$

(B) $H + 1 \leq N \leq 2^{H+1} - 1$

(C) $H \leq I \leq 2^{H} - 1$

(D) $H \leq L \leq 2^{H-1}$

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

SOLUTION

1. Take the height H as the number of edges on the longest root-to-leaf path. This is the convention under which the key holds. A single node has $H = 0$. A binary tree means every node has at most 2 children. Let $n_0 = L$, n_1, n_2 be the numbers of nodes with 0, 1 and 2 children. Then $I = n_1 + n_2$ and $N = n_0 + n_1 + n_2$.
2. Count the edges two ways. Every node except the root has exactly one parent edge, so there are $N - 1$ edges. Each node with k children contributes k edges, so there are also $n_1 + 2n_2$ edges. Hence $N - 1 = n_1 + 2n_2$. Substituting $N = n_0 + n_1 + n_2$ gives $n_0 = n_2 + 1$.
3. (A) $L = n_2 + 1 \leq (n_1 + n_2) + 1 = I + 1$ because $n_1 \geq 0$. Always true.
4. (B) The longest path alone contains $H + 1$ nodes, so $N \geq H + 1$, with equality for a chain. Level d holds at most 2^d nodes and levels run from 0 to H, so $N \leq \sum_{d=0}^{H} 2^d = 2^{H+1} - 1$, with equality for a full tree. Always true.
5. (C) Every node on the longest path except its last node (a leaf) has a child, so it is internal. That gives H internal nodes, so $I \geq H$. Internal nodes cannot lie on level H, because a node at level H has no children. So they lie on levels 0 to $H - 1$, and $I \leq \sum_{d=0}^{H-1} 2^d = 2^H - 1$. Always true, including $H = 0$, where $I = 0$.
6. (D) Take a chain of height $H = 2$: it has $L = 1 < H$, so $H \leq L$ fails. Also a full tree of height H has $L = 2^H > 2^{H-1}$, so $L \leq 2^{H-1}$ fails too. Not always true.
7. So the statements that are always true are A, B and C.

KEY IDEA

In a binary tree, $n_0 = n_2 + 1$ (leaves = degree-2 nodes + 1). Bound the node counts by the longest path, which gives the minimum, and by full levels, which gives the maximum.

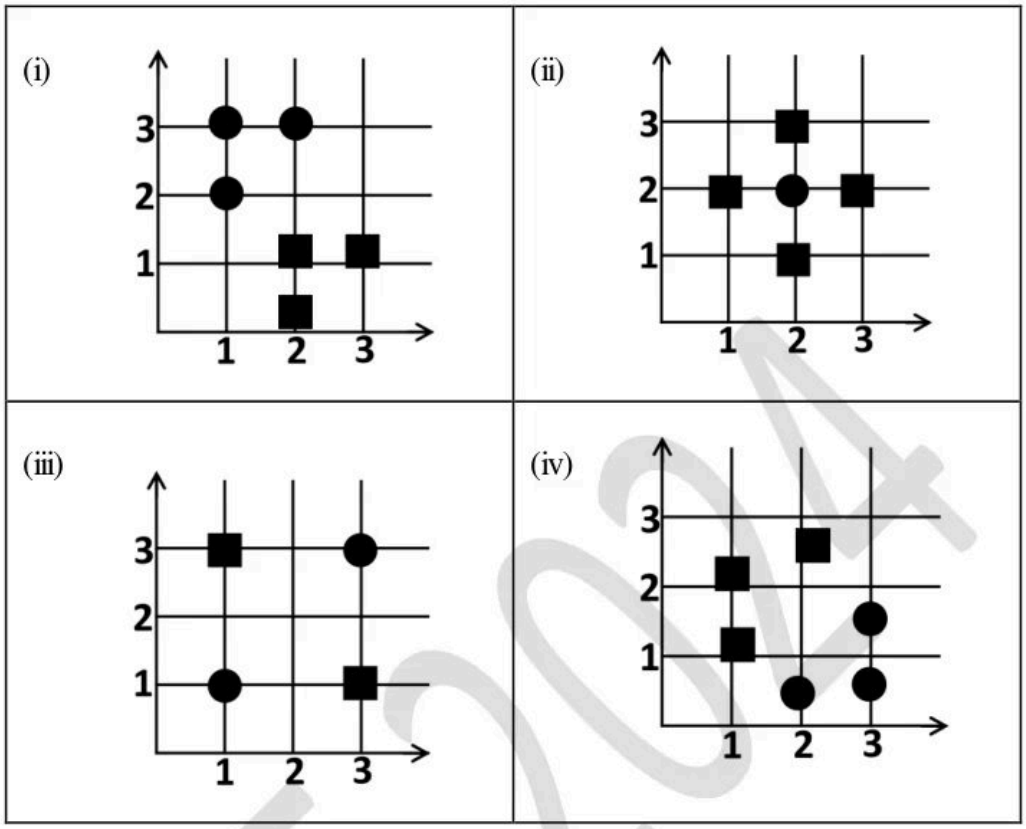
WATCH OUT FOR

Counting height in nodes instead of edges, which breaks the lower bounds in B and C. Another slip is assuming a tree of height H has at least H leaves. A chain has only one leaf.

Q.53 MSQ · 2 MARKS

Machine Learning › LDA, SVM & trees

Q.53 Consider the following figures representing datasets consisting of two-dimensional features with two classes denoted by circles and squares.



Which of the following is/are **TRUE**?

(A)	(i) is linearly separable.
(B)	(ii) is linearly separable.
(C)	(iii) is linearly separable.
(D)	(iv) is linearly separable.

Official answer: **A, D**

SOLUTION

1. Linear separability means one straight line (a hyperplane in 2-D) puts every circle on one side and every square on the other. Check each panel in turn.
2. (i) The circles are at $(1,3),(2,3),(1,2)$, all with $y \geq 2$. The squares are at $(2,1),(3,1),(2,0.3)$, all with $y \leq 1$. The horizontal line $y = 1.5$ separates them, so (i) is linearly separable.
3. (ii) The circle at $(2,2)$ is the centre of the four squares $(1,2),(3,2),(2,3),(2,1)$. It is the midpoint of the squares $(1,2)$ and $(3,2)$, so it lies in the convex hull of the squares. A line with all squares on one side

puts their whole hull, and so the circle, on that side too. So (ii) is not linearly separable.

4. (iii) This is the XOR layout. The squares are at $(1,3),(3,1)$ and the circles at $(3,3),(1,1)$. The segment joining the two squares and the segment joining the two circles both have midpoint $(2,2)$, so they cross and the convex hulls intersect. So (iii) is not linearly separable.
5. (iv) The squares are at $(1,2.2),(2,2.6),(1,1.2)$ and the circles at $(2,0.4),(3,1.5),(3,0.5)$. Take the line $y = 0.5x + 0.3$. At $x = 1, 2, 1$ its heights are $0.8,\ 1.3,\ 0.8$, and the squares' heights $2.2,\ 2.6,\ 1.2$ are all above these. At $x = 2, 3, 3$ its heights are $1.3,\ 1.8,\ 1.8$, and the circles' heights $0.4,\ 1.5,\ 0.5$ are all below these. The classes are separated, so (iv) is linearly separable.
6. True statements: (i) and (iv), so options A and D.

KEY IDEA

Two classes in the plane are linearly separable exactly when their convex hulls do not intersect. A point inside the other class's hull (panel ii) or crossing diagonals (the XOR layout, panel iii) rule out separation, and one explicit separating line proves it.

WATCH OUT FOR

Judging panel (iv) by eye from its tilted layout and calling it inseparable. Because the points are not on a grid, a slanted line is needed, and one such line separates the classes. The other slip is treating the XOR layout as separable because each class is symmetric.

Q.54 MSQ · 2 MARKS

Artificial Intelligence › Logic

Let $game(ball, rugby)$ be true if the ball is used in rugby and false otherwise. Let $shape(ball, round)$ be true if the ball is round and false otherwise. Consider the following logical sentences:

s1: $\forall ball\ \neg game(ball, rugby) \Rightarrow shape(ball, round)$
s2: $\forall ball\ \neg shape(ball, round) \Rightarrow game(ball, rugby)$
s3: $\forall ball\ game(ball, rugby) \Rightarrow \neg shape(ball, round)$
s4: $\forall ball\ shape(ball, round) \Rightarrow \neg game(ball, rugby)$

Which of the following choices is/are logical representations of the assertion, "All balls are round except balls used in rugby"?

(A) $s1 \wedge s3$

(B) $s1 \wedge s2$

(C) $s2 \wedge s3$

(D) $s3 \wedge s4$

Official answer: **A, C**

SOLUTION

1. Let G = game(ball, rugby) and R = shape(ball, round). The assertion "All balls are round except balls used in rugby" says two things: a ball that is not used in rugby is round, and a ball used in rugby is not round.
2. Formalise the two parts: $\forall ball\ (\neg G \rightarrow R)$ and $\forall ball\ (G \rightarrow \neg R)$. These are exactly $s1$ and $s3$.
3. Contrapositive: $\neg G \rightarrow R \equiv \neg R \rightarrow G$, so $s2$ is logically equivalent to $s1$.
4. Contrapositive: $G \rightarrow \neg R \equiv R \rightarrow \neg G$, so $s4$ is logically equivalent to $s3$.
5. Hence $s1 \wedge s3$ is the full assertion, so (A) is correct. Since $s2 \equiv s1$, we get $s2 \wedge s3 \equiv s1 \wedge s3$, so (C) is also correct.
6. $s1 \wedge s2 \equiv s1$ only says non-rugby balls are round. It never says rugby balls are not round, so (B) is incomplete. $s3 \wedge s4 \equiv s3$ only says rugby balls are not round. It says nothing about non-rugby balls, so (D) is incomplete.

KEY IDEA

"All X are Y except Z" means two implications: not-Z implies Y, and Z implies not-Y. Each implication is equivalent to its contrapositive, so $s1 \equiv s2$ and $s3 \equiv s4$. A pair is correct exactly when it contains one sentence from $\{s1, s2\}$ and one from $\{s3, s4\}$.

WATCH OUT FOR

Treating $s1 \wedge s2$ or $s3 \wedge s4$ as the full assertion. Each of these pairs holds two equivalent sentences, so it states only one half of the assertion. Another slip is to miss that $s2$ is the contrapositive of $s1$ and so drop option (C).

Q.55 MSQ · 2 MARKS

Database Management and Warehousing › Indexing

An OTT company is maintaining a large disk-based relational database of different movies with the following schema:

Movie(ID, CustomerRating)
Genre(ID, Name)
Movie_Genre(MovieID, GenreID)

(Primary keys underlined in print: Movie.ID; Genre.ID; Movie_Genre.MovieID and GenreID together.)

Consider the following SQL query on the relation database above:

```
SELECT *
FROM Movie, Genre, Movie_Genre
WHERE
   Movie.CustomerRating > 3.4 AND
   Genre.Name = "Comedy" AND
   Movie_Genre.MovieID = Movie.ID AND
   Movie_Genre.GenreID = Genre.ID;
```

This SQL query can be sped up using which of the following indexing options?

(A) B^+ tree on all the attributes.

(B) Hash index on `Genre.Name` and B^+ tree on the remaining attributes.

(C) Hash index on `Movie.CustomerRating` and B^+ tree on the remaining attributes.

(D) Hash index on all the attributes.

Official answer: **A, B**

SOLUTION

1. Classify each predicate by the access it needs. `Movie.CustomerRating > 3.4` is a **range** condition. `Genre.Name = "Comedy"` is an **equality** condition. The two join conditions, `Movie_Genre.MovieID = Movie.ID` and `Movie_Genre.GenreID = Genre.ID`, are equalities on key attributes.
2. A hash index answers only equality lookups, because hashing scatters keys and destroys their order. A B^+ tree keeps its leaves sorted and linked, so it answers both equality and range lookups.
3. Take an option to speed the query up when every predicate has an index type that can serve it. The range predicate on `CustomerRating` is the deciding one: it can use an index only if `Movie.CustomerRating` carries a B^+ tree. A hash index on it cannot answer "rating > 3.4", so `Movie` would still be scanned in full.
4. Option A: a B^+ tree on every attribute serves the range on `CustomerRating`, the equality on `Name`, and the join lookups on `ID`, `MovieID` and `GenreID`. Every predicate is served, so the query is sped up. **Correct.**
5. Option B: a hash index on `Genre.Name` is ideal for the single equality `Name = "Comedy"`. The remaining attributes carry B^+ trees, including `CustomerRating` (range) and the join keys. Every predicate again has a suitable index, so the query is sped up. **Correct.**

6. Option C: the hash index sits on CustomerRating, the one attribute that needs range access. The range filter on Movie therefore gets no usable index, which is exactly the mismatch that defeats the design for this query. The other attributes are indexed, but the option does not match the access pattern, so it is not accepted as a way to speed the query up. **Not correct.**
7. Option D: hash indexes on all attributes leave the range predicate on CustomerRating unsupported, for the same reason as C. **Not correct.**
8. Only A and B give every predicate a matching index type. The answer is A and B.

KEY IDEA
Match index type to predicate type. Hash indexes serve only equality, whereas B^+ trees serve equality and range. A range predicate such as `CustomerRating > 3.4` needs a B^+ tree on that attribute.

WATCH OUT FOR
Assuming a hash index is always faster than a B^+ tree and choosing a hash index on the attribute with the range condition (option C), or hashing everything (option D). Hashing cannot help with a $>$ comparison.

Q.56 NAT · 2 MARKS

Probability and Statistics › Continuous distributions

Let X be a random variable uniformly distributed in the interval $[1, 3]$ and Y be a random variable uniformly distributed in the interval $[2, 4]$. If X and Y are independent of each other, the probability $P(X \geq Y)$ is _____ (rounded off to **three** decimal places).

Official answer: **0.125**

SOLUTION

1. Independence gives the joint density $f(x, y) = \frac{1}{2} \cdot \frac{1}{2} = \frac{1}{4}$ on the rectangle $[1, 3] \times [2, 4]$, which has area $2 \times 2 = 4$.
2. So $P(X \geq Y) = \dfrac{\text{area of the part of the rectangle where } x \geq y}{4}$.
3. Since $y \geq 2$, the condition $x \geq y$ forces $x \geq 2$. So $x \in [2, 3]$ and, for each such x, $y \in [2, x]$.
4. The area is $\displaystyle\int_2^3 (x - 2)\, dx = \left[\frac{(x-2)^2}{2}\right]_2^3 = \frac{1}{2}$. This is a right triangle with legs of length 1.
5. Hence $P(X \geq Y) = \dfrac{1/2}{4} = \dfrac{1}{8} = 0.125$.
6. Check by conditioning on X: for $x \in [2, 3]$, $P(Y \leq x) = \frac{x-2}{2}$, and for $x \in [1, 2)$ it is 0. So $P(X \geq Y) = \int_2^3 \frac{1}{2} \cdot \frac{x-2}{2}\, dx = \frac{1}{4} \cdot \frac{1}{2} = \frac{1}{8}$.

KEY IDEA
For independent uniform variables the point (X, Y) is uniform on a rectangle, so the probability is the area of the region $x \geq y$ divided by the rectangle's area.

WATCH OUT FOR
Treating the two intervals as if they overlapped on $[2, 3]$ only and taking $P = \frac{1}{2}$ there, or integrating over $x \in [1, 3]$ and forgetting that $Y \geq 2$ makes $P(Y \leq x) = 0$ for $x < 2$. The area must be divided by the full rectangle area 4, not by 1.

Q.57 NAT · 2 MARKS

Probability and Statistics › Continuous distributions

Let X be a random variable exponentially distributed with parameter $\lambda > 0$. The probability density function of X is given by:

$$f_X(x) = \begin{cases} \lambda e^{-\lambda x}, & x \geq 0 \\ 0, & \text{otherwise} \end{cases}$$

If $5E(X) = \mathrm{Var}(X)$, where $E(X)$ and $\mathrm{Var}(X)$ indicate the expectation and variance of X, respectively, the value of λ is _____ (rounded off to one decimal place).

Official answer: **0.2**

SOLUTION

1. For an exponential random variable with density $f_X(x) = \lambda e^{-\lambda x}$ for $x \geq 0$, the mean is $E(X) = \int_0^\infty x\,\lambda e^{-\lambda x}\,dx = \dfrac{1}{\lambda}$ (integration by parts).
2. The second moment is $E(X^2) = \int_0^\infty x^2 \lambda e^{-\lambda x}\,dx = \dfrac{2}{\lambda^2}$, so $\mathrm{Var}(X) = E(X^2) - [E(X)]^2 = \dfrac{2}{\lambda^2} - \dfrac{1}{\lambda^2} = \dfrac{1}{\lambda^2}$.
3. Impose the condition $5E(X) = \mathrm{Var}(X)$, which gives $\dfrac{5}{\lambda} = \dfrac{1}{\lambda^2}$.
4. Multiply both sides by λ^2 (allowed since $\lambda > 0$): $5\lambda = 1$, so $\lambda = \dfrac{1}{5} = 0.2$.

KEY IDEA

For an exponential distribution with rate λ, the mean is $1/\lambda$ and the variance is $1/\lambda^2$, so the condition $5E(X) = \mathrm{Var}(X)$ reduces to $5/\lambda = 1/\lambda^2$.

WATCH OUT FOR

Using the mean and variance of an exponential written in terms of the scale (mean) parameter, or taking the variance as $1/\lambda$ instead of $1/\lambda^2$. That gives $5/\lambda = 1/\lambda$, which has no solution.

Q.58 NAT · 2 MARKS

Probability and Statistics › Conditional & Bayes

Consider two events T and S. Let $\bar{T}$ denote the complement of the event T. The probability associated with different events are given as follows: $P(\bar{T}) = 0.6$, $P(S \mid T) = 0.3$, $P(S \mid \bar{T}) = 0.6$. Then, $P(T \mid S)$ is _____ (rounded off to two decimal places).

Official answer: **0.25**

SOLUTION

1. Complement rule: $P(T) = 1 - P(\bar{T}) = 1 - 0.6 = 0.4$.
2. Total probability over the partition $\{T, \bar{T}\}$: $P(S) = P(S \mid T)\,P(T) + P(S \mid \bar{T})\,P(\bar{T})$.
3. $P(S) = 0.3 \times 0.4 + 0.6 \times 0.6 = 0.12 + 0.36 = 0.48$.
4. Bayes' theorem: $P(T \mid S) = \dfrac{P(S \mid T)\,P(T)}{P(S)} = \dfrac{0.12}{0.48}$.
5. $P(T \mid S) = \dfrac{0.12}{0.48} = 0.25$, which is already exact to two decimal places.

KEY IDEA

The question gives $P(\bar{T})$ and the likelihoods $P(S \mid T)$, $P(S \mid \bar{T})$. Get the prior $P(T)$ from the complement rule, get $P(S)$ by total probability over T and $\bar{T}$, then reverse the conditioning with Bayes' theorem.

WATCH OUT FOR

Using $P(\bar{T}) = 0.6$ as if it were $P(T)$. That gives $P(S) = 0.3 \times 0.6 + 0.6 \times 0.4 = 0.42$ and $P(T \mid S) = 0.18/0.42 \approx 0.43$. Another slip is dividing by $P(S \mid T)$ or by $P(T)$ instead of $P(S)$.

Q.59 NAT · 2 MARKS

Probability and Statistics › Random variables

Consider a joint probability density function of two random variables X and Y

$$f_{X,Y}(x,y) = \begin{cases} 2xy, & 0 < x < 2,\ 0 < y < x \\ 0, & \text{otherwise} \end{cases}$$

Then, $E[Y \mid X = 1.5]$ is _____.

Official answer: **Marks to all**

SOLUTION

1. The support is the triangle $0 < y < x < 2$, so once $X = x$ is fixed, Y ranges over $(0, x)$.
2. Check the total probability first.
$$\int_0^2 \int_0^x 2xy \, dy \, dx = \int_0^2 x \cdot x^2 \, dx = \int_0^2 x^3 \, dx = \frac{2^4}{4} = 4.$$
The total is 4, not 1, so the printed function is not a valid joint density as it stands.
3. The natural reading is that the density is proportional to $2xy$ on the triangle. Dividing by 4 gives the valid joint density $f_{X,Y}(x, y) = \dfrac{xy}{2}$ for $0 < y < x < 2$.
4. Marginal of X:
$$f_X(x) = \int_0^x \frac{xy}{2} \, dy = \frac{x}{2} \cdot \frac{x^2}{2} = \frac{x^3}{4}, \quad 0 < x < 2.$$
Check: $\int_0^2 \frac{x^3}{4} \, dx = \frac{16}{16} = 1$.
5. Conditional density of Y given $X = x$:
$$f_{Y|X}(y \mid x) = \frac{f_{X,Y}(x, y)}{f_X(x)} = \frac{xy/2}{x^3/4} = \frac{2y}{x^2}, \quad 0 < y < x.$$
It integrates to 1: $\int_0^x \frac{2y}{x^2} \, dy = \frac{x^2}{x^2} = 1$. The scaling constant cancels in this ratio, so the conditional density is the same whether or not the joint is normalised. That makes the answer well defined despite the misprint.
6. Conditional expectation:
$$E[Y \mid X = x] = \int_0^x y \cdot \frac{2y}{x^2} \, dy = \frac{2}{x^2} \cdot \frac{x^3}{3} = \frac{2x}{3}.$$
7. Substitute $x = 1.5$:
$$E[Y \mid X = 1.5] = \frac{2 \times 1.5}{3} = 1.$$

KEY IDEA

Given $X = x$, the conditional density is proportional to the joint density with x held fixed. On $0 < y < x$ that is proportional to y, so $f_{Y|X}(y \mid x) = 2y/x^2$ and $E[Y \mid X = x] = 2x/3$. Any constant in front of the joint density cancels in this ratio.

WATCH OUT FOR

One mistake is taking Y from 0 to 2 instead of 0 to x. Another is using the marginal mean $E[Y]$ instead of the conditional one. A third is to assume, without checking, that $2xy$ is a valid density: it integrates to 4 on this triangle, so x^3 is not the true marginal of X (the valid one is $x^3/4$). A fourth is to read the conditional as uniform on $(0, 1.5)$, which would give 0.75.

Q.60 NAT · 2 MARKS

Calculus and Optimization › Limits & continuity

Evaluate the following limit:

$$\lim_{x \to 0} \frac{\ln\left((x^2 + 1)\cos x\right)}{x^2} = ____.$$

Official answer: **0.5**

SOLUTION

1. Near $x = 0$ both $x^2 + 1$ and $\cos x$ are positive, so the logarithm of the product splits: $\ln\left((x^2 + 1)\cos x\right) = \ln(1 + x^2) + \ln(\cos x)$.

2. Divide by x^2: the limit is $\lim_{x\to 0}\dfrac{\ln(1+x^2)}{x^2}+\lim_{x\to 0}\dfrac{\ln(\cos x)}{x^2}$, valid provided both limits exist (shown below).
3. First term: put $t=x^2\to 0^+$. Then $\dfrac{\ln(1+t)}{t}\to 1$, since $\ln(1+t)=t-\dfrac{t^2}{2}+\cdots$ (equivalently, the derivative of $\ln(1+t)$ at $t=0$ is 1).
4. Second term: $\cos x=1-\dfrac{x^2}{2}+O(x^4)$. With $u=-\dfrac{x^2}{2}+O(x^4)\to 0$, $\ln(1+u)=u+O(u^2)$, so $\ln\cos x=-\dfrac{x^2}{2}+O(x^4)$. Therefore $\dfrac{\ln\cos x}{x^2}\to-\dfrac{1}{2}$.
5. Check by L'Hopital (the form is $0/0$): $\dfrac{d}{dx}\ln\cos x=-\tan x$ and $\dfrac{d}{dx}x^2=2x$, so the ratio is $-\dfrac{\tan x}{2x}\to -\dfrac{1}{2}$.
6. Both limits exist, so the sum is $1-\dfrac{1}{2}=\dfrac{1}{2}=0.5$.

KEY IDEA
Split the logarithm of the product into $\ln(1+x^2)+\ln\cos x$, then use the standard expansions $\ln(1+u)\approx u$ and $\cos x\approx 1-x^2/2$ so each piece divided by x^2 has a finite limit.

WATCH OUT FOR
Forgetting the $\ln\cos x$ contribution and answering 1 (treating $\cos x\to 1$ as if $\ln\cos x$ were negligible next to x^2), or getting its sign wrong. $\ln\cos x$ is of order x^2, so it changes the limit by $-\frac{1}{2}$.

Q.61 NAT · 2 MARKS

Linear Algebra › LU & SVD

Let $\mathbf{u}=\begin{bmatrix}1\\2\\3\\4\\5\end{bmatrix}$, and let $\sigma_1,\sigma_2,\sigma_3,\sigma_4,\sigma_5$ be the singular values of the matrix $\mathbf{M}=\mathbf{u}\mathbf{u}^T$ (where $\mathbf{u}^T$ is the transpose of $\mathbf{u}$). The value of $\sum_{i=1}^{5}\sigma_i$ is _____.

Official answer: **55**

SOLUTION

1. $\mathbf{M}=\mathbf{u}\mathbf{u}^T$ is the outer product of the nonzero vector $\mathbf{u}$ with itself, so every column of $\mathbf{M}$ is a multiple of $\mathbf{u}$ and $\operatorname{rank}(\mathbf{M})=1$.
2. A matrix of rank 1 has exactly one nonzero singular value; the other four singular values are 0. So $\sum_{i=1}^{5}\sigma_i=\sigma_1$.
3. $\mathbf{M}$ is symmetric, and for any $\mathbf{x}$ we have $\mathbf{x}^T\mathbf{M}\mathbf{x}=(\mathbf{u}^T\mathbf{x})^2\geq 0$, so it is positive semidefinite. For such a matrix the singular values equal the eigenvalues.
4. $\mathbf{M}\mathbf{u}=\mathbf{u}(\mathbf{u}^T\mathbf{u})=(\mathbf{u}^T\mathbf{u})\,\mathbf{u}$, so $\mathbf{u}$ is an eigenvector with eigenvalue $\mathbf{u}^T\mathbf{u}$. This is the nonzero eigenvalue, since the rank is 1.
5. $\mathbf{u}^T\mathbf{u}=1^2+2^2+3^2+4^2+5^2=1+4+9+16+25=55$.
6. Hence $\sigma_1=55$ and $\sigma_2=\cdots=\sigma_5=0$, giving $\sum_{i=1}^{5}\sigma_i=55$. Check: $\operatorname{trace}(\mathbf{M})=\sum_i u_i^2=55$, which equals the sum of the eigenvalues.

KEY IDEA
$\mathbf{u}\mathbf{u}^T$ has rank 1, so it has a single nonzero singular value, equal to $\mathbf{u}^T\mathbf{u}=\|\mathbf{u}\|^2$. The sum of the singular values is therefore the trace, $\sum u_i^2=55$.

WATCH OUT FOR
Taking the single singular value to be the norm $\|\mathbf{u}\|=\sqrt{55}\approx 7.42$ instead of $\|\mathbf{u}\|^2=55$. The matrix is $\mathbf{u}\mathbf{u}^T$, so the squared norm appears. Another slip is adding the entries of $\mathbf{u}$ (15).

Q.62 NAT · 2 MARKS

Machine Learning › LDA, SVM & trees

Q.62 Details of ten international cricket games between two teams "Green" and "Blue" are given in Table C. This table consists of matches played on different pitches, across formats along with their winners. The attribute *Pitch* can take one of two values: spin-friendly (represented as S) or pace-friendly (represented as F). The attribute *Format* can take one of two values: one-day match (represented as O) or test match (represented as T).

A cricket organization would like to use the information given in Table C to develop a decision-tree model to predict outcomes of future games between these two teams.

To develop such a model, the computed InformationGain(C, *Pitch*) with respect to the Target is _____ (rounded off to **two** decimal places).

Table C

Match Number	Pitch	Format	Winner (Target)
1	S	T	Green
2	S	T	Blue
3	F	O	Blue
4	S	O	Blue
5	F	T	Green
6	F	O	Blue
7	S	O	Green
8	F	T	Blue
9	F	O	Blue
10	S	O	Green

Official answer: **0.12 to 0.13**

SOLUTION

1. Count the target over all 10 matches. Green won matches 1, 5, 7 and 10, which is 4 matches. Blue won the other 6. So $P(\text{Green}) = 0.4$ and $P(\text{Blue}) = 0.6$.
2. Entropy of the target: $H(C) = -0.4\log_2 0.4 - 0.6\log_2 0.6 = 0.5288 + 0.4422 = 0.9710$.
3. Split on Pitch. Pitch S holds matches 1, 2, 4, 7 and 10 with winners G, B, B, G, G, so 3 Green and 2 Blue. Pitch F holds matches 3, 5, 6, 8 and 9 with winners B, G, B, B, B, so 1 Green and 4 Blue.
4. $H(S) = -0.6\log_2 0.6 - 0.4\log_2 0.4 = 0.9710$.
5. $H(F) = -0.2\log_2 0.2 - 0.8\log_2 0.8 = 0.4644 + 0.2575 = 0.7219$.
6. Each branch holds 5 of the 10 matches, so the weighted child entropy is $\frac{5}{10}(0.9710) + \frac{5}{10}(0.7219) = 0.8464$.
7. $$\text{IG}(C, \text{Pitch}) = H(C) - \sum_v \frac{|C_v|}{|C|} H(C_v) = 0.9710 - 0.8464 \approx 0.1245 \approx 0.12.$$

KEY IDEA

Information gain is the parent's entropy minus the size-weighted average entropy of the children. Count the class split in the whole table and in each Pitch branch, then compute the binary entropies.

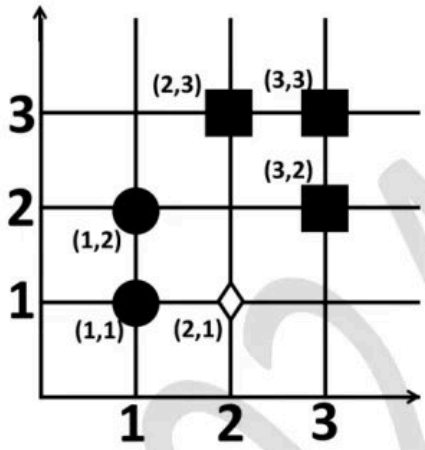
WATCH OUT FOR

Using the wrong class counts per branch, for example miscounting the Green wins on pace-friendly pitches. Another slip is forgetting to weight each branch by its share of the matches, or using natural log instead of base 2.

Q.63 NAT · 2 MARKS

Machine Learning › Logistic, k-NN, naive Bayes

Q.63 Given the two-dimensional dataset consisting of 5 data points from two classes (circles and squares) and assume that the Euclidean distance is used to measure the distance between two points. The minimum odd value of k in k-nearest neighbor algorithm for which the diamond ($\diamond$) shaped data point is assigned the label square is _____.



Official answer: **5**

SOLUTION

1. Read the points off the plot: circles at $(1,1)$ and $(1,2)$; squares at $(2,3)$, $(3,3)$ and $(3,2)$; the query (diamond) is at $(2,1)$.
2. Euclidean distances from $(2,1)$: to $(1,1)$ it is $\sqrt{1^2+0^2}=1$; to $(1,2)$ it is $\sqrt{1+1}=\sqrt{2}$; to $(3,2)$ it is $\sqrt{1+1}=\sqrt{2}$; to $(2,3)$ it is $\sqrt{0+4}=2$; to $(3,3)$ it is $\sqrt{1+4}=\sqrt{5}$.
3. Order by distance: circle $(1,1)$ at 1; circle $(1,2)$ and square $(3,2)$ tied at $\sqrt{2}\approx 1.41$; square $(2,3)$ at 2; square $(3,3)$ at ≈ 2.24.
4. $k=1$: the single nearest point is the circle at $(1,1)$, so the label is circle.
5. $k=3$: the three nearest points are $(1,1)$, $(1,2)$ and $(3,2)$. The tie at $\sqrt{2}$ is fully included, because both tied points fall inside the three, so there is no cut-off ambiguity. That gives 2 circles and 1 square, so the majority is circle.
6. $k=5$: all five points vote, giving 2 circles and 3 squares, so the majority is square.
7. The odd values $k=1$ and $k=3$ both give circle, and $k=5$ gives square, so the minimum odd k that assigns the label square is 5.

KEY IDEA

In k-NN, rank all training points by Euclidean distance from the query and take a majority vote among the k nearest. Compute the distances once, then test each odd k in turn.

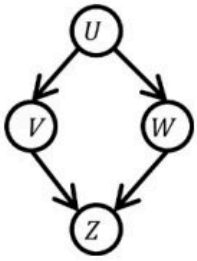
WATCH OUT FOR

Counting only the two squares nearest to the query, or assuming that the square at $(3, 2)$ is farther than the circle at $(1, 2)$. These two are tied at $\sqrt{2}$, so at $k = 3$ both are included, and the circles still win 2 to 1. With 5 points in total, the only way for squares to win a majority is to use all of them.

Q.64 NAT · 2 MARKS

Artificial Intelligence › Uncertainty

Q.64 Given the following Bayesian Network consisting of four Bernoulli random variables and the associated conditional probability tables:



	$P(\cdot)$
$U = 0$	0.5
$U = 1$	0.5

	$P(V = 0 \mid \cdot)$	$P(V = 1 \mid \cdot)$
$U = 0$	0.5	0.5
$U = 1$	0.5	0.5

	$P(W = 0 \mid \cdot)$	$P(W = 1 \mid \cdot)$
$U = 0$	1	0
$U = 1$	0	1

		$P(Z = 0 \mid \cdot)$	$P(Z = 1 \mid \cdot)$
$V = 0$	$W = 0$	0.5	0.5
$V = 0$	$W = 1$	1	0
$V = 1$	$W = 0$	1	0
$V = 1$	$W = 1$	0.5	0.5

The value of $P(U = 1, V = 1, W = 1, Z = 1) =$ _____ (rounded off to **three** decimal places).

Official answer: **0.125**

SOLUTION

1. The network has edges $U \to V$, $U \to W$, $V \to Z$, $W \to Z$, so the joint distribution factorises as

$$P(U,V,W,Z) = P(U)\,P(V \mid U)\,P(W \mid U)\,P(Z \mid V,W).$$

2. Prior: from the first table, $P(U=1) = 0.5$.
3. From the table for V, row $U=1$: $P(V=1 \mid U=1) = 0.5$.
4. From the table for W, row $U=1$: $P(W=1 \mid U=1) = 1$.
5. From the table for Z, row $V=1, W=1$: $P(Z=1 \mid V=1, W=1) = 0.5$.
6. Multiply the four factors:

$$P(U=1,V=1,W=1,Z=1) = 0.5 \times 0.5 \times 1 \times 0.5 = 0.125.$$

KEY IDEA

A Bayesian network's joint probability is the product of each node's conditional probability given its parents. Read the one row of each table that matches the assignment asked for, then multiply.

WATCH OUT FOR

Reading the wrong row of a table, for example using the $V=0, W=0$ row of the Z table. Another slip is to treat V and W as independent of U and drop the U conditioning. Here W is fully determined by U, so $P(W=1 \mid U=1) = 1$, not 0.5.

Q.65 NAT · 2 MARKS

Probability and Statistics › Descriptive statistics

Two fair coins are tossed independently. X is a random variable that takes a value of 1 if both tosses are heads and 0 otherwise. Y is a random variable that takes a value of 1 if at least one of the tosses is heads and 0 otherwise.

The value of the covariance of X and Y is _____ (rounded off to three decimal places).

Official answer: **0.062 to 0.063**

SOLUTION

1. Two fair coins give four equally likely outcomes: HH, HT, TH, TT, each with probability $\frac{1}{4}$.
2. $X = 1$ only for HH, so $E[X] = P(X=1) = \frac{1}{4}$.
3. $Y = 0$ only for TT, so $E[Y] = P(Y=1) = 1 - \frac{1}{4} = \frac{3}{4}$.
4. $XY = 1$ only when both $X=1$ and $Y=1$. Outcome HH has $X=1$ and also $Y=1$ (it has at least one head), so $XY=1$ exactly when the outcome is HH. Hence $E[XY] = P(HH) = \frac{1}{4}$.
5. Use $\mathrm{Cov}(X,Y) = E[XY] - E[X]\,E[Y]$.
6. $$\mathrm{Cov}(X,Y) = \frac{1}{4} - \frac{1}{4}\cdot\frac{3}{4} = \frac{1}{4} - \frac{3}{16} = \frac{1}{16} = 0.0625.$$
7. The exact value is 0.0625 (rounded to three decimals it is about 0.063).

KEY IDEA

Write each indicator's expectation as a probability, then use $\mathrm{Cov}(X,Y) = E[XY] - E[X]E[Y]$. Because $X=1$ forces $Y=1$, the product XY equals X.

WATCH OUT FOR

Treating X and Y as independent and answering 0. They are dependent, since "both heads" is a special case of "at least one head". Another slip is computing $E[XY]$ as $E[X]E[Y]$.

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 2024 question paper, organised by IISc Bengaluru; 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 IISc Bengaluru 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