CS Exit Ethiopia

Questions & Worked Answers

Computer Science Exit Examination

235 complete questions and 2 source variants

1 selected course | 18 November 2026

Selected courses

01. Introduction to Artificial Intelligence

https://csexitethiopia.com

abelasnake2003@gmail.com | Telegram: @Abechu7

Course contents

Each course begins on a new page. PDF bookmarks provide question, study-note section and revision-topic navigation.

Introduction to Artificial Intelligence 3

Reading the answers

Every question includes its answer, explanation, choice-by-choice contrasts and available evidence. Primary note proof quotes the study notes. Related reading and module passages are labelled separately; a module reference is not presented as primary note proof.

Introduction to Artificial Intelligence

Questions & Worked Answers

235 questions + 2 source variants

Question 1 | Q00001

Which one of the following is concerned with the meaning of a sentence in knowledge representation?

A. Structure

B. Computational aspect

C. Semantics

D. Syntax

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer C: Semantics

Semantics assigns meaning to sentences; syntax restricts their legal construction (PDF p. 43). The stem asks about meaning rather than representation layout or computational efficiency.

Why each choice is right or wrong

A - Incorrect: Structure concerns arrangement, not the stated meaning rule. See PDF p. 43.

B - Incorrect: Computational aspects concern processing rather than the interpretation definition. See PDF p. 43.

C - Correct: This is precisely the module definition of assigning meaning. See PDF p. 43.

D - Incorrect: Syntax decides legal symbols and sentence construction. See PDF p. 43.

Primary note proof

Study-note section AI24

> Knowledge representation encodes facts, objects, relations, events and rules so a machine can use them. Syntax defines legal expressions; semantics defines their meaning. A knowledge base stores assertions; an inference system derives consequences. Declarative knowledge says what is true; procedural knowledge specifies how to act.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 43 | lines 7-8

> o Semantic also involves assigning a meaning to each sentence.
> o It can be two logics : Propositional Logics and Predicate logics

Question sources

MOE Computer Science Exit Exam · 2016 · Question 1

Question 2 | Q00007

Which one of the following agents works by searching and matching the rules defined by percent and stored internal state?

A. Model-based reflex agent

B. Goal-based agent

C. Simple reflex agent

D. Utility based agent

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer A: Model-based reflex agent

Taking “percent” as a typographical error for percept, the stored internal state records percept history and the model specifies world evolution (PDF p. 14). That distinguishes model-based reflex selection from a simple current-percept rule.

Why each choice is right or wrong

A - Correct: Combines a model and retained internal state. See PDF p. 14, 14.

B - Incorrect: Goal-based selection needs a desired goal, absent here. See PDF p. 14, 14.

C - Incorrect: Ignores percept history and therefore cannot account for the stated memory. See PDF p. 14, 14.

D - Incorrect: Requires a utility measurement, absent here. See PDF p. 14, 14.

Primary note proof

Study-note section AI06

> A simple reflex agent matches condition-action rules against the current percept only: if dirty, clean. It has no persistent world-state memory. This is adequate when the current observation contains everything needed for its rule, but partial observability can make the same percept require different actions.

Study-note section AI06

> A model-based reflex agent maintains internal state using the previous state, previous action, current percept, and models of world change and action effects. It then applies rules to that state. Both kinds can use condition-action rules; memory and a world model distinguish them.

Study-note section AI07

> A goal-based agent considers future action sequences that reach a desired state. A utility-based agent compares degrees of desirability and trade-offs, choosing maximum expected utility: $EU(a)=\Sigma s\ P(s|a,evidence)U(s)$. Reaching the destination is a goal; balancing arrival time, safety and cost requires preferences among outcomes.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 14 | lines 29-30

> - Internal State: It is a representation of the current state based on percept history.
> • These agents have the model, "which is knowledge of the world" and based on the model they

Introduction to Artificial Intelligence module | PDF page 14 | lines 17-18

> • The Simple reflex agent works on Condition-action rule, which means it maps the current state to
> action.

Question sources

MOE Computer Science Exit Exam · 2016 · Question 7

Question 3 | Q00028

Assume the entire state space and transition model are known for every task below, despite any hidden current state. Which problem/environment coupling is incorrect?

A. Contingency problem -> partially observable, non-deterministic

B. Multiple state problem deterministic, partially observable

C. Exploration problem -> partially observable, deterministic

D. Single state problem-> fully observable, deterministic

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer C: Exploration problem -> partially observable, deterministic

A known deterministic model with incomplete observation can be handled by a belief state representing several possible current states. Unknown state spaces or transition models create the need for exploration; partial observability alone does not. Nondeterministic partial observation can require contingency branches, while full observation with known deterministic transitions supports a single-state formulation. The explicit knowledge assumption therefore makes the exploration pairing the incorrect one.

Why each choice is right or wrong

A - Incorrect: This is a valid coupling: partial observation and nondeterminism can require contingency branches, so it is not the incorrect pairing.

B - Incorrect: This is valid under the known-model assumption: incomplete observation can require a belief state, so it is not selected.

C - Correct: This is the incorrect coupling: hidden state with a known deterministic model does not by itself require exploration.

D - Incorrect: This is valid: full observation and known deterministic transitions support a single current state.

Primary note proof

Study-note section AI10

> With a known deterministic model and full observation, a single current state can be used. Without observation but with known deterministic transitions, a sensorless or multiple-state problem tracks a belief state: a set of possible states.

Study-note section AI10

> Nondeterminism or partial observation may require a contingency plan with branches based on later percepts. Exploration or online search is required when the state space or transition model is unknown; observation alone does not identify this category. These classifications can overlap. A pairing based only on “partially observable, deterministic” cannot distinguish a known sensorless problem from an unknown exploration problem.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 13 | lines 23-24

> ✓ Deterministic / Non-deterministic − If the next state of the environment is completely determined
> by the current state and the actions of the agent, then the environment is deterministic; otherwise it

Introduction to Artificial Intelligence module | PDF page 13 | lines 14-15

> ✓ Observable / Partially Observable − If it is possible to determine the complete state of the environment at each time point from the percepts it is observable; otherwise it is only partially

Question sources

MOE Computer Science Exit Exam · 2016 · Question 28

Question 4 | Q00033

Which one of the following is correct first order logic for the statement "There is someone standing at the door and is tall"?

A. ∀x (StandingAt (x; door)

B. ∀x (StandingAt (x; door)^ tall(x))

C. ∃x (StandingAt (x; door)^ tall(x))

D. ∃x (StandingAt (x; door) -> toll(x))

Hint

Translate the statement carefully, then check the scope of its quantifiers, connectives or inference rule. Distinguish guaranteed consequences from plausible generalizations or explanations.

Answer C: ∃x (StandingAt (x; door)^ tall(x))

“There is someone” requires ∃x, and “standing ... and ... tall” requires StandingAt(x,door) ∧ tall(x), so C asserts one witness with both properties (PDF pp. 46, 49). An implication does not assert its antecedent and therefore cannot assert that anyone is at the door.

Why each choice is right or wrong

A - Incorrect: Incomplete formula with a universal quantifier and no tall predicate. See PDF p. 49, 46.

B - Incorrect: Claims everyone has both properties rather than at least one person. See PDF p. 49, 46.

C - Correct: Existentially binds the same x in both conjuncts. See PDF p. 49, 46.

D - Incorrect: Uses implication instead of conjunction and also spells the predicate toll. See PDF p. 49, 46.

Primary note proof

Study-note section AI26

> Use ∀x(Human(x)→Mortal(x)) for all humans, and ∃x(Student(x)∧Tall(x)) for at least one tall student. These are useful translation patterns, not rules forbidding other connectives. ∀x∃y Likes(x,y) allows a different y for each x; ∃y∀x Likes(x,y) requires one shared y. Negation swaps quantifiers: ¬∀x P(x) ≡ ∃x¬P(x). Only occurrences within the correct quantifier’s scope are bound.

Study-note section AI25

> Propositional atoms represent whole true-or-false claims. ¬P negates; P∧Q requires both; P∨Q requires at least one; P→Q is false only when P is true and Q false; P↔Q requires matching truth values. Thus P→Q ≡ ¬P∨Q. De Morgan gives ¬(P∧Q) ≡ ¬P∨¬Q.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 49 | lines 23-24

> 2. Existential Quantifier: - are the type of quantifiers, which express that the statement within its scope
> is true for at least one instance of something.

Introduction to Artificial Intelligence module | PDF page 46 | lines 14-15

> - Conjunction: A sentence which has ∧ connective such as, P ∧ Q is called a conjunction.
> • Example: Rohan is intelligent and hardworking. It can be written as,

Question sources

MOE Computer Science Exit Exam · 2016 · Question 33

Question 5 | Q00057

Which statements are incorrect? Interpret rational action as maximizing expected performance given available information; do not assume a successful action always exists. Select all that apply.

A. A rational agent know the result of its action, but unrealistic

B. Rational agent take action but not 100% sure of success

C. Rational agent act to achieve utmost success based on information

D. An omniscience agent takes actions for 100% sure of success

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer A, D: A rational agent know the result of its action, but unrealistic; An omniscience agent takes actions for 100% sure of success

Rationality concerns choosing the action with the best expected performance given the agent's information. It does not mean knowing actual outcomes in advance, so the first claim incorrectly attributes omniscience to a rational agent. Even omniscience would not ensure success if all available actions fail. The remaining statements correctly allow uncertainty while requiring the best expected choice under the available evidence.

Why each choice is right or wrong

A - Correct: This is false: rationality does not imply knowing actual action outcomes in advance.

B - Incorrect: This is true: a rational action need not guarantee success.

C - Incorrect: This is true when utmost success means best expected performance given the available information.

D - Correct: This is false: complete outcome knowledge cannot guarantee that a successful action exists.

Primary note proof

Study-note section AI04

> A performance measure defines success externally: safety, useful deliveries, elapsed time, or cost. Rationality depends on that measure, prior knowledge, percept history, and available actions. A rational agent selects the action with the best expected performance given its information.

Study-note section AI04

> Omniscience would mean knowing actual action outcomes in advance. Rationality does not require omniscience or guaranteed success. Even complete outcome knowledge cannot create a successful action when every available action fails. Information gathering and exploration may therefore be rational actions.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 12 | lines 20-21

> o It knows the actual effects of its actions and it is impossible in real world.
> → A rational agent behaves according to its precepts and knowledge and attempts to maximize the

Introduction to Artificial Intelligence module | PDF page 12 | lines 22-23

> expected performance.
> 2.4 Structure of Intelligent Agents

Question sources

MOE Computer Science Exit Exam · 2016 · Question 57

Question 6 | Q00059

_____ finds the shortest path to the goal based on cost.

A. Depth first search

B. Iterative deepening

C. Uniform cost search

D. Breath first search

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer C: Uniform cost search

Uniform-cost search explicitly seeks the goal path with lowest cumulative cost and uses a priority queue (PDF p. 21). “Shortest” here means minimum cost, under the usual nonnegative-cost setting; minimum depth alone is not sufficient when edge costs vary.

Why each choice is right or wrong

A - Incorrect: Follows depth and is explicitly nonoptimal on p. 19. See PDF p. 21, 21.

B - Incorrect: Increases depth limits, not cumulative-cost priority. See PDF p. 21, 21.

C - Correct: Directly minimizes cumulative path cost. See PDF p. 21, 21.

D - Incorrect: Minimizes steps; equivalent to UCS only when all edge costs are equal. See PDF p. 21, 21.

Primary note proof

Study-note section AI15

UCS uses a priority queue ordered by actual accumulated path cost g(n). Remove a minimum-g node; accept a goal when removed, not when first generated. Relax each edge: if a cheaper path is found, update its frontier cost and parent.

Study-note section AI15

With finite branching and each step at least ε>0, UCS is complete and optimal. Finite graphs permit nonnegative edges with proper duplicate handling. Dijkstra uses the same minimum-distance principle, often computing all reachable distances; UCS commonly stops at a goal and generates states on demand. Queue initialization is an implementation choice, not a fundamental distinction. Negative edges invalidate the usual finalized-distance argument.

Study-note section AI13

BFS removes the oldest frontier entry using a FIFO queue. Start with the root, remove the head, test it, append unseen successors at the tail, and repeat. It explores depths 0,1,2,... and is complete for finite b and a finite-depth goal. It minimizes number of edges, so minimum cost requires equal positive step costs.

Study-note section AI16

DLS runs DFS to depth l and reports cutoff separately from failure. It misses goals deeper than l. Iterative deepening runs DLS with limits 0,1,2,...; for finite b it finds the shallowest solution with O(b^d) time and O(bd) frontier space, optimal for equal step costs. Repeated shallow work is small compared with the deepest level.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 21 | lines 20-21

• The primary goal of the uniform-cost search is to find a path to the goal node which has the lowest
cumulative cost.

Introduction to Artificial Intelligence module | PDF page 21 | lines 26-26

• Uniform cost search is equivalent to BFS algorithm if the path cost of all edges is the same.

Question sources

MOE Computer Science Exit Exam · 2016 · Question 59

Question 7 | Q00172

The data structure used in the standard implementation of Depth First Search is ?

A. Stack

B. Array

C. Linked List

D. Queue

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer A: Stack

DFS uses a stack (PDF p. 18): the newest pending state is explored first. Arrays or linked lists could implement a stack, but they do not themselves specify the required access order.

Why each choice is right or wrong

A - Correct: Explicitly named implementation structure. See PDF p. 18.

B - Incorrect: May store a stack but does not specify LIFO behavior. See PDF p. 18.

C - Incorrect: May implement a stack but is not the stated abstract structure. See PDF p. 18.

D - Incorrect: FIFO queue is used for BFS on p. 20. See PDF p. 18.

Primary note proof

Study-note section AI14

> DFS follows the deepest pending branch using a LIFO stack or recursion. FILO means the same order as LIFO. With the same children as BFS, a left-first DFS visits S,A,C before B; push rightmost children first to pop the leftmost first.

Study-note section AI13

> BFS removes the oldest frontier entry using a FIFO queue. Start with the root, remove the head, test it, append unseen successors at the tail, and repeat. It explores depths 0,1,2,... and is complete for finite b and a finite-depth goal. It minimizes number of edges, so minimum cost requires equal positive step costs.

Study-note section AI-S02

> Binary search exploits the order of a sorted sequence: a comparison narrows the remaining candidate interval instead of exploring arbitrary successor states. Array bisection therefore differs from breadth-first or depth-first traversal of a general state-space graph. A storage container such as an array does not by itself specify FIFO or LIFO frontier behavior.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 18 | lines 34-35

> • DFS uses a stack data structure for its implementation.
> • The process of the DFS algorithm is similar to the BFS algorithm.

Question sources

Addis Ababa University mock exam · Question 74

Question 8 | Q00178

Among the following, which kind of algorithm is used in the Game tree to make decisions of Win/Lose?

A. Heuristic Search Algorithm

B. Min/Max Algorithm

C. Depth First Search Algorithm

D. Greedy Search Algorithm

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer B: Min/Max Algorithm

Minimax recursively backs up utilities, choosing maxima on MAX turns and minima on MIN turns, so it selects moves against an optimal opponent (PDF pp. 33-34). DFS is its traversal mechanism rather than the adversarial decision rule.

Why each choice is right or wrong

A - Incorrect: A broad category, not the specified alternating-player rule. See PDF p. 33, 34.

B - Correct: Provides the game-tree decision procedure. See PDF p. 33, 34.

C - Incorrect: Traverses the tree but alone does not compare opponent utilities. See PDF p. 33, 34.

D - Incorrect: Chooses local heuristic promise rather than alternating MAX/MIN decisions. See PDF p. 33, 34.

Primary note proof

Study-note section AI22

> Minimax assigns terminal utilities, takes minima at MIN nodes and maxima at MAX nodes, then chooses a maximum-valued root action. It gives the best guaranteed outcome against optimal opposition within the modeled tree. For stochastic outcomes add chance nodes weighted by probabilities.

Study-note section AI23

> At a depth cutoff use an evaluation function estimating position desirability. Positive favors the designated MAX player, negative MIN; the scale is a convention, with +1/0/−1 often win/draw/loss. Material counts weighted pieces; mobility counts legal options; territory scores controlled regions. These are different features, often combined. Heuristics improve limited computation, not guarantee victory.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 33 | lines 9-10

Mini-max algorithm is a recursive or backtracking algorithm which is used in decision-making and
game theory. It provides an optimal move for the player assuming that opponent is also playing

Introduction to Artificial Intelligence module | PDF page 34 | lines 14-15

• If N is a Max move, take max of successors
• If N is a Min move, take min of successors

Question sources

Addis Ababa University mock exam · Question 80

Computer Science Exit Exam · Sene 2016 · Question 28-2

Computer Science Exit Exam practice collection · Question 205

Question 9 | Q00179

Agents' behavior can be best described by

A. Perception sequence

B. Agent function is performing

C. Sensors and Actuators

D. Environment in which agent

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer B: Agent function is performing

The module evaluates behavior and performance through the agent function, a mapping from perception history to action (PDF p. 11). B is grammatically incomplete but identifies that function.

Why each choice is right or wrong

A - Incorrect: Input history alone does not specify the selected actions. See PDF p. 11.

B - Correct: Names the function that describes action behavior. See PDF p. 11.

C - Incorrect: Describe input/output mechanisms, not the complete action mapping. See PDF p. 11.

D - Incorrect: Provides context, not the mapping. See PDF p. 11.

Primary note proof

Study-note section AI03

A percept is one observation; a percept sequence is the entire observation history. An agent function maps percept sequences to actions: f:P*→A. An agent program implements that mapping on an architecture; agent = architecture + program. Environment means the setting in

> which it senses and acts, not merely the search space. Autonomy means relying substantially on experience rather than only a designer's prior instructions.

Study-note section AI03

> An agent perceives an environment through sensors and acts through actuators or effectors. Cameras sense; motors move. Software sensors may be incoming messages and software actions may be API requests. Humans are agents but are not artificial agents.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 11 | lines 12-13

> • The behaviour and performance of intelligent agents have to be evaluated in terms of the agent function.

Question sources

Addis Ababa University mock exam · Question 81

Question 10 | Q00180

Which type of agent deals with happy and unhappy states?

A. Simple reflex agent

B. Model based agent

C. Utility based agent

D. Learning agent

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer C: Utility based agent

Happy/unhappy is an informal description of desirability. The utility function assigns each state a real number measuring how well actions achieve goals (PDF p. 15). The module does not literally use those emotional terms.

Why each choice is right or wrong

A - Incorrect: Selects current-percept rules, not ranked desirability. See PDF p. 15.

B - Incorrect: Tracks world state, without requiring desirability scores. See PDF p. 15.

C - Correct: Uses the state preference measure described in the stem. See PDF p. 15.

D - Incorrect: Improves through experience rather than defining state desirability. See PDF p. 15.

Primary note proof

Study-note section AI07

A goal-based agent considers future action sequences that reach a desired state. A utility-based agent compares degrees of desirability and trade-offs, choosing maximum expected utility: EU(a)=Σs P(s | a,evidence)U(s). Reaching the destination is a goal; balancing arrival time, safety and cost requires preferences among outcomes.

Study-note section AI06

A simple reflex agent matches condition-action rules against the current percept only: if dirty, clean. It has no persistent world-state memory. This is adequate when the current observation contains everything needed for its rule, but partial observability can make the same percept require different actions.

Study-note section AI08

The performance element selects actions. The learning element changes the agent to improve future performance. The critic evaluates behavior against a performance standard and supplies feedback. The problem generator suggests exploratory actions likely to yield useful experience.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 15 | lines 18-19

• The utility function maps each state to a real number to check how efficiently each action achieves
the goals.

Question sources

Addis Ababa University mock exam · Question 82

Question 11 | Q00181

What is the purpose of "Agent" in Artificial Intelligence?

A. Mapping of goal sequence to an action

B. Work without the direct interference of the people

C. Mapping of environment sequence to an action

D. Mapping of precept sequence to an action

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer D: Mapping of precept sequence to an action

The module defines Agent Function as a map from the percept sequence to an action (PDF p. 10). “Precept” is a spelling error for percept. D matches the definition in the current option order.

Why each choice is right or wrong

A - Incorrect: Goals are not the input sequence in the definition. See PDF p. 10.

B - Incorrect: Describes autonomy, not the function sought. See PDF p. 10.

C - Incorrect: The environment itself is not a percept sequence. See PDF p. 10.

D - Correct: Matches the stated percept-sequence-to-action mapping. See PDF p. 10.

Primary note proof

Study-note section AI03

> A percept is one observation; a percept sequence is the entire observation history. An agent function maps percept sequences to actions: f:P*→A. An agent program implements that mapping on an architecture; agent = architecture + program. Environment means the setting in which it senses and acts, not merely the search space. Autonomy means relying substantially on experience rather than only a designer’s prior instructions.

Study-note section AI04

> A performance measure defines success externally: safety, useful deliveries, elapsed time, or cost. Rationality depends on that measure, prior knowledge, percept history, and available actions. A rational agent selects the action with the best expected performance given its information.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 10 | lines 20-21

> • Agent Function − It is a map from the precept sequence to an action.
> • An agent is anything that can perceive its environment through sensors and acts upon that

Question sources

Addis Ababa University mock exam · Question 83

Artificial Intelligence practice questions · Question 25

Question 12 | Q00182

Here “exploratory learning” means discovering structure in data without target labels, oracle labeling or reward-driven interaction. Which learning regime is meant?

A. Supervised learning

B. Active learning

C. Unsupervised learning

D. Reinforcement learning

E. None of the Above

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer C: Unsupervised learning

Unsupervised learning seeks patterns such as clusters or associations without target labels. Active learning asks an oracle to label selected cases, supervised learning fits labeled input-target examples, and reinforcement learning improves behavior from reward-driven interaction. Specifying the feedback makes the intended use of exploratory precise. The term alone would have several usages, but the stated unlabeled structure-discovery task identifies a definite learning regime.

Why each choice is right or wrong

A - Incorrect: Supervised learning uses labeled targets, excluded by the stem.

B - Incorrect: Active learning obtains oracle labels, excluded by the stem.

C - Correct: Unsupervised learning discovers structure without target labels.

D - Incorrect: Reinforcement learning uses reward-driven interaction, excluded here.

E - Incorrect: The specified regime is present among the choices.

Primary note proof

Study-note section AI32

> Unsupervised learning seeks structure without target labels: clustering groups similar examples; association finds co-occurrence; dimensionality reduction compresses features. Good clusters have high within-cluster similarity and low between-cluster similarity under an appropriate distance. K-means alternates assignment to nearest centroids and centroid updates.

Study-note section AI32

> Semi-supervised learning combines some labels with additional unlabeled data. Active learning chooses informative cases for an oracle to label; it still uses labels. “Exploratory learning” is not an exclusive standard synonym: it can mean unsupervised exploration or active data acquisition, so define the feedback. A task named summarization or time-series analysis is not automatically supervised; its training targets determine that.

Study-note section AI33

> An RL agent observes state information, takes an action, receives reward and a new observation, and improves its policy to maximize expected cumulative reward. Rewards evaluate outcomes rather than supply the correct action label. Rewards may be delayed; not every correct-looking action receives immediate reward.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 69 | lines 7-8

> ⇒ Unsupervised learning means we are only given the x‘s. In either case, the goal is to estimate f. The system is supplied with a set of training examples consisting only of inputs and is required to

Introduction to Artificial Intelligence module | PDF page 69 | lines 15-16

> ⇒ Reinforcement learning is a feedback-based learning method, in which a learning agent gets a reward for each right action and gets a penalty for each wrong action.

Question sources

Addis Ababa University mock exam · Question 84

Question 13 | Q00183

Let F(x,y,t) mean that person x can fool person y at time t. Over nonempty person and time domains, which statement is equivalent to ∀x∃y∃t ¬F(x,y,t)?

A. Everyone can fool some person at some time

B. Every person has someone whom they cannot fool at any time.

C. No person can fool any person at any time.

D. No one can fool everyone all the time

Hint

Translate the statement carefully, then check the scope of its quantifiers, connectives or inference rule. Distinguish guaranteed consequences from plausible generalizations or explanations.

Answer D: No one can fool everyone all the time

For each person x, the formula requires a person y and a time t for which fooling fails. Moving negation across the quantifiers gives ¬∃x∀y∀t F(x,y,t): nobody can fool everyone all the time. The witness and time can depend on x. Requiring failure at every time, or forbidding all fooling events, would be stronger claims.

Why each choice is right or wrong

A - Incorrect: This asserts a successful fooling event for each person, rather than an unsuccessful event.

B - Incorrect: This requires one witness person who cannot be fooled at every time, a stronger time scope than the formula requires.

C - Incorrect: This forbids all fooling events, much stronger than requiring one unsuccessful event for each person.

D - Correct: Negating someone fooling everyone at every time yields, for each person, someone and some time when fooling fails.

Primary note proof

Study-note section AI26

> Use ∀x(Human(x)→Mortal(x)) for all humans, and ∃x(Student(x)∧Tall(x)) for at least one tall student. These are useful translation patterns, not rules forbidding other connectives. ∀x∃y Likes(x,y) allows a different y for each x; ∃y∀x Likes(x,y) requires one shared y. Negation swaps quantifiers: ¬∀x P(x) ≡ ∃x¬P(x). Only occurrences within the correct quantifier's scope are bound.

Study-note section AI44

> The formula ∀x∃y∃t¬F(x,y,t) means that for each person there is a person and time at which fooling fails. By quantifier negation this is ¬∃x∀y∀t F(x,y,t), so nobody fools everybody all the time. Forbidding every successful instance is a much stronger claim than this formula.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 49 | lines 12-13

> 1. Universal Quantifier: - is a symbol of logical representation, which specifies that the statement within its range is true for everything or every instance of a particular thing.

Introduction to Artificial Intelligence module | PDF page 49 | lines 23-24

> 2. Existential Quantifier: - are the type of quantifiers, which express that the statement within its scope
> is true for at least one instance of something.

Introduction to Artificial Intelligence module | PDF page 46 | lines 12-13

> - Negation: A sentence such as ¬ P is called negation of P. A literal can be either Positive literal or negative literal.

Question sources

Addis Ababa University mock exam · Question 85

Question 14 | Q00184

A language system must decide which interpretation is intended when a sentence permits more than one meaning or grammatical attachment. Which NLP challenge does this describe?

A. Handling Tokenization

B. Handling Ambiguity of Sentences

C. Handling POS-Tagging

D. Handling Semantics of Sentences

E. All of the mentioned

Hint

Match the input, output and purpose to the appropriate language, perception or physical-action task.

Answer B: Handling Ambiguity of Sentences

The difficulty described is ambiguity: a word or sentence can permit multiple meanings or structural attachments. Tokenization creates units, POS tagging assigns grammatical roles, and semantic analysis concerns meaning. These tasks can help resolve ambiguity, but they are not synonymous with the particular obstacle. The question therefore asks for the named challenge rather than ranking all NLP difficulties.

Why each choice is right or wrong

A - Incorrect: Tokenization decides processing units; it does not name the existence of multiple interpretations.

B - Correct: Ambiguity means that more than one interpretation is possible.

C - Incorrect: POS tagging assigns grammatical categories rather than naming the competing interpretations.

D - Incorrect: Semantics studies meaning generally; the specific obstacle described is multiplicity of interpretations.

E - Incorrect: This includes tokenization, tagging and semantic analysis in general, which do not all name the specific multiple-interpretation obstacle.

Primary note proof

Study-note section AI37

> Tokenization divides text into processing units; tokens may be words, punctuation or subwords. POS tagging assigns grammatical roles; named-entity recognition identifies people, organizations and locations. Parsing identifies grammatical structure; semantics concerns meaning; discourse links sentences; pragmatics uses context and speaker intent. Sentiment analysis predicts evaluative tone. Ambiguity can occur at every level: “bank” has multiple senses, and “I saw the man with a telescope” has multiple attachments.

Study-note section AI37

NLP interprets and generates human language: translation, dialogue, summarization, question answering and text generation. Speech recognition converts audio to text; speech synthesis generates speech. They overlap language processing but solve different input/output tasks.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 81 | lines 6-7

- NLP is difficult because Ambiguity and Uncertainty exist in the language.
- Ambiguity

Introduction to Artificial Intelligence module | PDF page 80 | lines 18-19

- Semantic analysis is concerned with the meaning representation. It mainly focuses on the literal meaning of words, phrases, and sentences.

Question sources

Addis Ababa University mock exam · Question 86

Question 15 | Q00199

A goal-directed search strategy uses a problem-specific estimate of remaining cost to order current frontier candidates. Which category is it?

A. Uninformed search

B. Informed search

C. Circular traversal without heuristic estimates

D. Binary search

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer B: Informed search

Informed search uses problem-specific heuristic information to estimate remaining cost or promise. Merely retaining the current state would not establish that category, because uninformed search also uses states, successors and known costs. The revised stem supplies the missing estimate. Binary search concerns an ordered interval rather than this state-space frontier, and circular traversal does not identify the heuristic-informed category being tested.

Why each choice is right or wrong

A - Incorrect: Uninformed search lacks the extra remaining-cost heuristic specified here.

B - Correct: Informed search uses problem-specific heuristic estimates to guide the frontier.

C - Incorrect: Traversal without heuristic estimates does not meet the specified remaining-cost-guided strategy.

D - Incorrect: Binary search needs an ordered search interval, not merely goal-directed state exploration.

Primary note proof

Study-note section AI17

> A heuristic h(n) estimates remaining cost or evaluates promise using problem-specific knowledge. It can reduce explored work, but computing it also costs time. Uninformed algorithms use only the problem definition and known costs; UCS remains uninformed despite using g.

Study-note section AI-S02

> Binary search exploits the order of a sorted sequence: a comparison narrows the remaining candidate interval instead of exploring arbitrary successor states. Array bisection therefore differs from breadth-first or depth-first traversal of a general state-space graph. A storage container such as an array does not by itself specify FIFO or LIFO frontier behavior.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 18 | lines 23-24

> • The uninformed search does not contain any domain knowledge such as closeness, the location of
> the goal. It is also called level order traversal.

Introduction to Artificial Intelligence module | PDF page 24 | lines 14-15

> • It uses domain knowledge. In an informed search, problem information is available which can guide
> the search.

Question sources

Mock exam 4 · Question 1

Question 16 | Q00200

An agent’s actions affect later decisions; identical state-action pairs can have different outcomes; and the world changes while the agent deliberates. Which list describes these three properties?

A. Deterministic, fully observable, static

B. Fully observable, sequential, deterministic

C. Partially observable, static, continuous

D. Sequential, non-deterministic, dynamic

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer D: Sequential, non-deterministic, dynamic

The scenario establishes three independent dimensions. Earlier actions affect later decisions, so it is sequential. Outcomes are not fixed by the state-action pair, so it is non-deterministic. Change during deliberation makes it dynamic. It gives no complete observation or continuity specification. Classifying the stated properties avoids an unsupported universal ranking of how difficult different environments must be.

Why each choice is right or wrong

A - Incorrect: Determinism and a static world contradict the uncertain outcomes and ongoing change.

B - Incorrect: Sequential fits, but deterministic outcomes contradict the scenario.

C - Incorrect: A static environment cannot change while the agent deliberates; this list also omits the stated action dependence.

D - Correct: Sequential captures dependent decisions, non-deterministic captures uncertain outcomes, and dynamic captures change during deliberation.

Primary note proof

Study-note section AI09

> Fully observable means sensors reveal relevant state; partially observable means relevant information is hidden or noisy. Deterministic means state and action determine the next state; stochastic means outcomes involve randomness. Strategic uncertainty about an opponent is distinct from dice-like randomness.

Study-note section AI09

> Episodic decisions are independent; sequential decisions affect later situations. Static worlds do not change while deliberating; dynamic ones do. A clock affecting score makes an otherwise static game semidynamic. Discrete quantities have distinct alternatives; continuous quantities vary along real-valued ranges. Known versus unknown concerns knowledge of the transition model, not observability.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 13 | lines 10-11

> 2.7 Properties of Task environments
> ✓ Discrete / Continuous − If there are a limited number of distinct, clearly defined, states of the

Introduction to Artificial Intelligence module | PDF page 13 | lines 26-27

> ✓ Episodic / Non-episodic − In an episodic environment, each episode consists of the agent perceiving
> and then acting. The quality of its action depends just on the episode itself. Subsequent episodes do

Variant 1 | Alternative version 1

An agent's actions affect later decisions; identical state-action pairs can have different outcomes; and the world changes while the agent deliberates. Which list describes these three properties?

A. Partially observable, non-deterministic, sequential, dynamic

B. Deterministic, sequential (non-episodic), dynamic, continuous

C. Static, episodic, continuous, partially observable

D. Fully observable, discrete, episodic, static

Answer A: Partially observable, non-deterministic, sequential, dynamic

Incomplete observations establish partial observability, uncertain transitions establish non-determinism, consequences for later decisions establish sequentiality, and change during deliberation establishes dynamics. The other lists contradict at least one of those conditions. No universal difficulty comparison is needed.

Why each choice is right or wrong

A - Correct: All four properties match the stated scenario.

B - Incorrect: Determinism contradicts uncertain transition outcomes.

C - Incorrect: Static and episodic contradict ongoing change and consequences for later decisions.

D - Incorrect: Full observability, episodicity and static conditions each contradict the scenario.

Primary note proof

Study-note section AI09

> Fully observable means sensors reveal relevant state; partially observable means relevant information is hidden or noisy. Deterministic means state and action determine the next state; stochastic means outcomes involve randomness. Strategic uncertainty about an opponent is distinct from dice-like randomness.

Study-note section AI09

> Episodic decisions are independent; sequential decisions affect later situations. Static worlds do not change while deliberating; dynamic ones do. A clock affecting score makes an otherwise static game semidynamic. Discrete quantities have distinct alternatives; continuous quantities vary along real-valued ranges. Known versus unknown concerns knowledge of the transition model, not observability.

Question sources

Mock exam 4 · Question 2

MOE Computer Science Exit Exam · Question 12-4

National Computer Science Exit Exam · Question 64

National Computer Science Exit Exam

MOE Computer Science Exit Exam

Computer Science Exit Exam practice collection · Question 196

Question 17 | Q00201

An agent which mainly operates based on a knowledge about "how things happen in the world" and representation of the current state based on percept history is

A. Simple reflex agent

B. Goal based agent

C. Utility based agent

D. Model based agent

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer D: Model based agent

The wording reproduces the two model-based factors: knowledge of how the world changes and a current-state representation based on percept history (PDF p. 14). No explicit goal or utility comparison is requested.

Why each choice is right or wrong

A - Incorrect: Ignores history. See PDF p. 14, 14.

B - Incorrect: Adds goals, which are not specified. See PDF p. 14, 14.

C - Incorrect: Adds a utility measure, which is not specified. See PDF p. 14, 14.

D - Correct: Both defining factors match directly. See PDF p. 14, 14.

Primary note proof

Study-note section AI06

> A model-based reflex agent maintains internal state using the previous state, previous action, current percept, and models of world change and action effects. It then applies rules to that state. Both kinds can use condition-action rules; memory and a world model distinguish them.

Study-note section AI07

> A goal-based agent considers future action sequences that reach a desired state. A utility-based agent compares degrees of desirability and trade-offs, choosing maximum expected utility: $EU(a)=\Sigma s\ P(s|a,evidence)U(s)$. Reaching the destination is a goal; balancing arrival time, safety and cost requires preferences among outcomes.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 14 | lines 27-28

- Model: It is knowledge about "how things happen in the world," so it is called a Model-based agent.

Introduction to Artificial Intelligence module | PDF page 14 | lines 29-30

- Internal State: It is a representation of the current state based on percept history.
• These agents have the model, "which is knowledge of the world" and based on the model they

Question sources

Mock exam 4 · Question 3

Question 18 | Q00202

An evaluation technique of classifier accuracy where we randomly partition the data into k mutually exclusive subsets, each approximately equal size and, use Di as test set and others as training set at i-th iteration is _____.

A. Random sampling

B. Cross-validation

C. Stratified cross-validation

D. Holdout method

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer B: k-fold cross-validation.

AI35 describes exactly the rotation in the stem: build k models, each trained on all folds except one, evaluate on that held-out fold, then average the scores. The stem does not say class proportions are preserved, so stratification is not established. A single holdout split or generic random sampling does not specify this complete rotation.

Why each choice is right or wrong

A - Incorrect: Random sampling alone does not specify the disjoint k-fold rotation of training and test subsets.

B - Correct: The described repeated use of each fold as the held-out subset is k-fold cross-validation.

C - Incorrect: Stratification additionally preserves label proportions, a condition not supplied by equal fold sizes alone.

D - Incorrect: A holdout method uses a fixed split rather than rotating through all k disjoint test folds.

Primary note proof

Study-note section AI35

Train data fits parameters; validation data selects models or settings; a held-out test set estimates final performance. In k-fold validation, train k models, each on k−1 folds and evaluate on the remaining fold, then average scores. Stratification approximately preserves label proportions. Fit preprocessing within each training fold. Time-ordered data needs splits respecting time; repeated tuning on the final test set leaks information.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 69 | lines 2-3

with a set of training examples consisting of inputs and corresponding outputs, and is required to
discover the relation or mapping between then, e.g. as a series of rules, or a neural network.
Agent

Question sources

Mock exam 4 · Question 4

Question 19 | Q00203

An intelligent agent is characterized by

A. Learning

B. Reasoning

C. All of them are answers

D. Perception

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer C: All of them are answers

Perception, reasoning and learning all appear in the intelligent-behavior list (PDF p. 6). C aggregates those three capabilities; it does not imply every individual agent must implement every technique.

Why each choice is right or wrong

A - Incorrect: Listed capability, included in C. See PDF p. 6, 6.

B - Incorrect: Listed capability, included in C. See PDF p. 6, 6.

C - Correct: Collects the three listed capabilities. See PDF p. 6, 6.

D - Incorrect: Listed capability, included in C. See PDF p. 6, 6.

Primary note proof

Study-note section AI01

> Artificial intelligence studies systems that perform tasks requiring intelligence: reasoning, learning, planning, perception and language. Applied AI builds useful solutions; cognitive AI uses computational models to investigate biological intelligence. A model should be compared with human behavior or experimental observations before claiming it explains human thought.

Study-note section AI03

> An agent perceives an environment through sensors and acts through actuators or effectors. Cameras sense; motors move. Software sensors may be incoming messages and software actions may be API requests. Humans are agents but are not artificial agents.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 6 | lines 15-16

> • Perception involving image recognition and computer vision
> • Reasoning

Introduction to Artificial Intelligence module | PDF page 6 | lines 3-5

> on the reasoning process. A rational agent is one that acts rationally, that is, is in the best possible
> manner.
> 1.2Typical AI problems

Question sources

Mock exam 4 · Question 5

Question 20 | Q00323

An agent type which has condition-action rule and memory to remind the current state of the world and its evolution is called

A. Simple reflex

B. Model based reflex

C. Goal based

D. Utility based

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer B: Model based reflex

The internal state records percept history; updating it requires how the world evolves and how actions affect it (PDF p. 14). A model-based reflex agent therefore matches the rule-plus-memory description.

Why each choice is right or wrong

A - Incorrect: Uses current percept and ignores the rest of the history. See PDF p. 14, 14.

B - Correct: Has the stated model and internal memory. See PDF p. 14, 14.

C - Incorrect: Would additionally require goal-directed lookahead. See PDF p. 14, 14.

D - Incorrect: Would additionally require utility measurement. See PDF p. 14, 14.

Primary note proof

Study-note section AI06

> A simple reflex agent matches condition-action rules against the current percept only: if dirty, clean. It has no persistent world-state memory. This is adequate when the current observation contains everything needed for its rule, but partial observability can make the same percept require different actions.

Study-note section AI06

> A model-based reflex agent maintains internal state using the previous state, previous action, current percept, and models of world change and action effects. It then applies rules to that state. Both kinds can use condition-action rules; memory and a world model distinguish them.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 14 | lines 29-30

> - Internal State: It is a representation of the current state based on percept history.
> • These agents have the model, "which is knowledge of the world" and based on the model they

Introduction to Artificial Intelligence module | PDF page 14 | lines 33-34

> - How the world evolves
> - How the agent's action affects the world.

Question sources

Computer Science model exam · Question 54

University of Gondar model exam · Question 54

Question 21 | Q00324

Assume that you are employed as a machine learning expert at the Ministry of Education (MoE). Which machine learning technique you would apply for the task of describing the existing higher education institutions as first generation, second generation and third generation.

A. Reinforcement Learning

B. Regression

C. Classification

D. Clustering

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer C: Classification

First, second and third generation are supplied category names. Treating the task as assigning institutions to these predefined labels makes it classification, which the module lists under supervised learning (PDF p. 69). If categories were to be discovered without labels, clustering would instead be relevant.

Why each choice is right or wrong

A - Incorrect: Requires action-dependent rewards, absent here. See PDF p. 69, 69.

B - Incorrect: A regression output is not the category assignment requested. See PDF p. 69, 69.

C - Correct: Matches assignment to the supplied categories. See PDF p. 69, 69.

D - Incorrect: Discovers groupings from inputs rather than assigning known labels. See PDF p. 69, 69.

Primary note proof

Study-note section AI31

> Supervised learning fits a mapping from labeled input-target examples. Classification predicts discrete labels, such as institution generations or spam/nonspam; regression predicts numeric quantities, such as prices. Predictive analytics uses historical patterns to forecast outcomes and can use either. Specify the target before choosing a method.

Study-note section AI32

> Unsupervised learning seeks structure without target labels: clustering groups similar examples; association finds co-occurrence; dimensionality reduction compresses features. Good clusters have high within-cluster similarity and low between-cluster similarity under an appropriate distance. K-means alternates assignment to nearest centroids and centroid updates.

Study-note section AI33

> An RL agent observes state information, takes an action, receives reward and a new observation, and improves its policy to maximize expected cumulative reward. Rewards evaluate outcomes rather than supply the correct action label. Rewards may be delayed; not every correct-looking action receives immediate reward.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 69 | lines 5-6

> o Classification
> o Regression

Introduction to Artificial Intelligence module | PDF page 69 | lines 10-11

o Most common task is clustering - e.g. taxi agent notices —bad traffic days‖
o Association rule learning

Question sources

Computer Science model exam · Question 55

University of Gondar model exam · Question 55

Question 22 | Q00325

Consider the definition of AI: "The effort to design and implement intelligent software that solves problems by identifying best action sequence to achieve the desired goal" can be views AI as

A. A system that acts humanly

B. A system that thinks humanly

C. A system that acts rationally

D. A system that thinks rationally

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer C: A system that acts rationally

The definition emphasizes the best action sequence for a goal. The acting-rationally view evaluates how a system acts and performs (PDF p. 6), rather than imitation of people or the internal laws of thought.

Why each choice is right or wrong

A - Incorrect: Would emphasize humanlike external behavior. See PDF p. 6.

B - Incorrect: Would model human thought processes. See PDF p. 6.

C - Correct: Matches goal-directed best possible action. See PDF p. 6.

D - Incorrect: Emphasizes inference and reasoning soundness rather than the action criterion. See PDF p. 6.

Primary note proof

Study-note section AI02

Thinking humanly models human mental processes; thinking rationally derives conclusions by valid logic. Acting humanly seeks humanlike behavior; acting rationally selects actions expected to achieve objectives. Thought processes and observed behavior are different evaluation targets.

Study-note section AI02

The conversational Turing test asks whether a judge can reliably distinguish a machine from a human through their dialogue. It requires language use, knowledge, reasoning and often learning. It tests an operational behavior criterion, not proof of consciousness. Cognitive models can be tested against response times, errors and psychological experiments.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 6 | lines 2-3

machines that act rationally. The focus is on how the system acts and performs, and not so much on the reasoning process. A rational agent is one that acts rationally, that is, is in the best possible

Question sources

Computer Science model exam · Question 56

University of Gondar model exam · Question 56

Question 23 | Q00326

What is completeness of search algorithms?

A. If there is a difference between the quality of goal states, the agent will find the best one

B. An estimate of how much information the algorithm needs to store for finding a solution

C. An estimate of how many search steps it takes to find a solution

D. Provided that a solution exists, the algorithm will find it

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer D: Provided that a solution exists, the algorithm will find it

Completeness is the guarantee of finding a solution when an appropriately reachable solution exists; the BFS paragraph illustrates it using a finite-depth goal (PDF p. 20). It is distinct from finding the lowest-cost solution and from time or memory requirements.

Why each choice is right or wrong

A - Incorrect: Defines optimality, not completeness. See PDF p. 20, 18.

B - Incorrect: Describes space complexity. See PDF p. 20, 18.

C - Incorrect: Describes time complexity. See PDF p. 20, 18.

D - Correct: Matches solution-finding guarantee. See PDF p. 20, 18.

Primary note proof

Study-note section AI12

Completeness means finding a solution if one exists under stated assumptions; optimality means returning the best solution. Time counts work, often generated nodes; space counts simultaneous storage. Let b be branching factor, d shallowest goal depth, m maximum search depth, and l a depth limit. Graph-search bounds in vertices V and edges E differ from implicit-tree bounds.

Study-note section AI13

BFS removes the oldest frontier entry using a FIFO queue. Start with the root, remove the head, test it, append unseen successors at the tail, and repeat. It explores depths 0,1,2,... and is complete for finite b and a finite-depth goal. It minimizes number of edges, so minimum cost requires equal positive step costs.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 20 | lines 26-27

Completeness: BFS is complete, which means if the shallowest goal node is at some finite depth, then
BFS will find a solution.

Introduction to Artificial Intelligence module | PDF page 18 | lines 20-21

Optimal Solution: If a solution has the lowest cost among all solutions.
3.5 Search Strategies

Question sources

Computer Science model exam · Question 57

University of Gondar model exam · Question 57

Question 24 | Q00327

Which one of the view of defining AI concerned with thought processing

A. Viewing AI as a system which has complete autonomous

B. Viewing AI as a system that think either humanly or think rationally

C. Viewing AI as a system that act either humanly or think rationally

D. Viewing AI as a system that act omnisciently to its environment

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer B: Viewing AI as a system that think either humanly or think rationally

Thinking humanly models human cognition; thinking rationally examines ideal inference (PDF p. 5). Both concern thought processing, so B is the complete pairing.

Why each choice is right or wrong

A - Incorrect: Autonomy alone does not specify thought processing. See PDF p. 5, 5.

B - Correct: Includes both thought-oriented approaches. See PDF p. 5, 5.

C - Incorrect: Mixes an action view with a thought view. See PDF p. 5, 5.

D - Incorrect: Omniscience is not one of the four views. See PDF p. 5, 5.

Primary note proof

Study-note section AI02

> Thinking humanly models human mental processes; thinking rationally derives conclusions by valid logic. Acting humanly seeks humanlike behavior; acting rationally selects actions expected to achieve objectives. Thought processes and observed behavior are different evaluation targets.

Study-note section AI04

> Omniscience would mean knowing actual action outcomes in advance. Rationality does not require omniscience or guaranteed success. Even complete outcome knowledge cannot create a successful action when every available action fails. Information gathering and exploration may therefore be rational actions.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 5 | lines 5-6

> - Are we looking at the thought process or reasoning ability of the system?
> - Or are we only interested in the final manifestations of the system in terms of its

Introduction to Artificial Intelligence module | PDF page 5 | lines 26-27

> think rationally
> 4. The fourth view of AI is that it is the study of rational agents. This view deals with building

Question sources

Computer Science model exam · Question 58

University of Gondar model exam · Question 58

Question 25 | Q00328

Which one of the following statement is false

A. Turing test is a mechanism to validate AI system that act humanly

B. Validating AI system that act rationally can be made by its performance

C. Validating AI system that thinks humanly is impossible

D. Validating AI system that think rationally can be made by using its reasoning logic

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer C: Validating AI system that thinks humanly is impossible

C makes an absolute impossibility claim. The module describes understanding and modeling human thought as the cognitive-science approach (PDF p. 5), and p. 7 explicitly says computers test theories of the human mind. That contradicts treating validation as inherently impossible.

Why each choice is right or wrong

A - Incorrect: Turing's imitation game tests humanlike behavior. See PDF p. 5, 5, 6.

B - Incorrect: Rational-agent view focuses on actions and performance. See PDF p. 5, 5, 6.

C - Correct: This is the false assertion requested: page7 explicitly says computers can test theories about how the human mind works, contradicting an absolute impossibility claim.

D - Incorrect: This statement accords with the rational-thought view on page5: inference mechanisms, reasoning, soundness and completeness are assessment criteria. It is therefore not the false choice.

Primary note proof

Study-note section AI02

> Thinking humanly models human mental processes; thinking rationally derives conclusions by valid logic. Acting humanly seeks humanlike behavior; acting rationally selects actions expected to achieve objectives. Thought processes and observed behavior are different evaluation targets.

Study-note section AI02

> The conversational Turing test asks whether a judge can reliably distinguish a machine from a human through their dialogue. It requires language use, knowledge, reasoning and often learning. It tests an operational behavior criterion, not proof of consciousness. Cognitive models can be tested against response times, errors and psychological experiments.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 5 | lines 12-13

> the human thought process. This view is the cognitive science approach to AI. -Think human like 2. The second approach is best embodied by the concept of the Turing Test. Turing held that in future

Introduction to Artificial Intelligence module | PDF page 5 | lines 15-16

> argument Turing put forward the idea of an 'imitation game', in which a human being and a computer would be interrogated under conditions where the interrogator would not know which

Introduction to Artificial Intelligence module | PDF page 6 | lines 2-3

> machines that act rationally. The focus is on how the system acts and performs, and not so much on the reasoning process. A rational agent is one that acts rationally, that is, is in the best possible

Introduction to Artificial Intelligence module | PDF page 7 | lines 9-10

> Cognitive AI: computers are used to test theories about how the human mind works--for example,
> theories about how we recognize faces and other objects, or about how we solve abstract problems.

Introduction to Artificial Intelligence module | PDF page 5 | lines 22-26

> 3. Logic and laws of thought deals with studies of ideal or rational thought process and inference. The emphasis in this case is on the inference mechanism, and its properties. That is how the system
> arrives at a conclusion, or the reasoning behind its selection of actions is very important in this point of view. The soundness and completeness of the inference mechanisms are important here.
> -
> think rationally

Question sources

Computer Science model exam · Question 59

University of Gondar model exam · Question 59

Question 26 | Q00534

Which one of the following is evaluating the degrees of success of an agent?

A. Performance measure

B. Action

C. Knowledge

D. Perception

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer A: Performance measure

The performance measure determines how successful an agent is (PDF p. 10). Actions, knowledge and percepts influence behavior, but evaluating that behavior needs the criterion. In the alternative option order the answer is B.

Why each choice is right or wrong

A - Correct: Explicit success criterion. See PDF p. 10.

B - Incorrect: A selected operation, not its evaluation criterion. See PDF p. 10.

C - Incorrect: Supports decisions but does not define success. See PDF p. 10.

D - Incorrect: Supplies sensed inputs rather than the measure. See PDF p. 10.

Primary note proof

Study-note section AI04

> A performance measure defines success externally: safety, useful deliveries, elapsed time, or cost. Rationality depends on that measure, prior knowledge, percept history, and available actions. A rational agent selects the action with the best expected performance given its information.

Study-note section AI03

> A percept is one observation; a percept sequence is the entire observation history. An agent function maps percept sequences to actions: f:P*→A. An agent program implements that mapping on an architecture; agent = architecture + program. Environment means the setting in which it senses and acts, not merely the search space. Autonomy means relying substantially on experience rather than only a designer's prior instructions.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 10 | lines 16-17

> • Performance Measure of Agent -It is the criteria, which determines how successful an agent is.
> • Behavior of Agent – It is the action that agent performs after any given sequence of percepts.

Variant 1 | Alternative version 1

Which one of the following is evaluating the degrees of success of an agent?

A. Its Perception

B. Its Performance measure

C. Its action

D. Its prior kowled

Answer B: Its Performance measure

Apply the same study-note rule to this alternative source set; the choice contrasts below show the derivation.

Why each choice is right or wrong

A - Incorrect: Perception supplies input, not the success criterion (PDF p. 10).

B - Correct: The performance measure determines success (PDF p. 10).

C - Incorrect: An action is assessed by the criterion (PDF p. 10).

D - Incorrect: Prior knowledge helps action selection, but is not the success criterion (PDF p. 12).

Primary note proof

Study-note section AI04

> A performance measure defines success externally: safety, useful deliveries, elapsed time, or cost. Rationality depends on that measure, prior knowledge, percept history, and available actions. A rational agent selects the action with the best expected performance given its information.

Study-note section AI03

> A percept is one observation; a percept sequence is the entire observation history. An agent function maps percept sequences to actions: f:P*→A. An agent program implements that mapping on an architecture; agent = architecture + program. Environment means the setting in which it senses and acts, not merely the search space. Autonomy means relying substantially on experience rather than only a designer's prior instructions.

Question sources

Computer Science model exam · 2015 · Question 3

MOE Computer Science Exit Exam · Question 14-4

National Computer Science Exit Exam · Question 65

National Computer Science Exit Exam

MOE Computer Science Exit Exam

Computer Science Exit Exam practice collection · Question 197

Question 27 | Q00541

The component of Expert system which provide interaction between user expert system and Expert system

A. Kowledge base

B. Inference engine

C. User interface

D. Human expert

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer C: User interface

The user interface explicitly provides interaction between the user and the expert system (PDF p. 64). Knowledge storage and rule application serve different roles.

Why each choice is right or wrong

A - Incorrect: Stores domain knowledge. See PDF p. 64.

B - Incorrect: Applies rules to derive conclusions. See PDF p. 64.

C - Correct: Directly provides user-system interaction. See PDF p. 64.

D - Incorrect: Provides expertise rather than the interaction component. See PDF p. 64.

Primary note proof

Study-note section AI40

> An expert system applies specialized knowledge to a domain such as equipment diagnosis. The domain expert supplies expertise; the knowledge engineer elicits and encodes it; the inference engine reasons over a knowledge base and working facts. The user interface handles interaction; an explanation facility presents why or how a conclusion was reached.

Study-note section AI24

> Knowledge representation encodes facts, objects, relations, events and rules so a machine can use them. Syntax defines legal expressions; semantics defines their meaning. A knowledge base stores assertions; an inference system derives consequences. Declarative knowledge says what is true; procedural knowledge specifies how to act.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 64 | lines 20-21

> • User interface provides interaction between user of the ES and the ES itself.
> • Natural Language Processing so as to be used by the user who is well-versed in the task domain

Question sources

Computer Science model exam · 2015 · Question 10

Question 28 | Q00557

Which Terminology of Natural language processing study about meaning full components of word

A. Phonetics

B. Phonology

C. Morphology

D. Pragmatics

Hint

Match the input, output and purpose to the appropriate language, perception or physical-action task.

Answer C: Morphology.

AI38 defines morphology through prefixes, stems and suffixes, such as un-happi-ness and cat plus plural -s. These are the meaningful components of words requested by the stem. Phonetics concerns physical sounds, phonology sound patterns, and pragmatics intended meaning in context; none names the study of internal meaningful word parts.

Why each choice is right or wrong

A - Incorrect: Phonetics studies the physical properties of speech sounds rather than meaningful word components.

B - Incorrect: Phonology studies a language's sound patterns and contrasts rather than morpheme structure.

C - Correct: Morphology analyzes meaningful internal parts such as prefixes, stems and suffixes.

D - Incorrect: Pragmatics interprets intended meaning in context, not the internal components of a word.

Primary note proof

Study-note section AI38

> Morphology studies meaningful word parts: un-happi-ness contains a prefix, stem and suffix; cats contains cat and plural -s. It differs from tokenization, which identifies processing boundaries. Phonetics studies physical speech sounds; phonology studies a language's sound patterns and contrasts.

Study-note section AI38

> Prosody concerns stress, rhythm, timing, prominence and intonation across syllables and phrases. A pitch rise can help signal a question; shifting stress can change emphasis. These cues assist dialogue and speech processing. Pragmatics concerns intended meaning in context, not the internal components of words.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 80 | lines 7-8

> 1. Lexical Analysis and Morphological
> - The first phase of NLP is the Lexical Analysis. This phase scans the source code as a

Question sources

Computer Science model exam · 2015 · Question 26

Question 29 | Q00577

The application area of AI which focus on recognize objects based on patterns in the same way as the human visual system does.

A. Expert System

B. Speech recogition

C. Computer vison

D. Robotics

Hint

Match the input, output and purpose to the appropriate language, perception or physical-action task.

Answer C: Computer vison

Computer vision extracts and interprets image information and explicitly includes object recognition (PDF p. 84). That is the application directly matching visual pattern recognition.

Why each choice is right or wrong

A - Incorrect: Uses domain expertise, not specifically image interpretation. See PDF p. 84, 84.

B - Incorrect: Converts spoken language rather than images. See PDF p. 84, 84.

C - Correct: Directly addresses visual objects. See PDF p. 84, 84.

D - Incorrect: Robotics may use vision but is the broader physical-agent domain. See PDF p. 84, 84.

Primary note proof

Study-note section AI39

> Computer vision interprets images or video: classification names image content, detection locates objects, segmentation labels pixels, and recognition identifies objects or people. Occlusion, lighting change, noise, unfamiliar conditions and incomplete context make open-world perception difficult even when a face detector works well.

Study-note section AI39

> Robotics combines mechanical, electrical, control and computing systems to sense, plan and act physically. Components include sensors, controller, power supply, actuators and an end effector. Locomotion may be legged, wheeled or tracked, with trade-offs in terrain, stability and efficiency. Software can act autonomously without being a robot. Humanoid shape is not evidence of general intelligence.

Study-note section AI37

> NLP interprets and generates human language: translation, dialogue, summarization, question answering and text generation. Speech recognition converts audio to text; speech synthesis generates speech. They overlap language processing but solve different input/output tasks.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 84 | lines 4-5

> o Computer vision automatically extracts, analyzes, and comprehends useful information from a single image or an array of images. This process involves development of algorithms to

Introduction to Artificial Intelligence module | PDF page 84 | lines 20-21

> o Object Recognition – they are installed in supermarkets, cameras, high-end cars such as BMW, GM, and Volvo.

Question sources

Computer Science model exam · 2015 · Question 46

Question 30 | Q00581

From steps of problem solving which one is concerned on looking various sequence of action which lead to the goal state, evaluate it and chose the sequence of optimal action?

A. Execute

B. Goal Formulation

C. problem formulation

D. Search

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer D: Search

The module defines a solution as an action sequence from start to goal and an optimal solution as the lowest-cost such solution (PDF p. 18). Searching is the process of finding/evaluating those sequences, so D matches the activity described.

Why each choice is right or wrong

A - Incorrect: Carries out a selected sequence rather than finding one. See PDF p. 18, 18.

B - Incorrect: Specifies a desired outcome rather than searching paths. See PDF p. 18, 18.

C - Incorrect: Sets up states/actions rather than choosing a path. See PDF p. 18, 18.

D - Correct: Finds the goal-reaching sequence. See PDF p. 18, 18.

Primary note proof

Study-note section AI11

First formulate the goal; then formulate the problem by specifying initial state, legal actions, transition model, goal test, and action costs. Search selects a sequence; execution applies it. A state abstracts away irrelevant world details.

Study-note section AI11

In route finding, states are cities, actions are roads, successors are neighboring cities, and path cost is accumulated distance or time. A solution is a legal path ending at a goal; an optimal solution has minimum cost. Automatic programming or program synthesis searches for a program satisfying a specification; conditional plans and loops require richer representations than a fixed action list.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 18 | lines 19-20

Solution: It is an action sequence which leads from the start node to the goal node.
Optimal Solution: If a solution has the lowest cost among all solutions.

Introduction to Artificial Intelligence module | PDF page 18 | lines 20-21

Optimal Solution: If a solution has the lowest cost among all solutions.
3.5 Search Strategies

Question sources

Computer Science model exam · 2015 · Question 50

Question 31 | Q00600

Which types of rule used in phonology to interpretation of the fluctuation in stress and intonation across the sentences?

A. Phonetic rule

B. Phonemic rule

C. Prosodic rule

D. Morpheme rule

Hint

Match the input, output and purpose to the appropriate language, perception or physical-action task.

Answer C: Prosodic rules.

AI38 identifies stress and intonation as prosodic phenomena, alongside rhythm, timing and prominence. That directly matches the sentence-level variations in the question. Physical sound production and phonemic contrasts describe other aspects of speech; morphemes are meaningful word parts. The phrase phonology in the stem does not change the specific prosodic level requested.

Why each choice is right or wrong

A - Incorrect: Phonetic rules concern physical speech sounds, while the stated stress and intonation patterns are specifically prosodic.

B - Incorrect: Phonemic contrasts distinguish sound categories; they are not the specific stress-and-intonation rule category requested.

C - Correct: Prosodic rules address stress, prominence, rhythm and intonation across syllables and phrases.

D - Incorrect: Morphemes are meaningful word parts, not the rule category for sentence intonation.

Primary note proof

Study-note section AI38

> Morphology studies meaningful word parts: un-happi-ness contains a prefix, stem and suffix; cats contains cat and plural -s. It differs from tokenization, which identifies processing boundaries. Phonetics studies physical speech sounds; phonology studies a language's sound patterns and contrasts.

Study-note section AI38

> Prosody concerns stress, rhythm, timing, prominence and intonation across syllables and phrases. A pitch rise can help signal a question; shifting stress can change emphasis. These cues assist dialogue and speech processing. Pragmatics concerns intended meaning in context, not the internal components of words.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 79 | lines 2-3

> 7.1 Natural language processing
> NLP stands for Natural Language Processing, which is a part of Computer Science, Human language,

Question sources

Computer Science model exam · 2015 · Question 69

Question 32 | Q00611

A system is evaluated by whether a judge can distinguish its natural-language conversation from a human's. Which of the four AI approaches supplies this evaluation criterion?

A. Thinking humanly

B. Acting Humanly

C. Thinking rationally

D. Acting rationally

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer B: Acting Humanly

The criterion is observable humanlike behavior: a judge compares the system's conversation with a human's. That is the acting-humanly approach underlying the conversational Turing test. Natural-language ability alone would not establish this classification; the evaluation criterion does. Thinking approaches concern internal processes, while rational approaches judge valid conclusions or effective actions rather than resemblance to humans.

Why each choice is right or wrong

A - Incorrect: Thinking humanly concerns internal mental processes, not only observed conversation.

B - Correct: Acting humanly evaluates humanlike observable behavior, as in the conversational Turing test.

C - Incorrect: Thinking rationally evaluates valid reasoning, not indistinguishability from human dialogue.

D - Incorrect: Acting rationally evaluates expected objective achievement, which need not resemble a human.

Primary note proof

Study-note section AI02

> Thinking humanly models human mental processes; thinking rationally derives conclusions by valid logic. Acting humanly seeks humanlike behavior; acting rationally selects actions expected to achieve objectives. Thought processes and observed behavior are different evaluation targets.

Study-note section AI02

> The conversational Turing test asks whether a judge can reliably distinguish a machine from a human through their dialogue. It requires language use, knowledge, reasoning and often learning. It tests an operational behavior criterion, not proof of consciousness. Cognitive models can be tested against response times, errors and psychological experiments.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 5 | lines 17-18

> was which, the communication being entirely by textual messages. Turing argued that if the interrogator could not distinguish them by questioning, then it would be unreasonable not to call

Introduction to Artificial Intelligence module | PDF page 6 | lines 10-11

> - Communicating (through natural language).
> - Navigating around obstacles on the streets

Question sources

Computer Science model exam · 2015 · Question 80

Question 33 | Q00619

From which components Intelligence is composed of

A. Reasoning

B. Problem solving

C. Linguistic perception

D. All

Hint

Match the input, output and purpose to the appropriate language, perception or physical-action task.

Answer D: All

Reasoning, learning, language understanding, problem solving and perception are all listed aspects of intelligent behavior (PDF p. 6). D aggregates the capabilities named by the question.

Why each choice is right or wrong

A - Incorrect: One listed capability. See PDF p. 6.

B - Incorrect: Another listed capability. See PDF p. 6.

C - Incorrect: Interpreted as language understanding/perception, both listed. See PDF p. 6.

D - Correct: Includes all three capabilities. See PDF p. 6.

Primary note proof

Study-note section AI01

> Artificial intelligence studies systems that perform tasks requiring intelligence: reasoning, learning, planning, perception and language. Applied AI builds useful solutions; cognitive AI uses computational models to investigate biological intelligence. A model should be compared with human behavior or experimental observations before claiming it explains human thought.

Study-note section AI37

> NLP interprets and generates human language: translation, dialogue, summarization, question answering and text generation. Speech recognition converts audio to text; speech synthesis generates speech. They overlap language processing but solve different input/output tasks.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 6 | lines 3-6

on the reasoning process. A rational agent is one that acts rationally, that is, is in the best possible
manner.
1.2Typical AI problems
While studying the typical range of tasks that we might expect an —intelligent entity‖ to perform,we

Question sources

Computer Science model exam · 2015 · Question 88

Question 34 | Q00622

An individual who involve in solving problems regarding a certain profession and are considered as knowledge providers in a certain field of study known as ----------

A. Domain Expert

B. Knowledge Engineer

C. User

D. All

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer A: Domain Expert

The knowledge engineer acquires information from a subject expert and organizes it into rules (PDF p. 63). The professional who supplies specialist problem-solving knowledge is therefore the domain expert.

Why each choice is right or wrong

A - Correct: Supplies the subject/domain expertise. See PDF p. 63.

B - Incorrect: Elicits and organizes knowledge from the expert. See PDF p. 63.

C - Incorrect: Uses the system rather than necessarily providing expert knowledge. See PDF p. 63.

D - Incorrect: Does not preserve the distinct roles. See PDF p. 63.

Primary note proof

Study-note section AI40

> An expert system applies specialized knowledge to a domain such as equipment diagnosis. The domain expert supplies expertise; the knowledge engineer elicits and encodes it; the inference engine reasons over a knowledge base and working facts. The user interface handles interaction; an explanation facility presents why or how a conclusion was reached.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 63 | lines 10-11

> o He acquires information from subject expert by recording, interviewing, and observing him at work,
> etc. He then categorizes and organizes the information in a meaningful way, in the form of IF-

Question sources

Computer Science model exam · 2015 · Question 91

Question 35 | Q00623

Which statement best describes a rational artificial agent?

A. An artificial agent that selects actions for the best expected performance given its percept history and knowledge.

B. It used only sensors and effectors to be effective

C. It is a combination of software and hardware agents

D. All

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer A: An artificial agent that selects actions for the best expected performance given its percept history and knowledge.

A rational artificial agent senses its environment and chooses actions expected to maximize its performance measure given its information. It need not know actual outcomes in advance. Sensors and effectors alone do not supply an action-selection policy, and combining two vaguely named agent categories is not the architecture-plus-program definition. The choice gives the intended rational-agent criterion without imposing omniscience.

Why each choice is right or wrong

A - Correct: This states the rational action-selection criterion without requiring omniscience.

B - Incorrect: Sensors and effectors provide interaction, but alone do not establish intelligent action selection.

C - Incorrect: Architecture plus program is not a combination of a software agent and a hardware agent.

D - Incorrect: The component statements are defective, so their aggregate is not a valid definition.

Primary note proof

Study-note section AI03

> An agent perceives an environment through sensors and acts through actuators or effectors. Cameras sense; motors move. Software sensors may be incoming messages and software actions may be API requests. Humans are agents but are not artificial agents.

Study-note section AI03

> A percept is one observation; a percept sequence is the entire observation history. An agent function maps percept sequences to actions: f:P*→A. An agent program implements that mapping on an architecture; agent = architecture + program. Environment means the setting in which it senses and acts, not merely the search space. Autonomy means relying substantially on experience rather than only a designer’s prior instructions.

Study-note section AI04

A performance measure defines success externally: safety, useful deliveries, elapsed time, or cost. Rationality depends on that measure, prior knowledge, percept history, and available actions. A rational agent selects the action with the best expected performance given its information.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 11 | lines 2-3

• An Intelligent Agent must sense, must act, and must be autonomous (to some extent). It also must
be rational.

Introduction to Artificial Intelligence module | PDF page 12 | lines 20-21

o It knows the actual effects of its actions and it is impossible in real world.
→ A rational agent behaves according to its precepts and knowledge and attempts to maximize the

Question sources

Computer Science model exam · 2015 · Question 92

Question 36 | Q00625

Which option gives the complete set of correct descriptions of robotics?

A. It is application area of AI emphasis of creation of intelligent and Efficient robot

B. It is a multi disciplinary area of AI

C. It focus on Electrical, Mechanical and software part of robot

D. All

Hint

Match the input, output and purpose to the appropriate language, perception or physical-action task.

Answer D: All

Robotics integrates mechanical, electrical, control and computing systems to sense, plan and act physically. Building intelligent robots is an AI application, and the required disciplines extend across hardware and software. Thus the three individual descriptions are true. Because this is a single-choice question requesting the complete set, the aggregate is selected and each individual description is an incomplete selection.

Why each choice is right or wrong

A - Incorrect: Creating intelligent, efficient robots is a valid application description, but this alone omits the other true descriptions.

B - Incorrect: Robotics is multidisciplinary, but this alone omits the other true descriptions.

C - Incorrect: Electrical, mechanical and software systems are part of robotics, but this alone omits the other true descriptions.

D - Correct: All three descriptions are compatible with robotics as an integrated discipline that builds physical sensing and acting systems.

Primary note proof

Study-note section AI39

> Robotics combines mechanical, electrical, control and computing systems to sense, plan and act physically. Components include sensors, controller, power supply, actuators and an end effector. Locomotion may be legged, wheeled or tracked, with trade-offs in terrain, stability and efficiency. Software can act autonomously without being a robot. Humanoid shape is not evidence of general intelligence.

Study-note section AI01

> Artificial intelligence studies systems that perform tasks requiring intelligence: reasoning, learning, planning, perception and language. Applied AI builds useful solutions; cognitive AI uses computational models to investigate biological intelligence. A model should be compared with human behavior or experimental observations before claiming it explains human thought.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 82 | lines 10-11

> Robotics is a domain in artificial intelligence that deals with the study of creating intelligent and efficient robots. Robots are the artificial agents acting in real world environment. Robots are aimed at

Introduction to Artificial Intelligence module | PDF page 83 | lines 23-24

> ⇒ Components of robotics
> o Power Supply − The robots are powered by batteries, solar power, hydraulic, or pneumatic

Question sources

Computer Science model exam · 2015 · Question 94

Question 37 | Q00646

Which option describes a conventional computer’s instruction-driven storage and retrieval of information, without making a universal comparison with human intelligence?

A. Machine can figure out complete objects even some part is missed

B. Machines are Naturally intelligent

C. Store and Recall information with set of instruction

D. Take precepts from environment with pattern

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer C: Store and Recall information with set of instruction

A conventional computer manipulates information according to program instructions, including instructions for storage and retrieval. This describes a mode of computation rather than a universal difference in intelligence between humans and machines. Programs can also implement pattern learning and object recognition. Those capabilities do not change which option names the particular storage-and-retrieval operation requested.

Why each choice is right or wrong

A - Incorrect: Recognizing an occluded object is a perception capability, not the definition of instruction-driven storage and retrieval.

B - Incorrect: A computer's programmed operations do not establish natural biological intelligence.

C - Correct: Storing and recalling information by executing instructions matches the operation described.

D - Incorrect: Processing perceptual patterns is a sensing or recognition task, not specifically storage and retrieval.

Primary note proof

Study-note section AI-S01

> A conventional computer manipulates digital information according to a program or sequence of instructions. Programs can implement storage, retrieval and learning operations. Neural networks can learn pattern recognition from examples, so executing instructions and learning patterns are not mutually exclusive descriptions. This provides no universal claim that every machine is better or worse than every human at completing missing information.

Study-note section AI34

> Feedforward networks pass signals toward outputs without recurrent state loops; recurrent networks carry state across sequence steps. Training adjusts weights, often through gradients and backpropagation. Backpropagation is not the same as a recurrent connection. Deep learning uses multiple representation layers and can learn features from images, sound and text. CNNs share filters across image locations. ANN is a brain-inspired abstraction, not a literal brain simulation; DL⊂ML⊂AI.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 9 | lines 13-14

> C) Humans recall and store information by patterns rather than algorithms
> D) Machines can‘t think out of the box but humans do

Introduction to Artificial Intelligence module | PDF page 9 | lines 11-12

A) Machines perceive by set of rules rather than patterns
B) Machines can label/figure out a missing part of a given object more effectively than humans

Question sources

Computer Science model exam · 2015 · Question 115

Question 38 | Q00651

Which types of searching strategies is not grouped under blind search?

A. Breadth first search

B. Uniform cost search

C. Constraint satisfaction search

D. Depth limited search

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer C: Constraint satisfaction search

BFS, UCS and depth-limited search appear in the uninformed-search sequence (PDF pp. 18-22). Constraint satisfaction is treated separately, so C is the item not grouped under blind search.

Why each choice is right or wrong

A - Incorrect: BFS is in the uninformed sequence. See PDF p. 18, 21, 36.

B - Incorrect: UCS is in the uninformed sequence. See PDF p. 18, 21, 36.

C - Correct: A separately presented constraint-based problem formulation/search topic. See PDF p. 18, 21, 36.

D - Incorrect: Depth-limited search is in the uninformed sequence. See PDF p. 18, 21, 36.

Primary note proof

Study-note section AI21

Backtracking assigns a variable, tries a legal value, and undoes the choice on failure. Choose the variable with fewest remaining values; break ties by degree; try values that rule out fewest neighbor choices. Forward checking removes incompatible neighbor values. Arc consistency removes a value when a binary constraint offers it no support in the other domain. CSP is a problem framework using search and propagation, not a universal fourth uninformed traversal.

Study-note section AI17

A heuristic h(n) estimates remaining cost or evaluates promise using problem-specific knowledge. It can reduce explored work, but computing it also costs time. Uninformed

algorithms use only the problem definition and known costs; UCS remains uninformed despite using g.

Study-note section AI13

BFS removes the oldest frontier entry using a FIFO queue. Start with the root, remove the head, test it, append unseen successors at the tail, and repeat. It explores depths 0,1,2,... and is complete for finite b and a finite-depth goal. It minimizes number of edges, so minimum cost requires equal positive step costs.

Study-note section AI16

DLS runs DFS to depth l and reports cutoff separately from failure. It misses goals deeper than l. Iterative deepening runs DLS with limits 0,1,2,...; for finite b it finds the shallowest solution with O(b^d) time and O(bd) frontier space, optimal for equal step costs. Repeated shallow work is small compared with the deepest level.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 18 | lines 22-23

3.5.1 Unformed Search Strategies
• The uninformed search does not contain any domain knowledge such as closeness, the location of

Introduction to Artificial Intelligence module | PDF page 21 | lines 17-18

4. Uniform cost search
• Uniform-cost search is a searching algorithm used for traversing a weighted tree or graph.

Introduction to Artificial Intelligence module | PDF page 36 | lines 7-8

3.7 Constraint Satisfaction Search
Constraint satisfaction is a technique where a problem is solved when its values satisfy certain

Question sources

Computer Science model exam · 2015 · Question 120

Question 39 | Q00655

Which statement correctly distinguishes propositional logic from first-order logic?

A. It evaluates whole propositions using a truth assignment, without quantifying over objects in a domain.

B. It is predicate logic

C. It assigns truth to open predicates without assigning values to their object variables.

D. It contains quantifier

Hint

Translate the statement carefully, then check the scope of its quantifiers, connectives or inference rule. Distinguish guaranteed consequences from plausible generalizations or explanations.

Answer A: It evaluates whole propositions using a truth assignment, without quantifying over objects in a domain.

Propositional logic treats atoms as whole true-or-false claims and evaluates compound formulas from a truth assignment. It does not quantify over individual objects. First-order logic adds predicates, object variables, interpretations and quantifiers. The distinction does not mean that propositional truth needs no assignment: the truth values of its atoms still matter.

Why each choice is right or wrong

A - Correct: Propositional evaluation uses atom truth values rather than object-domain quantification.

B - Incorrect: Propositional logic is not identical to first-order predicate logic.

C - Incorrect: Open predicates need interpretation and variable assignments; this is not a propositional truth rule.

D - Incorrect: Object quantifiers belong to first-order logic, not propositional syntax.

Primary note proof

Study-note section AI26

> First-order logic adds objects, predicates, functions and quantified object variables. Human(a) is a property; Likes(a,b) a relation; father(a) a function term. Propositional logic has no object quantifiers, but its truth still depends on the assigned atom values.

Study-note section AI25

> Propositional atoms represent whole true-or-false claims. ¬P negates; P∧Q requires both; P∨Q requires at least one; P→Q is false only when P is true and Q false; P↔Q requires matching truth values. Thus P→Q ≡ ¬P∨Q. De Morgan gives ¬(P∧Q) ≡ ¬P∨¬Q.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 45 | lines 16-17

> A proposition is a declarative statement which is either true or false. It is a technique of knowledge
> representation in logical and mathematical form.

Introduction to Artificial Intelligence module | PDF page 47 | lines 20-21

> • We cannot represent relations like ALL, some, or none with propositional logic. Example:
> • All the girls are intelligent.

Question sources

Computer Science model exam · 2015 · Question 124

Question 40 | Q00658

Which option names the complete set of dataset roles used across model fitting, model selection and final held-out evaluation?

A. Validation dataset

B. Train dataset

C. Test dataset

D. All

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer D: All

Training fits parameters, validation selects models or settings, and a held-out test set estimates final performance. The overall workflow uses all three roles, with the test set kept separate from fitting and repeated tuning. Each individual option is true about one role, but only the aggregate gives the complete set requested by this single-choice question.

Why each choice is right or wrong

A - Incorrect: Validation supports model selection, but alone does not cover fitting and final held-out evaluation.

B - Incorrect: Training fits model parameters, but alone does not cover model selection and final evaluation.

C - Incorrect: Held-out test data estimates final performance, but alone omits fitting and model selection.

D - Correct: Training, validation and test sets together cover the three distinct workflow roles.

Primary note proof

Study-note section AI35

> Train data fits parameters; validation data selects models or settings; a held-out test set estimates final performance. In k-fold validation, train k models, each on k−1 folds and evaluate on the remaining fold, then average scores. Stratification approximately preserves label proportions. Fit preprocessing within each training fold. Time-ordered data needs splits respecting time; repeated tuning on the final test set leaks information.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 69 | lines 2-3

> with a set of training examples consisting of inputs and corresponding outputs, and is required to

discover the relation or mapping between then, e.g. as a series of rules, or a neural network.
Agent

Question sources

Computer Science model exam · 2015 · Question 127

Question 41 | Q00721

algorithm traverses a graph in a Breadth-ward motion and uses a
queue to remember to get the next vertex to start a search, when a dead end occurs in
any iteration.

A. Breadth first search

B. Binary search

C. Depth first search

D. A&C

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer A: Breadth first search

Expanding a graph breadthwise with a FIFO queue is precisely BFS (PDF pp. 19-20). DFS follows a path deeply and uses a stack; binary search is not this traversal.

Why each choice is right or wrong

A - Correct: Matches breadthwise expansion and FIFO frontier. See PDF p. 19, 20.

B - Incorrect: Not the stated graph traversal. See PDF p. 19, 20.

C - Incorrect: DFS uses a stack and goes deeply first. See PDF p. 19, 20.

D - Incorrect: DFS does not meet the queue/breadthwise description. See PDF p. 19, 20.

Primary note proof

Study-note section AI13

BFS removes the oldest frontier entry using a FIFO queue. Start with the root, remove the head, test it, append unseen successors at the tail, and repeat. It explores depths 0,1,2,... and is complete for finite b and a finite-depth goal. It minimizes number of edges, so minimum cost requires equal positive step costs.

Study-note section AI14

DFS follows the deepest pending branch using a LIFO stack or recursion. FILO means the same order as LIFO. With the same children as BFS, a left-first DFS visits S,A,C before B; push rightmost children first to pop the leftmost first.

Study-note section AI-S02

> Binary search exploits the order of a sorted sequence: a comparison narrows the remaining candidate interval instead of exploring arbitrary successor states. Array bisection therefore differs from breadth-first or depth-first traversal of a general state-space graph. A storage container such as an array does not by itself specify FIFO or LIFO frontier behavior.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 19 | lines 25-26

> • BFS algorithm starts searching from the root node of the tree and expands all successor node at the
> current level before moving to nodes of next level.

Introduction to Artificial Intelligence module | PDF page 20 | lines 3-4

> • Breadth-first search implemented using FIFO queue data structure.
> • If branching factor (average number of child nodes for a given node) = b and depth = d, then number

Question sources

Computer Science model exam 2 · Question 62

Question 42 | Q00889

Which Al approach most directly supports understanding biological intelligence behavior as a scientific goal?

A. Weak AI

B. Applied AI

C. .
Cognitive Al

D. Strong AI

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer C: .
Cognitive Al

Cognitive AI tests theories of how minds recognize objects and solve problems; the science-based goal is understanding biological intelligence (PDF pp. 7-8). C links these two descriptions.

Why each choice is right or wrong

A - Incorrect: Simulating cognition does not by itself specify the scientific-theory-testing aim. See PDF p. 7, 8.

B - Incorrect: Commercially viable systems are the stated focus of applied AI. See PDF p. 7, 8.

C - Correct: Directly tests theories of mind and cognition. See PDF p. 7, 8.

D - Incorrect: Aims for genuine human-level reasoning rather than specifically theory testing. See PDF p. 7, 8.

Primary note proof

Study-note section AI01

> Artificial intelligence studies systems that perform tasks requiring intelligence: reasoning, learning, planning, perception and language. Applied AI builds useful solutions; cognitive AI uses computational models to investigate biological intelligence. A model should be compared with human behavior or experimental observations before claiming it explains human thought.

Study-note section AI01

> Narrow or weak AI addresses bounded tasks, such as face recognition, spam filtering or a voice assistant. General AI denotes broad competence transferable across domains; strong AI also has a philosophical meaning involving actual understanding. Do not equate a humanoid body, autonomous vehicle, or many features with general intelligence.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 7 | lines 9-10

> Cognitive AI: computers are used to test theories about how the human mind works--for example,
> theories about how we recognize faces and other objects, or about how we solve abstract problems.

Introduction to Artificial Intelligence module | PDF page 8 | lines 17-19

> Science based AI goal
> • Develop concepts, mechanisms and vocabulary to understand biological intelligent behaviours

Question sources

Computer Science Exit Exam · 2018 · Question 2

Question 43 | Q00895

Why is Greedy Best-First Search not guaranteed to be optimal or complete?

A. It ignores path cost g(n) and relies only on h(n)

B. It always chooses the shallowest node

C. .
It expands nodes in FIFO order

D. It requires full knowledge of the goal state in advance

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer A: It ignores path cost g(n) and relies only on h(n)

Greedy chooses by estimated remaining cost h(n), ignoring accumulated path cost g(n), so a promising estimate can hide an expensive route. The module also warns of loops and nonoptimality (PDF pp. 25-26). Completeness depends on implementation and state-space assumptions; the general guarantee is absent.

Why each choice is right or wrong

A - Correct: Its h-only choice gives the stated lack of a general guarantee. See PDF p. 25, 26.

B - Incorrect: Shallowest-first describes BFS. See PDF p. 25, 26.

C - Incorrect: FIFO order describes BFS, not heuristic priority. See PDF p. 25, 26.

D - Incorrect: Not the reason the module gives for greedy failures. See PDF p. 25, 26.

Primary note proof

Study-note section AI17

> Best-first is a family that orders OPEN by an evaluation function. Greedy best-first uses f=h and removes the smallest estimate, ignoring the cost already paid. It can choose an expensive route, loop without duplicate checks, or follow an infinite misleading branch. On a finite graph with full duplicate handling it is complete, but not generally optimal. Its implicit finite-tree worst case can be O(b^m).

Study-note section AI18

> A* orders by f(n)=g(n)+h(n). Admissibility means 0≤h(n)≤h*(n), the true cheapest remaining cost. Consistency means h(n)≤c(n,n′)+h(n′) on every edge, with h(goal)=0; it implies admissibility on states with goal paths. Consistency makes f nondecreasing along a path.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 25 | lines 21-22

> o Step 3: Remove the node n, from the OPEN list which has the lowest value of h(n), and places it in the CLOSED list.

Introduction to Artificial Intelligence module | PDF page 26 | lines 17-18

> o This algorithm is not optimal.
> • Time Complexity: The worst case time complexity of Greedy best first search is O(bm).

Question sources

Computer Science Exit Exam · 2018 · Question 8

Question 44 | Q00910

Which search strategy tries to expand the node that is closest to the goal, on the grounds that this is likely to lead to a solution quickly; Thus, it evaluates nodes by using just the heuristic function: f (n)=h (n).

A. Uniform-cost search

B. Greedy best-first search

C. ¢.
Depth-first search

D. A* search

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer B: Greedy best-first search

Selecting the node estimated closest to the goal uses h(n) alone, as the pure-heuristic and lowest-h OPEN-list rules state (PDF p. 25). The same page prints f(n)=g(n), an internal typo contradicted by those rules; visual inspection confirms the typo is in the PDF.

Why each choice is right or wrong

A - Incorrect: Uses g(n), accumulated cost. See PDF p. 25, 25.

B - Correct: Matches h-only priority. See PDF p. 25, 25.

C - Incorrect: Uses depth priority. See PDF p. 25, 25.

D - Incorrect: Uses g(n)+h(n) rather than h alone. See PDF p. 25, 25.

Primary note proof

Study-note section AI17

> Best-first is a family that orders OPEN by an evaluation function. Greedy best-first uses f=h and removes the smallest estimate, ignoring the cost already paid. It can choose an expensive route, loop without duplicate checks, or follow an infinite misleading branch. On a finite graph with full duplicate handling it is complete, but not generally optimal. Its implicit finite-tree worst case can be O(b^m).

Study-note section AI15

> UCS uses a priority queue ordered by actual accumulated path cost g(n). Remove a minimum-g node; accept a goal when removed, not when first generated. Relax each edge: if a cheaper path is found, update its frontier cost and parent.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 25 | lines 6-7

> • Pure Heuristic Search - Pure heuristic search is the simplest form of heuristic search algorithms. It
> expands nodes based on their heuristic value h(n).

Introduction to Artificial Intelligence module | PDF page 25 | lines 21-22

> o Step 3: Remove the node n, from the OPEN list which has the lowest value of h(n), and places it in the CLOSED list.

Question sources

Computer Science Exit Exam · 2018 · Question 23

Question 45 | Q00922

Which of the following best defines a heuristic in the context of AI search?

A. A function that estimates how close a state is to the goal

B. A guaranteed optimal path to the goal

C. .
A method that always avoids revisiting states

D. A brute-force enumeration of all possible states

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer A: A function that estimates how close a state is to the goal

The module says a heuristic estimates how close a state is to the goal (PDF p. 25). An estimate guides node ordering but is neither a path nor a guarantee against revisiting states.

Why each choice is right or wrong

A - Correct: Direct definition. See PDF p. 25.

B - Incorrect: An estimate does not itself specify a guaranteed optimal path. See PDF p. 25.

C - Incorrect: Repeated-state management is a separate search concern. See PDF p. 25.

D - Incorrect: Brute force is the uninformed contrast. See PDF p. 25.

Primary note proof

Study-note section AI17

> A heuristic h(n) estimates remaining cost or evaluates promise using problem-specific knowledge. It can reduce explored work, but computing it also costs time. Uninformed algorithms use only the problem definition and known costs; UCS remains uninformed despite using g.

Study-note section AI17

Best-first is a family that orders OPEN by an evaluation function. Greedy best-first uses f=h and removes the smallest estimate, ignoring the cost already paid. It can choose an expensive route, loop without duplicate checks, or follow an infinite misleading branch. On a finite graph with full duplicate handling it is complete, but not generally optimal. Its implicit finite-tree worst case can be O(b^m).

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 25 | lines 2-3

• Heuristic function estimates how close a state is to the goal. It is represented by h(n), and it calculates
the cost of an optimal path between the pair of states. Its value is always positive.

Question sources

Computer Science Exit Exam · 2018 · Question 35

Question 46 | Q00939

Which type of heuristic search strategy that evaluates nodes by combining g (n), the cost to reach the node, and h (n.), the cost to get from the node to the goal: f (n) = g (n) +h (n) Since g(n) gives the path cost from the start node to node n, and h(n) is the estimated cost of the cheapest path from n to the goal, we have f (n) = estimated cost of the cheapest solution through n ?

A. A* search

B. Greedy Best-First Search

C. Uniform-cost search

D. Depth-first search

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer A: A* search

A* uses f(n)=g(n)+h(n), combining the cost already paid with an estimate of the remaining cost. Greedy best-first uses only h; uniform-cost search uses only g; depth-first search follows its stack discipline. The sum is the distinguishing rule. The stray answer annotation attached to the depth-first option has been removed from the practice choices.

Why each choice is right or wrong

A - Correct: A* ranks by accumulated path cost plus estimated remaining cost.

B - Incorrect: Greedy best-first ranks by h alone and omits accumulated cost.

C - Incorrect: Uniform-cost search ranks by g alone and omits the heuristic estimate.

D - Incorrect: DFS follows a stack order rather than minimum g+h.

Primary note proof

Study-note section AI18

A* orders by f(n)=g(n)+h(n). Admissibility means 0≤h(n)≤h*(n), the true cheapest remaining cost. Consistency means h(n)≤c(n,n′)+h(n′) on every edge, with h(goal)=0; it implies admissibility on states with goal paths. Consistency makes f nondecreasing along a path.

Study-note section AI17

Best-first is a family that orders OPEN by an evaluation function. Greedy best-first uses f=h and removes the smallest estimate, ignoring the cost already paid. It can choose an expensive route, loop without duplicate checks, or follow an infinite misleading branch. On a finite graph with full duplicate handling it is complete, but not generally optimal. Its implicit finite-tree worst case can be O(b^m).

Study-note section AI15

UCS uses a priority queue ordered by actual accumulated path cost g(n). Remove a minimum-g node; accept a goal when removed, not when first generated. Relax each edge: if a cheaper path is found, update its frontier cost and parent.

Study-note section AI14

DFS follows the deepest pending branch using a LIFO stack or recursion. FILO means the same order as LIFO. With the same children as BFS, a left-first DFS visits S,A,C before B; push rightmost children first to pop the leftmost first.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 26 | lines 25-26

• A* search is the most commonly known form of best-first search.
• It uses heuristic function h(n), and cost to reach the node n from the start state g(n). It has combined

Introduction to Artificial Intelligence module | PDF page 28 | lines 4-5

the formula f(n)= g(n) + h(n), where g(n) is the cost to reach any node from start state.
• Here we will use OPEN and CLOSED list.

Question sources

Computer Science Exit Exam · 2018 · Question 52

Question 47 | Q00961

Why is perception still a challenge for AI-even though systems can recognize faces?

A. Because perception is unrelated to reasoning

B. Because AI lacks actuators

C. Because open-world inputs may contain occlusion, noise, unfamiliar conditions and incomplete context.

D. Because sensors are too expensive

Hint

Match the input, output and purpose to the appropriate language, perception or physical-action task.

Answer C: Because open-world inputs may contain occlusion, noise, unfamiliar conditions and incomplete context.

Recognizing faces in some settings does not solve all perception problems. Occlusion can hide relevant features, noise can corrupt observations, and unfamiliar conditions or missing context can make an interpretation unreliable. These are difficulties in interpreting sensory information. They do not require every perception task to use natural-language understanding, and they are distinct from actuator availability or sensor price.

Why each choice is right or wrong

A - Incorrect: Perception can use context and support reasoning; it is not unrelated to it.

B - Incorrect: Actuators implement actions, while the question concerns difficulty interpreting sensory input.

C - Correct: Occlusion, noise and unfamiliar or incomplete context can defeat otherwise successful recognition systems.

D - Incorrect: Cost is not the conceptual reason that open-world interpretation remains difficult.

Primary note proof

Study-note section AI39

> Computer vision interprets images or video: classification names image content, detection locates objects, segmentation labels pixels, and recognition identifies objects or people. Occlusion, lighting change, noise, unfamiliar conditions and incomplete context make open-world perception difficult even when a face detector works well.

Study-note section AI01

> Narrow or weak AI addresses bounded tasks, such as face recognition, spam filtering or a voice assistant. General AI denotes broad competence transferable across domains; strong AI also has a philosophical meaning involving actual understanding. Do not equate a humanoid body, autonomous vehicle, or many features with general intelligence.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 7 | lines 13-15

> knowledge. Common sense knowledge is needed to function in complex, open-ended worlds. Such a
> system also needs to understand unconstrained natural language. However these capabilities are not
> yet fully present in today‘s intelligent systems.

Question sources

Computer Science Exit Exam · 2018 · Question 74

Question 48 | Q00970

In which of the following Learning methods an agent receives feedback in the form of rewards or penalties to maximize cumulative reward?

A. Supervised learning

B. Reinforcement learning

C. .
Semi-supervised learning

D. Unsupervised learning

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer B: Reinforcement learning

Reinforcement learning gives rewards for right actions and penalties for wrong ones, with the goal of obtaining the most reward points (PDF p. 69). B matches both interaction and feedback.

Why each choice is right or wrong

A - Incorrect: Supervised learning uses input-target labels rather than action-reward feedback.

B - Correct: Reinforcement learning improves decisions to maximize cumulative reward from interaction.

C - Incorrect: Semi-supervised learning mixes labeled and unlabeled examples; that mixture does not define cumulative-reward learning.

D - Incorrect: Unsupervised learning discovers unlabeled structure rather than defining an action-reward objective.

Primary note proof

Study-note section AI33

An RL agent observes state information, takes an action, receives reward and a new observation, and improves its policy to maximize expected cumulative reward. Rewards evaluate outcomes rather than supply the correct action label. Rewards may be delayed; not every correct-looking action receives immediate reward.

Study-note section AI32

Semi-supervised learning combines some labels with additional unlabeled data. Active learning chooses informative cases for an oracle to label; it still uses labels. “Exploratory learning” is not an exclusive standard synonym: it can mean unsupervised exploration or active data acquisition, so define the feedback. A task named summarization or time-series analysis is not automatically supervised; its training targets determine that.

Study-note section AI31

Supervised learning fits a mapping from labeled input-target examples. Classification predicts discrete labels, such as institution generations or spam/nonspam; regression predicts numeric quantities, such as prices. Predictive analytics uses historical patterns to forecast outcomes and can use either. Specify the target before choosing a method.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 69 | lines 15-18

⇒ Reinforcement learning is a feedback-based learning method, in which a learning agent gets a reward for each right action and gets a penalty for each wrong action.
The agent learns automatically with these feedbacks and improves its performance. In reinforcement
learning, the agent interacts with the environment and explores it. The goal of an agent is to get the

Question sources

Computer Science Exit Exam · 2018 · Question 83

Question 49 | Q00995

Which one of the following is wrongly matched as AI environment property pairs?

A. Fully observable <-> Partially observable

B. Discrete <-> Continuous

C. Deterministic <-> Stochastic

D. Static <-> Stochastic

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer D: Static <-> Stochastic

The module pairs static with dynamic, and deterministic with nondeterministic (PDF p. 13). Thus static versus stochastic mixes two independent dimensions; D is the wrong pairing.

Why each choice is right or wrong

A - Incorrect: Opposite observability properties. See PDF p. 13, 13.

B - Incorrect: Opposite state/action granularity properties. See PDF p. 13, 13.

C - Incorrect: Stochastic is the question's term for nondeterministic transition uncertainty. See PDF p. 13, 13.

D - Correct: Incorrect pair: static's counterpart is dynamic. See PDF p. 13, 13.

Primary note proof

Study-note section AI09

> Fully observable means sensors reveal relevant state; partially observable means relevant information is hidden or noisy. Deterministic means state and action determine the next state; stochastic means outcomes involve randomness. Strategic uncertainty about an opponent is distinct from dice-like randomness.

Study-note section AI09

> Episodic decisions are independent; sequential decisions affect later situations. Static worlds do not change while deliberating; dynamic ones do. A clock affecting score makes an otherwise static game semidynamic. Discrete quantities have distinct alternatives; continuous quantities vary along real-valued ranges. Known versus unknown concerns knowledge of the transition model, not observability.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 13 | lines 17-18

> ✓ Static / Dynamic − If the environment does not change while an agent is acting, then it is static; otherwise it is dynamic.

Introduction to Artificial Intelligence module | PDF page 13 | lines 23-24

> ✓ Deterministic / Non-deterministic − If the next state of the environment is completely determined
> by the current state and the actions of the agent, then the environment is deterministic; otherwise it

Question sources

National Exit Exam practice questions · Question 9

Computer Science Exit Exam practice collection · Question 199

Question 50 | Q01007

Which option describes general data processing without AI learning or inference, rather than naming an AI-related field?

A. Machine learning

B. Robotics

C. Data analysis limited to fixed arithmetic summaries, without learning or inference.

D. Natural language processing

Hint

Match the input, output and purpose to the appropriate language, perception or physical-action task.

Answer C: Data analysis limited to fixed arithmetic summaries, without learning or inference.

Machine learning develops systems that improve from experience, natural language processing interprets or generates language, and robotics combines sensing, planning and physical action. Data analysis is broader: it may use AI, but fixed arithmetic summaries alone need not involve learning or inference. The option restricts that broad category to ordinary non-AI processing, avoiding the incorrect claim that all data analysis lies outside AI.

Why each choice is right or wrong

A - Incorrect: Machine learning is part of AI, and AI43 explicitly places learning architectures and regimes within it.

B - Incorrect: Robotics combines sensing, planning and physical action and is treated as an AI-related field in AI39.

C - Correct: The restricted fixed-summary task need not use AI; this does not exclude learning-based data analysis.

D - Incorrect: Natural language processing interprets and generates human language and is covered as an AI field in AI37.

Primary note proof

Study-note section AI01

> Artificial intelligence studies systems that perform tasks requiring intelligence: reasoning, learning, planning, perception and language. Applied AI builds useful solutions; cognitive AI uses computational models to investigate biological intelligence. A model should be compared with human behavior or experimental observations before claiming it explains human thought.

Study-note section AI39

> Robotics combines mechanical, electrical, control and computing systems to sense, plan and act physically. Components include sensors, controller, power supply, actuators and an end effector. Locomotion may be legged, wheeled or tracked, with trade-offs in terrain, stability and

efficiency. Software can act autonomously without being a robot. Humanoid shape is not evidence of general intelligence.

Study-note section AI37

NLP interprets and generates human language: translation, dialogue, summarization, question answering and text generation. Speech recognition converts audio to text; speech synthesis generates speech. They overlap language processing but solve different input/output tasks.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 6 | lines 12-23

1.3Intelligent behaviour
This discussion brings us back to the question of what constitutes intelligent behavior. Some ofthese
tasks and applications are:
• Perception involving image recognition and computer vision
• Reasoning
• Learning
• Understanding language involving natural language processing, speech processing
• Solving problems
• Robotics
• Thinking
• Acting- in the complex environment
• Knowledge - applying successfully in new situation

Introduction to Artificial Intelligence module | PDF page 67 | lines 29-30

• Data mining and knowledge discovery
• Speech/image/video (pattern) recognition

Question sources

Artificial Intelligence practice questions · Question 1

Question 51 | Q01008

Which techniques can improve learned parameters, candidate solutions or policies using training experience, fitness feedback or reward? Select all that apply. Treat D as a fixed knowledge-based expert system.

A. Artificial neural networks

B. Genetic algorithms

C. Reinforcement learning

D. Expert systems

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer A, B, C: Artificial neural networks; Genetic algorithms; Reinforcement learning

Neural-network training adjusts weights from experience, genetic algorithms evolve candidates using fitness and variation, and reinforcement learning improves policies using rewards from interaction. These are different improvement mechanisms that satisfy the broader capability requested. A fixed expert system applies encoded domain knowledge and does not inherently revise itself through experience. The fixed-system assumption distinguishes using knowledge from learning or optimizing with feedback.

Why each choice is right or wrong

A - Correct: Training adjusts neural weights from examples or feedback and can improve performance.

B - Correct: Genetic search improves a candidate population through fitness-based experience and variation.

C - Correct: RL improves policies from action outcomes and rewards.

D - Incorrect: A conventional fixed expert system applies encoded knowledge without inherently learning from experience.

Primary note proof

Study-note section AI43

> ANNs, genetic search and reinforcement learning all improve candidate behavior or performance through experience in different ways. Deep learning is an architecture family usable with supervised, unsupervised or reinforcement objectives. Decision trees and deep models may both be supervised. To distinguish learning regimes, specify labels, reward or unlabeled structure.

Study-note section AI40

> An expert system applies specialized knowledge to a domain such as equipment diagnosis. The domain expert supplies expertise; the knowledge engineer elicits and encodes it; the inference engine reasons over a knowledge base and working facts. The user interface handles interaction; an explanation facility presents why or how a conclusion was reached.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 67 | lines 18-19

> important feature of —intelligent‖. Machine learning is the subfield of AI concerned with intelligent
> systems that learn from experience and examples. It is the computational study of algorithms that

Introduction to Artificial Intelligence module | PDF page 69 | lines 15-16

> ⇒ Reinforcement learning is a feedback-based learning method, in which a learning agent gets a reward for each right action and gets a penalty for each wrong action.

Introduction to Artificial Intelligence module | PDF page 71 | lines 10-11

> working in unison to solve specific problems. ANNs, like people, learn by example. An ANN is configured for a specific application, such as pattern recognition or data classification, through a learning

Question sources

Artificial Intelligence practice questions · Question 2

Question 52 | Q01009

Which listed programming languages can implement AI or machine-learning algorithms? Select all that apply; this asks about implementation capability, not an exclusive popularity ranking.

A. C++

B. Java

C. Python

D. Ruby

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer A, B, C, D: C++; Java; Python; Ruby

All four languages can implement AI or machine-learning algorithms. The note explicitly identifies Python, C++ and Java as implementation languages. The supplementary primary Rumale project documentation supplies the missing Ruby fact by documenting a Ruby machine-learning library with classifiers and clustering algorithms. This establishes capability rather than a popularity ranking, and avoids treating one widely used language as the only valid implementation choice.

Why each choice is right or wrong

A - Correct: The note explicitly identifies C++ as an AI implementation language.

B - Correct: The note explicitly identifies Java as an AI implementation language.

C - Correct: The note identifies Python as widely used for AI implementation.

D - Correct: Rumale's primary project documentation demonstrates machine-learning algorithms implemented for Ruby.

Primary note proof

Study-note section AI43

Face recognition is both an applied task and a narrow capability. Broad machine-intelligence goals can include emulating human intelligence and building systems that learn. Humanlike problem-solving alone does not uniquely select expert systems. Python is widely used, but C++ and Java also implement AI; “commonly used” does not isolate one language. McCarthy is associated with naming AI, not its sole invention.

Study-note section AI-S01

AI algorithms are not exclusive to one programming language. Rumale is a Ruby machine-learning library with classification and clustering implementations, demonstrating that Ruby can implement learning algorithms as well as the general-purpose languages already discussed. Language choice and learning regime are different decisions.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 8 | lines 21-23

1.10 AI Techniques
Various techniques that have involved can be applied to a variety of AI tasks. The techniques are concerned with how we represent, manipulate and reason with knowledge in order to solve problem.

Question sources

Artificial Intelligence practice questions · Question 3

Question 53 | Q01010

Which listed goal specifically expresses the humanlike approaches to AI, rather than autonomy, general automation or replacement of workers?

A. To replace humans with machines

B. To develop autonomous systems

C. To simulate human intelligence

D. To automate tasks and improve efficiency

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer C: To simulate human intelligence

Humanlike AI approaches aim to model human thought or produce humanlike behavior. Simulating aspects of human intelligence expresses that goal. Autonomous action, automation and efficiency can also be useful AI objectives, but none requires human imitation. Replacing workers is a social or deployment choice rather than the field’s defining aim. This classification does not impose one primary goal on every AI system.

Why each choice is right or wrong

A - Incorrect: Replacing humans is a deployment choice, not the defining goal of humanlike AI.

B - Incorrect: Autonomy concerns independent action and need not emulate human intelligence.

C - Correct: Simulating human intelligence expresses the humanlike approach to thought or behavior.

D - Incorrect: Automation and efficiency can be achieved without modeling human intelligence.

Primary note proof

Study-note section AI01

> Artificial intelligence studies systems that perform tasks requiring intelligence: reasoning, learning, planning, perception and language. Applied AI builds useful solutions; cognitive AI uses computational models to investigate biological intelligence. A model should be compared with human behavior or experimental observations before claiming it explains human thought.

Study-note section AI02

> Thinking humanly models human mental processes; thinking rationally derives conclusions by valid logic. Acting humanly seeks humanlike behavior; acting rationally selects actions expected to achieve objectives. Thought processes and observed behavior are different evaluation targets.

Study-note section AI43

> Face recognition is both an applied task and a narrow capability. Broad machine-intelligence goals can include emulating human intelligence and building systems that learn. Humanlike problem-solving alone does not uniquely select expert systems. Python is widely used, but C++ and Java also implement AI; “commonly used” does not isolate one language. McCarthy is associated with naming AI, not its sole invention.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 8 | lines 9-13

> General AI goal
> • Replicate human intelligent
> • Solve knowledge intensive task
> • Make an intelligent connection between perception and action
> • Enhance human-computer interaction/ communication

Question sources

Artificial Intelligence practice questions · Question 4

Question 54 | Q01011

Which AI application is used to understand and interpret human language?

A. Computer vision

B. Speech recognition

C. Natural language processing

D. Expert systems

Hint

Match the input, output and purpose to the appropriate language, perception or physical-action task.

Answer C: Natural language processing

NLP is technology for machines to understand, analyze, manipulate and interpret human language (PDF p. 79). Speech recognition is narrower: converting spoken words into text (p. 82).

Why each choice is right or wrong

A - Incorrect: Works with image information. See PDF p. 79.

B - Incorrect: Transcribes spoken words; narrower than language interpretation. See PDF p. 79.

C - Correct: Matches the language-understanding definition. See PDF p. 79.

D - Incorrect: Uses domain expertise, not specifically language interpretation. See PDF p. 79.

Primary note proof

Study-note section AI37

> NLP interprets and generates human language: translation, dialogue, summarization, question answering and text generation. Speech recognition converts audio to text; speech synthesis generates speech. They overlap language processing but solve different input/output tasks.

Study-note section AI39

> Computer vision interprets images or video: classification names image content, detection locates objects, segmentation labels pixels, and recognition identifies objects or people. Occlusion, lighting change, noise, unfamiliar conditions and incomplete context make open-world perception difficult even when a face detector works well.

Study-note section AI40

> An expert system applies specialized knowledge to a domain such as equipment diagnosis. The domain expert supplies expertise; the knowledge engineer elicits and encodes it; the inference engine reasons over a knowledge base and working facts. The user interface handles interaction; an explanation facility presents why or how a conclusion was reached.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 79 | lines 4-5

and Artificial Intelligence. It is the technology that is used by machines to understand, analyze, manipulate, and interpret human's languages. It helps developers to organize knowledge for

Question sources

Artificial Intelligence practice questions · Question 5

Question 55 | Q01012

Which technique is defined by a specialized knowledge base of conventional domain facts and rules, with an inference engine applying them to solve domain problems?

A. Genetic algorithms

B. Fuzzy logic

C. Expert systems

D. Reinforcement learning

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer C: Expert systems

An expert system is characterized by specialized domain knowledge and an inference engine that reasons over its knowledge base and working facts. Merely saying that a method uses rules would not distinguish it from fuzzy systems, which may also use rules. Genetic algorithms instead evolve populations, and reinforcement learning uses reward feedback. The architecture in the question supplies the specific distinction.

Why each choice is right or wrong

A - Incorrect: Genetic algorithms evolve candidate populations through selection and variation rather than defining a knowledge-base inference architecture.

B - Incorrect: Fuzzy logic represents graded membership and can use rules, but rules alone do not make it the specialized architecture described.

C - Correct: Expert systems apply domain knowledge through an inference engine operating on a knowledge base and working facts.

D - Incorrect: Reinforcement learning improves decisions from rewards rather than defining the described knowledge-base architecture.

Primary note proof

Study-note section AI40

An expert system applies specialized knowledge to a domain such as equipment diagnosis. The domain expert supplies expertise; the knowledge engineer elicits and encodes it; the inference engine reasons over a knowledge base and working facts. The user interface handles interaction; an explanation facility presents why or how a conclusion was reached.

Study-note section AI30

A fuzzy controller fuzzifies measurements, applies graded IF-THEN rules, combines outputs and often defuzzifies to a numerical action. Fuzzy control suits linguistic rules about degrees; reinforcement learning suits reward-based adaptation. Merely saying an environment is complex or uncertain does not uniquely select either method.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 63 | lines 2-3

o It is in the form of IF-THEN-ELSE rules.
o Knowledge Acquisition

Introduction to Artificial Intelligence module | PDF page 63 | lines 20-21

o Applies rules repeatedly to the facts, which are obtained from earlier rule application.
o Adds new knowledge into the knowledge base if required.

Question sources

Artificial Intelligence practice questions · Question 6

Question 56 | Q01013

Which learning regime discovers structure in unlabeled data without action-reward interaction? Select the regime rather than a neural architecture family.

A. Supervised learning

B. Unsupervised learning

C. Reinforcement learning

D. Deep learning

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer B: Unsupervised learning

Unsupervised learning seeks clusters, associations or other structure without target labels. Reinforcement learning also need not receive correct-action labels, but it learns through actions and rewards, which this scenario excludes. Deep learning describes an architecture family that can use several learning regimes. Supervised learning requires labeled targets. Naming both the feedback and the requested category therefore isolates the intended regime.

Why each choice is right or wrong

A - Incorrect: Supervised learning requires labeled input-target examples.

B - Correct: Unsupervised learning discovers structure without target labels.

C - Incorrect: RL learns from action rewards, excluded by the stem.

D - Incorrect: Deep learning is an architecture family, not the specific unlabeled structure-discovery regime requested.

Primary note proof

Study-note section AI32

> Unsupervised learning seeks structure without target labels: clustering groups similar examples; association finds co-occurrence; dimensionality reduction compresses features. Good clusters have high within-cluster similarity and low between-cluster similarity under an appropriate distance. K-means alternates assignment to nearest centroids and centroid updates.

Study-note section AI33

> An RL agent observes state information, takes an action, receives reward and a new observation, and improves its policy to maximize expected cumulative reward. Rewards evaluate outcomes rather than supply the correct action label. Rewards may be delayed; not every correct-looking action receives immediate reward.

Study-note section AI43

> ANNs, genetic search and reinforcement learning all improve candidate behavior or performance through experience in different ways. Deep learning is an architecture family usable with supervised, unsupervised or reinforcement objectives. Decision trees and deep models may both be supervised. To distinguish learning regimes, specify labels, reward or unlabeled structure.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 69 | lines 7-8

> ⇒ Unsupervised learning means we are only given the x‘s. In either case, the goal is to estimate f. The system is supplied with a set of training examples consisting only of inputs and is required to

Introduction to Artificial Intelligence module | PDF page 76 | lines 7-8

> means there is no teacher as in supervised learning. After receiving the feedback, the network performs adjustments of the weights to get better critic information in future.

Question sources

Artificial Intelligence practice questions · Question 7

Question 57 | Q01014

What does the Turing test assess?

A. Machine learning capabilities

B. Expert system knowledge

C. Natural language processing proficiency

D. Artificial intelligence in simulating human conversation

Hint

Match the input, output and purpose to the appropriate language, perception or physical-action task.

Answer D: Artificial intelligence in simulating human conversation

Turing's test asks whether an interrogator can distinguish a computer from a human using textual questions (PDF p. 5). D captures the overall conversational behavioral test; NLP proficiency is a component rather than the full criterion.

Why each choice is right or wrong

A - Incorrect: Not a test specifically of learning. See PDF p. 5.

B - Incorrect: Not restricted to expert-system domain knowledge. See PDF p. 5.

C - Incorrect: A required conversational ability, not the complete assessment. See PDF p. 5.

D - Correct: Matches indistinguishability in textual human conversation. See PDF p. 5.

Primary note proof

Study-note section AI02

> The conversational Turing test asks whether a judge can reliably distinguish a machine from a human through their dialogue. It requires language use, knowledge, reasoning and often learning. It tests an operational behavior criterion, not proof of consciousness. Cognitive models can be tested against response times, errors and psychological experiments.

Study-note section AI37

> NLP interprets and generates human language: translation, dialogue, summarization, question answering and text generation. Speech recognition converts audio to text; speech synthesis generates speech. They overlap language processing but solve different input/output tasks.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 5 | lines 18-20

> interrogator could not distinguish them by questioning, then it would be unreasonable not to call the computer intelligent. Turing's 'imitation game' is now usually called 'the Turing test' for intelligence.

Question sources

Artificial Intelligence practice questions · Question 8

Question 58 | Q01015

Which AI technique aims to mimic the structure and function of the human brain?

A. Genetic algorithms

B. Neural networks

C. Expert systems

D. Reinforcement learning

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer B: Neural networks

Artificial neural networks are brain-inspired abstractions built from connected processing units with adjustable weights. This is inspiration rather than a literal recreation of the brain. Genetic algorithms use evolutionary population mechanisms, expert systems use encoded domain knowledge, and reinforcement learning defines a feedback regime that can be implemented with many model types, including neural networks.

Why each choice is right or wrong

A - Incorrect: Genetic algorithms are inspired by population selection and genetic variation, not a network of neural processing units.

B - Correct: Artificial neural networks are brain-inspired abstractions of interconnected weighted processing units.

C - Incorrect: Expert systems encode domain knowledge for inference rather than defining a brain-inspired network structure.

D - Incorrect: Reinforcement learning defines reward-based interaction and can use neural or non-neural implementations.

Primary note proof

Study-note section AI34

> A neuron computes z=Σiwixi+b and output a=φ(z). For x=(2,1), w=(0.5,−1), b=0.2, z=0.2. Linear activation returns z; sigmoid returns 1/(1+e^(−z)); tanh is bipolar; ReLU returns max(0,z). Nonlinearities let multiple layers represent nonlinear relationships.

Study-note section AI20

> A genetic algorithm maintains a population, evaluates fitness, selects parents, applies crossover and mutation, and repeats replacement. It searches without guaranteeing global optimality. Means-ends analysis selects a difference between current and goal states, chooses an operator

reducing it, and recursively solves its preconditions. Neither method is simply A* running backward.

Study-note section AI40

An expert system applies specialized knowledge to a domain such as equipment diagnosis. The domain expert supplies expertise; the knowledge engineer elicits and encodes it; the inference engine reasons over a knowledge base and working facts. The user interface handles interaction; an explanation facility presents why or how a conclusion was reached.

Study-note section AI33

An RL agent observes state information, takes an action, receives reward and a new observation, and improves its policy to maximize expected cumulative reward. Rewards evaluate outcomes rather than supply the correct action label. Rewards may be delayed; not every correct-looking action receives immediate reward.

Study-note section AI34

Feedforward networks pass signals toward outputs without recurrent state loops; recurrent networks carry state across sequence steps. Training adjusts weights, often through gradients and backpropagation. Backpropagation is not the same as a recurrent connection. Deep learning uses multiple representation layers and can learn features from images, sound and text. CNNs share filters across image locations. ANN is a brain-inspired abstraction, not a literal brain simulation; DL⊂ML⊂AI.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 71 | lines 6-8

An Artificial Neural Network (ANN) is an information processing paradigm that is inspired by the way
biological nervous systems, such as the brain, process information. The key element of this paradigm
is the novel structure of the information processing system.

Question sources

Artificial Intelligence practice questions · Question 9

Question 59 | Q01016

Which AI application is used to make predictions based on historical data?

A. Data mining

B. Pattern recognition

C. Predictive analytics

D. Genetic algorithms

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer C: Predictive analytics.

AI31 directly defines predictive analytics as using historical patterns to forecast outcomes. This names the application requested by the question. Data mining and pattern recognition are broader related activities that can contribute, while genetic algorithms perform evolutionary optimization. Their possible supporting roles do not make them the precise application term defined in the note.

Why each choice is right or wrong

A - Incorrect: Data mining is a broader discovery activity; the note names historical-outcome forecasting specifically as predictive analytics.

B - Incorrect: Recognizing patterns can support prediction, but does not name the defined historical-forecasting application as precisely.

C - Correct: Predictive analytics is explicitly defined as forecasting outcomes using historical patterns.

D - Incorrect: Genetic algorithms evolve candidate solutions through fitness and variation, rather than naming historical forecasting itself.

Primary note proof

Study-note section AI31

> Supervised learning fits a mapping from labeled input-target examples. Classification predicts discrete labels, such as institution generations or spam/nonspam; regression predicts numeric quantities, such as prices. Predictive analytics uses historical patterns to forecast outcomes and can use either. Specify the target before choosing a method.

Study-note section AI32

> Unsupervised learning seeks structure without target labels: clustering groups similar examples; association finds co-occurrence; dimensionality reduction compresses features. Good clusters have high within-cluster similarity and low between-cluster similarity under an appropriate distance. K-means alternates assignment to nearest centroids and centroid updates.

Study-note section AI20

> A genetic algorithm maintains a population, evaluates fitness, selects parents, applies crossover and mutation, and repeats replacement. It searches without guaranteeing global optimality. Means-ends analysis selects a difference between current and goal states, chooses an operator reducing it, and recursively solves its preconditions. Neither method is simply A* running backward.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 68 | lines 21-22

> • Machine learning constructs or uses the algorithms that learn from historical data.
> • Supervised

Introduction to Artificial Intelligence module | PDF page 71 | lines 18-19

> of information it has been given to analyse. This expert can then be used to provide projections given
> new situations of interest and answer "what if" questions. Other advantages include:

Question sources

Artificial Intelligence practice questions · Question 10

Question 60 | Q01017

Which descriptions explicitly establish a bounded task-specific AI capability, rather than only a body shape or an unspecified “general-purpose” label? Select all that apply.

A. Autonomous vehicles limited to driving tasks.

B. Virtual assistants like Siri or Alexa performing bounded assistant tasks.

C. Humanoid robots

D. Robots described only as general-purpose, with no capability details.

Hint

Match the input, output and purpose to the appropriate language, perception or physical-action task.

Answer A, B: Autonomous vehicles limited to driving tasks.; Virtual assistants like Siri or Alexa performing bounded assistant tasks.

Narrow AI addresses bounded tasks rather than broad competence transferable across domains. Driving-only automation and a bounded voice assistant therefore establish task-specific capabilities. A humanoid body reveals shape, not intellectual scope, while an unexplained general-purpose label supplies no demonstrated cross-domain competence. The question asks which descriptions establish narrowness, so the two insufficient descriptions are excluded without claiming that their particular robots could never be narrow.

Why each choice is right or wrong

A - Correct: Autonomous driving is a bounded capability and does not by itself establish broad cross-domain general intelligence.

B - Correct: Virtual voice assistants are the note’s direct examples of bounded narrow-AI capability.

C - Incorrect: Humanoid shape alone does not establish a bounded task-specific capability.

D - Incorrect: An undefined general-purpose label does not establish a specific bounded capability.

Primary note proof

Study-note section AI01

> Narrow or weak AI addresses bounded tasks, such as face recognition, spam filtering or a voice assistant. General AI denotes broad competence transferable across domains; strong AI also has a philosophical meaning involving actual understanding. Do not equate a humanoid body, autonomous vehicle, or many features with general intelligence.

Study-note section AI39

> Robotics combines mechanical, electrical, control and computing systems to sense, plan and act physically. Components include sensors, controller, power supply, actuators and an end effector. Locomotion may be legged, wheeled or tracked, with trade-offs in terrain, stability and efficiency. Software can act autonomously without being a robot. Humanoid shape is not evidence of general intelligence.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 7 | lines 3-5

> Weak AI: deals with the creation of some form of computer-based artificial intelligence that cannot
> truly reason and solve problems, but can act as if it were intelligent. Weak AI holds that suitably programmed machines can simulate human cognition.

Introduction to Artificial Intelligence module | PDF page 7 | lines 19-20

> • In Robotics, we have been able to make vehicles that are mostly autonomous.
> • In Natural language processing, we have systems that are capable of simple machine

Question sources

Artificial Intelligence practice questions · Question 11

Question 61 | Q01018

Which learning regime is defined by an agent choosing actions and improving its policy from environmental rewards or penalties, rather than by a model architecture?

A. Reinforcement learning

B. Supervised learning

C. Unsupervised learning

D. Deep learning

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer A: Reinforcement learning

Reinforcement learning is defined by interaction: the agent acts, receives reward feedback and improves its policy to increase expected cumulative return. Correct-action target labels are not required. Supervised and unsupervised learning use different feedback definitions. Deep learning describes a model architecture and may support reinforcement learning, but it does not name the particular feedback regime asked for.

Why each choice is right or wrong

A - Correct: Reinforcement learning uses action-reward interaction to improve expected return.

B - Incorrect: Supervised learning uses supplied input-target pairs rather than this reward-feedback definition.

C - Incorrect: Unsupervised learning seeks unlabeled structure rather than an action-reward policy.

D - Incorrect: Deep learning specifies multilayer neural representations and can implement several feedback regimes.

Primary note proof

Study-note section AI33

> An RL agent observes state information, takes an action, receives reward and a new observation, and improves its policy to maximize expected cumulative reward. Rewards evaluate outcomes rather than supply the correct action label. Rewards may be delayed; not every correct-looking action receives immediate reward.

Study-note section AI43

> ANNs, genetic search and reinforcement learning all improve candidate behavior or performance through experience in different ways. Deep learning is an architecture family usable with supervised, unsupervised or reinforcement objectives. Decision trees and deep models may both be supervised. To distinguish learning regimes, specify labels, reward or unlabeled structure.

Study-note section AI31

> Supervised learning fits a mapping from labeled input-target examples.

Study-note section AI32

> Unsupervised learning seeks structure without target labels: clustering groups similar examples; association finds co-occurrence; dimensionality reduction compresses features.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 69 | lines 16-18

> reward for each right action and gets a penalty for each wrong action.
> The agent learns automatically with these feedbacks and improves its performance. In reinforcement

learning, the agent interacts with the environment and explores it. The goal of an agent is to get the

Question sources

Artificial Intelligence practice questions · Question 12

Question 62 | Q01019

Which AI application is used to identify and classify objects in images?

A. Natural language processing

B. Speech recognition

C. Computer vision

D. Expert systems

Hint

Match the input, output and purpose to the appropriate language, perception or physical-action task.

Answer C: Computer vision

Computer vision automatically analyzes image information and includes object recognition (PDF p. 84), directly matching identifying/classifying image objects.

Why each choice is right or wrong

A - Incorrect: Processes language rather than image objects. See PDF p. 84, 84.

B - Incorrect: Transcribes speech. See PDF p. 84, 84.

C - Correct: The explicitly visual recognition application. See PDF p. 84, 84.

D - Incorrect: Domain expertise alone does not specify vision. See PDF p. 84, 84.

Primary note proof

Study-note section AI39

Computer vision interprets images or video: classification names image content, detection locates objects, segmentation labels pixels, and recognition identifies objects or people. Occlusion, lighting change, noise, unfamiliar conditions and incomplete context make open-world perception difficult even when a face detector works well.

Study-note section AI37

NLP interprets and generates human language: translation, dialogue, summarization, question answering and text generation. Speech recognition converts audio to text; speech synthesis generates speech. They overlap language processing but solve different input/output tasks.

Study-note section AI40

An expert system applies specialized knowledge to a domain such as equipment diagnosis. The domain expert supplies expertise; the knowledge engineer elicits and encodes it; the inference engine reasons over a knowledge base and working facts. The user interface handles interaction; an explanation facility presents why or how a conclusion was reached.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 84 | lines 4-5

o Computer vision automatically extracts, analyzes, and comprehends useful information from a single image or an array of images. This process involves development of algorithms to

Introduction to Artificial Intelligence module | PDF page 84 | lines 20-21

o Object Recognition – they are installed in supermarkets, cameras, high-end cars such as BMW, GM, and Volvo.

Question sources

Artificial Intelligence practice questions · Question 13

Question 63 | Q01020

Which approach represents imprecise linguistic categories such as warm by graded membership between zero and one, rather than treating membership as strictly true or false?

A. Genetic algorithms

B. Fuzzy logic

C. Neural networks

D. Reinforcement learning

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer B: Fuzzy logic

Fuzzy logic represents imprecise categories through graded membership, so an input can belong to a category to a degree rather than only completely or not at all. This differs from uncertainty about whether an event occurs. Genetic search, neural networks and reinforcement learning can be combined with fuzzy methods, but their defining mechanisms do not supply this membership interpretation.

Why each choice is right or wrong

A - Incorrect: Genetic algorithms evolve populations; graded category membership is not their defining mechanism.

B - Correct: Fuzzy logic represents degrees of membership in vague categories.

C - Incorrect: Neural networks adjust weighted representations and do not inherently define fuzzy membership.

D - Incorrect: Reinforcement learning uses rewards and policy improvement rather than defining graded category membership.

Primary note proof

Study-note section AI30

> Fuzzy sets assign graded membership μ(x) between 0 and 1 to concepts such as warm or tall. A temperature can belong to warm with degree 0.7. This is not a 70% probability that an otherwise crisp warm proposition is true. Boolean logic uses only true and false; probability describes uncertainty about events.

Study-note section AI20

> A genetic algorithm maintains a population, evaluates fitness, selects parents, applies crossover and mutation, and repeats replacement. It searches without guaranteeing global optimality. Means-ends analysis selects a difference between current and goal states, chooses an operator reducing it, and recursively solves its preconditions. Neither method is simply A* running backward.

Study-note section AI33

> An RL agent observes state information, takes an action, receives reward and a new observation, and improves its policy to maximize expected cumulative reward. Rewards evaluate outcomes rather than supply the correct action label. Rewards may be delayed; not every correct-looking action receives immediate reward.

Study-note section AI34

> Training adjusts weights, often through gradients and backpropagation.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 55 | lines 27-29

> • Probabilistic reasoning is a way of knowledge representation where we apply the concept of probability to indicate the uncertainty in knowledge.
> • In probabilistic reasoning, we combine probability theory with logic to handle the uncertainty.

Introduction to Artificial Intelligence module | PDF page 71 | lines 15-16

> Neural networks, with their remarkable ability to derive meaning from complicated or imprecise data,
> can be used to extract patterns and detect trends that are too complex to be noticed by either humans or

Question sources

Artificial Intelligence practice questions · Question 14

Question 64 | Q01021

Which AI application is used to generate human-like text or speech?

A. Natural language processing

B. Speech recognition

C. Computer vision

D. Expert systems

Hint

Match the input, output and purpose to the appropriate language, perception or physical-action task.

Answer A: Natural language processing

NLG is the NLP component for generating language from computerized representations, including text and sentence planning (PDF pp. 79-80). Thus NLP is the offered application matching generation; the source does not provide a full speech-synthesis treatment.

Why each choice is right or wrong

A - Correct: Includes the explicit language-generation component. See PDF p. 79, 80.

B - Incorrect: Recognizes spoken input rather than generating output. See PDF p. 79, 80.

C - Incorrect: Works with image understanding. See PDF p. 79, 80.

D - Incorrect: Generates domain conclusions, not necessarily natural-language output. See PDF p. 79, 80.

Primary note proof

Study-note section AI37

> NLP interprets and generates human language: translation, dialogue, summarization, question answering and text generation. Speech recognition converts audio to text; speech synthesis generates speech. They overlap language processing but solve different input/output tasks.

Study-note section AI39

> Computer vision interprets images or video: classification names image content, detection locates objects, segmentation labels pixels, and recognition identifies objects or people. Occlusion, lighting change, noise, unfamiliar conditions and incomplete context make open-world perception difficult even when a face detector works well.

Study-note section AI40

> An expert system applies specialized knowledge to a domain such as equipment diagnosis. The domain expert supplies expertise; the knowledge engineer elicits and encodes it; the inference engine reasons over a knowledge base and working facts. The user interface handles interaction; an explanation facility presents why or how a conclusion was reached.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 79 | lines 31-33

> - Natural Language Generation (NLG) acts as a translator that converts the computerized data into natural language representation. It mainly involves Text planning, Sentence planning, and Text Realization.

Introduction to Artificial Intelligence module | PDF page 80 | lines 3-4

> - NLG is the process of writing or generating language.
> - NLG - It produces non-linguistic outputs from natural language inputs.

Question sources

Artificial Intelligence practice questions · Question 15

Question 65 | Q01022

Which technique involves evolving solutions using principles from genetics and natural selection?

A. Genetic algorithms

B. Neural networks

C. Reinforcement learning

D. Expert systems

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer A: Genetic algorithms

AI20 describes a genetic algorithm as maintaining a population, evaluating fitness, selecting parents and applying crossover and mutation. Those operations embody the genetics and natural-selection principles in the stem. Neural networks, reward-based learning and rule-based expert systems use different organizing mechanisms, even where hybrid systems combine them with evolutionary search.

Why each choice is right or wrong

A - Correct: Genetic algorithms use fitness, parent selection, crossover and mutation to evolve candidate populations.

B - Incorrect: Neural networks learn weighted representations; evolutionary population replacement is not required.

C - Incorrect: Reinforcement learning improves policies from environmental rewards rather than requiring genetic variation.

D - Incorrect: Expert systems reason with specialized encoded knowledge rather than requiring evolutionary populations.

Primary note proof

Study-note section AI20

A genetic algorithm maintains a population, evaluates fitness, selects parents, applies crossover and mutation, and repeats replacement. It searches without guaranteeing global optimality. Means-ends analysis selects a difference between current and goal states, chooses an operator reducing it, and recursively solves its preconditions. Neither method is simply A* running backward.

Study-note section AI34

A neuron computes $z=\Sigma iwixi+b$ and output $a=\varphi(z)$. For $x=(2,1)$, $w=(0.5,-1)$, $b=0.2$, $z=0.2$. Linear activation returns z; sigmoid returns $1/(1+e^{(-z)})$; tanh is bipolar; ReLU returns $\max(0,z)$. Nonlinearities let multiple layers represent nonlinear relationships.

Study-note section AI33

An RL agent observes state information, takes an action, receives reward and a new observation, and improves its policy to maximize expected cumulative reward. Rewards evaluate outcomes rather than supply the correct action label. Rewards may be delayed; not every correct-looking action receives immediate reward.

Study-note section AI40

An expert system applies specialized knowledge to a domain such as equipment diagnosis. The domain expert supplies expertise; the knowledge engineer elicits and encodes it; the inference engine reasons over a knowledge base and working facts. The user interface handles interaction; an explanation facility presents why or how a conclusion was reached.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 8 | lines 36-37

o Genetic Algorithms
1. Which one of the following is correctly matched with the definition and approaches of AI

Question sources

Artificial Intelligence practice questions · Question 16

Question 66 | Q01023

Which technique directly represents degrees of linguistic concepts such as warm and applies graded IF-THEN rules to choose an action?

A. Fuzzy logic

B. Reinforcement learning

C. Expert systems

D. Genetic algorithms

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer A: Fuzzy logic

Fuzzy control represents degrees of membership in concepts such as warm and combines graded rules to produce an action. This handles vague category boundaries rather than requiring probabilistic uncertainty or reward-based adaptation. Reinforcement learning improves policies through rewards, expert systems apply specialized knowledge, and genetic algorithms evolve populations. The added representation and rule requirements distinguish fuzzy control from other methods that may also address complex decisions.

Why each choice is right or wrong

A - Correct: Fuzzy logic uses graded membership and fuzzy IF-THEN rules.

B - Incorrect: RL uses reward-driven adaptation rather than defining graded linguistic membership.

C - Incorrect: An ordinary expert system need not represent degrees of membership in linguistic categories.

D - Incorrect: Genetic algorithms evolve candidates through fitness and variation rather than fuzzy rule evaluation.

Primary note proof

Study-note section AI30

> Fuzzy sets assign graded membership μ(x) between 0 and 1 to concepts such as warm or tall. A temperature can belong to warm with degree 0.7. This is not a 70% probability that an otherwise crisp warm proposition is true. Boolean logic uses only true and false; probability describes uncertainty about events.

Study-note section AI30

> A fuzzy controller fuzzifies measurements, applies graded IF-THEN rules, combines outputs and often defuzzifies to a numerical action. Fuzzy control suits linguistic rules about degrees; reinforcement learning suits reward-based adaptation. Merely saying an environment is complex or uncertain does not uniquely select either method.

Study-note section AI33

> An RL agent observes state information, takes an action, receives reward and a new observation, and improves its policy to maximize expected cumulative reward. Rewards evaluate outcomes rather than supply the correct action label. Rewards may be delayed; not every correct-looking action receives immediate reward.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 12 | lines 22-23

> expected performance.
> 2.4 Structure of Intelligent Agents

Introduction to Artificial Intelligence module | PDF page 69 | lines 18-19

> learning, the agent interacts with the environment and explores it. The goal of an agent is to get the
> most reward points, and hence, it improves its performance. The robotic dog, which automatically

Introduction to Artificial Intelligence module | PDF page 60 | lines 21-22

> • Instructing and assisting human in decision
> making

Question sources

Artificial Intelligence practice questions · Question 17

Question 67 | Q01024

Which AI application is used to simulate human conversation?

A. Natural language processing

B. Speech recognition

C. Computer vision

D. Expert systems

Hint

Match the input, output and purpose to the appropriate language, perception or physical-action task.

Answer A: Natural language processing

The module lists chatbots as an important NLP application used for chat services (PDF p. 82). That directly supports NLP for simulated conversation.

Why each choice is right or wrong

A - Correct: Includes the chatbot application. See PDF p. 82.

B - Incorrect: Transcription alone is not conversation. See PDF p. 82.

C - Incorrect: Image comprehension does not define conversation. See PDF p. 82.

D - Incorrect: Expert conclusions may be communicated, but chat is specifically listed under NLP. See PDF p. 82.

Primary note proof

Study-note section AI37

> NLP interprets and generates human language: translation, dialogue, summarization, question answering and text generation. Speech recognition converts audio to text; speech synthesis generates speech. They overlap language processing but solve different input/output tasks.

Study-note section AI02

> The conversational Turing test asks whether a judge can reliably distinguish a machine from a human through their dialogue. It requires language use, knowledge, reasoning and often learning. It tests an operational behavior criterion, not proof of consciousness. Cognitive models can be tested against response times, errors and psychological experiments.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 82 | lines 5-6

> - Chatbot - is one of the important applications of NLP. It is used by many companies to provide the customer's chat services.

Question sources

Artificial Intelligence practice questions · Question 18

Question 68 | Q01025

A collection of documents has no target labels. The task is to group similar documents, with no action-reward interaction. Which learning regime describes this objective, regardless of the architecture used?

A. Supervised learning

B. Unsupervised learning

C. Reinforcement learning

D. Deep learning

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer B: Unsupervised learning

Grouping similar documents without target labels is clustering, an unsupervised learning objective. The data modality does not determine the feedback regime, and a neural architecture could still be used for this task. Supervised learning requires targets; reinforcement learning requires an action-and-reward objective. The question asks for the learning regime rather than which architecture might process documents.

Why each choice is right or wrong

A - Incorrect: Supervised learning requires target labels, absent from the described task.

B - Correct: Unsupervised learning discovers structure, including clusters, without supplied targets.

C - Incorrect: No actions or reward feedback are present, so this is not the reinforcement objective.

D - Incorrect: Deep learning describes a multilayer architecture that might implement clustering, not the learning regime.

Primary note proof

Study-note section AI32

> Unsupervised learning seeks structure without target labels: clustering groups similar examples; association finds co-occurrence; dimensionality reduction compresses features. Good clusters have high within-cluster similarity and low between-cluster similarity under an appropriate distance. K-means alternates assignment to nearest centroids and centroid updates.

Study-note section AI43

> ANNs, genetic search and reinforcement learning all improve candidate behavior or performance through experience in different ways. Deep learning is an architecture family usable with supervised, unsupervised or reinforcement objectives. Decision trees and deep models may both be supervised. To distinguish learning regimes, specify labels, reward or unlabeled structure.

Study-note section AI31

> Supervised learning fits a mapping from labeled input-target examples.

Study-note section AI33

> An RL agent observes state information, takes an action, receives reward and a new observation, and improves its policy to maximize expected cumulative reward.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 69 | lines 7-10

> ⇒ Unsupervised learning means we are only given the x‘s. In either case, the goal is to estimate f. The system is supplied with a set of training examples consisting only of inputs and is required to

discover for itself what appropriate outputs should be,
o Most common task is clustering - e.g. taxi agent notices —bad traffic days‖

Question sources

Artificial Intelligence practice questions · Question 19

Question 69 | Q01026

Which AI technique involves training a model on a large dataset and multiple layers of artificial neurons?

A. Genetic algorithms

B. Neural networks

C. Expert systems

D. Reinforcement learning

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer B: Neural networks

AI34 describes artificial neurons with learned weights and nonlinear outputs, then defines deep learning through multiple representation layers. This matches the large-data, multilayer neural architecture in the stem. Reinforcement learning defines reward-based feedback rather than a specific architecture; genetic algorithms and expert systems do not inherently consist of layered artificial neurons.

Why each choice is right or wrong

A - Incorrect: Genetic algorithms evolve populations and are not defined by layered artificial neurons.

B - Correct: A layered collection of artificial neurons is a neural-network architecture whose weights are adjusted during training.

C - Incorrect: Expert systems define knowledge-base reasoning rather than layers of trainable neurons.

D - Incorrect: Reinforcement learning specifies the reward objective, not a required layered-neuron architecture.

Primary note proof

Study-note section AI34

A neuron computes $z=\Sigma iwixi+b$ and output $a=\varphi(z)$. For $x=(2,1)$, $w=(0.5,-1)$, $b=0.2$, $z=0.2$. Linear activation returns z; sigmoid returns $1/(1+e^{(-z)})$; tanh is bipolar; ReLU returns $\max(0,z)$. Nonlinearities let multiple layers represent nonlinear relationships.

Study-note section AI20

A genetic algorithm maintains a population, evaluates fitness, selects parents, applies crossover and mutation, and repeats replacement. It searches without guaranteeing global optimality. Means-ends analysis selects a difference between current and goal states, chooses an operator reducing it, and recursively solves its preconditions. Neither method is simply A* running backward.

Study-note section AI40

An expert system applies specialized knowledge to a domain such as equipment diagnosis. The domain expert supplies expertise; the knowledge engineer elicits and encodes it; the inference engine reasons over a knowledge base and working facts. The user interface handles interaction; an explanation facility presents why or how a conclusion was reached.

Study-note section AI33

An RL agent observes state information, takes an action, receives reward and a new observation, and improves its policy to maximize expected cumulative reward. Rewards evaluate outcomes rather than supply the correct action label. Rewards may be delayed; not every correct-looking action receives immediate reward.

Study-note section AI34

Feedforward networks pass signals toward outputs without recurrent state loops; recurrent networks carry state across sequence steps. Training adjusts weights, often through gradients and backpropagation. Backpropagation is not the same as a recurrent connection. Deep learning uses multiple representation layers and can learn features from images, sound and text. CNNs share filters across image locations. ANN is a brain-inspired abstraction, not a literal brain simulation; DL⊂ML⊂AI.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 71 | lines 9-10

It is composed of large number of highly interconnected simple processing elements (neurons) working in unison to solve specific problems. ANNs, like people, learn by example. An ANN is

Introduction to Artificial Intelligence module | PDF page 74 | lines 4-6

- Network Topology
- Adjustments of Weights or Learning
- Activation Functions

Question sources

Artificial Intelligence practice questions · Question 20

Question 70 | Q01027

What is Artificial Intelligence?

A. Artificial Intelligence is a field that aims to make humans more intelligent

B. Artificial Intelligence is a field that aims to improve the security

C. Artificial Intelligence is a field that aims to develop intelligent machines

D. Artificial Intelligence is a field that aims to mine the data

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer C: Artificial Intelligence is a field that aims to develop intelligent machines

The module defines AI as designing intelligence in an artificial device (PDF p. 4). Developing intelligent machines is the matching general description, rather than an individual application.

Why each choice is right or wrong

A - Incorrect: Making people smarter is not the stated field definition. See PDF p. 4.

B - Incorrect: Security is a possible application, not the field definition. See PDF p. 4.

C - Correct: Matches intelligence in an artificial device. See PDF p. 4.

D - Incorrect: Data mining is one learning application, not the field's definition. See PDF p. 4.

Primary note proof

Study-note section AI01

> Artificial intelligence studies systems that perform tasks requiring intelligence: reasoning, learning, planning, perception and language. Applied AI builds useful solutions; cognitive AI uses computational models to investigate biological intelligence. A model should be compared with human behavior or experimental observations before claiming it explains human thought.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 4 | lines 26-27

> is concerned with the design of intelligence in an artificial device. The term was coined by McCarthy
> in 1956. There are two ideas in the definition. i.e Intelligence, and Artificial device.

Question sources

Artificial Intelligence practice questions · Question 21

Question 71 | Q01028

Who is credited with coining the term artificial intelligence in the proposal for the Dartmouth summer research project?

A. Geoffrey Hinton

B. Andrew Ng

C. John McCarthy

D. Jürgen Schmidhuber

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer C: John McCarthy

John McCarthy is credited with coining the term in the 1955 proposal for the Dartmouth summer research project. The workshop itself took place in 1956 and helped establish the research field. This naming attribution does not make McCarthy the sole inventor of every AI idea or method. The question now asks for the documented attribution rather than an oversimplified invention claim.

Why each choice is right or wrong

A - Incorrect: The historical naming attribution in the note is to McCarthy, not Hinton.

B - Incorrect: The historical naming attribution in the note is to McCarthy, not Ng.

C - Correct: McCarthy is credited with introducing the term in the 1955 Dartmouth proposal.

D - Incorrect: The historical naming attribution in the note is to McCarthy, not Schmidhuber.

Primary note proof

Study-note section AI-S04

> John McCarthy is credited with coining the term artificial intelligence in the 1955 proposal for the Dartmouth summer research project. The Dartmouth workshop took place in 1956 and helped establish the field. Naming the field does not mean that one person invented all its ideas or methods.

Study-note section AI43

> Face recognition is both an applied task and a narrow capability. Broad machine-intelligence goals can include emulating human intelligence and building systems that learn. Humanlike problem-solving alone does not uniquely select expert systems. Python is widely used, but C++ and Java also implement AI; “commonly used” does not isolate one language. McCarthy is associated with naming AI, not its sole invention.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 4 | lines 26-27

is concerned with the design of intelligence in an artificial device. The term was coined by McCarthy
in 1956. There are two ideas in the definition. i.e Intelligence, and Artificial device.

Question sources

Artificial Intelligence practice questions · Question 22

Question 72 | Q01029

Which activity, as stated, directly requires an intelligence-related capability, rather than merely naming a computing operation that may or may not use AI?

A. It helps to exploit vulnerabilities to secure the firm

B. Language understanding and problem-solving (Text analytics and NLP)

C. Easy to create a website

D. It helps to deploy applications on the cloud

Hint

Match the input, output and purpose to the appropriate language, perception or physical-action task.

Answer B: Language understanding and problem-solving (Text analytics and NLP)

AI concerns capabilities such as reasoning, learning and language understanding. The language-understanding and problem-solving option explicitly names such capabilities. Website creation, cloud deployment and security operations can incorporate AI, but the activity names alone do not establish its use. The distinction is between an explicit intelligence-related task and a computing context that could host many kinds of software.

Why each choice is right or wrong

A - Incorrect: Security testing or exploitation can use ordinary programmed tools; the activity alone does not require intelligent learning or reasoning.

B - Correct: Understanding human language and solving problems directly name intelligence-related capabilities.

C - Incorrect: Creating a website can use conventional software tools without intelligent behavior.

D - Incorrect: Deploying an application to a cloud platform does not itself require an intelligent application.

Primary note proof

Study-note section AI01

Artificial intelligence studies systems that perform tasks requiring intelligence: reasoning, learning, planning, perception and language. Applied AI builds useful solutions; cognitive AI uses computational models to investigate biological intelligence. A model should be compared with human behavior or experimental observations before claiming it explains human thought.

Study-note section AI37

NLP interprets and generates human language: translation, dialogue, summarization, question answering and text generation. Speech recognition converts audio to text; speech synthesis generates speech. They overlap language processing but solve different input/output tasks.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 6 | lines 18-19

• Understanding language involving natural language processing, speech processing
• Solving problems

Introduction to Artificial Intelligence module | PDF page 6 | lines 19-20

• Solving problems
• Robotics

Question sources

Artificial Intelligence practice questions · Question 23

Question 73 | Q01030

Which of the following is a component of Artificial Intelligence?

A. Learning

B. Training

C. Designing

D. Puzzling

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer A: Learning

Learning appears explicitly among intelligent behaviors (PDF p. 6), and is defined as improving performance from experience on p. 67. The other words are not offered as parallel components in the module.

Why each choice is right or wrong

A - Correct: Explicitly named intelligent capability. See PDF p. 6.

B - Incorrect: Training is a model-building activity, not the named component here. See PDF p. 6.

C - Incorrect: Designing is not presented as the matching intelligent capability. See PDF p. 6.

D - Incorrect: Puzzling is not a defined component. See PDF p. 6.

Primary note proof

Study-note section AI08

> The performance element selects actions. The learning element changes the agent to improve future performance. The critic evaluates behavior against a performance standard and supplies feedback. The problem generator suggests exploratory actions likely to yield useful experience.

Study-note section AI08

> Learning is improvement on task T, measured by P, through experience E. For spam filtering, T is classifying email, E is past labeled messages, and P might be recall at a required precision. Exploration trades immediate reward for information; exploitation uses the best action currently known. Learning can improve reflex, goal-based or utility-based agents.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 6 | lines 17-19

> • Learning
> • Understanding language involving natural language processing, speech processing
> • Solving problems

Question sources

Artificial Intelligence practice questions · Question 24

Question 74 | Q01032

Which of the following is not a type of Artificial Intelligence agent?

A. Learning AI agent

B. Goal-based AI agent

C. Simple reflex AI agent

D. Unity-based AI agent

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer D: Unity-based AI agent

The module lists simple reflex, model-based reflex, goal-based, utility-based and learning agents (PDF pp. 14-15). “Unity-based” is not “utility-based” and is not one of those classes.

Why each choice is right or wrong

A - Incorrect: Learning is a listed class. See PDF p. 14, 15.

B - Incorrect: Goal-based is listed. See PDF p. 14, 15.

C - Incorrect: Simple reflex is listed. See PDF p. 14, 15.

D - Correct: Unity-based is not a listed class. See PDF p. 14, 15.

Primary note proof

Study-note section AI42

> The common teaching list contains five agent classes when learning agents are included, or four decision architectures before learning is added. State the convention when counting agent types. “Target-based” can mean goal-based; “unity-based” is not utility-based.

Study-note section AI06

> A simple reflex agent matches condition-action rules against the current percept only: if dirty, clean. It has no persistent world-state memory. This is adequate when the current observation contains everything needed for its rule, but partial observability can make the same percept require different actions.

Study-note section AI08

> The performance element selects actions. The learning element changes the agent to improve future performance. The critic evaluates behavior against a performance standard and supplies feedback. The problem generator suggests exploratory actions likely to yield useful experience.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 14 | lines 10-11

> Agents can be grouped into five classes based on their degree of perceived intelligence and capability.
> All these agents can improve their performance and generate better action over the time.

Introduction to Artificial Intelligence module | PDF page 15 | lines 20-21

> 2.9.5 Learning agent
> • Learning allows an agent to operate in initially unknown environments.

Question sources

Artificial Intelligence practice questions · Question 26

Question 75 | Q01033

Which agent architecture applies fixed condition-action rules directly to the current percept, without retaining an internal world-state model?

A. Simple reflex agent

B. Model-based agent

C. Goal-based agent

D. Utility-based agent

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer A: Simple reflex agent

The absence of persistent internal state is the key distinction. A simple reflex agent maps the current percept to an action using condition-action rules. Model-based reflex agents may also use rules, but they maintain a world-state representation. Goal-based and utility-based agents use goals or outcome preferences to guide selection, rather than being defined by current-percept-only rule matching.

Why each choice is right or wrong

A - Correct: A simple reflex agent chooses from the current percept without persistent world-state memory.

B - Incorrect: A model-based reflex agent adds internal state and applies rules to that modeled state.

C - Incorrect: Goal-based selection evaluates actions or future sequences against a desired state.

D - Incorrect: Utility-based selection ranks outcomes by desirability or expected utility.

Primary note proof

Study-note section AI06

> A simple reflex agent matches condition-action rules against the current percept only: if dirty, clean. It has no persistent world-state memory. This is adequate when the current observation contains everything needed for its rule, but partial observability can make the same percept require different actions.

Study-note section AI06

> A model-based reflex agent maintains internal state using the previous state, previous action, current percept, and models of world change and action effects. It then applies rules to that state. Both kinds can use condition-action rules; memory and a world model distinguish them.

Study-note section AI07

> A goal-based agent considers future action sequences that reach a desired state. A utility-based agent compares degrees of desirability and trade-offs, choosing maximum expected utility: $EU(a)=\Sigma s\ P(s|a,evidence)U(s)$. Reaching the destination is a goal; balancing arrival time, safety and cost requires preferences among outcomes.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 14 | lines 17-18

> • The Simple reflex agent works on Condition-action rule, which means it maps the current state to
> action.

Introduction to Artificial Intelligence module | PDF page 14 | lines 29-30

- Internal State: It is a representation of the current state based on percept history.
• These agents have the model, "which is knowledge of the world" and based on the model they

Question sources

Artificial Intelligence practice questions · Question 27

Question 76 | Q01034

Which type of agent maintains an internal model of the world to make decisions?

A. Simple reflex agent

B. Model-based agent

C. Goal-based agent

D. Utility-based agent

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer B: Model-based agent

An internal state based on percept history and a model of how the world changes define the model-based category. Goal and utility agents can extend that architecture, so the requested defining type is B. Evidence: PDF 14.

Why each choice is right or wrong

A - Incorrect: Current-percept rules omit history. See PDF p. 14.

B - Correct: Matches the internal model definition. See PDF p. 14.

C - Incorrect: Adds explicit goal information to a model. See PDF p. 14.

D - Incorrect: Adds a numerical utility criterion. See PDF p. 14.

Primary note proof

Study-note section AI06

A model-based reflex agent maintains internal state using the previous state, previous action, current percept, and models of world change and action effects. It then applies rules to that state. Both kinds can use condition-action rules; memory and a world model distinguish them.

Study-note section AI07

A goal-based agent considers future action sequences that reach a desired state. A utility-based agent compares degrees of desirability and trade-offs, choosing maximum expected utility: $EU(a)=\Sigma s\ P(s|a,evidence)U(s)$. Reaching the destination is a goal; balancing arrival time, safety and cost requires preferences among outcomes.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 14 | lines 29-31

- Internal State: It is a representation of the current state based on percept history.
• These agents have the model, "which is knowledge of the world" and based on the model they perform actions.

Question sources

Artificial Intelligence practice questions · Question 28

Question 77 | Q01035

Which type of agent selects actions based on their expected utility or desirability?

A. Simple reflex agent

B. Model-based agent

C. Goal-based agent

D. Utility-based agent

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer D: Utility-based agent

A utility function assigns real numbers to states to compare how efficiently actions achieve goals; this is precisely desirability-based selection. Evidence: PDF 15.

Why each choice is right or wrong

A - Incorrect: Uses condition-action rules. See PDF p. 15.

B - Incorrect: A world model alone supplies no preference scale. See PDF p. 15.

C - Incorrect: Goal achievement alone does not rank desirable outcomes. See PDF p. 15.

D - Correct: Numerical utility ranks alternatives. See PDF p. 15.

Primary note proof

Study-note section AI07

> A goal-based agent considers future action sequences that reach a desired state. A utility-based agent compares degrees of desirability and trade-offs, choosing maximum expected utility: EU(a)=Σs P(s|a,evidence)U(s). Reaching the destination is a goal; balancing arrival time, safety and cost requires preferences among outcomes.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 15 | lines 18-20

> • The utility function maps each state to a real number to check how efficiently each action achieves
> the goals.
> 2.9.5 Learning agent

Question sources

Artificial Intelligence practice questions · Question 29

Question 78 | Q01036

Which type of agent represents goals and selects actions that maximize the achievement of those goals?

A. Simple reflex agent

B. Model-based agent

C. Goal-based agent

D. Utility-based agent

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer C: Goal-based agent

Explicit desirable goal situations and choosing actions to reach them characterize goal-based agents. Evidence: PDF 15.

Why each choice is right or wrong

A - Incorrect: Responds to current percept. See PDF p. 15.

B - Incorrect: Tracks state without necessarily representing a goal. See PDF p. 15.

C - Correct: Explicit goal information supplies the requested criterion. See PDF p. 15.

D - Incorrect: Adds utility rankings beyond mere goal achievement. See PDF p. 15.

Primary note proof

Study-note section AI07

> A goal-based agent considers future action sequences that reach a desired state. A utility-based agent compares degrees of desirability and trade-offs, choosing maximum expected utility: EU(a)=Σs P(s|a,evidence)U(s). Reaching the destination is a goal; balancing arrival time, safety and cost requires preferences among outcomes.

Study-note section AI07

> Updating an internal state after selecting an action predicts the action’s effect; it does not itself modify the external environment. A stored variable named goal is insufficient to prove a goal-based policy if BEST_ACTION never uses it. Model-based, goal-based, utility-based and learning capabilities can coexist.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 15 | lines 6-8

> • Goal-based agents expand the capabilities of the model-based agent by having the "goal" information.
> • They choose an action, so that they can achieve the goal.

Question sources

Artificial Intelligence practice questions · Question 30

Question 79 | Q01037

Which type of agent improves its performance over time through learning from experience?

A. Simple reflex agent

B. Model-based agent

C. Goal-based agent

D. Learning agent

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer D: Learning agent

Learning agents adapt using past experience and a learning element that changes performance. Evidence: PDF 15.

Why each choice is right or wrong

A - Incorrect: Current rules do not themselves explain improvement. See PDF p. 15.

B - Incorrect: Tracking state is not itself learning. See PDF p. 15.

C - Incorrect: Goal pursuit need not change performance over time. See PDF p. 15.

D - Correct: Explicitly learns from experience. See PDF p. 15.

Primary note proof

Study-note section AI08

> The performance element selects actions. The learning element changes the agent to improve future performance. The critic evaluates behavior against a performance standard and supplies feedback. The problem generator suggests exploratory actions likely to yield useful experience.

Study-note section AI08

> Learning is improvement on task T, measured by P, through experience E. For spam filtering, T is classifying email, E is past labeled messages, and P might be recall at a required precision. Exploration trades immediate reward for information; exploitation uses the best action currently known. Learning can improve reflex, goal-based or utility-based agents.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 15 | lines 24-26

• A learning agent in AI is the type of agent which can learn from its past experiences, or it has learning
capabilities.
• It starts to act with basic knowledge and then able to act and adapt automatically through learning.

Question sources

Artificial Intelligence practice questions · Question 31

Question 80 | Q01038

A face-recognition system is built to control entry to a building, without trying to explain biological recognition processes. Which approach classifies this practical engineering purpose, rather than its breadth of capability?

A. Weak AI approach

B. Applied AI approach

C. Cognitive AI approach

D. Strong AI approach

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer B: Applied AI approach

Applied AI builds useful solutions to practical tasks, so building-access recognition fits that purpose. The same system is also narrow or weak AI in capability scope. These labels answer different classification questions and are compatible. Cognitive AI would investigate biological intelligence, while strong-AI claims concern broader competence or actual understanding. The explicit purpose axis isolates the intended selection.

Why each choice is right or wrong

A - Incorrect: Weak AI describes bounded capability, which also fits face recognition, but does not classify the purpose requested.

B - Correct: Applied AI builds a useful practical solution, matching the entry-control purpose.

C - Incorrect: Cognitive AI investigates biological intelligence; that is not the stated purpose.

D - Incorrect: Strong AI concerns actual understanding or broad intelligence claims, not the application purpose described.

Primary note proof

Study-note section AI01

Artificial intelligence studies systems that perform tasks requiring intelligence: reasoning, learning, planning, perception and language. Applied AI builds useful solutions; cognitive AI uses

computational models to investigate biological intelligence. A model should be compared with human behavior or experimental observations before claiming it explains human thought.

Study-note section AI01

Narrow or weak AI addresses bounded tasks, such as face recognition, spam filtering or a voice assistant. General AI denotes broad competence transferable across domains; strong AI also has a philosophical meaning involving actual understanding. Do not equate a humanoid body, autonomous vehicle, or many features with general intelligence.

Study-note section AI43

Face recognition is both an applied task and a narrow capability. Broad machine-intelligence goals can include emulating human intelligence and building systems that learn. Humanlike problem-solving alone does not uniquely select expert systems. Python is widely used, but C++ and Java also implement AI; “commonly used” does not isolate one language. McCarthy is associated with naming AI, not its sole invention.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 7 | lines 6-8

Applied AI: aims to produce commercially viable smart system. For example, a security system that
is able to recognize the faces of people who are permitted to enter a particular building. Applied AI
has already enjoyed considerable success.

Question sources

Artificial Intelligence practice questions · Question 32

Question 81 | Q01039

Which single option gives the complete set of possible AI benefits in suitable, evaluated applications? Interpret the choices operationally, without assuming consciousness or guaranteed improvement.

A. Can reduce search work and time with an effective heuristic.

B. Can support protective filtering by detecting unwanted spam messages.

C. Can automate reasoning and learning to help solve tasks.

D. All of the above

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer D: All of the above

Useful AI methods can reduce search work when a heuristic guides exploration effectively. Learned detection of unwanted spam can support protective filtering, and computational reasoning or learning can automate parts of a task. These are possible operational benefits rather than guarantees for every deployment or claims of consciousness. Each individual choice gives only one benefit, so the inclusive response supplies the complete set.

Why each choice is right or wrong

A - Incorrect: Effective heuristics can reduce explored work, but this is only one of the listed possible benefits.

B - Incorrect: Spam classification can support protective filtering, but this alone is incomplete.

C - Incorrect: Automating computational reasoning and learning can help tasks, but this alone is incomplete.

D - Correct: All three are possible operational benefits under the stated suitable-application scope.

Primary note proof

Study-note section AI17

> A heuristic h(n) estimates remaining cost or evaluates promise using problem-specific knowledge. It can reduce explored work, but computing it also costs time. Uninformed algorithms use only the problem definition and known costs; UCS remains uninformed despite using g.

Study-note section AI08

> Learning is improvement on task T, measured by P, through experience E. For spam filtering, T is classifying email, E is past labeled messages, and P might be recall at a required precision. Exploration trades immediate reward for information; exploitation uses the best action currently known. Learning can improve reflex, goal-based or utility-based agents.

Study-note section AI01

> Artificial intelligence studies systems that perform tasks requiring intelligence: reasoning, learning, planning, perception and language. Applied AI builds useful solutions; cognitive AI uses computational models to investigate biological intelligence. A model should be compared with human behavior or experimental observations before claiming it explains human thought.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 7 | lines 6-8

> Applied AI: aims to produce commercially viable smart system. For example, a security system that
> is able to recognize the faces of people who are permitted to enter a particular building. Applied AI
> has already enjoyed considerable success.

Introduction to Artificial Intelligence module | PDF page 79 | lines 11-12

- It is very time efficient.
- Most of the companies use NLP to improve the efficiency of documentation processes,

Introduction to Artificial Intelligence module | PDF page 7 | lines 3-5

Weak AI: deals with the creation of some form of computer-based artificial intelligence that cannot
truly reason and solve problems, but can act as if it were intelligent. Weak AI holds that suitably programmed machines can simulate human cognition.

Question sources

Artificial Intelligence practice questions · Question 33

Question 82 | Q01040

Which of the following is/are the composition for AI agents?

A. Program only

B. Architecture only

C. Both Program and Architecture

D. None of the mentioned

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer C: Both Program and Architecture

The structural equation adds architecture and agent program; neither component alone is the complete agent. Evidence: PDF 12.

Why each choice is right or wrong

A - Incorrect: Omits executing machinery. See PDF p. 12.

B - Incorrect: Omits implementing program. See PDF p. 12.

C - Correct: Contains both terms in the equation. See PDF p. 12.

D - Incorrect: Contradicts the explicit equation. See PDF p. 12.

Primary note proof

Study-note section AI03

A percept is one observation; a percept sequence is the entire observation history. An agent function maps percept sequences to actions: f:P*→A. An agent program implements that mapping on an architecture; agent = architecture + program. Environment means the setting in which it senses and acts, not merely the search space. Autonomy means relying substantially on experience rather than only a designer’s prior instructions.

Study-note section AI-S03

> An AI task is modeled as an agent interacting with its environment: observations enter through sensors and selected actions affect the environment through actuators. This task-level description differs from an agent implementation, which combines a program with a computing architecture.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 12 | lines 25-27

> • Agent = Architecture + Agent Program
> • Architecture = the machinery that an agent executes on.
> • Agent Program = an implementation of an agent function.

Question sources

Artificial Intelligence practice questions · Question 34

Question 83 | Q01041

Which are examples of artificial agents, assuming the spacecraft and robot use constructed controllers that sense and select actions autonomously? Select all that apply.

A. Autonomous Spacecraft

B. Human

C. Robot

D. All of the mentioned

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer A, C: Autonomous Spacecraft; Robot

An artificial agent uses a constructed sensing and action-selection system to interact with its environment. Under the explicit control assumption, both the autonomous spacecraft and robot qualify. Humans also perceive and act, so they are agents, but they are not artificial agents. The inclusive choice fails because it includes humans. This classification concerns artificial agency and does not assert that a robot has general intelligence.

Why each choice is right or wrong

A - Correct: An autonomous spacecraft uses constructed sensing and action-selection mechanisms to interact with its environment.

B - Incorrect: A human is an agent but is not an artificial agent.

C - Correct: An AI-controlled robot senses and acts physically and is an artificial-agent example.

D - Incorrect: Including humans makes this aggregate false under the question's artificial-agent wording.

Primary note proof

Study-note section AI03

> An agent perceives an environment through sensors and acts through actuators or effectors. Cameras sense; motors move. Software sensors may be incoming messages and software actions may be API requests. Humans are agents but are not artificial agents.

Study-note section AI39

> Robotics combines mechanical, electrical, control and computing systems to sense, plan and act physically. Components include sensors, controller, power supply, actuators and an end effector. Locomotion may be legged, wheeled or tracked, with trade-offs in terrain, stability and efficiency. Software can act autonomously without being a robot. Humanoid shape is not evidence of general intelligence.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 12 | lines 2-4

> • Human: sensor (eye, ears, skins, taste). effectors (Hand, fingers, legs, mouth)
> • Robot: sensor (camera, infrared,..), effectors (gripper, wheel, speaker,.)
> • Software agent: function as sensor and actuator

Introduction to Artificial Intelligence module | PDF page 82 | lines 11-13

> efficient robots. Robots are the artificial agents acting in real world environment. Robots are aimed at
> manipulating the objects by perceiving, picking, moving, modifying the physical properties of object,
> destroying it, or to have an effect thereby freeing manpower from doing repetitive functions without

Question sources

Artificial Intelligence practice questions · Question 35

Question 84 | Q01042

Which of the following is an expansion of Artificial Intelligence application?

A. Game Playing

B. Planning and Scheduling

C. Diagnosis

D. All of the mentioned

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer D: All of the mentioned

The application list includes medical diagnosis, telescope planning/scheduling, and grand-master game play. All three categories are present. Evidence: PDF 7, 7.

Why each choice is right or wrong

A - Incorrect: Games are explicitly listed. See PDF p. 7, 7.

B - Incorrect: Planning and scheduling are explicitly listed. See PDF p. 7, 7.

C - Incorrect: Medical diagnosis is explicitly listed. See PDF p. 7, 7.

D - Correct: Combines all three documented applications. See PDF p. 7, 7.

Primary note proof

Study-note section AI-S03

> Planning constructs actions to reach a goal. Scheduling additionally assigns activities to times and resources while respecting constraints such as resource conflicts and precedence. It is an AI application when search or constraint reasoning chooses feasible or useful assignments; automated telescope experiment scheduling is an example.

Study-note section AI22

> Adversarial search models agents with competing objectives. A game specifies states, player to move, legal actions, successor function, terminal test and utility. Perfect information means previous moves and current state are visible; it does not reveal future choices or imply absence of chance. Zero-sum utility makes one player’s gain the other’s loss.

Study-note section AI40

> An expert system applies specialized knowledge to a domain such as equipment diagnosis. The domain expert supplies expertise; the knowledge engineer elicits and encodes it; the inference engine reasons over a knowledge base and working facts. The user interface handles interaction; an explanation facility presents why or how a conclusion was reached.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 7 | lines 22-24

> • Today‘s Expert systems can carry out medical diagnosis in a narrow domain
> • Speech understanding systems are capable of recognizing several thousand words continuous speech

Introduction to Artificial Intelligence module | PDF page 7 | lines 25-29

> • Planning and scheduling systems had been employed in scheduling experiments with the Hubble Telescope

• The Learning systems are capable of doing text categorization into about a 1000 topics
• In Games, AI systems can play at the Grand Master level in chess (world champion), checkers, etc.

Question sources

Artificial Intelligence practice questions · Question 36

Question 85 | Q01043

Which option gives the complete set of valid general descriptions or examples of artificial agents?

A. Receives input from its surroundings and selects operations that act on those surroundings.

B. A line-following robot controller that reads sensors and commands the robot's motors.

C. Perceives its environment through sensors and acting upon that environment through actuators

D. All of the mentioned

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer D: All of the mentioned

An agent perceives its environment and acts on it. The input-and-operation paraphrase describes that interaction, the sensor-driven line-following controller is an implementation example, and the sensor/actuator statement supplies the general definition. All are valid under the stated descriptions. Since the question requests the complete set in one selection, the aggregate is selected; each individual statement is true but incomplete.

Why each choice is right or wrong

A - Incorrect: This is a valid input-and-action paraphrase, but it alone omits the other valid descriptions.

B - Incorrect: The controller senses and acts through the robot, so it is a valid example, but alone is not the complete set.

C - Incorrect: This gives the general sensing-and-acting definition, but alone omits the valid paraphrase and example.

D - Correct: The paraphrase, robot-controller example and sensor/actuator definition all fit artificial agency.

Primary note proof

Study-note section AI03

An agent perceives an environment through sensors and acts through actuators or effectors. Cameras sense; motors move. Software sensors may be incoming messages and software actions may be API requests. Humans are agents but are not artificial agents.

Study-note section AI03

A percept is one observation; a percept sequence is the entire observation history. An agent function maps percept sequences to actions: f:P*→A. An agent program implements that mapping on an architecture; agent = architecture + program. Environment means the setting in which it senses and acts, not merely the search space. Autonomy means relying substantially on experience rather than only a designer's prior instructions.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 10 | lines 21-23

• An agent is anything that can perceive its environment through sensors and acts upon that environment through effectors. An agent perceives its environment through sensors.
• The complete set of inputs at a given time is called a percept.

Question sources

Artificial Intelligence practice questions · Question 37

Question 86 | Q01044

What is the name of Artificial Intelligence which allows machines to handle vague information with a deftness that mimics human intuition?

A. Human intelligence

B. Boolean logic

C. Functional logic

D. Fuzzy logic

Hint

Translate the statement carefully, then check the scope of its quantifiers, connectives or inference rule. Distinguish guaranteed consequences from plausible generalizations or explanations.

Answer D: Fuzzy logic.

AI30 represents vague categories with degrees between zero and one and explains fuzzy control using graded rules. This supplies the concrete mechanism behind the stem's loose phrase about human intuition. Boolean logic uses only true or false. Human intelligence is not the named machine-reasoning formalism here, and functional logic is not the provided graded-membership method.

Why each choice is right or wrong

A - Incorrect: Human intelligence names biological capability, not the specified artificial graded-membership technique.

B - Incorrect: Boolean logic uses crisp true/false values rather than degrees of membership in vague categories.

C - Incorrect: Functional logic is not the graded-membership and fuzzy-rule formalism described in AI30.

D - Correct: Fuzzy logic handles vague linguistic categories through graded membership and graded rules.

Primary note proof

Study-note section AI30

> Fuzzy sets assign graded membership μ(x) between 0 and 1 to concepts such as warm or tall. A temperature can belong to warm with degree 0.7. This is not a 70% probability that an otherwise crisp warm proposition is true. Boolean logic uses only true and false; probability describes uncertainty about events.

Study-note section AI25

> Propositional atoms represent whole true-or-false claims. ¬P negates; P∧Q requires both; P∨Q requires at least one; P→Q is false only when P is true and Q false; P↔Q requires matching truth values. Thus P→Q ≡ ¬P∨Q. De Morgan gives ¬(P∧Q) ≡ ¬P∨¬Q.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 45 | lines 19-21

> • Propositional/Boolean logic as it works on 0 and 1.
> • In PL, we use symbolic variables to represent the logic, and we can use any symbol for a representing
> a proposition, such A, B, C, P, Q, R, etc.

Question sources

Artificial Intelligence practice questions · Question 38

Question 87 | Q01045

An AI system is composed of?

A. agent

B. environment

C. Both A and B

D. None of the Above

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer C: Both A and B

The explicit system composition is agent plus environment. Evidence: PDF 11.

Why each choice is right or wrong

A - Incorrect: Only one component. See PDF p. 11.

B - Incorrect: Only the other component. See PDF p. 11.

C - Correct: Exactly both documented components. See PDF p. 11.

D - Incorrect: Conflicts with the explicit composition. See PDF p. 11.

Primary note proof

Study-note section AI-S03

> An AI task is modeled as an agent interacting with its environment: observations enter through sensors and selected actions affect the environment through actuators. This task-level description differs from an agent implementation, which combines a program with a computing architecture.

Study-note section AI03

> A percept is one observation; a percept sequence is the entire observation history. An agent function maps percept sequences to actions: f:P*→A. An agent program implements that mapping on an architecture; agent = architecture + program. Environment means the setting in which it senses and acts, not merely the search space. Autonomy means relying substantially on experience rather than only a designer’s prior instructions.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 11 | lines 21-23

> • An AI system is composed of an agent and its environment.
> • The agents act in their environment. The environment may contain other agents.
> • An agent is anything that can perceive its environment through sensors and acts upon that

Question sources

Artificial Intelligence practice questions · Question 39

Question 88 | Q01046

What are Agent?

A. An agent perceives its environment through sensors and acts on it through actuators.

B. An agent is anything that can change its environment through sensors

C. An agent is anything that can control its environment through sensors

D. None of the Above

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer A: An agent perceives its environment through sensors and acts on it through actuators.

An agent is an entity that perceives an environment through sensors and acts on it through actuators or effectors. Sensing alone is only part of this definition, so the first option now includes acting. The other descriptions incorrectly give sensors the action or control role. An agent need not control everything in its environment or guarantee a desired outcome.

Why each choice is right or wrong

A - Correct: This states both parts of the agent interaction: sensing and acting.

B - Incorrect: Sensors supply observations; they are not the general components that execute environmental changes.

C - Incorrect: Sensors do not generally control the environment; action uses actuators or effectors.

D - Incorrect: A complete agent definition is supplied by the sensing-and-acting statement.

Primary note proof

Study-note section AI03

> An agent perceives an environment through sensors and acts through actuators or effectors. Cameras sense; motors move. Software sensors may be incoming messages and software actions may be API requests. Humans are agents but are not artificial agents.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 10 | lines 21-23

> • An agent is anything that can perceive its environment through sensors and acts upon that environment through effectors. An agent perceives its environment through sensors.
> • The complete set of inputs at a given time is called a percept.

Question sources

Artificial Intelligence practice questions · Question 40

Question 89 | Q01047

Which instruments are used for perceiving and acting upon the environment?

A. Sensors and Actuators

B. Sensors

C. Perceiver

D. Perceiver and Sensor

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer A: Sensors and Actuators

Sensors receive environmental input and actuators deliver output, so both are needed. Evidence: PDF 13.

Why each choice is right or wrong

A - Correct: Includes perception and action interfaces. See PDF p. 13.

B - Incorrect: Omits action interface. See PDF p. 13.

C - Incorrect: Does not identify the two instruments. See PDF p. 13.

D - Incorrect: Two perception terms omit actuators. See PDF p. 13.

Primary note proof

Study-note section AI03

> An agent perceives an environment through sensors and acts through actuators or effectors. Cameras sense; motors move. Software sensors may be incoming messages and software actions may be API requests. Humans are agents but are not artificial agents.

Study-note section AI05

> PEAS means Performance measure, Environment, Actuators, Sensors. Specify all four before designing an agent. For a taxi: performance includes safety and travel time; environment includes roads, traffic and passengers; actuators include steering, brakes and displays; sensors include cameras, positioning and speed readings.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 13 | lines 6-8

> • The agent takes input from the environment through sensors and delivers the output to the environment through actuators.
> • The most famous artificial environment is the Turing Test environment, in which one real and

Question sources

Artificial Intelligence practice questions · Question 41

Question 90 | Q01048

In the teaching taxonomy of decision architectures before learning agents are added as a separate class, how many agent types are counted?

A. 1

B. 2

C. 3

D. 4

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer D: 4

The note distinguishes a four-architecture convention from a five-class list that also includes learning agents. The specified convention counts simple reflex, model-based reflex, goal-based and utility-based architectures. Learning can improve several of those architectures, but it is excluded as a separate class for this count. The explicit taxonomy therefore determines the number without claiming a universal count for every possible classification.

Why each choice is right or wrong

A - Incorrect: One does not match either teaching convention described in AI42.

B - Incorrect: Two does not match the stated architecture or learning-inclusive list.

C - Incorrect: Three is not either count given here.

D - Correct: The stated pre-learning architecture taxonomy contains four types.

Primary note proof

Study-note section AI42

> The common teaching list contains five agent classes when learning agents are included, or four decision architectures before learning is added. State the convention when counting agent types. “Target-based” can mean goal-based; “unity-based” is not utility-based.

Study-note section AI06

> A simple reflex agent matches condition-action rules against the current percept only: if dirty, clean. It has no persistent world-state memory. This is adequate when the current observation contains everything needed for its rule, but partial observability can make the same percept require different actions.

Study-note section AI07

A goal-based agent considers future action sequences that reach a desired state. A utility-based agent compares degrees of desirability and trade-offs, choosing maximum expected utility: EU(a)=Σs P(s | a,evidence)U(s). Reaching the destination is a goal; balancing arrival time, safety and cost requires preferences among outcomes.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 14 | lines 10-12

Agents can be grouped into five classes based on their degree of perceived intelligence and capability.
All these agents can improve their performance and generate better action over the time.
2.9.1 Simple reflex agent

Question sources

Artificial Intelligence practice questions · Question 42

Question 91 | Q01049

Which of the following is not an agent decision architecture?

A. Model based

B. Utility based

C. Simple reflex

D. Search frontier

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer D: Search frontier

Model-based, utility-based and simple reflex describe ways to organize agent decision making. A search frontier instead stores generated nodes awaiting selection or expansion during a search. It is an algorithm component rather than a decision architecture. The overlapping target-based distractor has been replaced because target-based can reasonably mean goal-based, which is a valid agent category.

Why each choice is right or wrong

A - Incorrect: A model-based architecture maintains an internal representation used in action selection.

B - Incorrect: A utility-based architecture selects using outcome desirability.

C - Incorrect: Simple reflex is a current-percept condition-action architecture.

D - Correct: A search frontier is the collection of generated candidates awaiting selection, not an agent decision architecture.

Primary note proof

Study-note section AI42

The common teaching list contains five agent classes when learning agents are included, or four decision architectures before learning is added. State the convention when counting agent types. “Target-based” can mean goal-based; “unity-based” is not utility-based.

Study-note section AI12

A search node records a state plus parent, action, depth and path cost. Expanding a node generates its successors. OPEN or frontier stores discovered candidates awaiting expansion; CLOSED records expanded states. Reaching the same state through different paths does not imply the same node or cost.

Study-note section AI06

A simple reflex agent matches condition-action rules against the current percept only: if dirty, clean.

Study-note section AI06

A model-based reflex agent maintains internal state using the previous state, previous action, current percept, and models of world change and action effects.

Study-note section AI07

A goal-based agent considers future action sequences that reach a desired state. A utility-based agent compares degrees of desirability and trade-offs, choosing maximum expected utility: $EU(a)=\Sigma s\ P(s|a,evidence)U(s)$. Reaching the destination is a goal; balancing arrival time, safety and cost requires preferences among outcomes.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 14 | lines 9-11

2.9 Agent Types
Agents can be grouped into five classes based on their degree of perceived intelligence and capability.
All these agents can improve their performance and generate better action over the time.

Introduction to Artificial Intelligence module | PDF page 15 | lines 2-3

2.9.3 Goal-based agent
• The knowledge of the current state environment is not always sufficient to decide for an agent to

Introduction to Artificial Intelligence module | PDF page 15 | lines 12-13

2.9.4 Utility-based agent
• These agents are similar to the goal-based agent but provide an extra component of utility

Question sources

Artificial Intelligence practice questions · Question 43

Question 92 | Q01050

Which is used to improve the agents performance?

A. Perceiving

B. Observing

C. Learning

D. Sequence

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer C: Learning

The learning element makes improvements from the environment; perception and observation supply data but are not the improvement mechanism. Evidence: PDF 15.

Why each choice is right or wrong

A - Incorrect: Supplies input without necessarily changing policy. See PDF p. 15.

B - Incorrect: Observation alone need not improve performance. See PDF p. 15.

C - Correct: Named improvement mechanism. See PDF p. 15.

D - Incorrect: A history is data, not the learning operation. See PDF p. 15.

Primary note proof

Study-note section AI08

> The performance element selects actions. The learning element changes the agent to improve future performance. The critic evaluates behavior against a performance standard and supplies feedback. The problem generator suggests exploratory actions likely to yield useful experience.

Study-note section AI08

> Learning is improvement on task T, measured by P, through experience E. For spam filtering, T is classifying email, E is past labeled messages, and P might be recall at a required precision. Exploration trades immediate reward for information; exploitation uses the best action currently known. Learning can improve reflex, goal-based or utility-based agents.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 15 | lines 28-30

> - Learning element is for making improvements by learning from environment
> - Critic Learning element takes feedback from critic which describes that how well the agent is doing with respect to a fixed performance standard.

Question sources

Artificial Intelligence practice questions · Question 44

Question 93 | Q01051

Rationality is?

A. being reasonable

B. being sensible

C. having good sense of judgment.

D. All of the above

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer D: All of the above

The module defines rationality with all three phrases: reasonable, sensible, and good judgment. Evidence: PDF 11.

Why each choice is right or wrong

A - Incorrect: Part of the definition. See PDF p. 11.

B - Incorrect: Part of the definition. See PDF p. 11.

C - Incorrect: Part of the definition. See PDF p. 11.

D - Correct: Includes every listed phrase. See PDF p. 11.

Primary note proof

Study-note section AI-S03

> In ordinary usage, rationality means reasonable, sensible judgment. For an AI agent, the operational criterion is the action with the best expected performance given its percept history and available knowledge, rather than guaranteed success or knowledge of the actual future.

Study-note section AI04

> A performance measure defines success externally: safety, useful deliveries, elapsed time, or cost. Rationality depends on that measure, prior knowledge, percept history, and available actions. A rational agent selects the action with the best expected performance given its information.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 11 | lines 27-29

• Rationalityis nothing but status of being reasonable, sensible, and having good sense of judgment.
• Rationality is concerned with expected actions and results depending upon what the agent has perceived.

Question sources

Artificial Intelligence practice questions · Question 45

Question 94 | Q01052

Rationality of an agent does not depends on?

A. performance measures

B. Percept Sequence

C. reaction

D. actions

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer C: reaction

The explicit determinants are performance measures, percept history, environmental knowledge, and available actions. Reaction is not a separate determinant in that list. Evidence: PDF 12, 12.

Why each choice is right or wrong

A - Incorrect: Explicit determinant. See PDF p. 12, 12.

B - Incorrect: Explicit determinant. See PDF p. 12, 12.

C - Correct: Not a separate listed factor; may colloquially refer to action. See PDF p. 12, 12.

D - Incorrect: Available actions explicitly matter. See PDF p. 12, 12.

Primary note proof

Study-note section AI04

A performance measure defines success externally: safety, useful deliveries, elapsed time, or cost. Rationality depends on that measure, prior knowledge, percept history, and available actions. A rational agent selects the action with the best expected performance given its information.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 12 | lines 10-12

• Rationality of an agent depends on the following –
o The performance measures, which determine the degree of success.
o Agent‘s Percept Sequence till now.

Introduction to Artificial Intelligence module | PDF page 12 | lines 11-14

o The performance measures, which determine the degree of success.
o Agent‘s Percept Sequence till now.
o The agent‘s posterior knowledge about the environment.
o The actions that the agent can carry out.

Question sources

Artificial Intelligence practice questions · Question 46

Question 95 | Q01053

What is the full form of PEAS?

A. Perceptual Measure, Environment, Actuators, and Sensors

B. Performance Measure, Environment, Actuators, and Sensors

C. Performance Measure, Entity, Actuators, and Sensors

D. Performance Measure, Environment, Agent Function, and Sensors

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer B: Performance Measure, Environment, Actuators, and Sensors

PEAS comprises Performance measure, Environment, Actuators, Sensors, spanning the page boundary. Evidence: PDF 13, 14.

Why each choice is right or wrong

A - Incorrect: Perceptual replaces Performance incorrectly. See PDF p. 13, 14.

B - Correct: All four components match. See PDF p. 13, 14.

C - Incorrect: Entity replaces Environment. See PDF p. 13, 14.

D - Incorrect: Agent Function replaces Actuators. See PDF p. 13, 14.

Primary note proof

Study-note section AI05

PEAS means Performance measure, Environment, Actuators, Sensors. Specify all four before designing an agent. For a taxi: performance includes safety and travel time; environment includes roads, traffic and passengers; actuators include steering, brakes and displays; sensors include cameras, positioning and speed readings.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 13 | lines 30-32

> 2.8 PEAS examples
> • The are four elements to take account when designing an agent to solve a particular problem
> • PEAS

Introduction to Artificial Intelligence module | PDF page 14 | lines 2-3

> o Actuator: agent acts with in its environment
> o Sensor: the agent senses its environment

Question sources

Artificial Intelligence practice questions · Question 47

Question 96 | Q01054

Chess is example of which properties?

A. Discrete

B. Continuous

C. Episodic

D. Non-deterministic

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer A: Discrete

Chess is the explicit discrete-environment example: finitely many distinct states/actions. Evidence: PDF 13, 13.

Why each choice is right or wrong

A - Correct: Explicit chess classification. See PDF p. 13, 13.

B - Incorrect: Driving, not chess, is the continuous example. See PDF p. 13, 13.

C - Incorrect: Moves affect future positions rather than independent episodes. See PDF p. 13, 13.

D - Incorrect: Board transitions are determined by state and selected move. See PDF p. 13, 13.

Primary note proof

Study-note section AI09

> Episodic decisions are independent; sequential decisions affect later situations. Static worlds do not change while deliberating; dynamic ones do. A clock affecting score makes an otherwise static game semidynamic. Discrete quantities have distinct alternatives; continuous quantities

vary along real-valued ranges. Known versus unknown concerns knowledge of the transition model, not observability.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 13 | lines 11-13

✓ Discrete / Continuous − If there are a limited number of distinct, clearly defined, states of the environment, the environment is discrete (For example, chess); otherwise it is continuous (For example, driving).

Introduction to Artificial Intelligence module | PDF page 13 | lines 23-28

✓ Deterministic / Non-deterministic − If the next state of the environment is completely determined
by the current state and the actions of the agent, then the environment is deterministic; otherwise it
is non-deterministic.
✓ Episodic / Non-episodic − In an episodic environment, each episode consists of the agent perceiving
and then acting. The quality of its action depends just on the episode itself. Subsequent episodes do
not depend on the actions in the previous episodes. Episodic environments are much simpler because

Question sources

Artificial Intelligence practice questions · Question 48

Question 97 | Q01055

For an untimed crossword puzzle whose clues and grid do not change while the solver deliberates, which environment property follows from that lack of change?

A. Dynamic

B. Static

C. Semi Dynamic

D. Continuous

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer B: Static

A static environment remains unchanged while the agent deliberates. The untimed, fixed-grid puzzle meets that condition. A running clock that affects performance could make an otherwise fixed task semidynamic, but the stem excludes it. Dynamic change and continuous-valued state or action descriptions are separate properties and do not follow from the stated puzzle setting.

Why each choice is right or wrong

A - Incorrect: Dynamic means the environment can change during deliberation, excluded by the premise.

B - Correct: Static means the relevant environment does not change while the agent deliberates.

C - Incorrect: Semidynamic describes changing performance, such as a running clock, despite an otherwise fixed environment; none is present.

D - Incorrect: Continuity concerns real-valued states or actions, not whether the puzzle changes during deliberation.

Primary note proof

Study-note section AI09

> Episodic decisions are independent; sequential decisions affect later situations. Static worlds do not change while deliberating; dynamic ones do. A clock affecting score makes an otherwise static game semidynamic. Discrete quantities have distinct alternatives; continuous quantities vary along real-valued ranges. Known versus unknown concerns knowledge of the transition model, not observability.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 13 | lines 17-19

> ✓ Static / Dynamic – If the environment does not change while an agent is acting, then it is static; otherwise it is dynamic.
> ✓ Single agent / Multiple agents – The environment may contain other agents which may be of the

Question sources

Artificial Intelligence practice questions · Question 49

Question 98 | Q01056

Which learning regime improves an agent's action choices through interaction with an environment that supplies rewards or penalties?

A. Supervised learning

B. Reinforcement learning

C. Semi-supervised learning

D. Unsupervised learning

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer B: Reinforcement learning

Interaction alone is too broad to identify a learning regime. Here the environment supplies rewards or penalties for the agent's actions, and the agent improves its choices from that feedback. This defines reinforcement learning. Supervised and semi-supervised methods depend on target labels; unsupervised methods discover unlabeled structure. None of those label-based descriptions specifies this action-and-reward loop.

Why each choice is right or wrong

A - Incorrect: Supervised learning supplies correct target outputs rather than defining action-reward interaction.

B - Correct: Reinforcement learning improves decisions through environmental reward feedback.

C - Incorrect: Semi-supervised learning mixes labeled and unlabeled examples, not a defining action-reward objective.

D - Incorrect: Unsupervised learning seeks unlabeled structure without the stated reward-feedback objective.

Primary note proof

Study-note section AI33

> An RL agent observes state information, takes an action, receives reward and a new observation, and improves its policy to maximize expected cumulative reward. Rewards evaluate outcomes rather than supply the correct action label. Rewards may be delayed; not every correct-looking action receives immediate reward.

Study-note section AI32

> Semi-supervised learning combines some labels with additional unlabeled data. Active learning chooses informative cases for an oracle to label; it still uses labels. "Exploratory learning" is not an exclusive standard synonym: it can mean unsupervised exploration or active data acquisition, so define the feedback. A task named summarization or time-series analysis is not automatically supervised; its training targets determine that.

Study-note section AI31

> Supervised learning fits a mapping from labeled input-target examples.

Study-note section AI32

> Unsupervised learning seeks structure without target labels: clustering groups similar examples; association finds co-occurrence; dimensionality reduction compresses features.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 69 | lines 15-17

> ⇒ Reinforcement learning is a feedback-based learning method, in which a learning agent gets a reward for each right action and gets a penalty for each wrong action.
> The agent learns automatically with these feedbacks and improves its performance. In reinforcement

Question sources

Artificial Intelligence practice questions · Question 50

Question 99 | Q01057

Which is reinforcement learning's immediate parent field in the AI → machine learning → reinforcement learning hierarchy?

A. Artifical Intelligence

B. Data Analytics

C. Machine Learning

D. None of the above

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer C: Machine Learning

Reinforcement learning is a branch of machine learning, and machine learning is within artificial intelligence. Artificial intelligence therefore also contains reinforcement learning, but it is the broader ancestor rather than the immediate parent requested. Data analytics is an application context rather than the stated parent field. The explicit hierarchy makes the nearer containing field the unique taxonomic selection.

Why each choice is right or wrong

A - Incorrect: AI contains RL indirectly through machine learning, but is not the immediate parent requested.

B - Incorrect: Data analytics is not the stated parent-field classification of RL in the note.

C - Correct: Machine learning is the immediate parent of reinforcement learning in the stated hierarchy.

D - Incorrect: The immediate parent is present among the choices.

Primary note proof

Study-note section AI33

An RL agent observes state information, takes an action, receives reward and a new observation, and improves its policy to maximize expected cumulative reward. Rewards evaluate outcomes rather than supply the correct action label. Rewards may be delayed; not every correct-looking action receives immediate reward.

Study-note section AI43

ANNs, genetic search and reinforcement learning all improve candidate behavior or performance through experience in different ways. Deep learning is an architecture family usable with supervised, unsupervised or reinforcement objectives. Decision trees and deep models may both be supervised. To distinguish learning regimes, specify labels, reward or unlabeled structure.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 68 | lines 16-18

6.2 Types of learning (machine learning)
• Machine learning enables a machine to automaticallylearn from data, improve performance from
experiences, and predict things without being explicitly programmed.

Introduction to Artificial Intelligence module | PDF page 67 | lines 18-20

important feature of —intelligent‖. Machine learning is the subfield of AI concerned with intelligent
systems that learn from experience and examples. It is the computational study of algorithms that
improve performance based on experience

Question sources

Artificial Intelligence practice questions · Question 51

Question 100 | Q01058

Which type of feedback does an agent in reinforcement learning receive?

A. Predictions

B. Labels

C. Clusters

D. Rewards or penalties

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer D: Rewards or penalties

Rewards and penalties are the stated reinforcement feedback; supervised targets and unsupervised grouping differ. Evidence: PDF 69.

Why each choice is right or wrong

A - Incorrect: A prediction is output, not the reward signal. See PDF p. 69.

B - Incorrect: Labels belong to supervised examples. See PDF p. 69.

C - Incorrect: Clusters are an unsupervised result. See PDF p. 69.

D - Correct: Matches the explicit feedback description. See PDF p. 69.

Primary note proof

Study-note section AI33

> An RL agent observes state information, takes an action, receives reward and a new observation, and improves its policy to maximize expected cumulative reward. Rewards evaluate outcomes rather than supply the correct action label. Rewards may be delayed; not every correct-looking action receives immediate reward.

Study-note section AI31

> Supervised learning fits a mapping from labeled input-target examples. Classification predicts discrete labels, such as institution generations or spam/nonspam; regression predicts numeric quantities, such as prices. Predictive analytics uses historical patterns to forecast outcomes and can use either. Specify the target before choosing a method.

Study-note section AI32

> Unsupervised learning seeks structure without target labels: clustering groups similar examples; association finds co-occurrence; dimensionality reduction compresses features. Good clusters have high within-cluster similarity and low between-cluster similarity under an appropriate distance. K-means alternates assignment to nearest centroids and centroid updates.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 69 | lines 16-18

> reward for each right action and gets a penalty for each wrong action.
> The agent learns automatically with these feedbacks and improves its performance. In reinforcement
> learning, the agent interacts with the environment and explores it. The goal of an agent is to get the

Question sources

Artificial Intelligence practice questions · Question 52

Question 101 | Q01059

What is the main task of a problem-solving agent?

A. Solve the given problem and reach to goal

B. To find out which sequence of action will get it to the goal state

C. Both A and B

D. None of the Above

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer C: Both A and B

A solution is an action sequence leading from start to goal; this supports both general goal achievement and finding that sequence. Evidence: PDF 18.

Why each choice is right or wrong

A - Incorrect: Describes the overall objective. See PDF p. 18.

B - Incorrect: Gives the operational search task. See PDF p. 18.

C - Correct: Both are supported by the solution definition. See PDF p. 18.

D - Incorrect: Contradicts the defined objective. See PDF p. 18.

Primary note proof

Study-note section AI11

> First formulate the goal; then formulate the problem by specifying initial state, legal actions, transition model, goal test, and action costs. Search selects a sequence; execution applies it. A state abstracts away irrelevant world details.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 18 | lines 19-21

> Solution: It is an action sequence which leads from the start node to the goal node.
> Optimal Solution: If a solution has the lowest cost among all solutions.
> 3.5 Search Strategies

Question sources

Artificial Intelligence practice questions · Question 53

Question 102 | Q01060

For DFS on an implicit tree with branching factor b and maximum depth m, storing pending siblings, which expression gives the standard worst-case space bound?

A. b

B. b^2

C. b^b

D. O(bm)

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer D: O(bm)

Depth-first search retains the active branch and pending alternatives rather than an entire breadth frontier. With up to b pending entries per level over depth m, the standard stored-frontier space bound is O(bm). Path-only on-demand generation can use O(m), but that is a different implementation convention. The expression gives the depth-dependent storage bound rather than an exponential time bound.

Why each choice is right or wrong

A - Incorrect: b omits growth with search depth m in the general DFS storage bound.

B - Incorrect: b^2 is not the stated depth-dependent worst-case DFS storage expression.

C - Incorrect: b^b is neither of the note's DFS space bounds.

D - Correct: With pending siblings, at most a branching-factor-sized set per depth gives the standard O(bm) space bound.

Primary note proof

Study-note section AI14

> DFS follows the deepest pending branch using a LIFO stack or recursion. FILO means the same order as LIFO. With the same children as BFS, a left-first DFS visits S,A,C before B; push rightmost children first to pop the leftmost first.

Study-note section AI14

> Implicit-tree worst-case time is O(b^m), and space is O(bm) when pending siblings are stored, or O(m) for path-only recursive generation. DFS is not generally optimal and can fail on infinite paths or cycles. On a finite graph, visited-state tracking makes traversal complete but needs O(V) storage. Do not confuse the O(b^m) worst-case time bound with the DFS space bounds above.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 19 | lines 20-22

• Space Complexity: DFS algorithm needs to store only single path from the root node, hence space
complexity of DFS is equivalent to the size of the fringe set, which is O(bm).
• Optimal: DFS search algorithm is non-optimal, as it may generate a large number of steps or high

Question sources

Artificial Intelligence practice questions · Question 54

Question 103 | Q01061

For BFS on an implicit tree with branching factor b>1 and shallowest goal depth d, using generation-time goal testing, which expression gives the standard worst-case generated-node order?

A. b

B. b^d

C. b^2

D. b^b

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer B: b^d

BFS generates levels in order. Under generation-time goal testing, the generated nodes through the shallowest goal depth d are bounded by the geometric sum 1+b+...+b^d, whose dominant term is b^d when b>1. Testing only after expanding deeper nodes can change the conventional bound by another factor of b, so the question states the goal-test convention.

Why each choice is right or wrong

A - Incorrect: b counts one level and does not account for arbitrary goal depth.

B - Correct: The sum 1+b+...+b^d has order b^d for b>1.

C - Incorrect: b^2 only describes a fixed depth-two case rather than general depth d.

D - Incorrect: b^b wrongly replaces solution depth d with the branching factor.

Primary note proof

Study-note section AI13

The tree contains 1+b+...+b^d nodes through depth d, giving O(b^d) under generation-time goal testing. Expanding depth-d nodes before testing can generate depth d+1 and yields O(b^(d+1)); specify the convention. Explicit finite graph traversal with visited-state tracking takes O(V+E) time and O(V) space.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 20 | lines 20-22

> Time Complexity: is obtained by the number of nodes traversed in BFS until the shallowest Node.
> T
> (b) =O(bd
>) where d= depth of shallowest solution and b is a node at every state.

Question sources

Artificial Intelligence practice questions · Question 55

Question 104 | Q01062

Which labels describe the recursive DFS stack’s removal order? Select all that apply.

A. LIFO

B. LILO

C. FIFO

D. FILO

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer A, D: LIFO; FILO

Recursive depth-first search returns first from the most recently entered unfinished call. That is last-in-first-out stack order, also called first-in-last-out. The two labels are synonyms rather than distinct data structures. FIFO or LILO describes the opposite end-to-end order and does not characterize the DFS call stack. The multiple-selection instruction therefore retains both correct names for the same discipline.

Why each choice is right or wrong

A - Correct: LIFO is the stack order used by recursion and depth-first exploration.

B - Incorrect: LILO does not express the required last-in-first-out removal order.

C - Incorrect: FIFO is queue order, associated with breadth-first rather than depth-first selection.

D - Correct: FILO is another name for LIFO and therefore also describes the DFS stack.

Primary note proof

Study-note section AI14

> DFS follows the deepest pending branch using a LIFO stack or recursion. FILO means the same order as LIFO. With the same children as BFS, a left-first DFS visits S,A,C before B; push rightmost children first to pop the leftmost first.

Study-note section AI13

BFS removes the oldest frontier entry using a FIFO queue. Start with the root, remove the head, test it, append unseen successors at the tail, and repeat. It explores depths 0,1,2,... and is complete for finite b and a finite-depth goal. It minimizes number of edges, so minimum cost requires equal positive step costs.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 18 | lines 34-35

- DFS uses a stack data structure for its implementation.
- The process of the DFS algorithm is similar to the BFS algorithm.

Question sources

Artificial Intelligence practice questions · Question 56

Question 105 | Q01063

Which of the following search algorithm searches forward from initial state and backward from goal state till both meet to identify a common state?

A. Uniform Cost Search

B. Iterative Deepening Depth-First Search

C. Bidirectional Search

D. None of the Above

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer C: Bidirectional Search

The algorithm starts one search at the initial state and another at the goal and stops at their intersection. Evidence: PDF 23.

Why each choice is right or wrong

A - Incorrect: Orders by path cost rather than two directions. See PDF p. 23.

B - Incorrect: Repeats depth-limited search from the start. See PDF p. 23.

C - Correct: Exactly the two-front definition. See PDF p. 23.

D - Incorrect: C explicitly fits. See PDF p. 23.

Primary note proof

Study-note section AI16

Bidirectional search grows from start and goal until frontiers connect. With suitable reverse actions, meeting tests and equal-cost BFS, time and space may fall to O(b^(d/2)). Weighted variants require a sound stopping bound; the first intersection need not be cheapest. IDA* deepens an f=g+h threshold rather than only a depth limit.

Study-note section AI15

UCS uses a priority queue ordered by actual accumulated path cost g(n). Remove a minimum-g node; accept a goal when removed, not when first generated. Relax each edge: if a cheaper path is found, update its frontier cost and parent.

Study-note section AI16

DLS runs DFS to depth l and reports cutoff separately from failure. It misses goals deeper than l. Iterative deepening runs DLS with limits 0,1,2,...; for finite b it finds the shallowest solution with O(b^d) time and O(bd) frontier space, optimal for equal step costs. Repeated shallow work is small compared with the deepest level.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 23 | lines 16-18

• Bidirectional search algorithm runs two simultaneous searches, one form initial state called as forward-search and other from goal node called as backward-search, to find the goal node.
• Bidirectional search replaces one single search graph with two small sub-graphs in which one starts

Question sources

Artificial Intelligence practice questions · Question 57

Question 106 | Q01064

When is breadth-first search is optimal?

A. When there is less number of nodes

B. When all step costs are equal

C. When all step costs are unequal

D. All of the above

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer B: When all step costs are equal

Equal positive step costs make total path cost proportional to depth; BFS first reaches a shallowest solution and hence a cheapest one. Evidence: PDF 20.

Why each choice is right or wrong

A - Incorrect: Node count alone says nothing about cost. See PDF p. 20.

B - Correct: Equal step costs ensure depth order is cost order. See PDF p. 20.

C - Incorrect: Unequal costs can make a deeper path cheaper. See PDF p. 20.

D - Incorrect: A and C are not sufficient conditions. See PDF p. 20.

Primary note proof

Study-note section AI13

> BFS removes the oldest frontier entry using a FIFO queue. Start with the root, remove the head, test it, append unseen successors at the tail, and repeat. It explores depths 0,1,2,... and is complete for finite b and a finite-depth goal. It minimizes number of edges, so minimum cost requires equal positive step costs.

Study-note section AI13

> The tree contains 1+b+...+b^d nodes through depth d, giving O(b^d) under generation-time goal testing. Expanding depth-d nodes before testing can generate depth d+1 and yields O(b^(d+1)); specify the convention. Explicit finite graph traversal with visited-state tracking takes O(V+E) time and O(V) space.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 20 | lines 28-30

> Optimality: BFS is optimal if path cost is a non-decreasing function of the depth of the node.
> 3. Depth limited search
> A depth-limited search algorithm is similar to depth-first search with a predetermined limit.
> Depth-

Question sources

Artificial Intelligence practice questions · Question 58

Question 107 | Q01065

What is the heuristic function of greedy best-first search?

A. f(n) != h(n)

B. f(n) < h(n)

C. f(n) = h(n)

D. f(n) > h(n)

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer C: f(n) = h(n)

The operational rule removes the node with lowest h(n), so the greedy evaluation is f(n)=h(n). The printed f(n)=g(n) on this same page conflicts with its own algorithm and nearby h definition; the operational evidence supports C. Evidence: PDF 25.

Why each choice is right or wrong

A - Incorrect: Inequality would not express that rule. See PDF p. 25.

B - Incorrect: Strictly smaller is not the definition. See PDF p. 25.

C - Correct: Matches ranking exclusively by h. See PDF p. 25.

D - Incorrect: Strictly greater is not the definition. See PDF p. 25.

Primary note proof

Study-note section AI17

> Best-first is a family that orders OPEN by an evaluation function. Greedy best-first uses f=h and removes the smallest estimate, ignoring the cost already paid. It can choose an expensive route, loop without duplicate checks, or follow an infinite misleading branch. On a finite graph with full duplicate handling it is complete, but not generally optimal. Its implicit finite-tree worst case can be O(b^m).

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 25 | lines 21-23

> o Step 3: Remove the node n, from the OPEN list which has the lowest value of h(n), and places it in the CLOSED list.
> o Step 4: Expand the node n, and generate the successors of node n.

Question sources

Artificial Intelligence practice questions · Question 59

Question 108 | Q01066

Greedy best-first search selects which descriptions of the next node? Interpret “closest” as estimated remaining cost h, not actual distance. Select all that apply.

A. Shallowest

B. Deepest

C. The one closest to the goal node

D. Minimum heuristic cost

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer C, D: The one closest to the goal node; Minimum heuristic cost

Greedy best-first search orders the frontier by h alone, choosing the smallest estimated remaining cost and ignoring cost already paid. Under the explicit interpretation, estimated closest to the goal and minimum heuristic cost express that same rule. Shallowest-first uses depth as in breadth-first search, while deepest-first follows depth-first ordering. Both heuristic descriptions are therefore valid and neither depth description is selected.

Why each choice is right or wrong

A - Incorrect: Shallowest-first is BFS-style depth ordering rather than greedy heuristic selection.

B - Incorrect: Deepest-first is DFS-style selection rather than minimum estimated remaining cost.

C - Correct: Read closest as estimated by h; it is the same greedy criterion as minimum heuristic cost.

D - Correct: Greedy best-first selects a node with minimum h, ignoring already-paid path cost.

Primary note proof

Study-note section AI17

> Best-first is a family that orders OPEN by an evaluation function. Greedy best-first uses f=h and removes the smallest estimate, ignoring the cost already paid. It can choose an expensive route, loop without duplicate checks, or follow an infinite misleading branch. On a finite graph with full duplicate handling it is complete, but not generally optimal. Its implicit finite-tree worst case can be O(b^m).

Study-note section AI13

> BFS removes the oldest frontier entry using a FIFO queue. Start with the root, remove the head, test it, append unseen successors at the tail, and repeat. It explores depths 0,1,2,... and is complete for finite b and a finite-depth goal. It minimizes number of edges, so minimum cost requires equal positive step costs.

Study-note section AI14

DFS follows the deepest pending branch using a LIFO stack or recursion. FILO means the same order as LIFO. With the same children as BFS, a left-first DFS visits S,A,C before B; push rightmost children first to pop the leftmost first.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 25 | lines 14-16

• In the best first search algorithm, we expand the node which is closest to the goal node and the closest
cost is estimated by heuristic function-
• f(n)= g(n). Were, h(n)= estimated cost from node n to the goal.

Introduction to Artificial Intelligence module | PDF page 25 | lines 21-22

o Step 3: Remove the node n, from the OPEN list which has the lowest value of h(n), and places it in the CLOSED list.

Question sources

Artificial Intelligence practice questions · Question 60

Question 109 | Q01067

Which general search family contains A* because it selects frontier nodes by a minimum evaluation value?

A. Breadth-First-Search

B. Depth-First-Search

C. Uniform Cost Search

D. Best-First-Search

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer D: Best-First-Search

Best-first search is the general family that selects a frontier node according to an evaluation value. A* is one member, using f=g+h. Uniform-cost search is another member, using g alone. Breadth-first and depth-first names describe depth-oriented selection disciplines. The question asks for the containing family rather than claiming A* is identical to one simpler algorithm.

Why each choice is right or wrong

A - Incorrect: BFS selects by depth/FIFO order rather than the general evaluation-priority family.

B - Incorrect: DFS follows stack/deepest-branch order rather than minimum evaluation priority.

C - Incorrect: UCS is another best-first method using g alone; it is not the general family containing A*.

D - Correct: Best-first search is the family that selects by an evaluation, with A* using g+h.

Primary note proof

Study-note section AI18

A* orders by f(n)=g(n)+h(n). Admissibility means 0≤h(n)≤h*(n), the true cheapest remaining cost. Consistency means h(n)≤c(n,n′)+h(n′) on every edge, with h(goal)=0; it implies admissibility on states with goal paths. Consistency makes f nondecreasing along a path.

Study-note section AI15

UCS uses a priority queue ordered by actual accumulated path cost g(n). Remove a minimum-g node; accept a goal when removed, not when first generated. Relax each edge: if a cheaper path is found, update its frontier cost and parent.

Study-note section AI17

Best-first is a family that orders OPEN by an evaluation function. Greedy best-first uses f=h and removes the smallest estimate, ignoring the cost already paid. It can choose an expensive route, loop without duplicate checks, or follow an infinite misleading branch. On a finite graph with full duplicate handling it is complete, but not generally optimal. Its implicit finite-tree worst case can be O(b^m).

Study-note section AI13

BFS removes the oldest frontier entry using a FIFO queue.

Study-note section AI14

DFS follows the deepest pending branch using a LIFO stack or recursion.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 26 | lines 25-27

• A* search is the most commonly known form of best-first search.
• It uses heuristic function h(n), and cost to reach the node n from the start state g(n). It has combined
features of UCS and greedy best-first search, by which it solves the problem efficiently.

Question sources

Artificial Intelligence practice questions · Question 61

Question 110 | Q01188

Function MY_AGENT (PERCEPT) return action
static state, a description of the current world state
goal, a description of the goal to achieve may be in terms of state
state ← UPDATE_STATE(state, percept)
actionSet ← POSSIBLE_ACTIONS(state)
action ← BEST_ACTION(actionSet)
state ← UPDATE_STATE(state, action)
return action
What is the purpose of the final instruction state ← UPDATE_STATE(state, action), after BEST_ACTION has selected an action?

A. Modify the state of the environment if necessary

B. Update the state information to be like the one after executing the best action chosen

C. Update the environment so that the agent can perceive the next percept

D. Modify the state of the environment once the decision of which action to choose is made.

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer B: Update the state information to be like the one after executing the best action chosen

The final update uses the chosen action to predict what the agent's internal state representation should look like after that action. It is distinct from the earlier update using the current percept. Updating memory does not execute an actuator or directly modify the external environment. Naming the exact instruction resolves the missing referent while preserving the model-update interpretation of the original prototype.

Why each choice is right or wrong

A - Incorrect: Updating an internal state representation does not itself modify the external environment.

B - Correct: The final call predicts and records the internal state after the selected action.

C - Incorrect: Neither internal update executes an environmental action or produces the next percept.

D - Incorrect: Selecting an action and updating memory are distinct from changing the external world through execution.

Primary note proof

Study-note section AI07

Updating an internal state after selecting an action predicts the action's effect; it does not itself modify the external environment. A stored variable named goal is insufficient to prove a goal-based policy if BEST_ACTION never uses it. Model-based, goal-based, utility-based and learning capabilities can coexist.

Study-note section AI06

A model-based reflex agent maintains internal state using the previous state, previous action, current percept, and models of world change and action effects. It then applies rules to that state. Both kinds can use condition-action rules; memory and a world model distinguish them.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 14 | lines 32-34

• Updating the agent state requires information about:
- How the world evolves
- How the agent's action affects the world.

Question sources

Computer Science model exam 2 · Question 9

Question 111 | Q01194

To implement greedy best-first search, let the evaluation be the estimated remaining cost h(n). Which OPEN-list insertion policy supplies its node-selection order?

A. Append newly generated nodes at the beginning of the open list

B. Append newly generated nodes at the end of the open list

C. Insert newly generated nodes in the open list according to their path cost (lowest values first)

D. Insert newly generated nodes in the open list according to their evaluation function (lowest values first)

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer D: Insert newly generated nodes in the open list according to their evaluation function (lowest values first)

Greedy best-first search orders its frontier by h alone, the estimate of remaining cost. A priority-ordered OPEN list with the smallest evaluation first therefore implements the rule when that evaluation is h. Ordering by g would instead implement uniform-cost search. Unconditionally prepending or appending nodes gives stack-like or FIFO behavior without the required heuristic ranking.

Why each choice is right or wrong

A - Incorrect: Prepending every new node gives a stack-like policy rather than ordering by remaining-cost estimates.

B - Incorrect: Appending every new node gives FIFO breadth order rather than heuristic priority.

C - Incorrect: Ordering by accumulated path cost g implements uniform-cost selection, not greedy h-only selection.

D - Correct: Ordering by the stated evaluation h puts the smallest estimated remaining cost first.

Primary note proof

Study-note section AI17

> Best-first is a family that orders OPEN by an evaluation function. Greedy best-first uses f=h and removes the smallest estimate, ignoring the cost already paid. It can choose an expensive route, loop without duplicate checks, or follow an infinite misleading branch. On a finite graph with full duplicate handling it is complete, but not generally optimal. Its implicit finite-tree worst case can be O(b^m).

Study-note section AI15

> UCS uses a priority queue ordered by actual accumulated path cost g(n). Remove a minimum-g node; accept a goal when removed, not when first generated. Relax each edge: if a cheaper path is found, update its frontier cost and parent.

Study-note section AI13

> BFS removes the oldest frontier entry using a FIFO queue. Start with the root, remove the head, test it, append unseen successors at the tail, and repeat. It explores depths 0,1,2,... and is complete for finite b and a finite-depth goal. It minimizes number of edges, so minimum cost requires equal positive step costs.

Study-note section AI14

> DFS follows the deepest pending branch using a LIFO stack or recursion. FILO means the same order as LIFO. With the same children as BFS, a left-first DFS visits S,A,C before B; push rightmost children first to pop the leftmost first.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 25 | lines 21-23

> o Step 3: Remove the node n, from the OPEN list which has the lowest value of h(n), and places it in the CLOSED list.
> o Step 4: Expand the node n, and generate the successors of node n.

Question sources

Computer Science model exam 2 · Question 15

Question 112 | Q01199

In classical propositional logic over a nonempty set of interpretations, which statements are false? In D, “models” means all interpretations. Select all that apply.

A. A sentence is said to be invalid iff it is unsatisfiable.

B. A sentence is valid if it is satisfiable.

C. A sentence which is not invalid is satisfiable.

D. A sentence is unsatisfiable if it is false in all models.

Hint

Translate the statement carefully, then check the scope of its quantifiers, connectives or inference rule. Distinguish guaranteed consequences from plausible generalizations or explanations.

Answer A, B: A sentence is said to be invalid iff it is unsatisfiable.; A sentence is valid if it is satisfiable.

Validity requires truth in every interpretation, while satisfiability requires truth in at least one. A contingent atom can therefore be satisfiable and invalid, refuting both the proposed invalid-unsatisfiable equivalence and the claim that satisfiability implies validity. Not invalid means valid and hence satisfiable in the stated nonempty setting. Unsatisfiable means false in every interpretation, so the remaining two statements are true.

Why each choice is right or wrong

A - Correct: False: a satisfiable contingent sentence can be invalid without being unsatisfiable.

B - Correct: False: truth in one interpretation does not imply truth in every interpretation.

C - Incorrect: True: not invalid means valid, which entails satisfiability in the ordinary model setting.

D - Incorrect: True: unsatisfiable means there is no interpretation making the sentence true.

Primary note proof

Study-note section AI25

> A model is an interpretation making a sentence true. Valid means true in every interpretation; satisfiable means true in at least one; unsatisfiable means true in none. Invalid means not valid, and may still be satisfiable. P∨¬P is valid; P is satisfiable but not valid; P∧¬P is unsatisfiable. Entailment KB⊨α means every model of KB also satisfies α.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 45 | lines 27-29

> • A proposition formula which is always true is called tautology, and it is also called a valid sentence.
> • A proposition formula which is always false is called Contradiction.
> • A proposition formula which has both true and false values is called Contingency.

Question sources

Computer Science model exam 2 · Question 20

Question 113 | Q01202

Consider this agent prototype:
Function MY_AGENT(percept) returns action
static state, goal
state ← UPDATE_STATE(state, percept)
actionSet ← POSSIBLE_ACTIONS(state)
action ← BEST_ACTION(actionSet, state, goal)
state ← UPDATE_STATE(state, action)
return action
BEST_ACTION uses lookahead over future action sequences and chooses by whether the specified desired state is reached; it does not rank successful outcomes by utility or learn from experience. Which category best describes this action-selection criterion?

A. Goal based

B. Model based reflex

C. Simple reflex

D. Utility based

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer A: Goal based

The action selector explicitly consults the desired state and looks ahead over sequences that can reach it. That is the defining goal-based decision criterion. Persistent state also gives the agent a model-based capability, but memory alone does not identify its selection policy. A reflex policy would apply condition-action rules instead of goal-directed lookahead, and utility-based choice would rank degrees of desirability among outcomes.

Why each choice is right or wrong

A - Correct: Goal-directed lookahead chooses sequences reaching a desired state, matching goal-based action selection.

B - Incorrect: The agent maintains a model, but the stated policy uses goal-directed lookahead rather than reflex rules.

C - Incorrect: Persistent state and future-sequence evaluation contradict a current-percept-only reflex policy.

D - Incorrect: No utility ranking is used; the selector evaluates attainment of the desired state.

Primary note proof

Study-note section AI07

A goal-based agent considers future action sequences that reach a desired state. A utility-based agent compares degrees of desirability and trade-offs, choosing maximum expected utility: EU(a)=Σs P(s | a,evidence)U(s). Reaching the destination is a goal; balancing arrival time, safety and cost requires preferences among outcomes.

Study-note section AI07

Updating an internal state after selecting an action predicts the action's effect; it does not itself modify the external environment. A stored variable named goal is insufficient to prove a goal-based policy if BEST_ACTION never uses it. Model-based, goal-based, utility-based and learning capabilities can coexist.

Study-note section AI06

A simple reflex agent matches condition-action rules against the current percept only: if dirty, clean. It has no persistent world-state memory. This is adequate when the current observation contains everything needed for its rule, but partial observability can make the same percept require different actions.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 15 | lines 6-8

• Goal-based agents expand the capabilities of the model-based agent by having the "goal" information.
• They choose an action, so that they can achieve the goal.

Introduction to Artificial Intelligence module | PDF page 14 | lines 29-31

- Internal State: It is a representation of the current state based on percept history.
• These agents have the model, "which is knowledge of the world" and based on the model they perform actions.

Question sources

Computer Science model exam 2 · Question 23

Question 114 | Q01220

Which are general terms for the components through which an arbitrary AI agent acts on its environment, rather than sensing components or one particular robot subtype? Select all that apply.

A. Sensors

B. Effectors

C. Actuators

D. Manipulators

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer B, C: Effectors; Actuators

Actuators and effectors are equivalent general names for the agent's acting interface here. Sensors instead supply observations. A manipulator can be a particular physical acting device, but it is not the general term covering every agent, including software agents. The scope allows the two synonymous general terms without incorrectly denying that robotic manipulators can carry out actions.

Why each choice is right or wrong

A - Incorrect: Sensors receive observations rather than enact the agent's selected environmental changes.

B - Correct: Effectors are an accepted name for the components through which the agent acts.

C - Correct: Actuators are the note's equivalent general term for acting components.

D - Incorrect: A manipulator is a particular physical device, not the general acting-interface term requested.

Primary note proof

Study-note section AI03

> An agent perceives an environment through sensors and acts through actuators or effectors. Cameras sense; motors move. Software sensors may be incoming messages and software actions may be API requests. Humans are agents but are not artificial agents.

Study-note section AI42

> A generic rule does not distinguish simple from model-based reflex without specifying memory. A goal declaration unused by the action policy does not establish goal-based behavior. "Sensors and effectors" describes interaction, not omniscience or intelligence by itself. Define the mechanical distinction when comparing effectors and actuators as acting components.

Study-note section AI39

> Robotics combines mechanical, electrical, control and computing systems to sense, plan and act physically. Components include sensors, controller, power supply, actuators and an end effector. Locomotion may be legged, wheeled or tracked, with trade-offs in terrain, stability and efficiency. Software can act autonomously without being a robot. Humanoid shape is not evidence of general intelligence.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 10 | lines 25-27

- The agent can change the environment through actuators or effectors.
- An operation involving an effector is called an action.
- Actions can be grouped into action sequences.

Question sources

Computer Science model exam 2 · Question 41

Question 115 | Q01252

A control agent has incomplete observations, uncertain transition outcomes, actions that affect later decisions, a world that changes during deliberation, and real-valued measurements and controls. Which list matches all these properties?

A. Fully observable, discrete, episodic, static.

B. Static, episodic, continuous, partially observable.

C. Partially observable, non-deterministic, non-episodic, dynamic, continuous.

D. Deterministic, sequential (non-episodic), dynamic, continuous.

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer C: Partially observable, non-deterministic, non-episodic, dynamic, continuous.

The scenario establishes each environment dimension directly. Incomplete observations give partial observability; uncertain outcomes give non-determinism; effects on later decisions give sequentiality; ongoing change gives dynamics; and real-valued quantities give continuity. Only one list matches the whole combination. This classification avoids claiming that these dimensions establish a universal numerical ranking of task difficulty.

Why each choice is right or wrong

A - Incorrect: Full observability, episodic independence and a static discrete setting contradict the scenario.

B - Incorrect: Partial observability and continuity fit, but static and episodic do not.

C - Correct: Partial observability, non-determinism, sequentiality, dynamics and continuity all match.

D - Incorrect: The other stated properties fit, but determinism contradicts uncertain transition outcomes.

Primary note proof

Study-note section AI09

Fully observable means sensors reveal relevant state; partially observable means relevant information is hidden or noisy. Deterministic means state and action determine the next state;

stochastic means outcomes involve randomness. Strategic uncertainty about an opponent is distinct from dice-like randomness.

Study-note section AI09

Episodic decisions are independent; sequential decisions affect later situations. Static worlds do not change while deliberating; dynamic ones do. A clock affecting score makes an otherwise static game semidynamic. Discrete quantities have distinct alternatives; continuous quantities vary along real-valued ranges. Known versus unknown concerns knowledge of the transition model, not observability.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 16 | lines 17-19

✓ The most challenging environments are partially observable, stochastic, sequential, dynamic, and
continuous, and contain multiple intelligent agents.
2.11 Questions

Question sources

Computer Science model exam 2 · Question 73

Question 116 | Q01279

Which described training tasks are supervised learning? Select all that apply.

A. Summarization trained on document/reference-summary pairs.

B. Time-series forecasting trained on history/next-value target pairs.

C. Sequence discovery without target labels.

D. Clustering unlabeled examples.

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer A, B: Summarization trained on document/reference-summary pairs.; Time-series forecasting trained on history/next-value target pairs.

Supervised training uses paired inputs and target outputs. Document/reference-summary pairs therefore make the summarization task supervised, and history/next-value pairs make the forecasting task supervised. Sequence discovery and clustering without targets instead seek structure in unlabeled data. The labels in the descriptions determine the regime; task names alone would not establish whether summarization or time-series analysis is supervised.

Why each choice is right or wrong

A - Correct: The reference summary is a target output paired with each input document.

B - Correct: The next value is a numeric target paired with each input history.

C - Incorrect: Without target labels, sequence discovery seeks unlabeled structure.

D - Incorrect: Clustering unlabeled examples is unsupervised structure discovery.

Primary note proof

Study-note section AI31

> Supervised learning fits a mapping from labeled input-target examples. Classification predicts discrete labels, such as institution generations or spam/nonspam; regression predicts numeric quantities, such as prices. Predictive analytics uses historical patterns to forecast outcomes and can use either. Specify the target before choosing a method.

Study-note section AI32

> Unsupervised learning seeks structure without target labels: clustering groups similar examples; association finds co-occurrence; dimensionality reduction compresses features. Good clusters have high within-cluster similarity and low between-cluster similarity under an appropriate distance. K-means alternates assignment to nearest centroids and centroid updates.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 69 | lines 5-7

> o Classification
> o Regression
> ⇒ Unsupervised learning means we are only given the x‘s. In either case, the goal is to estimate f.

Question sources

Computer Science model exam 2 · Question 100

Question 117 | Q01485

A common name used to describe search-based optimization technique based on the principles of Genetics and Natural Selection is ____.

A. Genetic algorithms

B. Heuristic algorithm

C. Hill climbing

D. Simulated annealing

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer A: Genetic algorithms.

AI20 names the search method with population, fitness, parent selection, crossover and mutation as a genetic algorithm. That is the specific common name requested by the genetics/natural-selection description. Heuristic algorithm is a broad umbrella rather than this name; hill climbing follows improved neighbors, and simulated annealing is not the population-selection mechanism described.

Why each choice is right or wrong

A - Correct: This is the specific name of the evolutionary search technique based on fitness, selection and genetic variation.

B - Incorrect: Heuristic algorithm is a broad category and does not specifically name the genetics-based population method.

C - Incorrect: Hill climbing improves a candidate through neighbors rather than maintaining the described genetic population.

D - Incorrect: Simulated annealing is not the named selection/crossover/mutation method defined in the note.

Primary note proof

Study-note section AI20

> A genetic algorithm maintains a population, evaluates fitness, selects parents, applies crossover and mutation, and repeats replacement. It searches without guaranteeing global optimality. Means-ends analysis selects a difference between current and goal states, chooses an operator reducing it, and recursively solves its preconditions. Neither method is simply A* running backward.

Study-note section AI20

> Local search keeps a complete candidate configuration and improves its objective; the path may be irrelevant. Hill climbing chooses better neighbors. First-improvement accepts the first improvement; steepest-ascent examines all neighbors; stochastic variants choose among improvements. Local maxima, plateaus and ridges can stop progress. Random restarts or limited sideways steps can help without guaranteeing success.

Study-note section AI-S05

> Simulated annealing follows a current candidate and randomly samples a neighboring move. It accepts improvements and sometimes accepts worse moves according to a temperature that decreases over time. Unlike genetic search, its defining mechanism is not a population undergoing parent selection, crossover and mutation.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 8 | lines 36-38

o Genetic Algorithms
1. Which one of the following is correctly matched with the definition and approaches of AI
A)Acting humanly - The rational agent approach

Question sources

Mock exam 1 · Question 1

Question 118 | Q01486

A complete history of everything that the agent has ever perceived is

A. Agent program

B. Agent function

C. Percept

D. Percept sequence

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer D: Percept sequence

Percept sequence is explicitly the history of all the agent has perceived, unlike a single current percept. Evidence: PDF 10.

Why each choice is right or wrong

A - Incorrect: Implementation of agent function, not history. See PDF p. 10.

B - Incorrect: Maps history to actions, not history itself. See PDF p. 10.

C - Incorrect: One perceptual instant. See PDF p. 10.

D - Correct: Exact history definition. See PDF p. 10.

Primary note proof

Study-note section AI03

> A percept is one observation; a percept sequence is the entire observation history. An agent function maps percept sequences to actions: f:P*→A. An agent program implements that mapping on an architecture; agent = architecture + program. Environment means the setting in which it senses and acts, not merely the search space. Autonomy means relying substantially on experience rather than only a designer’s prior instructions.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 10 | lines 19-21

> • Percept Sequence − It is the history of all that an agent has perceived till date.
> • Agent Function − It is a map from the precept sequence to an action.
> • An agent is anything that can perceive its environment through sensors and acts upon that

Question sources

Mock exam 1 · Question 2

Question 119 | Q01487

A good cluster is characterized by

A. Equal number of instances

B. Equal distance among cluster members

C. High inter-class similarity

D. High intra-class similarity

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer D: High intra-class similarity

The module groups similar input vectors into clusters. Good within-cluster (intra-class) similarity follows from that grouping principle, with no requirement of equal sizes or equal pairwise distances. Evidence: PDF 75.

Why each choice is right or wrong

A - Incorrect: No equal-size condition given. See PDF p. 75.

B - Incorrect: Similarity does not require identical distances. See PDF p. 75.

C - Incorrect: Between-cluster similarity is not the grouping objective. See PDF p. 75.

D - Correct: Matches similarity inside a group. See PDF p. 75.

Primary note proof

Study-note section AI32

> Unsupervised learning seeks structure without target labels: clustering groups similar examples; association finds co-occurrence; dimensionality reduction compresses features. Good clusters have high within-cluster similarity and low between-cluster similarity under an appropriate distance. K-means alternates assignment to nearest centroids and centroid updates.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 75 | lines 19-21

> similar type are combined to form clusters. When a new input pattern is applied, then the neural network gives an output response indicating the class to which the input pattern belongs. There is no feedback from the environment as to what should be the desired output and if it is

Question sources

Mock exam 1 · Question 3

Question 120 | Q01569

A problem in supervised machine learning model due to too much attention to the particular data set it is trained on, causing it to perform poorly on unseen data is ____

A. Under fitting

B. Over fitting

C. Learning curve

D. Variance

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer B: Overfitting.

AI35 directly defines overfitting as fitting training-specific details so performance drops on unseen data. This is the complete failure pattern described in the stem. Underfitting is insufficient fit even on training data; a learning curve is a plot used to inspect performance. Variance may describe sensitivity but does not name the training-versus-generalization failure asked for.

Why each choice is right or wrong

A - Incorrect: Underfitting means the model is too limited or poorly fitted even on the training data.

B - Correct: Overfitting captures training-specific details at the expense of performance on unseen examples.

C - Incorrect: A learning curve plots performance against training size or progress; it is not the failure itself.

D - Incorrect: Variance is related to sensitivity, but overfitting is the precise failure pattern defined by the stem.

Primary note proof

Study-note section AI35

> Overfitting means fitting training-specific details so performance drops on unseen data; underfitting means the model is too limited or poorly fitted even on training data. A learning curve plots performance against training size or training progress. Complexity control, regularization, pruning and more representative data can help.

Study-note section AI-S06

> Support-vector machines include supervised classification and regression methods. They use support vectors in their decision function rather than being defined by a next-state conditional-dependence assumption. In estimation, variance measures how predictions change across different training samples; overfitting instead names poor generalization caused by fitting training-specific details.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 68 | lines 3-5

> tasks T and performance measure P, if its performance at tasks in T, as measured by P, improves with experience E.
> Thus a learning system is characterized by:

Question sources

Mock exam 3 · Question 1

Question 121 | Q01570

Which general AI search category uses the problem definition and known costs without a problem-specific estimate of remaining cost to a goal?

A. Informed search

B. Uninformed search

C. Greedy best-first search

D. Binary search

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer B: Uninformed search

Uninformed search uses the problem definition, including legal actions, goal tests and known action costs, without a domain-specific remaining-cost estimate. Informed methods add such guidance, and greedy best-first uses it directly as h. Binary search is a different ordered-lookup algorithm for sorted sequences. The formerly undefined circular-search distractor is replaced with a meaningful search method.

Why each choice is right or wrong

A - Incorrect: Informed search adds problem-specific guidance beyond the basic problem definition.

B - Correct: Uninformed search uses the problem definition without a heuristic estimate of remaining cost.

C - Incorrect: Greedy best-first explicitly uses h, an estimate of remaining cost.

D - Incorrect: Binary search narrows a sorted sequence by comparisons; it is a specific ordered-lookup algorithm, not the broad AI search category requested.

Primary note proof

Study-note section AI17

A heuristic h(n) estimates remaining cost or evaluates promise using problem-specific knowledge. It can reduce explored work, but computing it also costs time. Uninformed algorithms use only the problem definition and known costs; UCS remains uninformed despite using g.

Study-note section AI-S02

Binary search exploits the order of a sorted sequence: a comparison narrows the remaining candidate interval instead of exploring arbitrary successor states. Array bisection therefore differs from breadth-first or depth-first traversal of a general state-space graph. A storage container such as an array does not by itself specify FIFO or LIFO frontier behavior.

Study-note section AI17

Best-first is a family that orders OPEN by an evaluation function. Greedy best-first uses f=h and removes the smallest estimate, ignoring the cost already paid. It can choose an expensive route, loop without duplicate checks, or follow an infinite misleading branch. On a finite graph with full duplicate handling it is complete, but not generally optimal. Its implicit finite-tree worst case can be O(b^m).

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 18 | lines 23-25

• The uninformed search does not contain any domain knowledge such as closeness, the location of
the goal. It is also called level order traversal.
• It operates in a brute-force way as it only includes information about how to traverse the tree and

Question sources

Mock exam 3 · Question 2

Question 122 | Q01571

Which named test uses a judge’s inability to distinguish a machine from a human through conversation as its operational criterion?

A. The Chinese room

B. A psychological response-time experiment

C. Turing test

D. Black box testing

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer C: Turing test

The conversational Turing test evaluates whether a judge can reliably distinguish a machine's dialogue from a human's. It supplies an observable behavior criterion rather than proof of understanding. The Chinese-room argument addresses understanding philosophically, response-time experiments test cognitive models empirically, and black-box testing concerns externally visible software functionality. Those are different questions or procedures.

Why each choice is right or wrong

A - Incorrect: The Chinese-room argument concerns whether formal symbol manipulation suffices for understanding; it is not the judge-and-conversation test.

B - Incorrect: A response-time experiment evaluates a cognitive model against human measurements rather than applying conversational indistinguishability.

C - Correct: The conversational Turing test uses this judge-and-hidden-participant criterion.

D - Incorrect: Black-box testing checks external software functionality without inspecting internals; it does not inherently require humanlike conversation.

Primary note proof

Study-note section AI02

> The conversational Turing test asks whether a judge can reliably distinguish a machine from a human through their dialogue. It requires language use, knowledge, reasoning and often learning. It tests an operational behavior criterion, not proof of consciousness. Cognitive models can be tested against response times, errors and psychological experiments.

Study-note section AI-S07

> Searle's Chinese-room argument challenges whether executing formal symbol-manipulation rules is sufficient for understanding. It is a philosophical argument about understanding, not the conversational judge-and-hidden-participant procedure of the Turing test.

Study-note section AI-S07

> Black-box software testing examines externally visible functionality without inspecting implementation internals. A specification may concern any software behavior; this testing method does not by itself require conversational indistinguishability from a human.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 5 | lines 15-17

> argument Turing put forward the idea of an 'imitation game', in which a human being and a computer would be interrogated under conditions where the interrogator would not know which
> was which, the communication being entirely by textual messages. Turing argued that if the

Question sources

Mock exam 3 · Question 3

Question 123 | Q01612

A learning agent has mainly four conceptual components, which component is responsible for suggesting actions that will lead to new and informative experiences.

A. Problem generator

B. Performance element

C. Learning element

D. Critic

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer A: Problem generator

Problem generator proposes actions producing new informative experience, while the performance element chooses external action, learning element improves it, and critic evaluates it. Evidence: PDF 15.

Why each choice is right or wrong

A - Correct: Matches the explicit exploration role. See PDF p. 15.

B - Incorrect: Selects external action, different component role. See PDF p. 15.

C - Incorrect: Changes performance based on feedback. See PDF p. 15.

D - Incorrect: Evaluates against a performance standard. See PDF p. 15.

Primary note proof

Study-note section AI08

> The performance element selects actions. The learning element changes the agent to improve future performance. The critic evaluates behavior against a performance standard and supplies feedback. The problem generator suggests exploratory actions likely to yield useful experience.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 15 | lines 32-33

> - Problem generator is responsible for suggesting actions that will lead to new and informative experiences.

Question sources

Mock exam 2 · Question 1

Question 124 | Q01613

A learning scheme in which agent learns from a series of rewards & punishments is ____

A. Opinion mining

B. Unsupervised learning

C. Supervised learning

D. Reinforcement learning

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer D: Reinforcement learning

Rewards and penalties define the module reinforcement scheme. Evidence: PDF 69.

Why each choice is right or wrong

A - Incorrect: Opinion mining is a language application. See PDF p. 69.

B - Incorrect: Input-only pattern discovery lacks reward feedback. See PDF p. 69.

C - Incorrect: Uses paired desired outputs. See PDF p. 69.

D - Correct: Exactly reward/penalty feedback. See PDF p. 69.

Primary note proof

Study-note section AI33

> An RL agent observes state information, takes an action, receives reward and a new observation, and improves its policy to maximize expected cumulative reward. Rewards evaluate outcomes rather than supply the correct action label. Rewards may be delayed; not every correct-looking action receives immediate reward.

Study-note section AI37

> Tokenization divides text into processing units; tokens may be words, punctuation or subwords. POS tagging assigns grammatical roles; named-entity recognition identifies people, organizations and locations. Parsing identifies grammatical structure; semantics concerns meaning; discourse links sentences; pragmatics uses context and speaker intent. Sentiment analysis predicts evaluative tone. Ambiguity can occur at every level: “bank” has multiple senses, and “I saw the man with a telescope” has multiple attachments.

Study-note section AI31

> Supervised learning fits a mapping from labeled input-target examples. Classification predicts discrete labels, such as institution generations or spam/nonspam; regression predicts numeric quantities, such as prices. Predictive analytics uses historical patterns to forecast outcomes and can use either. Specify the target before choosing a method.

Study-note section AI32

Unsupervised learning seeks structure without target labels: clustering groups similar examples; association finds co-occurrence; dimensionality reduction compresses features. Good clusters have high within-cluster similarity and low between-cluster similarity under an appropriate distance. K-means alternates assignment to nearest centroids and centroid updates.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 69 | lines 16-18

reward for each right action and gets a penalty for each wrong action.
The agent learns automatically with these feedbacks and improves its performance. In reinforcement
learning, the agent interacts with the environment and explores it. The goal of an agent is to get the

Question sources

Mock exam 2 · Question 2

Question 125 | Q01614

A learning scheme in which agent learns patterns in the input without any explicit feedback is _____.

A. Supervised learning

B. Unsupervised learning

C. Opinion mining

D. Reinforcement learning

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer B: Unsupervised learning

The unsupervised network must discover features without feedback about correct outputs. Evidence: PDF 75.

Why each choice is right or wrong

A - Incorrect: Has desired outputs supplied. See PDF p. 75.

B - Correct: Discovers patterns without teacher feedback. See PDF p. 75.

C - Incorrect: Opinion mining is an application, not this learning scheme. See PDF p. 75.

D - Incorrect: Uses evaluative environmental feedback. See PDF p. 75.

Primary note proof

Study-note section AI32

Unsupervised learning seeks structure without target labels: clustering groups similar examples; association finds co-occurrence; dimensionality reduction compresses features. Good clusters have high within-cluster similarity and low between-cluster similarity under an appropriate distance. K-means alternates assignment to nearest centroids and centroid updates.

Study-note section AI31

Supervised learning fits a mapping from labeled input-target examples. Classification predicts discrete labels, such as institution generations or spam/nonspam; regression predicts numeric quantities, such as prices. Predictive analytics uses historical patterns to forecast outcomes and can use either. Specify the target before choosing a method.

Study-note section AI33

An RL agent observes state information, takes an action, receives reward and a new observation, and improves its policy to maximize expected cumulative reward. Rewards evaluate outcomes rather than supply the correct action label. Rewards may be delayed; not every correct-looking action receives immediate reward.

Study-note section AI37

Tokenization divides text into processing units; tokens may be words, punctuation or subwords. POS tagging assigns grammatical roles; named-entity recognition identifies people, organizations and locations. Parsing identifies grammatical structure; semantics concerns meaning; discourse links sentences; pragmatics uses context and speaker intent. Sentiment analysis predicts evaluative tone. Ambiguity can occur at every level: “bank” has multiple senses, and “I saw the man with a telescope” has multiple attachments.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 75 | lines 21-23

There is no feedback from the environment as to what should be the desired output and if it is correct or incorrect. Hence, in this type of learning, the network itself must discover the patterns and features from the input data, and the relation for the input data over the output.

Question sources

Mock exam 2 · Question 3

Question 126 | SM0010

Suppose you have developed a self-driving car equipped with relevant sensors and cameras. Its sensors are able to detect the environment but not all the aspects. Hence, in which environment is your car
acting?

A. Stochastic

B. Discrete

C. Fully observable

D. Partially observable

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer D: Partially observable

The sensors cannot determine complete state; this is precisely partial observability. It says nothing about transition randomness. Evidence: PDF 13.

Why each choice is right or wrong

A - Incorrect: Stochastic describes transitions, not hidden aspects. See PDF p. 13.

B - Incorrect: Discrete describes state structure, not visibility. See PDF p. 13.

C - Incorrect: Contradicted by incomplete detection. See PDF p. 13.

D - Correct: Matches missing state information. See PDF p. 13.

Primary note proof

Study-note section AI09

> Fully observable means sensors reveal relevant state; partially observable means relevant information is hidden or noisy. Deterministic means state and action determine the next state; stochastic means outcomes involve randomness. Strategic uncertainty about an opponent is distinct from dice-like randomness.

Study-note section AI09

> Episodic decisions are independent; sequential decisions affect later situations. Static worlds do not change while deliberating; dynamic ones do. A clock affecting score makes an otherwise static game semidynamic. Discrete quantities have distinct alternatives; continuous quantities vary along real-valued ranges. Known versus unknown concerns knowledge of the transition model, not observability.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 13 | lines 14-16

✓ Observable / Partially Observable – If it is possible to determine the complete state of the environment at each time point from the percepts it is observable; otherwise it is only partially observable.

Question sources

MOE Computer Science Exit Exam · Question 3-2

National Computer Science Exit Exam · Question 31

National Computer Science Exit Exam

MOE Computer Science Exit Exam

Computer Science Exit Exam practice collection · Question 193

Question 127 | SM0020

Which quantifier can be used for some portion of the universe?

A. For all (universal)

B. Implies

C. There exists (existential)

D. And

Hint

Translate the statement carefully, then check the scope of its quantifiers, connectives or inference rule. Distinguish guaranteed consequences from plausible generalizations or explanations.

Answer C: There exists (existential)

Some or at least one is existential quantification; universal applies to all. Evidence: PDF 49.

Why each choice is right or wrong

A - Incorrect: All instances rather than some. See PDF p. 49.

B - Incorrect: A connective, not a quantifier. See PDF p. 49.

C - Correct: Matches some/at least one. See PDF p. 49.

D - Incorrect: A connective, not a quantifier. See PDF p. 49.

Primary note proof

Study-note section AI26

First-order logic adds objects, predicates, functions and quantified object variables. Human(a) is a property; Likes(a,b) a relation; father(a) a function term. Propositional logic has no object quantifiers, but its truth still depends on the assigned atom values.

Study-note section AI25

Propositional atoms represent whole true-or-false claims. ¬P negates; P∧Q requires both; P∨Q requires at least one; P→Q is false only when P is true and Q false; P↔Q requires matching truth values. Thus P→Q ≡ ¬P∨Q. De Morgan gives ¬(P∧Q) ≡ ¬P∨¬Q.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 49 | lines 11-13

• Existential quantifier, (for some, at least one).
1. Universal Quantifier: - is a symbol of logical representation, which specifies that the statement within its range is true for everything or every instance of a particular thing.

Question sources

MOE Computer Science Exit Exam · Question 5-4

National Computer Science Exit Exam · Question 21

National Computer Science Exit Exam

MOE Computer Science Exit Exam

Computer Science Exit Exam practice collection · Question 194

Question 128 | SM0037

Which one of the following is a step in problem solving that lists the actions and states that the agent has to consider provided a goal?

A. Problem formulation

B. Execute

C. Search

D. Goal formulation

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer A: Problem formulation.

AI11 says the goal is formulated first, then the problem specifies initial state, legal actions, transitions, goal test and action costs. Given a goal, listing the relevant actions and states belongs to that second step. Search chooses a sequence within the formulated problem, and execution applies it; neither defines the alternatives initially.

Why each choice is right or wrong

A - Correct: Problem formulation specifies the states, available actions and model the agent must consider for its goal.

B - Incorrect: Execution applies a selected action sequence after formulation and search.

C - Incorrect: Search selects a solution using the already specified states, actions, model and goal test.

D - Incorrect: Goal formulation chooses the objective; the stem already provides a goal before listing actions and states.

Primary note proof

Study-note section AI11

> First formulate the goal; then formulate the problem by specifying initial state, legal actions, transition model, goal test, and action costs. Search selects a sequence; execution applies it. A state abstracts away irrelevant world details.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 18 | lines 16-18

> Actions: It gives the description of all the available actions to the agent.
> Transition model: A description of what each action do, can be represented as a transition model.
> Path Cost: It is a function which assigns a numeric cost to each path.

Question sources

MOE Computer Science Exit Exam · Question 11-3

National Computer Science Exit Exam · Question 98

National Computer Science Exit Exam

MOE Computer Science Exit Exam

Computer Science Exit Exam practice collection · Question 195

Question 129 | SM0061

Which term names the entire ordered history of observations an agent has received so far?

A. Performance measure

B. Percept sequence

C. Knowledge

D. Action

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer B: Percept sequence

A percept is one observation, while a percept sequence is the entire observation history. A performance measure defines success instead of storing observations. Knowledge may include prior facts and inferences beyond the observation history, and an action is selected behavior rather than observed input. Replacing the overly broad perception option supplies the precise established term without changing the agent-history concept.

Why each choice is right or wrong

A - Incorrect: A performance measure defines success, not what the agent has observed.

B - Correct: A percept sequence is the entire observation history.

C - Incorrect: Knowledge includes prior and inferred information, not precisely the observation history.

D - Incorrect: An action is selected behavior rather than the agent's accumulated observations.

Primary note proof

Study-note section AI03

> A percept is one observation; a percept sequence is the entire observation history. An agent function maps percept sequences to actions: f:P*→A. An agent program implements that mapping on an architecture; agent = architecture + program. Environment means the setting in which it senses and acts, not merely the search space. Autonomy means relying substantially on experience rather than only a designer's prior instructions.

Study-note section AI04

> A performance measure defines success externally: safety, useful deliveries, elapsed time, or cost. Rationality depends on that measure, prior knowledge, percept history, and available actions. A rational agent selects the action with the best expected performance given its information.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 10 | lines 19-21

> • Percept Sequence − It is the history of all that an agent has perceived till date.
> • Agent Function − It is a map from the precept sequence to an action.
> • An agent is anything that can perceive its environment through sensors and acts upon that

Question sources

MOE Computer Science Exit Exam · Question 17-4

National Computer Science Exit Exam · Question 51

National Computer Science Exit Exam

Computer Science Exit Exam practice collection · Question 198

Question 130 | SM0211

Every training example has an input and its desired target output; there is no extra unlabeled training pool or interaction reward. Which learning regime fits this setup?

A. Semi-supervised learning

B. Unsupervised learning

C. Supervised learning

D. Reinforcement learning

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer C: Supervised learning

The training data consist entirely of input-target pairs, so the task is supervised learning. The desired target provides the learning signal. Semi-supervised learning would add an unlabeled pool, unsupervised learning would seek structure without targets, and reinforcement learning would use rewards from interaction. Stating the complete data setup distinguishes the regimes without implying that any use of labels excludes semi-supervised methods.

Why each choice is right or wrong

A - Incorrect: Semi-supervised learning requires additional unlabeled training data alongside labeled examples, absent here.

B - Incorrect: Unsupervised learning discovers structure without the supplied target labels.

C - Correct: Supervised learning fits the mapping from the labeled input-target pairs described.

D - Incorrect: Reinforcement learning uses evaluative interaction rewards, not the described desired-output pairs.

Primary note proof

Study-note section AI31

> Supervised learning fits a mapping from labeled input-target examples. Classification predicts discrete labels, such as institution generations or spam/nonspam; regression predicts numeric quantities, such as prices. Predictive analytics uses historical patterns to forecast outcomes and can use either. Specify the target before choosing a method.

Study-note section AI32

> Semi-supervised learning combines some labels with additional unlabeled data. Active learning chooses informative cases for an oracle to label; it still uses labels. “Exploratory learning” is not an exclusive standard synonym: it can mean unsupervised exploration or active data acquisition, so define the feedback. A task named summarization or time-series analysis is not automatically supervised; its training targets determine that.

Study-note section AI33

An RL agent observes state information, takes an action, receives reward and a new observation, and improves its policy to maximize expected cumulative reward. Rewards evaluate outcomes rather than supply the correct action label. Rewards may be delayed; not every correct-looking action receives immediate reward.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 68 | lines 25-25

⇒ Supervised learning implies we are given a set of (x, y) pairs by a "teacher." The system is supplied

Introduction to Artificial Intelligence module | PDF page 69 | lines 2-4

with a set of training examples consisting of inputs and corresponding outputs, and is required to
discover the relation or mapping between then, e.g. as a series of rules, or a neural network. Agent
learns a function from observing example input-output pairs

Question sources

Computer Science Exit Exam · Sene 2016 · Question 4-1

Computer Science Exit Exam practice collection · Question 207

Question 131 | SM0222

The simple strategy in which the root node is expanded first, then all successors are expanded next, layer by layer, refers to:

A. Greedy best-first search

B. Depth-first search (DFS)

C. Depth limited search (DLS)

D. Breadth-first search (BFS)

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer D: Breadth-first search (BFS)

Expanding every node at one level before the next is the stated BFS order. Evidence: PDF 19.

Why each choice is right or wrong

A - Incorrect: Ranks by estimated goal distance. See PDF p. 19.

B - Incorrect: Follows a path deeply before siblings. See PDF p. 19.

C - Incorrect: Restricts DFS to a bound. See PDF p. 19.

D - Correct: Exact layer-order expansion. See PDF p. 19.

Primary note proof

Study-note section AI13

BFS removes the oldest frontier entry using a FIFO queue. Start with the root, remove the head, test it, append unseen successors at the tail, and repeat. It explores depths 0,1,2,... and is complete for finite b and a finite-depth goal. It minimizes number of edges, so minimum cost requires equal positive step costs.

Study-note section AI14

DFS follows the deepest pending branch using a LIFO stack or recursion. FILO means the same order as LIFO. With the same children as BFS, a left-first DFS visits S,A,C before B; push rightmost children first to pop the leftmost first.

Study-note section AI16

DLS runs DFS to depth l and reports cutoff separately from failure. It misses goals deeper than l. Iterative deepening runs DLS with limits 0,1,2,...; for finite b it finds the shallowest solution with O(b^d) time and O(bd) frontier space, optimal for equal step costs. Repeated shallow work is small compared with the deepest level.

Study-note section AI17

Best-first is a family that orders OPEN by an evaluation function. Greedy best-first uses f=h and removes the smallest estimate, ignoring the cost already paid. It can choose an expensive route, loop without duplicate checks, or follow an infinite misleading branch. On a finite graph with full duplicate handling it is complete, but not generally optimal. Its implicit finite-tree worst case can be O(b^m).

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 19 | lines 25-26

• BFS algorithm starts searching from the root node of the tree and expands all successor node at the
current level before moving to nodes of next level.

Question sources

Computer Science Exit Exam · Sene 2016 · Question 7-3

Computer Science Exit Exam practice collection · Question 206

Question 132 | SM0253

Which statement is not a correct characterization of reasoning?

A. Using knowledge to derive conclusions that can guide what a system does.

B. A process of thinking, logically arguing, and drawing inference

C. Theory of blind searching

D. The act of deriving a conclusion from certain premises using a given methodology

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer C: Theory of blind searching

Reasoning uses knowledge and a specified methodology to derive conclusions from premises. Logical argument and inference are therefore appropriate descriptions. Blind search is instead a search strategy that explores alternatives without heuristic guidance; it does not define reasoning as a whole. Clarifying the incomplete knowledge-system option into a claim about deriving conclusions leaves the intended search-versus-inference distinction as the unique exception.

Why each choice is right or wrong

A - Incorrect: Using knowledge to derive conclusions is a valid characterization, so it is not the exception.

B - Incorrect: Logical argument and drawing inferences describe reasoning, so this is not a false characterization for the negative question.

C - Correct: Blind search is one search strategy rather than a definition of reasoning.

D - Incorrect: Deriving a conclusion from premises using a methodology describes reasoning and does not satisfy the negative question.

Primary note proof

Study-note section AI24

> Knowledge representation encodes facts, objects, relations, events and rules so a machine can use them. Syntax defines legal expressions; semantics defines their meaning. A knowledge base stores assertions; an inference system derives consequences. Declarative knowledge says what is true; procedural knowledge specifies how to act.

Study-note section AI17

> A heuristic h(n) estimates remaining cost or evaluates promise using problem-specific knowledge. It can reduce explored work, but computing it also costs time. Uninformed algorithms use only the problem definition and known costs; UCS remains uninformed despite using g.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 53 | lines 18-20

> - Reasoning - is the mental process of deriving logical conclusion and making predictions from available knowledge, facts, and beliefs. "Reasoning is a way to infer facts from existing data." It is a general process of thinking rationally, to find valid conclusions. In artificial intelligence, the

Question sources

Computer Science Exit Exam · Sene 2016 · Question 17-2

Computer Science Exit Exam practice collection · Question 208

Question 133 | SM0261

Which reasoning technique applies valid inference rules so that true premises guarantee a true conclusion?

A. Abductive reasoning

B. Analogical reasoning

C. Deductive reasoning

D. Inductive reasoning

Hint

Translate the statement carefully, then check the scope of its quantifiers, connectives or inference rule. Distinguish guaranteed consequences from plausible generalizations or explanations.

Answer C: Deductive reasoning

Deductive reasoning uses truth-preserving inference: when the premises are true and the rule is valid, the conclusion must be true. Rules in a broad sense are not sufficient to identify deduction, since many reasoning methods use procedures. Induction generalizes from cases, abduction proposes explanations, and analogy transfers conclusions across similarities; these do not supply the stated logical guarantee.

Why each choice is right or wrong

A - Incorrect: Abduction proposes plausible explanatory hypotheses that can still be false.

B - Incorrect: Analogy transfers a conclusion from a similar case without a truth-preservation guarantee.

C - Correct: Valid deductive inference preserves truth from premises to conclusion.

D - Incorrect: Induction generalizes beyond the observed cases and remains revisable.

Primary note proof

Study-note section AI28

> Deduction applies truth-preserving rules: all people are mortal; Ada is a person; therefore Ada is mortal. Induction generalizes observations to a broader rule and can be overturned by new

examples. Abduction proposes a plausible explanation: wet ground suggests rain, although sprinklers could explain it. Analogy transfers a pattern between similar situations.

Study-note section AI28

Classical logical entailment is monotonic: adding premises preserves previous consequences. Default reasoning can be nonmonotonic: infer a bird usually flies, then withdraw this when told it is a penguin. Common-sense defaults are useful with incomplete knowledge, but are not guaranteed truths. Blind search is a search strategy, not a definition of reasoning.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 53 | lines 23-25

A. Deductive reasoning - it deducing new information from logically related known information.
o It is the form of valid reasoning, which means the argument's conclusion must be true when the
premises are true.

Question sources

Computer Science Exit Exam · Sene 2016 · Question 20-1

Computer Science Exit Exam practice collection · Question 204

Question 134 | SM0269

For A* tree search with an admissible heuristic and positive step costs, what is the intended benefit of using the heuristic alongside accumulated edge costs?

A. To seek a least-cost path while potentially reducing expansions through heuristic guidance.

B. To find the path with the maximum distance from the start to the goal node

C. To find paths with non-optimal searches from node to the goal node

D. To find the path with the maximum cost from the start to the goal node

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer A: To seek a least-cost path while potentially reducing expansions through heuristic guidance.

A* adds accumulated cost g to an admissible remaining-cost estimate h. Under the stated tree-search conditions, this supports least-cost solutions, while useful heuristic guidance can reduce expansions. It does not guarantee a speed improvement for every problem or heuristic. Maximizing distance or cost is a different objective, and deliberately seeking nonoptimal paths is not the benefit described.

Why each choice is right or wrong

A - Correct: An admissible remaining-cost estimate supports optimal tree search and can guide exploration toward useful candidates.

B - Incorrect: A* minimizes total path cost; it does not seek maximum path length.

C - Incorrect: Nonoptimality is not the purpose of admissible A* tree search.

D - Incorrect: A* seeks minimum, not maximum, total cost.

Primary note proof

Study-note section AI18

A* orders by f(n)=g(n)+h(n). Admissibility means 0≤h(n)≤h*(n), the true cheapest remaining cost. Consistency means h(n)≤c(n,n′)+h(n′) on every edge, with h(goal)=0; it implies admissibility on states with goal paths. Consistency makes f nondecreasing along a path.

Study-note section AI18

Admissibility suffices for optimal A* tree search; graph search that never reopens CLOSED states normally requires consistency. With an admissible inconsistent heuristic, reopen states when g improves. Test a goal when removed at minimum f. Completeness also needs finite branching and positive lower-bounded costs in infinite spaces. Worst-case time and memory can be exponential; no heuristic universally guarantees speed.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 27 | lines 2-4

• A* search algorithm finds the shortest path through the search space using the heuristic function.
This search algorithm expands less search tree and provides optimal result faster. A* algorithm is
similar to UCS except that it uses g(n)+h(n) instead of g(n).

Introduction to Artificial Intelligence module | PDF page 28 | lines 14-19

• Optimal: A* search algorithm is optimal if it follows below two conditions:
o Admissible: the first condition requires for optimality is that h(n) should be an admissible heuristic for A* tree search. An admissible heuristic is optimistic in nature.
o Consistency: Second required condition is consistency for only A* graph-search.
• If the heuristic function is admissible, then A* tree search will always find the least cost path.
• Time Complexity: The time complexity of A* search algorithm depends on heuristic function, and

Question sources

Computer Science Exit Exam · Sene 2016 · Question 22-3

Computer Science Exit Exam practice collection · Question 209

Question 135 | SM0299

What role does machine learning play in understanding human perception?

A. Machine learning helps simulate human perception through neural networks

B. Machine learning has no relevance to understanding human perception

C. Machine learning is not capable of analyzing perceptual data

D. Machine learning only focuses on logical reasoning and decision-making

SYSTEMS

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer A: Machine learning helps simulate human perception through neural networks

ANNs model brain-inspired processing and learn pattern recognition; combined with p67 recognition applications this supports the perception-modeling role. Evidence: PDF 71, 71, 67.

Why each choice is right or wrong

A - Correct: Matches brain-inspired models applied to recognition. See PDF p. 71, 71, 67.

B - Incorrect: Contradicts documented pattern-recognition applications. See PDF p. 71, 71, 67.

C - Incorrect: Contradicts recognition/classification capabilities. See PDF p. 71, 71, 67.

D - Incorrect: Machine learning also handles sensory patterns, not only logic. See PDF p. 71, 71, 67.

Primary note proof

Study-note section AI34

> Feedforward networks pass signals toward outputs without recurrent state loops; recurrent networks carry state across sequence steps. Training adjusts weights, often through gradients and backpropagation. Backpropagation is not the same as a recurrent connection. Deep learning uses multiple representation layers and can learn features from images, sound and text. CNNs share filters across image locations. ANN is a brain-inspired abstraction, not a literal brain simulation; DL⊂ML⊂AI.

Study-note section AI39

> Computer vision interprets images or video: classification names image content, detection locates objects, segmentation labels pixels, and recognition identifies objects or people. Occlusion, lighting change, noise, unfamiliar conditions and incomplete context make open-world perception difficult even when a face detector works well.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 71 | lines 6-8

> An Artificial Neural Network (ANN) is an information processing paradigm that is inspired by the way
> biological nervous systems, such as the brain, process information. The key element of this paradigm
> is the novel structure of the information processing system.

Introduction to Artificial Intelligence module | PDF page 71 | lines 11-13

> configured for a specific application, such as pattern recognition or data classification, through a learning
> process. Learning in biological systems involves adjustments to the synaptic connections that exist
> between the neurons. This is true of ANNs as well. ANNs can process information at a great speed

Introduction to Artificial Intelligence module | PDF page 67 | lines 30-31

> • Speech/image/video (pattern) recognition
> • Adaptive control

Question sources

Computer Science Exit Exam · Sene 2016 · Question 32-2

Computer Science Exit Exam practice collection · Question 210

Question 136 | SS100

An informed search keeps one priority frontier ordered by g(n)+h(n), permanently closes states, and tests goals on removal. With a consistent heuristic, finite branching and positive lower-bounded step costs, which algorithm is this?

A. Iterative Deepening A* Search

B. Greedy Best-First Search

C. Uniform Cost Search

D. A* Search

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer D: A* Search

A* uses a best-first priority frontier ordered by g+h. Consistency provides the standard guarantee for graph search that permanently closes states, with the stated cost and goal-test conditions. IDA* also uses f values but organizes repeated depth-first threshold searches instead. Greedy search omits g, and UCS omits h. The frontier organization distinguishes the algorithms rather than a worst-case complexity shared by several methods.

Why each choice is right or wrong

A - Incorrect: IDA* repeats depth-first searches under increasing f thresholds rather than retaining this single priority frontier.

B - Incorrect: Greedy best-first uses h alone and does not have the stated cost-aware priority.

C - Incorrect: UCS uses g alone, without the heuristic term.

D - Correct: A* uses g+h; consistency supports optimality when graph states are closed without reopening.

Primary note proof

Study-note section AI18

A* orders by f(n)=g(n)+h(n). Admissibility means 0≤h(n)≤h*(n), the true cheapest remaining cost. Consistency means h(n)≤c(n,n′)+h(n′) on every edge, with h(goal)=0; it implies admissibility on states with goal paths. Consistency makes f nondecreasing along a path.

Study-note section AI18

Admissibility suffices for optimal A* tree search; graph search that never reopens CLOSED states normally requires consistency. With an admissible inconsistent heuristic, reopen states when g improves. Test a goal when removed at minimum f. Completeness also needs finite branching and positive lower-bounded costs in infinite spaces. Worst-case time and memory can be exponential; no heuristic universally guarantees speed.

Study-note section AI16

Bidirectional search grows from start and goal until frontiers connect. With suitable reverse actions, meeting tests and equal-cost BFS, time and space may fall to O(b^(d/2)). Weighted variants require a sound stopping bound; the first intersection need not be cheapest. IDA* deepens an f=g+h threshold rather than only a depth limit.

Study-note section AI17

Best-first is a family that orders OPEN by an evaluation function. Greedy best-first uses f=h and removes the smallest estimate, ignoring the cost already paid. It can choose an expensive route, loop without duplicate checks, or follow an infinite misleading branch. On a finite graph with full duplicate handling it is complete, but not generally optimal. Its implicit finite-tree worst case can be O(b^m).

Study-note section AI15

UCS uses a priority queue ordered by actual accumulated path cost g(n).

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 28 | lines 19-21

• Time Complexity: The time complexity of A* search algorithm depends on heuristic function, and
the number of nodes expanded is exponential to the depth of solution d. So the time complexity

is
O(bd), where b is the branching factor.

Introduction to Artificial Intelligence module | PDF page 28 | lines 14-19

• Optimal: A* search algorithm is optimal if it follows below two conditions:
o Admissible: the first condition requires for optimality is that h(n) should be an admissible heuristic for A* tree search. An admissible heuristic is optimistic in nature.
o Consistency: Second required condition is consistency for only A* graph-search.
• If the heuristic function is admissible, then A* tree search will always find the least cost path.
• Time Complexity: The time complexity of A* search algorithm depends on heuristic function, and

Introduction to Artificial Intelligence module | PDF page 26 | lines 17-18

o This algorithm is not optimal.
• Time Complexity: The worst case time complexity of Greedy best first search is O(bm).

Question sources

Computer Science model exam collection · Question 7

Question 137 | SS143

Which general term denotes an entity that perceives its environment through sensors and acts on it through effectors?

A. API

B. Expert System

C. Intelligence

D. Agent

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer D: Agent

An agent is defined by perceiving an environment and acting on it through sensors and effectors. The acting mechanism may be physical or software-based. An API can expose operations used by an agent, and an expert system can contribute domain reasoning, but neither names the general entity defined here. Intelligence is a capability rather than the entity itself.

Why each choice is right or wrong

A - Incorrect: An API exposes software operations; it may support an agent's actions but is not the general perceiving-and-acting entity.

B - Incorrect: An expert system is a specialized knowledge-based application, not the general term for every sensing-and-acting entity.

C - Incorrect: Intelligence is a capability, not the name of the defined entity.

D - Correct: Agent is the general term for an entity with this perception-and-action interaction.

Primary note proof

Study-note section AI03

> An agent perceives an environment through sensors and acts through actuators or effectors. Cameras sense; motors move. Software sensors may be incoming messages and software actions may be API requests. Humans are agents but are not artificial agents.

Study-note section AI40

> An expert system applies specialized knowledge to a domain such as equipment diagnosis. The domain expert supplies expertise; the knowledge engineer elicits and encodes it; the inference engine reasons over a knowledge base and working facts. The user interface handles interaction; an explanation facility presents why or how a conclusion was reached.

Study-note section AI-S06

> An application programming interface (API) exposes defined operations through which software requests functionality. It can be an interface used by an agent, but the interface alone is not the general entity that perceives an environment and selects actions.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 10 | lines 21-23

> • An agent is anything that can perceive its environment through sensors and acts upon that environment through effectors. An agent perceives its environment through sensors.
> • The complete set of inputs at a given time is called a percept.

Question sources

Computer Science model exam collection · Question 50

Question 138 | SS171

A new dataset has no target labels. The objective is to discover its groups or structure, with neither interaction rewards nor reuse of a pretrained model. Which learning regime applies?

A. Reinforcement learning

B. Unsupervised learning

C. Transfer learning

D. Supervised learning

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer B: Unsupervised learning

The objective is structure discovery without target labels, which defines unsupervised learning. The stem separately excludes action-reward interaction and reuse of a pretrained model. Those exclusions distinguish reinforcement feedback and transfer learning from this task. Transfer learning describes a reuse strategy rather than a label-availability regime and can be combined with different learning objectives.

Why each choice is right or wrong

A - Incorrect: Reinforcement learning requires a reward-based interaction objective, which is excluded.

B - Correct: Unsupervised learning discovers structure in data without supplied target labels.

C - Incorrect: Transfer learning concerns reuse of prior learned representations or models; no such reuse is present.

D - Incorrect: Supervised learning requires target outputs paired with the inputs.

Primary note proof

Study-note section AI32

> Unsupervised learning seeks structure without target labels: clustering groups similar examples; association finds co-occurrence; dimensionality reduction compresses features. Good clusters have high within-cluster similarity and low between-cluster similarity under an appropriate distance. K-means alternates assignment to nearest centroids and centroid updates.

Study-note section AI31

> Supervised learning fits a mapping from labeled input-target examples. Classification predicts discrete labels, such as institution generations or spam/nonspam; regression predicts numeric quantities, such as prices. Predictive analytics uses historical patterns to forecast outcomes and can use either. Specify the target before choosing a method.

Study-note section AI33

> An RL agent observes state information, takes an action, receives reward and a new observation, and improves its policy to maximize expected cumulative reward. Rewards evaluate outcomes rather than supply the correct action label. Rewards may be delayed; not every correct-looking action receives immediate reward.

Study-note section AI-S06

> Transfer learning reuses representations or a model learned on an earlier task for a new task, for example adapting pretrained image features to new image classes. This is a reuse strategy, not a specification that the new task has no labels or rewards.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 69 | lines 7-9

⇒ Unsupervised learning means we are only given the x‘s. In either case, the goal is to estimate f. The system is supplied with a set of training examples consisting only of inputs and is required to discover for itself what appropriate outputs should be,

Question sources

Computer Science model exam collection · Question 78

Question 139 | SS175

One of the following statement is correct about Artificial Intelligence (Al)?

A. It refers to the development of machines that can perform tasks only in a controlled laboratory setting.

B. It refers to the development of machines that can only perform simple, repetitive tasks.

C. Itrefers to the development of machines that can only perform physical tasks.

D. It refers to the development of machines that can perform tasks that typically require human intelligence.

Hint

Match the input, output and purpose to the appropriate language, perception or physical-action task.

Answer D: It refers to the development of machines that can perform tasks that typically require human intelligence.

AI designs intelligence in artificial devices and includes language, learning, reasoning and real-world action, not the restrictive only-claims. Evidence: PDF 4, 6.

Why each choice is right or wrong

A - Incorrect: Real-world robots contradict laboratory-only. See PDF p. 4, 6.

B - Incorrect: Reasoning and complex problem solving exceed repetition. See PDF p. 4, 6.

C - Incorrect: Language/reasoning are not merely physical. See PDF p. 4, 6.

D - Correct: Matches design of artificial intelligence. See PDF p. 4, 6.

Primary note proof

Study-note section AI01

Artificial intelligence studies systems that perform tasks requiring intelligence: reasoning, learning, planning, perception and language. Applied AI builds useful solutions; cognitive AI uses computational models to investigate biological intelligence. A model should be compared with human behavior or experimental observations before claiming it explains human thought.

Study-note section AI39

Robotics combines mechanical, electrical, control and computing systems to sense, plan and act physically. Components include sensors, controller, power supply, actuators and an end effector.

Locomotion may be legged, wheeled or tracked, with trade-offs in terrain, stability and efficiency. Software can act autonomously without being a robot. Humanoid shape is not evidence of general intelligence.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 4 | lines 26-28

is concerned with the design of intelligence in an artificial device. The term was coined by McCarthy
in 1956. There are two ideas in the definition. i.e Intelligence, and Artificial device.
What is intelligence?

Introduction to Artificial Intelligence module | PDF page 6 | lines 12-21

1.3Intelligent behaviour
This discussion brings us back to the question of what constitutes intelligent behavior. Some ofthese
tasks and applications are:
• Perception involving image recognition and computer vision
• Reasoning
• Learning
• Understanding language involving natural language processing, speech processing
• Solving problems
• Robotics
• Thinking

Question sources

Computer Science model exam collection · Question 82

Question 140 | SS187

A techniques for generating plans with conditionals and loops are almost identical to those for generating programs from logical specifications is
called__

A. Automatic learning

B. Automatic goal testing

C. Automatic monitoring

D. Automatic programming

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer D: Automatic programming

Automatic programming, also called program synthesis, constructs a program that satisfies a specification. Generating conditional plans or loops is analogous because the output must implement the required behavior. Learning changes performance through experience, goal testing checks whether a candidate succeeds, and monitoring checks execution. None of those by itself names the specification-to-program construction task.

Why each choice is right or wrong

A - Incorrect: Learning improves performance through experience rather than specifically synthesizing a program from logical requirements.

B - Incorrect: A goal test checks whether a candidate meets the goal; it does not generate a program from its specification.

C - Incorrect: Monitoring checks an executing system rather than generating a program from a specification.

D - Correct: Automatic programming or program synthesis constructs a program meeting a specification, including conditional or looping behavior.

Primary note proof

Study-note section AI11

> In route finding, states are cities, actions are roads, successors are neighboring cities, and path cost is accumulated distance or time. A solution is a legal path ending at a goal; an optimal solution has minimum cost. Automatic programming or program synthesis searches for a program satisfying a specification; conditional plans and loops require richer representations than a fixed action list.

Study-note section AI08

> Learning is improvement on task T, measured by P, through experience E. For spam filtering, T is classifying email, E is past labeled messages, and P might be recall at a required precision. Exploration trades immediate reward for information; exploitation uses the best action currently known. Learning can improve reflex, goal-based or utility-based agents.

Study-note section AI11

> First formulate the goal; then formulate the problem by specifying initial state, legal actions, transition model, goal test, and action costs. Search selects a sequence; execution applies it. A state abstracts away irrelevant world details.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 64 | lines 6-8

> • Backward Chaining
> o With this strategy, an expert system finds out the answer to the question, —Why this happened?‖

o On the basis of what has already happened, the Inference Engine tries to find out which conditions

Question sources

Computer Science model exam collection · Question 94

Question 141 | AM0624

Which choices describe broad intelligence-related goals of AI rather than implementation means or hardware capacity? Select all that apply.

A. Emulating human intelligence

B. Creating machines that can think and learn

C. Developing advanced algorithms

D. Enhancing computer processing power

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer A, B: Emulating human intelligence; Creating machines that can think and learn

AI can aim to emulate aspects of human intelligence or to build systems that reason and learn. Those are compatible broad goals, although different AI approaches evaluate humanlike behavior and rational performance differently. Advanced algorithms are means that may serve many fields, while additional processing power is hardware capacity rather than intelligence itself. The distinction therefore permits both intelligence-related objectives.

Why each choice is right or wrong

A - Correct: Humanlike thought and behavior are established AI approaches and fit this broad goal.

B - Correct: Building systems with intelligent reasoning and learning is another supported broad AI objective.

C - Incorrect: Advanced algorithms are implementation means, not the defining intelligence goal.

D - Incorrect: More processing power alone does not establish intelligent behavior or learning.

Primary note proof

Study-note section AI01

Artificial intelligence studies systems that perform tasks requiring intelligence: reasoning, learning, planning, perception and language. Applied AI builds useful solutions; cognitive AI uses computational models to investigate biological intelligence. A model should be compared with human behavior or experimental observations before claiming it explains human thought.

Study-note section AI43

> Face recognition is both an applied task and a narrow capability. Broad machine-intelligence goals can include emulating human intelligence and building systems that learn. Humanlike problem-solving alone does not uniquely select expert systems. Python is widely used, but C++ and Java also implement AI; “commonly used” does not isolate one language. McCarthy is associated with naming AI, not its sole invention.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 8 | lines 9-11

> General AI goal
> • Replicate human intelligent
> • Solve knowledge intensive task

Introduction to Artificial Intelligence module | PDF page 8 | lines 14-16

> Engineering based AI goal
> • Develop concepts, theory and practice of building intelligent machines
> • Emphasis is on system building

Question sources

Computer Science MCQ practice collection · Question 624

Question 142 | AM0625

What is knowledge representation in artificial intelligence?

A. The process of storing information in a computer

B. The ability of machines to reason and make decisions

C. The study of how humans acquire knowledge

D. The process of encoding knowledge for machine use

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer D: The process of encoding knowledge for machine use

Knowledge representation organizes and formalizes knowledge in a machine-usable knowledge base. Storage alone omits the encoding/formalization purpose. Evidence: PDF 62.

Why each choice is right or wrong

A - Incorrect: Too broad: mere storage need not represent usable knowledge. See PDF p. 62.

B - Incorrect: Reasoning uses representations but is not representation itself. See PDF p. 62.

C - Incorrect: Human acquisition is a different subject. See PDF p. 62.

D - Correct: Best match to formal machine-usable encoding. See PDF p. 62.

Primary note proof

Study-note section AI24

> Knowledge representation encodes facts, objects, relations, events and rules so a machine can use them. Syntax defines legal expressions; semantics defines their meaning. A knowledge base stores assertions; an inference system derives consequences. Declarative knowledge says what is true; procedural knowledge specifies how to act.

Study-note section AI40

> An expert system applies specialized knowledge to a domain such as equipment diagnosis. The domain expert supplies expertise; the knowledge engineer elicits and encodes it; the inference engine reasons over a knowledge base and working facts. The user interface handles interaction; an explanation facility presents why or how a conclusion was reached.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 62 | lines 27-27

> o It is the method used to organize and formalize the knowledge in the knowledge base.

Question sources

Computer Science MCQ practice collection · Question 625

Question 143 | AM0626

Which representation system matches IF-THEN production rules against working memory and applies a selected rule to derive a conclusion or action?

A. Semantic networks

B. Neural networks

C. Production systems

D. Bayesian networks

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer C: Production systems

A production system uses IF-THEN rules whose conditions are matched against working memory; an applicable rule is selected and applied. Semantic networks encode concept relationships, neural networks learn weighted representations, and Bayesian networks encode probabilistic dependencies. All can support reasoning in broader systems, so the distinctive feature here is the explicit production-rule cycle rather than inference in general.

Why each choice is right or wrong

A - Incorrect: Semantic networks represent concepts and labeled relationships, not this working-memory production-rule cycle.

B - Incorrect: Neural networks compute through weighted units and learned representations rather than the specified rule-matching architecture.

C - Correct: Production systems match condition-action or condition-conclusion rules against working memory.

D - Incorrect: Bayesian networks represent conditional probabilistic dependencies rather than the stated production-rule mechanism.

Primary note proof

Study-note section AI24

A semantic network connects concept nodes by labeled relations such as is-a. Frames describe stereotyped entities through slots, values and defaults. Production systems use IF conditions THEN conclusions or actions, matching rules against working memory and selecting among applicable rules. A graph suits relationships such as frequently bought together; frames suit attributes. No representation is always best for recommendation without task requirements.

Study-note section AI29

A Bayesian network is a directed acyclic graph with conditional probability tables; its joint distribution factors as the product of each node's probability given its parents. A Markov model assumes the next state depends on the current state rather than the entire past. These represent uncertainty; fuzzy membership represents a different issue.

Study-note section AI34

Feedforward networks pass signals toward outputs without recurrent state loops; recurrent networks carry state across sequence steps. Training adjusts weights, often through gradients and backpropagation. Backpropagation is not the same as a recurrent connection. Deep learning uses multiple representation layers and can learn features from images, sound and text. CNNs share filters across image locations. ANN is a brain-inspired abstraction, not a literal brain simulation; DL⊂ML⊂AI.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 63 | lines 20-22

o Applies rules repeatedly to the facts, which are obtained from earlier rule application.
o Adds new knowledge into the knowledge base if required.
o Resolves rules conflict when multiple rules are applicable to a particular case

Introduction to Artificial Intelligence module | PDF page 64 | lines 5-6

and production rule systems.
• Backward Chaining

Question sources

Computer Science MCQ practice collection · Question 626

Question 144 | AM0627

Which reasoning technique involves drawing conclusions from specific observations or examples?

A. Deductive reasoning

B. Inductive reasoning

C. Abductive reasoning

D. Probabilistic reasoning

Hint

Translate the statement carefully, then check the scope of its quantifiers, connectives or inference rule. Distinguish guaranteed consequences from plausible generalizations or explanations.

Answer B: Inductive reasoning

Specific facts generalized into a conclusion define induction. The conclusion is supported rather than logically guaranteed by the finite examples. Evidence: PDF 53.

Why each choice is right or wrong

A - Incorrect: Typically general premises to specific conclusion. See PDF p. 53.

B - Correct: Exact specific-to-general pattern. See PDF p. 53.

C - Incorrect: Seeks most likely explanation for observations. See PDF p. 53.

D - Incorrect: Uses probability to express uncertainty, not this defining direction. See PDF p. 53.

Primary note proof

Study-note section AI28

> Deduction applies truth-preserving rules: all people are mortal; Ada is a person; therefore Ada is mortal. Induction generalizes observations to a broader rule and can be overturned by new examples. Abduction proposes a plausible explanation: wet ground suggests rain, although sprinklers could explain it. Analogy transfers a pattern between similar situations.

Study-note section AI29

> Probability describes uncertainty about events. $P(\neg A)=1-P(A)$; $P(A|B)=P(A\cap B)/P(B)$, for $P(B)>0$. Bayes gives $P(H|E)=P(E|H)P(H)/P(E)$: posterior equals likelihood times prior divided by evidence probability. For mutually exclusive exhaustive hypotheses, $P(E)=\Sigma_i P(E|Hi)P(Hi)$.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 53 | lines 32-33

o It starts with the series of specific facts or data and reaches to a general statement or conclusion.

Question sources

Computer Science MCQ practice collection · Question 627

Question 145 | AM0628

A classifier is trained on labeled examples and predicts by following attribute tests from a root to a class-labeled leaf. Which learning technique matches this representation?

A. Reinforcement learning

B. Unsupervised learning

C. Deep learning

D. Decision tree learning

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer D: Decision tree learning

A decision tree predicts by testing an attribute at each internal node, following a branch and returning the leaf prediction. Labeled training makes this a supervised classification task. Deep learning can also be supervised, so supervision alone would not distinguish it. The representation in the stem provides the distinction; reinforcement and unsupervised learning describe different feedback regimes.

Why each choice is right or wrong

A - Incorrect: Reinforcement learning specifies environmental reward feedback rather than this attribute-test representation.

B - Incorrect: Unsupervised learning seeks structure without supplied target labels.

C - Incorrect: Deep learning uses multiple neural representation layers; it can be supervised but is not the attribute-test-and-leaf model described.

D - Correct: Decision tree learning produces attribute-test branches with predictions at leaves.

Primary note proof

Study-note section AI31

Supervised learning fits a mapping from labeled input-target examples. Classification predicts discrete labels, such as institution generations or spam/nonspam; regression predicts numeric quantities, such as prices. Predictive analytics uses historical patterns to forecast outcomes and can use either. Specify the target before choosing a method.

Study-note section AI31

A decision tree tests an attribute at each internal node, follows the corresponding branch and predicts at a leaf. ID3 chooses high information gain: $H(S)=-\Sigma p_i \log_2 p_i$; gain is parent entropy minus weighted child entropy. Gini impurity is $1-\Sigma p_i^2$, zero for a pure class. Pruning limits overfitting. Neural networks, support-vector machines and trees can all be supervised.

Study-note section AI43

ANNs, genetic search and reinforcement learning all improve candidate behavior or performance through experience in different ways. Deep learning is an architecture family usable with supervised, unsupervised or reinforcement objectives. Decision trees and deep models may both be supervised. To distinguish learning regimes, specify labels, reward or unlabeled structure.

Study-note section AI34

Feedforward networks pass signals toward outputs without recurrent state loops; recurrent networks carry state across sequence steps. Training adjusts weights, often through gradients and backpropagation. Backpropagation is not the same as a recurrent connection. Deep learning uses multiple representation layers and can learn features from images, sound and text. CNNs share filters across image locations. ANN is a brain-inspired abstraction, not a literal brain simulation; DL⊂ML⊂AI.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 69 | lines 21-23

A. Decision tree algorithm
⇒ It is a classification and prediction tool having a tree like structure, where each internal node denotes
a test on an attribute, each branch represents an outcome of the test, and each leaf node (terminal

Introduction to Artificial Intelligence module | PDF page 69 | lines 5-7

o Classification
o Regression
⇒ Unsupervised learning means we are only given the x‘s. In either case, the goal is to estimate f.

Question sources

Computer Science MCQ practice collection · Question 628

Question 146 | AM0629

What is the primary objective of natural language processing (NLP) in artificial intelligence?

A. Understanding and generating human language by machines

B. Analyzing and processing large datasets

C. Enhancing computer hardware capabilities

D. Simulating human-like decision-making

Hint

Match the input, output and purpose to the appropriate language, perception or physical-action task.

Answer A: Understanding and generating human language by machines

Understanding/interpretation is defined in NLP, and NLG turns data into language, supporting both understanding and generation. Evidence: PDF 79, 79.

Why each choice is right or wrong

A - Correct: Matches understanding plus NLG. See PDF p. 79, 79.

B - Incorrect: Large-data processing is broader than language. See PDF p. 79, 79.

C - Incorrect: Not a hardware discipline. See PDF p. 79, 79.

D - Incorrect: Too broad; omits the defining language objective. See PDF p. 79, 79.

Primary note proof

Study-note section AI37

> NLP interprets and generates human language: translation, dialogue, summarization, question answering and text generation. Speech recognition converts audio to text; speech synthesis generates speech. They overlap language processing but solve different input/output tasks.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 79 | lines 4-6

> and Artificial Intelligence. It is the technology that is used by machines to understand, analyze, manipulate, and interpret human's languages. It helps developers to organize knowledge for performing tasks such as translation, automatic summarization, Named Entity Recognition (NER),

Introduction to Artificial Intelligence module | PDF page 79 | lines 31-33

> - Natural Language Generation (NLG) acts as a translator that converts the computerized data into natural language representation. It mainly involves Text planning, Sentence planning, and Text Realization.

Question sources

Computer Science MCQ practice collection · Question 629

Question 147 | AM0630

Which model family assumes that the distribution of the next state depends on the current state rather than the entire earlier history?

A. Decision trees

B. Markov models

C. Support vector machines

D. Reinforcement learning

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer B: Markov models

The Markov assumption makes the current state sufficient for the next-state distribution, rather than requiring the full preceding history. That identifies the model family. Decision trees and support-vector machines are prediction methods with other defining structures. Reinforcement learning is a reward-based learning framework that can use a Markov model, but it is not itself the particular conditional-dependence assumption.

Why each choice is right or wrong

A - Incorrect: Decision trees are defined by attribute tests and leaf predictions rather than this conditional transition assumption.

B - Correct: A Markov model conditions the next state on the current state instead of the full previous history.

C - Incorrect: Support-vector machines use support vectors for prediction; a next-state conditional-dependence assumption does not define the family.

D - Incorrect: Reinforcement learning specifies improvement from rewards and may use different model families.

Primary note proof

Study-note section AI29

> A Bayesian network is a directed acyclic graph with conditional probability tables; its joint distribution factors as the product of each node's probability given its parents. A Markov model assumes the next state depends on the current state rather than the entire past. These represent uncertainty; fuzzy membership represents a different issue.

Study-note section AI31

> A decision tree tests an attribute at each internal node, follows the corresponding branch and predicts at a leaf. ID3 chooses high information gain: $H(S)=-\Sigma p_i \log_2 p_i$; gain is parent entropy minus weighted child entropy. Gini impurity is $1-\Sigma p_i^2$, zero for a pure class. Pruning limits overfitting. Neural networks, support-vector machines and trees can all be supervised.

Study-note section AI33

An RL agent observes state information, takes an action, receives reward and a new observation, and improves its policy to maximize expected cumulative reward. Rewards evaluate outcomes rather than supply the correct action label. Rewards may be delayed; not every correct-looking action receives immediate reward.

Study-note section AI-S06

Support-vector machines include supervised classification and regression methods. They use support vectors in their decision function rather than being defined by a next-state conditional-dependence assumption. In estimation, variance measures how predictions change across different training samples; overfitting instead names poor generalization caused by fitting training-specific details.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 55 | lines 27-29

• Probabilistic reasoning is a way of knowledge representation where we apply the concept of probability to indicate the uncertainty in knowledge.
• In probabilistic reasoning, we combine probability theory with logic to handle the uncertainty.

Question sources

Computer Science MCQ practice collection · Question 630

Question 148 | AM0631

Which learning regime is defined by an agent trying actions, receiving environmental rewards and improving a policy to maximize expected cumulative return?

A. Reinforcement learning

B. Unsupervised learning

C. Deep learning

D. Genetic algorithms

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer A: Reinforcement learning

Reinforcement learning improves an action policy through rewards received during interaction, with expected cumulative return as the objective. Generic trial and error is too broad because evolutionary search also evaluates and improves candidates. Genetic algorithms are identified by population selection and variation, while deep learning identifies architecture. Unsupervised learning concerns unlabeled structure rather than this sequential reward objective.

Why each choice is right or wrong

A - Correct: Reinforcement learning is defined by this action-reward policy-improvement objective.

B - Incorrect: Unsupervised learning discovers unlabeled structure rather than defining cumulative-return optimization.

C - Incorrect: Deep learning specifies multilayer neural representations and can be used within several learning regimes.

D - Incorrect: Genetic algorithms improve candidate populations using selection and genetic variation; trial and error alone does not make them the stated action-reward learning regime.

Primary note proof

Study-note section AI33

> An RL agent observes state information, takes an action, receives reward and a new observation, and improves its policy to maximize expected cumulative reward. Rewards evaluate outcomes rather than supply the correct action label. Rewards may be delayed; not every correct-looking action receives immediate reward.

Study-note section AI43

> ANNs, genetic search and reinforcement learning all improve candidate behavior or performance through experience in different ways. Deep learning is an architecture family usable with supervised, unsupervised or reinforcement objectives. Decision trees and deep models may both be supervised. To distinguish learning regimes, specify labels, reward or unlabeled structure.

Study-note section AI20

> A genetic algorithm maintains a population, evaluates fitness, selects parents, applies crossover and mutation, and repeats replacement. It searches without guaranteeing global optimality. Means-ends analysis selects a difference between current and goal states, chooses an operator reducing it, and recursively solves its preconditions. Neither method is simply A* running backward.

Study-note section AI32

> Unsupervised learning seeks structure without target labels: clustering groups similar examples; association finds co-occurrence; dimensionality reduction compresses features.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 69 | lines 18-20

> learning, the agent interacts with the environment and explores it. The goal of an agent is to get the
> most reward points, and hence, it improves its performance. The robotic dog, which automatically
> learns the movement of his arms, is an example of Reinforcement learning.

Question sources

Computer Science MCQ practice collection · Question 631

Question 149 | AM0632

Which type of learning algorithm aims to mimic the structure and function of the human brain?

A. Supervised learning

B. Unsupervised learning

C. Deep learning

D. Reinforcement learning

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer C: Deep learning, as a multilayer brain-inspired neural architecture.

AI34 calls ANNs brain-inspired abstractions and describes deep learning as using multiple representation layers. AI43 stresses that deep models can use supervised, unsupervised or reinforcement objectives. Thus deep learning is the offered architecture-based answer, while the other choices classify feedback regimes rather than requiring brain-inspired structure. This is inspiration, not a claim of literal brain simulation.

Why each choice is right or wrong

A - Incorrect: Supervised learning specifies labeled targets and can use non-neural models such as trees or SVMs.

B - Incorrect: Unsupervised learning specifies unlabeled structure discovery and does not require a neural architecture.

C - Correct: Deep learning uses multilayer neural representations within the brain-inspired ANN family.

D - Incorrect: Reinforcement learning specifies reward-based interaction and need not use brain-inspired neural structure.

Primary note proof

Study-note section AI34

> Feedforward networks pass signals toward outputs without recurrent state loops; recurrent networks carry state across sequence steps. Training adjusts weights, often through gradients and backpropagation. Backpropagation is not the same as a recurrent connection. Deep learning uses multiple representation layers and can learn features from images, sound and text. CNNs share filters across image locations. ANN is a brain-inspired abstraction, not a literal brain simulation; DL⊂ML⊂AI.

Study-note section AI43

ANNs, genetic search and reinforcement learning all improve candidate behavior or performance through experience in different ways. Deep learning is an architecture family usable with supervised, unsupervised or reinforcement objectives. Decision trees and deep models may both be supervised. To distinguish learning regimes, specify labels, reward or unlabeled structure.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 71 | lines 7-9

biological nervous systems, such as the brain, process information. The key element of this paradigm
is the novel structure of the information processing system.
It is composed of large number of highly interconnected simple processing elements (neurons)

Introduction to Artificial Intelligence module | PDF page 74 | lines 15-17

B. Multilayer feedforward network − The concept is of eedforward ANN having more than one weighted layer. As this network has one or more layers between the input and the output layer, it
is called hidden layers.

Question sources

Computer Science MCQ practice collection · Question 632

Question 150 | AM0633

What is the primary purpose of expert systems in artificial intelligence?

A. Emulate human expertise in a specific domain

B. Perform complex mathematical calculations

C. Process and analyze large datasets

D. Optimize computer hardware performance

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer A: Emulate human expertise in a specific domain

Expert systems solve complex problems in a particular domain at the level of human expertise. Evidence: PDF 60.

Why each choice is right or wrong

A - Correct: Exact domain-expertise purpose. See PDF p. 60.

B - Incorrect: May be a tool, not defining purpose. See PDF p. 60.

C - Incorrect: May be a workload, not defining purpose. See PDF p. 60.

D - Incorrect: Not the defining expert-system purpose. See PDF p. 60.

Primary note proof

Study-note section AI40

> An expert system applies specialized knowledge to a domain such as equipment diagnosis. The domain expert supplies expertise; the knowledge engineer elicits and encodes it; the inference engine reasons over a knowledge base and working facts. The user interface handles interaction; an explanation facility presents why or how a conclusion was reached.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 60 | lines 16-18

> The expert systems are the computer applications developed to solve complex problems in a particular
> domain, at the level of extra-ordinary human intelligence and expertise. Some ccharacteristics of expert system are High performance, Understandable, Reliable, and Highly responsive.

Question sources

Computer Science MCQ practice collection · Question 633

Question 151 | AM0634

Within the cognitive approach to AI, which listed objective expresses the use of computational models to investigate human intelligence and perception?

A. Replicate human intelligence and perception

B. Enhance computer processing power

C. Analyze large datasets efficiently

D. Improve natural language processing

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer A: Replicate human intelligence and perception

Cognitive AI uses computational models to investigate biological intelligence, including human perception and thought. Replicating relevant aspects of these processes expresses that modeling objective, provided the model is compared with observations. Greater processing power and data efficiency are resources or methods. Language processing is a narrower task and need not aim to explain biological intelligence.

Why each choice is right or wrong

A - Correct: Modeling or replicating aspects of intelligence and perception is the cognitive investigation objective.

B - Incorrect: Faster hardware enables computation but does not itself supply a model of intelligence.

C - Incorrect: Efficient dataset analysis is a method or workload, not the specific cognitive-modeling objective.

D - Incorrect: NLP is one application area and does not by itself express the broader intelligence-and-perception investigation.

Primary note proof

Study-note section AI01

> Artificial intelligence studies systems that perform tasks requiring intelligence: reasoning, learning, planning, perception and language. Applied AI builds useful solutions; cognitive AI uses computational models to investigate biological intelligence. A model should be compared with human behavior or experimental observations before claiming it explains human thought.

Study-note section AI02

> Thinking humanly models human mental processes; thinking rationally derives conclusions by valid logic. Acting humanly seeks humanlike behavior; acting rationally selects actions expected to achieve objectives. Thought processes and observed behavior are different evaluation targets.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 8 | lines 10-12

- Replicate human intelligent
- Solve knowledge intensive task
- Make an intelligent connection between perception and action

Question sources

Computer Science MCQ practice collection · Question 634

Question 152 | AM0635

Which AI approach uses computational models to investigate human thought processes and compares those models with human behavior?

A. Machine learning,

B. Natural language processing

C. Cognitive AI (cognitive modeling)

D. Robotics

Hint

Match the input, output and purpose to the appropriate language, perception or physical-action task.

Answer C: Cognitive AI (cognitive modeling)

Cognitive AI investigates biological intelligence through computational models, testing those models against human behavior or experimental observations. This is a scientific modeling purpose rather than simply useful automation. Machine learning, NLP and robotics may contribute tools to cognitive research, but none is defined by that purpose. The option now uses the explicit approach name taught in the note.

Why each choice is right or wrong

A - Incorrect: Machine learning improves performance from experience and need not model human thought processes.

B - Incorrect: NLP focuses on language interpretation and generation rather than necessarily explaining human cognition.

C - Correct: Cognitive AI uses computational models to investigate biological intelligence and human thought.

D - Incorrect: Robotics integrates physical sensing and action; it need not aim to explain human thought.

Primary note proof

Study-note section AI01

Artificial intelligence studies systems that perform tasks requiring intelligence: reasoning, learning, planning, perception and language. Applied AI builds useful solutions; cognitive AI uses computational models to investigate biological intelligence. A model should be compared with human behavior or experimental observations before claiming it explains human thought.

Study-note section AI02

Thinking humanly models human mental processes; thinking rationally derives conclusions by valid logic. Acting humanly seeks humanlike behavior; acting rationally selects actions expected to achieve objectives. Thought processes and observed behavior are different evaluation targets.

Study-note section AI37

NLP interprets and generates human language: translation, dialogue, summarization, question answering and text generation. Speech recognition converts audio to text; speech synthesis generates speech. They overlap language processing but solve different input/output tasks.

Study-note section AI39

Robotics combines mechanical, electrical, control and computing systems to sense, plan and act physically. Components include sensors, controller, power supply, actuators and an end effector. Locomotion may be legged, wheeled or tracked, with trade-offs in terrain, stability and efficiency. Software can act autonomously without being a robot. Humanoid shape is not evidence of general intelligence.

Study-note section AI08

Learning is improvement on task T, measured by P, through experience E.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 5 | lines 12-14

the human thought process. This view is the cognitive science approach to AI. -Think human like
2. The second approach is best embodied by the concept of the Turing Test. Turing held that in future
computers can be programmed to acquire abilities rivaling human intelligence.As part of his

Introduction to Artificial Intelligence module | PDF page 7 | lines 9-10

Cognitive AI: computers are used to test theories about how the human mind works--for example,
theories about how we recognize faces and other objects, or about how we solve abstract problems.

Question sources

Computer Science MCQ practice collection · Question 635

Question 153 | AM0636

Which area of AI research focuses on understanding and replicating human vision?

A. Natural language processing

B. Computer vision

C. Robotics

D. Expert systems

Hint

Match the input, output and purpose to the appropriate language, perception or physical-action task.

Answer B: Computer vision

Computer vision extracts and comprehends image information, making it the specific vision field. Evidence: PDF 84.

Why each choice is right or wrong

A - Incorrect: Language rather than image interpretation. See PDF p. 84.

B - Correct: Exact visual-comprehension field. See PDF p. 84.

C - Incorrect: Robots may use vision but robotics is broader. See PDF p. 84.

D - Incorrect: Domain expertise rather than vision specifically. See PDF p. 84.

Primary note proof

Study-note section AI39

> Computer vision interprets images or video: classification names image content, detection locates objects, segmentation labels pixels, and recognition identifies objects or people. Occlusion, lighting change, noise, unfamiliar conditions and incomplete context make open-world perception difficult even when a face detector works well.

Study-note section AI37

> NLP interprets and generates human language: translation, dialogue, summarization, question answering and text generation. Speech recognition converts audio to text; speech synthesis generates speech. They overlap language processing but solve different input/output tasks.

Study-note section AI40

> An expert system applies specialized knowledge to a domain such as equipment diagnosis. The domain expert supplies expertise; the knowledge engineer elicits and encodes it; the inference engine reasons over a knowledge base and working facts. The user interface handles interaction; an explanation facility presents why or how a conclusion was reached.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 84 | lines 6-8

accomplish automatic visual comprehension.
o Hardware of Computer Vision System - This involves –
• Power supply

Question sources

Computer Science MCQ practice collection · Question 636

Question 154 | AM0637

Which AI technique is used to extract insights and patterns from large datasets?

A. Natural language processing

B. Machine learning

C. Expert systems

D. Robotics

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer B: Machine learning

Machine learning applications include data mining/knowledge discovery and learning patterns; ANN example p71 specifically extracts complex patterns and trends. Evidence: PDF 67, 71.

Why each choice is right or wrong

A - Incorrect: Language-specific tools do not define general data learning. See PDF p. 67, 71.

B - Correct: Matches learning/discovery from data. See PDF p. 67, 71.

C - Incorrect: Knowledge-based domain expertise rather than general learned patterns. See PDF p. 67, 71.

D - Incorrect: Physical agents rather than the broad data-analysis technique. See PDF p. 67, 71.

Primary note proof

Study-note section AI08

Learning is improvement on task T, measured by P, through experience E. For spam filtering, T is classifying email, E is past labeled messages, and P might be recall at a required precision. Exploration trades immediate reward for information; exploitation uses the best action currently known. Learning can improve reflex, goal-based or utility-based agents.

Study-note section AI31

Supervised learning fits a mapping from labeled input-target examples. Classification predicts discrete labels, such as institution generations or spam/nonspam; regression predicts numeric quantities, such as prices. Predictive analytics uses historical patterns to forecast outcomes and can use either. Specify the target before choosing a method.

Study-note section AI32

Unsupervised learning seeks structure without target labels: clustering groups similar examples; association finds co-occurrence; dimensionality reduction compresses features. Good clusters have high within-cluster similarity and low between-cluster similarity under an appropriate distance. K-means alternates assignment to nearest centroids and centroid updates.

Study-note section AI40

An expert system applies specialized knowledge to a domain such as equipment diagnosis. The domain expert supplies expertise; the knowledge engineer elicits and encodes it; the inference engine reasons over a knowledge base and working facts. The user interface handles interaction; an explanation facility presents why or how a conclusion was reached.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 67 | lines 29-31

• Data mining and knowledge discovery
• Speech/image/video (pattern) recognition
• Adaptive control

Introduction to Artificial Intelligence module | PDF page 71 | lines 16-18

can be used to extract patterns and detect trends that are too complex to be noticed by either humans or
other computer techniques. A trained neural network can be thought of as an "expert" in the category
of information it has been given to analyse. This expert can then be used to provide projections given

Question sources

Computer Science MCQ practice collection · Question 637

Question 155 | AM0638

Which concept in AI refers to the ability of machines to understand and interpret human language?

A. Natural language processing

B. Machine learning

C. Expert systems

D. Robotics

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer A: Natural language processing

NLP is expressly machine understanding, analysis and interpretation of human languages. Evidence: PDF 79.

Why each choice is right or wrong

A - Correct: Exact language-understanding concept. See PDF p. 79.

B - Incorrect: Can support NLP but is broader learning from data. See PDF p. 79.

C - Incorrect: Domain expertise is not language understanding generally. See PDF p. 79.

D - Incorrect: Physical embodied action is a different focus. See PDF p. 79.

Primary note proof

Study-note section AI37

> NLP interprets and generates human language: translation, dialogue, summarization, question answering and text generation. Speech recognition converts audio to text; speech synthesis generates speech. They overlap language processing but solve different input/output tasks.

Study-note section AI08

> Learning is improvement on task T, measured by P, through experience E. For spam filtering, T is classifying email, E is past labeled messages, and P might be recall at a required precision. Exploration trades immediate reward for information; exploitation uses the best action currently known. Learning can improve reflex, goal-based or utility-based agents.

Study-note section AI40

> An expert system applies specialized knowledge to a domain such as equipment diagnosis. The domain expert supplies expertise; the knowledge engineer elicits and encodes it; the inference engine reasons over a knowledge base and working facts. The user interface handles interaction; an explanation facility presents why or how a conclusion was reached.

Study-note section AI39

> Robotics combines mechanical, electrical, control and computing systems to sense, plan and act physically. Components include sensors, controller, power supply, actuators and an end effector. Locomotion may be legged, wheeled or tracked, with trade-offs in terrain, stability and efficiency. Software can act autonomously without being a robot. Humanoid shape is not evidence of general intelligence.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 79 | lines 5-7

> manipulate, and interpret human's languages. It helps developers to organize knowledge for performing tasks such as translation, automatic summarization, Named Entity Recognition (NER),
> speech recognition, relationship extraction, and topic segmentation.

Question sources

Computer Science MCQ practice collection · Question 638

Question 156 | AM0639

Which AI technique is used to simulate human-like conversation and interaction?

A. Machine learning

B. Natural language processing

C. Robotics

D. Expert systems

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer B: Natural language processing

The module explicitly identifies chatbots providing customer chat as an NLP application. Evidence: PDF 82.

Why each choice is right or wrong

A - Incorrect: May power a chatbot, but not the specific application field. See PDF p. 82.

B - Correct: Explicit conversation application field. See PDF p. 82.

C - Incorrect: Embodiment not required for conversation. See PDF p. 82.

D - Incorrect: Expert advice is not the defining conversation technique. See PDF p. 82.

Primary note proof

Study-note section AI37

> NLP interprets and generates human language: translation, dialogue, summarization, question answering and text generation. Speech recognition converts audio to text; speech synthesis generates speech. They overlap language processing but solve different input/output tasks.

Study-note section AI08

> Learning is improvement on task T, measured by P, through experience E. For spam filtering, T is classifying email, E is past labeled messages, and P might be recall at a required precision. Exploration trades immediate reward for information; exploitation uses the best action currently known. Learning can improve reflex, goal-based or utility-based agents.

Study-note section AI40

> An expert system applies specialized knowledge to a domain such as equipment diagnosis. The domain expert supplies expertise; the knowledge engineer elicits and encodes it; the inference engine reasons over a knowledge base and working facts. The user interface handles interaction; an explanation facility presents why or how a conclusion was reached.

Study-note section AI39

> Robotics combines mechanical, electrical, control and computing systems to sense, plan and act physically. Components include sensors, controller, power supply, actuators and an end effector. Locomotion may be legged, wheeled or tracked, with trade-offs in terrain, stability and efficiency. Software can act autonomously without being a robot. Humanoid shape is not evidence of general intelligence.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 82 | lines 5-7

> - Chatbot - is one of the important applications of NLP. It is used by many companies to provide the customer's chat services.
> - Information extraction - Information extraction is one of the most important applications of

Question sources

Computer Science MCQ practice collection · Question 639

Question 157 | AM0640

What is the specifically computational-modeling contribution of cognitive AI to studying human thought, as distinct from collecting measurements or running psychological experiments alone?

A. By analyzing human brain activity

B. By replicating human thought processes

C. By conducting psychological experiments

D. By improving computer hardware capabilities

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer B: By replicating human thought processes

Cognitive AI contributes computational models that implement proposed mechanisms of human thought. Their predictions can then be compared with behavior, response times, errors or other observations. Brain measurements and psychological experiments can provide valuable evidence, but collecting evidence alone is distinct from constructing the computational model. Faster hardware similarly enables work without supplying a cognitive explanation.

Why each choice is right or wrong

A - Incorrect: Brain measurements can provide evidence, but measurement alone does not construct the computational thought model requested.

B - Correct: Replicating aspects of thought in a computational model provides a theory whose behavior can be tested.

C - Incorrect: Psychological experiments can test a model, but experiments alone do not supply the computational representation asked for.

D - Incorrect: More capable hardware supplies resources without itself representing a theory of human thought.

Primary note proof

Study-note section AI01

> Artificial intelligence studies systems that perform tasks requiring intelligence: reasoning, learning, planning, perception and language. Applied AI builds useful solutions; cognitive AI uses computational models to investigate biological intelligence. A model should be compared with human behavior or experimental observations before claiming it explains human thought.

Study-note section AI02

> Thinking humanly models human mental processes; thinking rationally derives conclusions by valid logic. Acting humanly seeks humanlike behavior; acting rationally selects actions expected to achieve objectives. Thought processes and observed behavior are different evaluation targets.

Study-note section AI02

> The conversational Turing test asks whether a judge can reliably distinguish a machine from a human through their dialogue. It requires language use, knowledge, reasoning and often learning. It tests an operational behavior criterion, not proof of consciousness. Cognitive models can be tested against response times, errors and psychological experiments.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 7 | lines 9-11

> Cognitive AI: computers are used to test theories about how the human mind works--for example,
> theories about how we recognize faces and other objects, or about how we solve abstract problems.
> 1.6Limits of AI Today

Introduction to Artificial Intelligence module | PDF page 5 | lines 11-12

> This view involves trying to understand human thought and an effort to build machines that emulate
> the human thought process. This view is the cognitive science approach to AI. -Think human like

Question sources

Computer Science MCQ practice collection · Question 640

Question 158 | AM0641

Which field builds integrated physical agents that sense their surroundings and act autonomously through motors or other actuators?

A. Machine learning

B. Robotics

C. Natural language processing

D. Expert systems

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer B: Robotics

Robotics builds physical systems that sense, plan and act, combining mechanical, electrical, control and computing components. The requirement for integrated physical sensing and actuation distinguishes it from software autonomy in general. Machine learning, NLP and expert systems can be components of a robot, but none by itself names the field responsible for the complete physical agent.

Why each choice is right or wrong

A - Incorrect: Machine learning provides experience-based improvement but does not itself define construction of physical sensing-and-acting systems.

B - Correct: Robotics integrates the mechanical, electrical, control and computing systems of physical agents.

C - Incorrect: NLP processes human language and does not require a physical body.

D - Incorrect: Expert systems supply specialized knowledge-based reasoning and need not control physical actuators.

Primary note proof

Study-note section AI39

> Robotics combines mechanical, electrical, control and computing systems to sense, plan and act physically. Components include sensors, controller, power supply, actuators and an end effector. Locomotion may be legged, wheeled or tracked, with trade-offs in terrain, stability and efficiency. Software can act autonomously without being a robot. Humanoid shape is not evidence of general intelligence.

Study-note section AI08

> Learning is improvement on task T, measured by P, through experience E. For spam filtering, T is classifying email, E is past labeled messages, and P might be recall at a required precision. Exploration trades immediate reward for information; exploitation uses the best action currently known. Learning can improve reflex, goal-based or utility-based agents.

Study-note section AI37

NLP interprets and generates human language: translation, dialogue, summarization, question answering and text generation. Speech recognition converts audio to text; speech synthesis generates speech. They overlap language processing but solve different input/output tasks.

Study-note section AI40

An expert system applies specialized knowledge to a domain such as equipment diagnosis. The domain expert supplies expertise; the knowledge engineer elicits and encodes it; the inference engine reasons over a knowledge base and working facts. The user interface handles interaction; an explanation facility presents why or how a conclusion was reached.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 82 | lines 10-12

Robotics is a domain in artificial intelligence that deals with the study of creating intelligent and efficient robots. Robots are the artificial agents acting in real world environment. Robots are aimed at
manipulating the objects by perceiving, picking, moving, modifying the physical properties of object,

Question sources

Computer Science MCQ practice collection · Question 641

Question 159 | AM0642

When studying how people recognize objects, what is the computational-modeling contribution of cognitive AI, rather than building a sensor or gathering biological and behavioral measurements alone?

A. By simulating human sensory organs

B. By analyzing human brain waves

C. By replicating human perceptual processes

D. By conducting psychological experiments

Hint

Match the input, output and purpose to the appropriate language, perception or physical-action task.

Answer C: By replicating human perceptual processes

Cognitive AI can implement proposed perceptual processes and compare the resulting recognition behavior with human observations. This is distinct from only building an input sensor or collecting brain-wave and behavioral measurements. Those activities can inform or test the model, but they do not alone provide its computational account of perception. The scope is modeling the recognition process, not claiming that other research methods lack value.

Why each choice is right or wrong

A - Incorrect: A sensory-organ simulation models an input mechanism, but does not by itself explain how observed patterns are recognized.

B - Incorrect: Brain-wave measurements provide data rather than the computational recognition model requested.

C - Correct: Replicating aspects of perceptual processing supplies a model whose recognition behavior can be tested.

D - Incorrect: Psychological experiments provide tests or observations; by themselves they do not construct the computational perception model.

Primary note proof

Study-note section AI01

> Artificial intelligence studies systems that perform tasks requiring intelligence: reasoning, learning, planning, perception and language. Applied AI builds useful solutions; cognitive AI uses computational models to investigate biological intelligence. A model should be compared with human behavior or experimental observations before claiming it explains human thought.

Study-note section AI02

> The conversational Turing test asks whether a judge can reliably distinguish a machine from a human through their dialogue. It requires language use, knowledge, reasoning and often learning. It tests an operational behavior criterion, not proof of consciousness. Cognitive models can be tested against response times, errors and psychological experiments.

Study-note section AI39

> Computer vision interprets images or video: classification names image content, detection locates objects, segmentation labels pixels, and recognition identifies objects or people. Occlusion, lighting change, noise, unfamiliar conditions and incomplete context make open-world perception difficult even when a face detector works well.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 7 | lines 10-12

> theories about how we recognize faces and other objects, or about how we solve abstract problems.
> 1.6Limits of AI Today
> Today‘s successful AI systems operate in well-defined domains and employ narrow, specialized

Question sources

Computer Science MCQ practice collection · Question 642

Question 160 | AM0643

Which AI technique solves specialized domain problems by applying expert knowledge encoded in a knowledge base through an inference engine?

A. Machine learning

B. Expert systems

C. Natural language processing

D. Robotics

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer B: Expert systems

An expert system applies specialized domain knowledge through an inference engine operating on a knowledge base and working facts. This is more specific than merely exhibiting intelligent problem solving. Machine learning emphasizes improvement from experience, NLP concerns language, and robotics combines sensing, planning and physical action. Those methods may be combined with expert systems, but they are not the technique defined by the stated knowledge-based architecture.

Why each choice is right or wrong

A - Incorrect: Machine learning fits behavior from experience rather than defining this expert-knowledge architecture.

B - Correct: Expert systems use a knowledge base and inference engine for specialized domain reasoning.

C - Incorrect: NLP concerns human language rather than this domain-expertise architecture.

D - Incorrect: Robotics concerns physical sensing and action rather than the stated knowledge-base technique.

Primary note proof

Study-note section AI40

> An expert system applies specialized knowledge to a domain such as equipment diagnosis. The domain expert supplies expertise; the knowledge engineer elicits and encodes it; the inference engine reasons over a knowledge base and working facts. The user interface handles interaction; an explanation facility presents why or how a conclusion was reached.

Study-note section AI08

> Learning is improvement on task T, measured by P, through experience E. For spam filtering, T is classifying email, E is past labeled messages, and P might be recall at a required precision. Exploration trades immediate reward for information; exploitation uses the best action currently known. Learning can improve reflex, goal-based or utility-based agents.

Study-note section AI37

NLP interprets and generates human language: translation, dialogue, summarization, question answering and text generation. Speech recognition converts audio to text; speech synthesis generates speech. They overlap language processing but solve different input/output tasks.

Study-note section AI39

Robotics combines mechanical, electrical, control and computing systems to sense, plan and act physically. Components include sensors, controller, power supply, actuators and an end effector. Locomotion may be legged, wheeled or tracked, with trade-offs in terrain, stability and efficiency. Software can act autonomously without being a robot. Humanoid shape is not evidence of general intelligence.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 60 | lines 17-19

domain, at the level of extra-ordinary human intelligence and expertise. Some ccharacteristics of expert system are High performance, Understandable, Reliable, and Highly responsive. Capabilities of Expert Systems include

Introduction to Artificial Intelligence module | PDF page 67 | lines 19-21

systems that learn from experience and examples. It is the computational study of algorithms that
improve performance based on experience
Machine learning is particularly attractive in several real- l i f e problems because ofthe following

Question sources

Computer Science MCQ practice collection · Question 643

Question 161 | AM0644

What is a heuristic in the context of artificial intelligence (AI)?

A. A predefined set of rules for decision making

B. A mathematical model used to solve search problems

C. An intelligent agent capable of learning from experience

D. A strategy or rule of thumb used to guide problem-solving

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer D: A strategy or rule of thumb used to guide problem-solving

A heuristic guides search toward promising states using estimated closeness; it is a problem-solving guide rather than a complete learning agent or guaranteed exact solver. Evidence: PDF 24, 24.

Why each choice is right or wrong

A - Incorrect: A predefined rule could implement a heuristic, but not all decision rules are heuristics. See PDF p. 24, 24.

B - Incorrect: A heuristic need not be a complete mathematical search model. See PDF p. 24, 24.

C - Incorrect: A learning agent is an entity, not the guiding strategy. See PDF p. 24, 24.

D - Correct: Closest to the module's guiding-way definition. See PDF p. 24, 24.

Primary note proof

Study-note section AI17

> A heuristic h(n) estimates remaining cost or evaluates promise using problem-specific knowledge. It can reduce explored work, but computing it also costs time. Uninformed algorithms use only the problem definition and known costs; UCS remains uninformed despite using g.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 24 | lines 17-19

> • Informed search is also called a Heuristic search. A heuristic is a way which might not always be
> guaranteed for best solutions but guaranteed to find a good solution in reasonable time.
> • Informed search can solve much complex problem which could not be solved in another way.

Introduction to Artificial Intelligence module | PDF page 24 | lines 22-24

> • It takes the current state of the agent as its input and produces the estimation of how close agent is
> from the goal.
> • The heuristic method, however, might not always give the best solution, but it guaranteed to find a

Question sources

Computer Science MCQ practice collection · Question 644

Question 162 | AM0645

How are heuristics useful in search problems?

A. They guarantee finding the optimal solution.

B. They reduce the search space and speed up the search process.

C. They eliminate the need for search algorithms.

D. They ensure a systematic exploration of the search space.

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer B: They reduce the search space and speed up the search process.

Problem knowledge guides search toward promising states, often reducing explored nodes and time. It does not universally guarantee optimality or replace a search algorithm. Evidence: PDF 24.

Why each choice is right or wrong

A - Incorrect: Module explicitly warns heuristic may not give best solution. See PDF p. 24.

B - Correct: Matches guided efficient exploration. See PDF p. 24.

C - Incorrect: A heuristic is used within informed search. See PDF p. 24.

D - Incorrect: Systematic exploration alone does not distinguish heuristic guidance. See PDF p. 24.

Primary note proof

Study-note section AI17

> A heuristic h(n) estimates remaining cost or evaluates promise using problem-specific knowledge. It can reduce explored work, but computing it also costs time. Uninformed algorithms use only the problem definition and known costs; UCS remains uninformed despite using g.

Study-note section AI18

> Admissibility suffices for optimal A* tree search; graph search that never reopens CLOSED states normally requires consistency. With an admissible inconsistent heuristic, reopen states when g improves. Test a goal when removed at minimum f. Completeness also needs finite branching and positive lower-bounded costs in infinite spaces. Worst-case time and memory can be exponential; no heuristic universally guarantees speed.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 24 | lines 16-18

> • Informed search strategies can find a solution more efficiently than an uninformed search strategy.

• Informed search is also called a Heuristic search. A heuristic is a way which might not always be
guaranteed for best solutions but guaranteed to find a good solution in reasonable time.

Question sources

Computer Science MCQ practice collection · Question 645

Question 163 | AM0646

Which general characterization says that search uses problem-specific knowledge to estimate remaining cost, without itself requiring either non-overestimation or an edge-by-edge consistency inequality?

A. Admissible heuristic

B. Consistent heuristic

C. Informed heuristic

D. Dominance comparison between heuristics

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer C: Informed heuristic

Informed search adds problem-specific knowledge to guide exploration, often through a remaining-cost estimate. That general characterization does not guarantee admissibility or consistency, which impose additional mathematical conditions. Dominance instead compares two estimates pointwise. The question asks for the broad knowledge-use property rather than one of those stronger conditions or comparisons.

Why each choice is right or wrong

A - Incorrect: Admissibility adds a non-overestimation bound, stronger than the general knowledge-use characterization requested.

B - Incorrect: Consistency adds an edge-by-edge inequality, stronger than the general property requested.

C - Correct: Informed heuristic guidance uses problem-specific knowledge to estimate remaining cost or promise.

D - Incorrect: Dominance compares the numerical estimates of two heuristics rather than naming general use of problem knowledge.

Primary note proof

Study-note section AI17

A heuristic h(n) estimates remaining cost or evaluates promise using problem-specific knowledge. It can reduce explored work, but computing it also costs time. Uninformed algorithms use only the problem definition and known costs; UCS remains uninformed despite using g.

Study-note section AI18

A* orders by f(n)=g(n)+h(n). Admissibility means $0 \le h(n) \le h^*(n)$, the true cheapest remaining cost. Consistency means $h(n) \le c(n,n') + h(n')$ on every edge, with h(goal)=0; it implies admissibility on states with goal paths. Consistency makes f nondecreasing along a path.

Study-note section AI-S05

For two admissible heuristics, h_2 dominates h_1 when $h_2(n) \ge h_1(n)$ for every node. Dominance compares estimates; it is not itself a condition ensuring that an arbitrary heuristic never overestimates. Both estimates must separately satisfy the relevant admissibility assumptions.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 24 | lines 14-16

• It uses domain knowledge. In an informed search, problem information is available which can guide
the search.
• Informed search strategies can find a solution more efficiently than an uninformed search strategy.

Question sources

Computer Science MCQ practice collection · Question 646

Question 164 | AM0647

Which term names the heuristic property defined by $0 \le h(n) \le h^*(n)$, where $h^*(n)$ is the true cheapest remaining cost?

A. Admissible heuristic

B. Consistent heuristic

C. Informed heuristic

D. Dominance comparison between heuristics

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer A: Admissible heuristic

The displayed non-overestimation inequality defines admissibility. Consistency is a different edge condition; with the usual nonnegative heuristic estimates and a zero goal estimate, it implies admissibility for states with a goal path. Thus a consistent heuristic can also satisfy the bound, but consistency is not the name of the property being defined. Informedness and dominance do not independently establish the bound.

Why each choice is right or wrong

A - Correct: Admissibility is exactly the non-overestimation property stated by this inequality.

B - Incorrect: Consistency is an edge-by-edge inequality with goal normalization, rather than the name of the displayed non-overestimation bound.

C - Incorrect: Being informed only means using problem knowledge and does not impose this bound.

D - Incorrect: Dominance compares two heuristic estimates and does not by itself bound either one by the true remaining cost.

Primary note proof

Study-note section AI18

> A* orders by f(n)=g(n)+h(n). Admissibility means 0≤h(n)≤h*(n), the true cheapest remaining cost. Consistency means h(n)≤c(n,n′)+h(n′) on every edge, with h(goal)=0; it implies admissibility on states with goal paths. Consistency makes f nondecreasing along a path.

Study-note section AI-S05

> For two admissible heuristics, h_2 dominates h_1 when $h_2(n) \geq h_1(n)$ for every node. Dominance compares estimates; it is not itself a condition ensuring that an arbitrary heuristic never overestimates. Both estimates must separately satisfy the relevant admissibility assumptions.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 25 | lines 4-6

> • Admissibility of the heuristic function is given as: h(n) <= h*(n) Here h(n) is heuristic cost, and h*(n) is the estimated cost. Hence heuristic cost should be less than or equal to the estimated cost.
> • Pure Heuristic Search - Pure heuristic search is the simplest form of heuristic search algorithms. It

Introduction to Artificial Intelligence module | PDF page 28 | lines 17-18

> o Consistency: Second required condition is consistency for only A* graph-search.
> • If the heuristic function is admissible, then A* tree search will always find the least cost path.

Question sources

Computer Science MCQ practice collection · Question 647

Question 165 | AM0648

Which forward best-first algorithm ranks frontier nodes by accumulated path cost plus a heuristic estimate of remaining cost?

A. Depth-first search

B. Breadth-first search

C. A* search algorithm

D. Hill climbing algorithm

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer C: A* search algorithm

A* orders frontier nodes by f(n) = g(n) + h(n), combining the cost already paid with an estimate of the cost still needed. Estimating future cost is not a backward search frontier. DFS and BFS use depth-related removal orders, while hill climbing improves neighboring complete candidates. Replacing the mistaken backward-search phrase preserves the intended heuristic-and-cost algorithm identification.

Why each choice is right or wrong

A - Incorrect: DFS follows a stack order and does not inherently combine a heuristic with backward search.

B - Incorrect: BFS uses breadth order without requiring heuristics or a backward frontier.

C - Correct: A* is forward best-first search using f = g + h.

D - Incorrect: Hill climbing selects improving neighbors rather than ordering a frontier by accumulated plus remaining path cost.

Primary note proof

Study-note section AI18

> A* orders by f(n)=g(n)+h(n). Admissibility means 0≤h(n)≤h*(n), the true cheapest remaining cost. Consistency means h(n)≤c(n,n′)+h(n′) on every edge, with h(goal)=0; it implies admissibility on states with goal paths. Consistency makes f nondecreasing along a path.

Study-note section AI13

> BFS removes the oldest frontier entry using a FIFO queue. Start with the root, remove the head, test it, append unseen successors at the tail, and repeat. It explores depths 0,1,2,... and is complete for finite b and a finite-depth goal. It minimizes number of edges, so minimum cost requires equal positive step costs.

Study-note section AI14

DFS follows the deepest pending branch using a LIFO stack or recursion. FILO means the same order as LIFO. With the same children as BFS, a left-first DFS visits S,A,C before B; push rightmost children first to pop the leftmost first.

Study-note section AI20

Local search keeps a complete candidate configuration and improves its objective; the path may be irrelevant. Hill climbing chooses better neighbors. First-improvement accepts the first improvement; steepest-ascent examines all neighbors; stochastic variants choose among improvements. Local maxima, plateaus and ridges can stop progress. Random restarts or limited sideways steps can help without guaranteeing success.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 26 | lines 25-27

• A* search is the most commonly known form of best-first search.
• It uses heuristic function h(n), and cost to reach the node n from the start state g(n). It has combined
features of UCS and greedy best-first search, by which it solves the problem efficiently.

Introduction to Artificial Intelligence module | PDF page 27 | lines 20-25

o Place the starting node in the OPEN list.
o Check if the OPEN list is empty or not, if the list is empty then return failure and stops.
o Select the node from the OPEN list which has the smallest value of evaluation function (g+h), if node n is goal node then return success and stop, otherwise
o Expand node n and generate all of its successors, and put n into the closed list. For each successor n', check whether n' is already in the OPEN or CLOSED list, if not then compute

Introduction to Artificial Intelligence module | PDF page 23 | lines 17-19

forward-search and other from goal node called as backward-search, to find the goal node.
• Bidirectional search replaces one single search graph with two small sub-graphs in which one starts
the search from an initial vertex and other starts from goal vertex.

Question sources

Computer Science MCQ practice collection · Question 648

Question 166 | AM0649

What is the main advantage of using heuristics in game playing AI systems?

A. They guarantee victory in every game.

B. They make the game unpredictable and exciting.

C. They allow the AI to learn from its opponents.

D. They improve the AI's decision-making and gameplay.

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer D: They improve the AI's decision-making and gameplay.

Game search computes candidate moves with a heuristic and chooses the best backed-up outcome, supporting improved decision quality rather than automatic victory or learning. Evidence: PDF 33, 34.

Why each choice is right or wrong

A - Incorrect: Optimal opponent may prevent a win. See PDF p. 33, 34.

B - Incorrect: Unpredictability is not the stated purpose. See PDF p. 33, 34.

C - Incorrect: Heuristic evaluation is not necessarily learning. See PDF p. 33, 34.

D - Correct: Matches heuristic move evaluation and selection. See PDF p. 33, 34.

Primary note proof

Study-note section AI23

> At a depth cutoff use an evaluation function estimating position desirability. Positive favors the designated MAX player, negative MIN; the scale is a convention, with +1/0/−1 often win/draw/loss. Material counts weighted pieces; mobility counts legal options; territory scores controlled regions. These are different features, often combined. Heuristics improve limited computation, not guarantee victory.

Study-note section AI22

> Adversarial search models agents with competing objectives. A game specifies states, player to move, legal actions, successor function, terminal test and utility. Perfect information means previous moves and current state are visible; it does not reveal future choices or imply absence of chance. Zero-sum utility makes one player’s gain the other’s loss.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 33 | lines 35-35

> • Compute the move using a heuristic function.

Introduction to Artificial Intelligence module | PDF page 34 | lines 3-4

> • Make the best move from the choices.
> • For example, in the above figure, the two players MAX and MIN are there. MAX starts the game

Question sources

Computer Science MCQ practice collection · Question 649

Question 167 | AM0650

Which component is the scoring mapping that assigns a heuristic desirability value to a game state, rather than a search, pruning or sampling procedure?

A. Minimax algorithm

B. Alpha-beta pruning

C. Monte Carlo simulation

D. Evaluation function

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer D: Evaluation function

An evaluation function assigns a desirability score to a game state, especially when search stops at a depth cutoff. Minimax uses scores to choose moves, and alpha-beta avoids branches that cannot alter the backed-up result. Monte Carlo sampling can estimate values from possible continuations, but the sampling procedure is distinct from the mapping whose output is a state score.

Why each choice is right or wrong

A - Incorrect: Minimax backs values up through competing choices; it uses state scores rather than being the score mapping itself.

B - Incorrect: Alpha-beta pruning skips branches that cannot affect a minimax decision; it is not the scoring mapping.

C - Incorrect: Monte Carlo simulation samples continuations whose outcomes may estimate value; sampling is a procedure, not the score mapping itself.

D - Correct: An evaluation function maps a state to its estimated desirability score.

Primary note proof

Study-note section AI23

> At a depth cutoff use an evaluation function estimating position desirability. Positive favors the designated MAX player, negative MIN; the scale is a convention, with +1/0/−1 often win/draw/loss. Material counts weighted pieces; mobility counts legal options; territory scores controlled regions. These are different features, often combined. Heuristics improve limited computation, not guarantee victory.

Study-note section AI23

> Alpha is MAX’s best lower bound; beta is MIN’s best upper bound. Pass bounds down, update alpha at MAX and beta at MIN, and prune when alpha≥beta. Pruning preserves the minimax value for the same searched tree. Good move ordering increases pruning: worst-case O(b^m), ideal ordering about O(b^(m/2)).

Study-note section AI22

Adversarial search models agents with competing objectives. A game specifies states, player to move, legal actions, successor function, terminal test and utility. Perfect information means previous moves and current state are visible; it does not reveal future choices or imply absence of chance. Zero-sum utility makes one player's gain the other's loss.

Study-note section AI-S05

A Monte Carlo game simulation samples possible continuations and uses their outcomes to estimate position value. An evaluation function is the scoring mapping itself; sampling is one way of obtaining an estimate, while minimax backs values up and alpha-beta removes branches that cannot affect that backed-up decision.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 33 | lines 6-8

• Games are modeled as a Search problem and heuristic evaluation function, and these are the two
main factors which help to model and solve games in AI.
1. Minimax algorithm

Question sources

Computer Science MCQ practice collection · Question 650

Question 168 | AM0651

What is the purpose of using heuristics in game tree search algorithms?

A. To guarantee a win in every game.

B. To eliminate the need for search algorithms.

C. To reduce the search space and improve efficiency.

D. To ensure a systematic exploration of the search tree.

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer C: To reduce the search space and improve efficiency.

Pruning removes nodes that cannot affect the decision and speeds game search; heuristic domain-informed ordering improves pruning effectiveness. This supports efficiency, not guaranteed wins or eliminating search. Evidence: PDF 35, 35.

Why each choice is right or wrong

A - Incorrect: A better search can still lose to stronger positions. See PDF p. 35, 35.

B - Incorrect: Pruning/search guidance still requires search. See PDF p. 35, 35.

C - Correct: Matches reducing irrelevant exploration. See PDF p. 35, 35.

D - Incorrect: Systematic traversal does not describe the efficiency purpose. See PDF p. 35, 35.

Primary note proof

Study-note section AI23

> Alpha is MAX's best lower bound; beta is MIN's best upper bound. Pass bounds down, update alpha at MAX and beta at MIN, and prune when alpha≥beta. Pruning preserves the minimax value for the same searched tree. Good move ordering increases pruning: worst-case O(b^m), ideal ordering about O(b^(m/2)).

Study-note section AI23

> At a depth cutoff use an evaluation function estimating position desirability. Positive favors the designated MAX player, negative MIN; the scale is a convention, with +1/0/−1 often win/draw/loss. Material counts weighted pieces; mobility counts legal options; territory scores controlled regions. These are different features, often combined. Heuristics improve limited computation, not guarantee victory.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 35 | lines 8-10

> • It removes all the nodes which are not really affecting the final decision but making algorithm slow.
> Hence by pruning these nodes, it makes the algorithm fast.
> • The main condition which required for alpha-beta pruning is α>=β

Introduction to Artificial Intelligence module | PDF page 35 | lines 32-33

> • Use domain knowledge while finding the best move.
> • We can bookkeep the states, as there is a possibility that states may repeat.

Question sources

Computer Science MCQ practice collection · Question 651

Question 169 | AM0652

Which game-state heuristic evaluates the weighted number and value of pieces held by each player, rather than legal move count or controlled territory?

A. Material heuristic

B. Mobility heuristic

C. Distance heuristic

D. Tactical heuristic

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer A: Material heuristic

Material evaluation counts pieces with weights reflecting their values. Mobility instead measures legal options, while territory evaluates controlled regions. Restricting the stem to weighted pieces removes the original conflation of material and territory. Distance and tactical features may inform other evaluations, but they do not name the weighted piece-count feature requested. The distinction concerns what the evaluator measures.

Why each choice is right or wrong

A - Correct: Material is the weighted piece-count/value feature.

B - Incorrect: Mobility measures available legal moves rather than the amount of material or territory controlled.

C - Incorrect: Distance evaluates spatial separation, not the piece-count or controlled-region feature stated here.

D - Incorrect: This does not name weighted piece holdings; the feature described is material evaluation.

Primary note proof

Study-note section AI23

> At a depth cutoff use an evaluation function estimating position desirability. Positive favors the designated MAX player, negative MIN; the scale is a convention, with +1/0/−1 often win/draw/loss. Material counts weighted pieces; mobility counts legal options; territory scores controlled regions. These are different features, often combined. Heuristics improve limited computation, not guarantee victory.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 33 | lines 6-8

> • Games are modeled as a Search problem and heuristic evaluation function, and these are the two
> main factors which help to model and solve games in AI.
> 1. Minimax algorithm

Question sources

Computer Science MCQ practice collection · Question 652

Question 170 | AM0653

Which heuristic evaluates a game state by considering the available moves and options for each player?

A. Material heuristic

B. Mobility heuristic

C. Distance heuristic

D. Tactical heuristic

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer B: The mobility heuristic.

AI23 explicitly says mobility counts legal options. Available moves and options for each player are exactly that feature, making mobility the supported answer. Material measures weighted pieces, while distance and tactics refer to different aspects of a position. The heuristic may be combined with other features, but the specific feature requested is unambiguous.

Why each choice is right or wrong

A - Incorrect: Material evaluates the amount or value of pieces, not the number of available moves.

B - Correct: Mobility directly counts or scores the legal move options available to a player.

C - Incorrect: Distance concerns separation or proximity, not the legal-option count described in the question.

D - Incorrect: Tactical evaluation concerns concrete threats or sequences rather than the specific mobility feature requested.

Primary note proof

Study-note section AI23

> At a depth cutoff use an evaluation function estimating position desirability. Positive favors the designated MAX player, negative MIN; the scale is a convention, with +1/0/−1 often win/draw/loss. Material counts weighted pieces; mobility counts legal options; territory scores controlled regions. These are different features, often combined. Heuristics improve limited computation, not guarantee victory.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 33 | lines 6-8

> • Games are modeled as a Search problem and heuristic evaluation function, and these are the two

main factors which help to model and solve games in AI.
1. Minimax algorithm

Question sources

Computer Science MCQ practice collection · Question 653

Question 171 | CS0200

_____ is a set of procedures that uses representational language to conclude new facts from existing ones.

A. Inference Engine

B. Explanatory facility

C. Interface

D. Knowledge Base

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer A: Inference Engine

Inference engine manipulates knowledge and repeatedly applies rules to facts to derive new knowledge. The knowledge base stores facts, while the user interface handles interaction/explanation. Evidence: PDF 63, 63, 64.

Why each choice is right or wrong

A - Correct: Matches procedural derivation of new facts. See PDF p. 63, 63, 64.

B - Incorrect: Explaining an outcome differs from deriving facts. See PDF p. 63, 63, 64.

C - Incorrect: Interaction layer differs from inference procedures. See PDF p. 63, 63, 64.

D - Incorrect: Stores knowledge rather than the inference procedure. See PDF p. 63, 63, 64.

Primary note proof

Study-note section AI40

An expert system applies specialized knowledge to a domain such as equipment diagnosis. The domain expert supplies expertise; the knowledge engineer elicits and encodes it; the inference engine reasons over a knowledge base and working facts. The user interface handles interaction; an explanation facility presents why or how a conclusion was reached.

Study-note section AI24

Knowledge representation encodes facts, objects, relations, events and rules so a machine can use them. Syntax defines legal expressions; semantics defines their meaning. A knowledge base

stores assertions; an inference system derives consequences. Declarative knowledge says what is true; procedural knowledge specifies how to act.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 63 | lines 17-19

• In case of knowledge-based ES, the Inference Engine acquires and manipulates the knowledge from
the knowledge base to arrive at a particular solution.
• In case of rule-based ES, it –

Introduction to Artificial Intelligence module | PDF page 63 | lines 20-22

o Applies rules repeatedly to the facts, which are obtained from earlier rule application.
o Adds new knowledge into the knowledge base if required.
o Resolves rules conflict when multiple rules are applicable to a particular case

Introduction to Artificial Intelligence module | PDF page 64 | lines 20-22

• User interface provides interaction between user of the ES and the ES itself.
• Natural Language Processing so as to be used by the user who is well-versed in the task domain
• It explains how the ES has arrived at a particular recommendation.

Question sources

Computer Science Exit Exam practice collection · Question 200

Question 172 | CS0201

Which one is a space where an agent perceives and acts, so that an intelligent agent designer has to understand the type of it?

A. Environment

B. Actuator

C. Sensor

D. Search space

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer A: Environment

Environment is the surrounding space in which an agent operates and senses/acts; sensors and actuators are interfaces to it. Evidence: PDF 13.

Why each choice is right or wrong

A - Correct: Exact surrounding operating context. See PDF p. 13.

B - Incorrect: Action-output component. See PDF p. 13.

C - Incorrect: Perceptual-input component. See PDF p. 13.

D - Incorrect: Abstract candidate-solution set, not necessarily the operating surroundings. See PDF p. 13.

Primary note proof

Study-note section AI03

> An agent perceives an environment through sensors and acts through actuators or effectors. Cameras sense; motors move. Software sensors may be incoming messages and software actions may be API requests. Humans are agents but are not artificial agents.

Study-note section AI03

> A percept is one observation; a percept sequence is the entire observation history. An agent function maps percept sequences to actions: f:P*→A. An agent program implements that mapping on an architecture; agent = architecture + program. Environment means the setting in which it senses and acts, not merely the search space. Autonomy means relying substantially on experience rather than only a designer's prior instructions.

Study-note section AI12

> A search node records a state plus parent, action, depth and path cost. Expanding a node generates its successors. OPEN or frontier stores discovered candidates awaiting expansion; CLOSED records expanded states. Reaching the same state through different paths does not imply the same node or cost.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 13 | lines 4-6

> • The environment is where agent lives, operate and provide the agent with something to sense and
> act upon it.
> • The agent takes input from the environment through sensors and delivers the output to the

Question sources

Computer Science Exit Exam practice collection · Question 201

Question 173 | CS0202

Which one of the following is an uninformed searching method?

A. A* search

B. Iterative improvement

C. Breadth first search

D. Greedy search

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer C: Breadth first search

The listed uninformed strategies include BFS. A* and best-first are separately listed informed strategies; local iterative improvement is not the cited uninformed breadthwise procedure. Evidence: PDF 17, 17.

Why each choice is right or wrong

A - Incorrect: Informed heuristic method. See PDF p. 17, 17.

B - Incorrect: Local improvement is not the listed uninformed method. See PDF p. 17, 17.

C - Correct: Explicit uninformed method. See PDF p. 17, 17.

D - Incorrect: Informed heuristic best-first method. See PDF p. 17, 17.

Primary note proof

Study-note section AI13

> BFS removes the oldest frontier entry using a FIFO queue. Start with the root, remove the head, test it, append unseen successors at the tail, and repeat. It explores depths 0,1,2,... and is complete for finite b and a finite-depth goal. It minimizes number of edges, so minimum cost requires equal positive step costs.

Study-note section AI17

> A heuristic h(n) estimates remaining cost or evaluates promise using problem-specific knowledge. It can reduce explored work, but computing it also costs time. Uninformed algorithms use only the problem definition and known costs; UCS remains uninformed despite using g.

Study-note section AI20

> Local search keeps a complete candidate configuration and improves its objective; the path may be irrelevant. Hill climbing chooses better neighbors. First-improvement accepts the first improvement; steepest-ascent examines all neighbors; stochastic variants choose among improvements. Local maxima, plateaus and ridges can stop progress. Random restarts or limited sideways steps can help without guaranteeing success.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 17 | lines 35-37

> o Unformed Search Strategies
> - DFS, BFS, Uniform cost search, DLS, Iterative Deeping DFS, and Bi-directional search.

Introduction to Artificial Intelligence module | PDF page 17 | lines 38-41

o Informed Search Strategies
- A star search, Best first search
o Local Search Strategies
- Hill climbing search, Means ends analysis, Simulated annealing

Question sources

Computer Science Exit Exam practice collection · Question 202

Question 174 | CS0203

Which one of the following searching strategies best explains: 'The searching agent is based on problem-specific knowledge and the agent knows how far away a different state is from the goal'?

A. Blind search

B. Informed search

C. Uninformed search

D. Not informed search

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer B: Informed search

Problem-specific knowledge of estimated distance to the goal is the module's informed/heuristic-search distinction. Blind search lacks such guidance. Evidence: PDF 24, 18.

Why each choice is right or wrong

A - Incorrect: Blind search omits goal-closeness knowledge. See PDF p. 24, 18.

B - Correct: Exact informed-search distinction. See PDF p. 24, 18.

C - Incorrect: Lacks the stated domain guidance. See PDF p. 24, 18.

D - Incorrect: Synonymous with uninformed/blind search. See PDF p. 24, 18.

Primary note proof

Study-note section AI17

A heuristic h(n) estimates remaining cost or evaluates promise using problem-specific knowledge. It can reduce explored work, but computing it also costs time. Uninformed algorithms use only the problem definition and known costs; UCS remains uninformed despite using g.

Module passages - related reading

These source passages provide context; the primary note proof is identified separately above.

Introduction to Artificial Intelligence module | PDF page 24 | lines 14-16

• It uses domain knowledge. In an informed search, problem information is available which can guide
the search.
• Informed search strategies can find a solution more efficiently than an uninformed search strategy.

Introduction to Artificial Intelligence module | PDF page 18 | lines 23-25

• The uninformed search does not contain any domain knowledge such as closeness, the location of
the goal. It is also called level order traversal.
• It operates in a brute-force way as it only includes information about how to traverse the tree and

Question sources

Computer Science Exit Exam practice collection · Question 203

Question 175 | NS0001

Which type of AI is capable of performing any intellectual task that a human can do?

A. Narrow AI

B. General AI

C. Machine Learning

D. Expert Systems

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer B: General AI.

AI01 defines general AI by broad competence transferable across domains, matching the hypothetical ability to perform any human intellectual task. Narrow AI is bounded; machine learning names a family of methods and expert systems supply specialized domain knowledge. Neither automatically gives universal intellectual capability. The statement describes a category, not a claim that such a system presently exists.

Why each choice is right or wrong

A - Incorrect: Narrow AI is limited to bounded tasks rather than broad intellectual transfer.

B - Correct: General AI is the broad cross-domain competence category described.

C - Incorrect: Machine learning is a method family and does not itself imply general intellectual competence.

D - Incorrect: Expert systems apply specialized domain knowledge, not arbitrary human intellectual capability.

Primary note proof

Study-note section AI01

> Narrow or weak AI addresses bounded tasks, such as face recognition, spam filtering or a voice assistant. General AI denotes broad competence transferable across domains; strong AI also has a philosophical meaning involving actual understanding. Do not equate a humanoid body, autonomous vehicle, or many features with general intelligence.

Study-note section AI08

> Learning is improvement on task T, measured by P, through experience E. For spam filtering, T is classifying email, E is past labeled messages, and P might be recall at a required precision. Exploration trades immediate reward for information; exploitation uses the best action currently known. Learning can improve reflex, goal-based or utility-based agents.

Study-note section AI40

> An expert system applies specialized knowledge to a domain such as equipment diagnosis. The domain expert supplies expertise; the knowledge engineer elicits and encodes it; the inference engine reasons over a knowledge base and working facts. The user interface handles interaction; an explanation facility presents why or how a conclusion was reached.

Question sources

Artificial Intelligence review questions · Question 1

Question 176 | NS0002

What is the primary goal of Machine Learning?

A. To hard-code rules into a machine

B. To allow systems to learn and improve from experience without being explicitly programmed

C. To simulate physical environments

D. To build robotic hardware

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer B: Improve task performance through experience.

AI08 defines learning as improvement on a task, measured by performance, through experience. B captures that adaptive objective; without explicitly programmed means the learned behavior is not simply an exhaustive hand-coded rule list, not that no algorithm is programmed. Simulation and robot construction can use ML but do not define its primary goal.

Why each choice is right or wrong

A - Incorrect: Hard-coding all task rules is not the experience-based improvement defining machine learning.

B - Correct: The system improves measured task performance from experience rather than only fixed task rules.

C - Incorrect: Physical simulation is one application context, not the definition or primary learning objective.

D - Incorrect: Building hardware is not required for machine learning, which also operates in software.

Primary note proof

Study-note section AI08

> Learning is improvement on task T, measured by P, through experience E. For spam filtering, T is classifying email, E is past labeled messages, and P might be recall at a required precision. Exploration trades immediate reward for information; exploitation uses the best action currently known. Learning can improve reflex, goal-based or utility-based agents.

Study-note section AI-S01

> A conventional computer manipulates digital information according to a program or sequence of instructions. Programs can implement storage, retrieval and learning operations. Neural networks can learn pattern recognition from examples, so executing instructions and learning patterns are not mutually exclusive descriptions. This provides no universal claim that every machine is better or worse than every human at completing missing information.

Question sources

Artificial Intelligence review questions · Question 2

Question 177 | NS0003

Which reasoning type guarantees that if premises are true, the conclusion is also true?

A. Inductive Reasoning

B. Abductive Reasoning

C. Deductive Reasoning

D. Probabilistic Reasoning

Hint

Translate the statement carefully, then check the scope of its quantifiers, connectives or inference rule. Distinguish guaranteed consequences from plausible generalizations or explanations.

Answer C: Deductive reasoning.

AI28 defines deduction by truth-preserving rules and contrasts induction's revisable generalizations and abduction's plausible explanations. With valid deductive inference and true premises, the conclusion must be true. Probabilistic support expresses uncertainty rather than this logical guarantee, so deductive is the relevant reasoning type.

Why each choice is right or wrong

A - Incorrect: Induction generalizes beyond observations and can be overturned by new examples.

B - Incorrect: Abduction proposes plausible explanations without guaranteeing that they are true.

C - Correct: Valid deductive rules preserve truth from true premises to the conclusion.

D - Incorrect: Probability can quantify uncertain conclusions rather than provide the stated deductive guarantee.

Primary note proof

Study-note section AI28

> Deduction applies truth-preserving rules: all people are mortal; Ada is a person; therefore Ada is mortal. Induction generalizes observations to a broader rule and can be overturned by new examples. Abduction proposes a plausible explanation: wet ground suggests rain, although sprinklers could explain it. Analogy transfers a pattern between similar situations.

Study-note section AI29

> Probability describes uncertainty about events. $P(\neg A)=1-P(A)$; $P(A|B)=P(A\cap B)/P(B)$, for $P(B)>0$. Bayes gives $P(H|E)=P(E|H)P(H)/P(E)$: posterior equals likelihood times prior divided by evidence probability. For mutually exclusive exhaustive hypotheses, $P(E)=\Sigma i P(E|Hi)P(Hi)$.

Question sources

Artificial Intelligence review questions · Question 3

Question 178 | NS0004

A self-driving car's environment is best classified as:

A. Fully Observable, Deterministic, Static

B. Partially Observable, Stochastic, Dynamic

C. Fully Observable, Stochastic, Static

D. Partially Observable, Deterministic, Dynamic

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer B: Partially observable, stochastic and dynamic.

AI09 describes driving as partially observable, sequential and dynamic. Traffic participants and noisy observations make outcomes uncertain rather than fully determined by the car's chosen action. B matches that ordinary driving setting. Neither full visibility nor a static world fits roads whose relevant conditions change while the agent deliberates.

Why each choice is right or wrong

A - Incorrect: Driving has hidden/noisy information and a changing world rather than full observability and static conditions.

B - Correct: Hidden state, uncertain outcomes and ongoing traffic changes fit the ordinary driving task.

C - Incorrect: Full observability and static conditions do not fit the stated driving environment.

D - Incorrect: Partial observation and dynamics fit, but deterministic outcomes ignore the uncertainty in ordinary driving.

Primary note proof

Study-note section AI09

> Fully observable means sensors reveal relevant state; partially observable means relevant information is hidden or noisy. Deterministic means state and action determine the next state; stochastic means outcomes involve randomness. Strategic uncertainty about an opponent is distinct from dice-like randomness.

Study-note section AI09

> Episodic decisions are independent; sequential decisions affect later situations. Static worlds do not change while deliberating; dynamic ones do. A clock affecting score makes an otherwise static game semidynamic. Discrete quantities have distinct alternatives; continuous quantities vary along real-valued ranges. Known versus unknown concerns knowledge of the transition model, not observability.

Question sources

Artificial Intelligence review questions · Question 4

Question 179 | NS0005

In the A* algorithm, what does the evaluation function f(n) represent?

A. Only the heuristic cost to goal

B. Only the cost from start to node n

C. The sum of cost from start to n (g(n)) and estimated cost from n to goal (h(n))

D. The product of g(n) and h(n)

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer C: f(n)=g(n)+h(n).

AI18 directly defines A* evaluation as actual start-to-node cost g plus estimated node-to-goal cost h. This estimates a solution cost through the node. Using only h is greedy best-first, using only g is uniform-cost selection, and multiplication is not the A* rule.

Why each choice is right or wrong

A - Incorrect: Using h alone is greedy best-first rather than A* evaluation.

B - Incorrect: Using g alone is uniform-cost ordering and omits the remaining-cost estimate.

C - Correct: A* adds the accumulated path cost and estimated remaining cost.

D - Incorrect: The defined evaluation is a sum, not a product.

Primary note proof

Study-note section AI18

> A* orders by f(n)=g(n)+h(n). Admissibility means 0≤h(n)≤h*(n), the true cheapest remaining cost. Consistency means h(n)≤c(n,n′)+h(n′) on every edge, with h(goal)=0; it implies admissibility on states with goal paths. Consistency makes f nondecreasing along a path.

Study-note section AI15

> UCS uses a priority queue ordered by actual accumulated path cost g(n). Remove a minimum-g node; accept a goal when removed, not when first generated. Relax each edge: if a cheaper path is found, update its frontier cost and parent.

Study-note section AI17

> Best-first is a family that orders OPEN by an evaluation function. Greedy best-first uses f=h and removes the smallest estimate, ignoring the cost already paid. It can choose an expensive route, loop without duplicate checks, or follow an infinite misleading branch. On a finite graph with full duplicate handling it is complete, but not generally optimal. Its implicit finite-tree worst case can be O(b^m).

Question sources

Artificial Intelligence review questions · Question 5

Question 180 | NS0006

Which reflex architecture maintains an internal world state and applies condition-action rules to that modeled state, rather than using only the current percept?

A. Simple Reflex Agent

B. Goal-Based Agent

C. Model-Based Reflex Agent

D. Learning Agent

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer C: Model-Based Reflex Agent

Model-based reflex agents update internal state from previous state, action and current observation, using models of world change and action effects. They then apply rules to that modeled state. Simple reflex agents lack persistent world-state memory. Goal-based and learning capabilities may coexist with a model, but the question specifically asks for the reflex architecture defined by state maintenance and rule application.

Why each choice is right or wrong

A - Incorrect: Simple reflex uses only the current percept without persistent world-state memory.

B - Incorrect: Goal-based choice uses future sequences reaching desired states, whereas the requested architecture applies reflex rules to modeled state.

C - Correct: Model-based reflex maintains internal state and applies rules to that state.

D - Incorrect: Learning concerns improving the agent from experience rather than naming this reflex architecture.

Primary note proof

Study-note section AI06

> A simple reflex agent matches condition-action rules against the current percept only: if dirty, clean. It has no persistent world-state memory. This is adequate when the current observation contains everything needed for its rule, but partial observability can make the same percept require different actions.

Study-note section AI06

> A model-based reflex agent maintains internal state using the previous state, previous action, current percept, and models of world change and action effects. It then applies rules to that state. Both kinds can use condition-action rules; memory and a world model distinguish them.

Study-note section AI07

A goal-based agent considers future action sequences that reach a desired state. A utility-based agent compares degrees of desirability and trade-offs, choosing maximum expected utility: $EU(a)=\Sigma s\ P(s|a,evidence)U(s)$. Reaching the destination is a goal; balancing arrival time, safety and cost requires preferences among outcomes.

Study-note section AI08

The performance element selects actions. The learning element changes the agent to improve future performance. The critic evaluates behavior against a performance standard and supplies feedback. The problem generator suggests exploratory actions likely to yield useful experience.

Question sources

Artificial Intelligence review questions · Question 6

Question 181 | NS0007

What distinguishes a Rational Agent from an Omniscient Agent?

A. Rational agents always make perfect decisions

B. Rational agents act on best available information but can make mistakes; omniscient agents know all outcomes

C. Omniscient agents are a common AI type

D. Rational agents do not use information

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer B: Rational choice uses available information; omniscience knows actual outcomes.

AI04 defines rationality by best expected performance using information available, while omniscience means knowing actual outcomes in advance. Rational actions can fail and do not require perfect predictions. The contrast is epistemic, not whether an agent uses any information or whether omniscient systems are common implementations.

Why each choice is right or wrong

A - Incorrect: Rationality does not guarantee perfect outcomes or foreknowledge.

B - Correct: This distinguishes information-conditioned expected performance from knowing actual outcomes.

C - Incorrect: The note treats omniscience as an idealized outcome-knowledge condition, not a common AI implementation.

D - Incorrect: Rationality explicitly depends on percepts, prior knowledge and available actions.

Primary note proof

Study-note section AI04

> A performance measure defines success externally: safety, useful deliveries, elapsed time, or cost. Rationality depends on that measure, prior knowledge, percept history, and available actions. A rational agent selects the action with the best expected performance given its information.

Study-note section AI04

> Omniscience would mean knowing actual action outcomes in advance. Rationality does not require omniscience or guaranteed success. Even complete outcome knowledge cannot create a successful action when every available action fails. Information gathering and exploration may therefore be rational actions.

Question sources

Artificial Intelligence review questions · Question 7

Question 182 | NS0008

Which search algorithm is guaranteed to find the shortest path in an unweighted graph?

A. Depth-First Search

B. Depth-Limited Search

C. Breadth-First Search

D. Best-First Search

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer C: Breadth-first search.

AI13 says BFS explores levels in increasing depth and minimizes the number of edges. An unweighted graph gives equal step costs, making that shortest-path property applicable with visited-state tracking. DFS or a fixed depth limit does not guarantee minimum depth, and generic best-first needs its evaluation rule specified.

Why each choice is right or wrong

A - Incorrect: DFS can reach a longer path first and is not generally optimal.

B - Incorrect: A depth limit can miss a solution or return a nonshortest path within its bound.

C - Correct: BFS explores increasing edge counts and finds shortest paths in an unweighted graph.

D - Incorrect: Best-first is a family; an unspecified evaluation rule does not guarantee shortest paths.

Primary note proof

Study-note section AI13

> BFS removes the oldest frontier entry using a FIFO queue. Start with the root, remove the head, test it, append unseen successors at the tail, and repeat. It explores depths 0,1,2,... and is complete for finite b and a finite-depth goal. It minimizes number of edges, so minimum cost requires equal positive step costs.

Study-note section AI14

> DFS follows the deepest pending branch using a LIFO stack or recursion. FILO means the same order as LIFO. With the same children as BFS, a left-first DFS visits S,A,C before B; push rightmost children first to pop the leftmost first.

Study-note section AI16

> DLS runs DFS to depth l and reports cutoff separately from failure. It misses goals deeper than l. Iterative deepening runs DLS with limits 0,1,2,...; for finite b it finds the shallowest solution with O(b^d) time and O(bd) frontier space, optimal for equal step costs. Repeated shallow work is small compared with the deepest level.

Question sources

Artificial Intelligence review questions · Question 8

Question 183 | NS0009

Why is Uniform Cost Search (UCS) NOT considered a heuristic approach?

A. It doesn't consider costs

B. It uses cost as a known factor but employs no domain-specific knowledge or estimates

C. It uses heuristic functions

D. It only works on trees, not graphs

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer B: UCS uses known accumulated cost without a heuristic estimate.

AI15 orders UCS by actual path cost g, and AI17 explicitly calls it uninformed despite using g. Known action costs come from the problem definition; a heuristic adds an estimate of future cost. UCS can operate on graphs with duplicate handling, so neither ignoring costs nor tree-only operation explains the distinction.

Why each choice is right or wrong

A - Incorrect: UCS explicitly orders its frontier by accumulated path cost.

B - Correct: Known costs are used, but no problem-specific estimate of remaining cost is required.

C - Incorrect: UCS does not require an h estimate; g is actual accumulated cost.

D - Incorrect: UCS works on graphs with proper relaxation and duplicate handling.

Primary note proof

Study-note section AI15

> UCS uses a priority queue ordered by actual accumulated path cost g(n). Remove a minimum-g node; accept a goal when removed, not when first generated. Relax each edge: if a cheaper path is found, update its frontier cost and parent.

Study-note section AI17

> A heuristic h(n) estimates remaining cost or evaluates promise using problem-specific knowledge. It can reduce explored work, but computing it also costs time. Uninformed algorithms use only the problem definition and known costs; UCS remains uninformed despite using g.

Question sources

Artificial Intelligence review questions · Question 9

Question 184 | NS0010

Which data structure supplies the defining next-node removal order of a depth-first search frontier?

A. Queue (FIFO)

B. Stack (LIFO)

C. Priority Queue

D. Hash Table

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer B: Stack (LIFO)

DFS follows the deepest pending branch using a LIFO stack or recursion. That removal discipline distinguishes it from breadth-first FIFO selection and cost- or heuristic-priority selection. A hash table can still be used alongside DFS to record visited states. The question therefore specifies frontier removal rather than broadly asking which structures may track nodes.

Why each choice is right or wrong

A - Incorrect: A FIFO queue removes the oldest entry, giving breadth-first rather than depth-first frontier order.

B - Correct: A LIFO stack processes the most recently pending branch first, implementing DFS.

C - Incorrect: A priority queue removes according to a key rather than the defining DFS LIFO order.

D - Incorrect: A hash table may store visited states, but does not supply the frontier's next-node removal order.

Primary note proof

Study-note section AI14

> DFS follows the deepest pending branch using a LIFO stack or recursion. FILO means the same order as LIFO. With the same children as BFS, a left-first DFS visits S,A,C before B; push rightmost children first to pop the leftmost first.

Study-note section AI13

> BFS removes the oldest frontier entry using a FIFO queue. Start with the root, remove the head, test it, append unseen successors at the tail, and repeat. It explores depths 0,1,2,... and is complete for finite b and a finite-depth goal. It minimizes number of edges, so minimum cost requires equal positive step costs.

Study-note section AI15

UCS uses a priority queue ordered by actual accumulated path cost g(n). Remove a minimum-g node; accept a goal when removed, not when first generated. Relax each edge: if a cheaper path is found, update its frontier cost and parent.

Question sources

Artificial Intelligence review questions · Question 10

Question 185 | NS0011

What is the main limitation of Depth-First Search regarding completeness?

A. It uses too much memory

B. It can get stuck in cycles or infinite paths and fail to find a solution

C. It always finds suboptimal solutions

D. It only works for trees

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer B: Cycles or infinite branches can prevent finding an existing solution.

AI14 says DFS can fail on infinite paths or cycles and is not generally optimal. Without suitable restrictions or duplicate handling, it may never return to a pending branch containing a goal. Finite visited-state graph traversal changes that result; the limitation is therefore not that DFS only works on trees or always returns a bad path.

Why each choice is right or wrong

A - Incorrect: DFS is usually memory-efficient; excess storage is not its defining completeness failure.

B - Correct: It may follow an infinite branch or cycle indefinitely while a solution waits elsewhere.

C - Incorrect: A returned path can be optimal by coincidence; always suboptimal is too strong.

D - Incorrect: DFS also traverses graphs, especially with visited-state tracking.

Primary note proof

Study-note section AI14

> DFS follows the deepest pending branch using a LIFO stack or recursion. FILO means the same order as LIFO. With the same children as BFS, a left-first DFS visits S,A,C before B; push rightmost children first to pop the leftmost first.

Study-note section AI14

> Implicit-tree worst-case time is O(b^m), and space is O(bm) when pending siblings are stored, or O(m) for path-only recursive generation. DFS is not generally optimal and can fail on infinite paths or cycles. On a finite graph, visited-state tracking makes traversal complete but needs O(V) storage. Do not confuse the O(b^m) worst-case time bound with the DFS space bounds above.

Question sources

Artificial Intelligence review questions · Question 11

Question 186 | NS0012

Which options can be used without a training set of labeled input-target examples? Select all that apply. For D, consider whether some deep-learning objectives permit unlabeled or reward feedback.

A. Supervised Learning

B. Reinforcement Learning

C. Unsupervised Learning

D. Deep Learning

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer B, C, D: Reinforcement Learning; Unsupervised Learning; Deep Learning

Unsupervised learning discovers structure without target labels. Reinforcement learning can learn from rewards rather than correct-action target examples. Deep learning is an architecture family and can be trained with unsupervised or reinforcement objectives, so it also need not require labeled input-target data. Supervised learning is specifically defined by those labeled pairs. The capability wording permits both regimes and architecture families to be judged accurately.

Why each choice is right or wrong

A - Incorrect: Supervised learning uses labeled examples by definition.

B - Correct: RL learns through rewards and interaction, without requiring labeled correct actions.

C - Correct: Unsupervised learning explicitly seeks structure without target labels.

D - Correct: Deep learning can use unlabeled or reward objectives, so labels are not inherently required.

Primary note proof

Study-note section AI31

> Supervised learning fits a mapping from labeled input-target examples. Classification predicts discrete labels, such as institution generations or spam/nonspam; regression predicts numeric quantities, such as prices. Predictive analytics uses historical patterns to forecast outcomes and can use either. Specify the target before choosing a method.

Study-note section AI32

> Unsupervised learning seeks structure without target labels: clustering groups similar examples; association finds co-occurrence; dimensionality reduction compresses features. Good clusters have high within-cluster similarity and low between-cluster similarity under an appropriate distance. K-means alternates assignment to nearest centroids and centroid updates.

Study-note section AI33

An RL agent observes state information, takes an action, receives reward and a new observation, and improves its policy to maximize expected cumulative reward. Rewards evaluate outcomes rather than supply the correct action label. Rewards may be delayed; not every correct-looking action receives immediate reward.

Study-note section AI43

ANNs, genetic search and reinforcement learning all improve candidate behavior or performance through experience in different ways. Deep learning is an architecture family usable with supervised, unsupervised or reinforcement objectives. Decision trees and deep models may both be supervised. To distinguish learning regimes, specify labels, reward or unlabeled structure.

Question sources

Artificial Intelligence review questions · Question 12

Question 187 | NS0013

For A* tree search with finite branching, step costs bounded below by a positive constant, goal testing on removal and h(goal)=0, which listed heuristic condition is the weakest sufficient guarantee of optimality? Treat monotone as synonymous with consistent.

A. Consistent

B. Admissible (never overestimates the actual cost)

C. Monotone

D. Domain-independent

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer B: Admissible (never overestimates the actual cost)

Admissibility requires that a heuristic never overestimate the true cheapest remaining cost. It suffices for optimal A* tree search under the stated cost and goal-testing assumptions. Consistency, also called monotonicity here, is stronger and is particularly important when graph search permanently closes states. Domain independence is unrelated to the optimality condition. Asking for the weakest sufficient listed condition distinguishes admissibility from the stronger alternatives.

Why each choice is right or wrong

A - Incorrect: Consistency is sufficient but stronger than the weakest listed condition for tree A*.

B - Correct: Admissibility is sufficient for optimal tree A* under the stated assumptions.

C - Incorrect: Monotonicity is consistency here and is stronger than necessary for the tree-search guarantee.

D - Incorrect: Domain independence is not a heuristic optimality condition.

Primary note proof

Study-note section AI18

> A* orders by f(n)=g(n)+h(n). Admissibility means 0≤h(n)≤h*(n), the true cheapest remaining cost. Consistency means h(n)≤c(n,n′)+h(n′) on every edge, with h(goal)=0; it implies admissibility on states with goal paths. Consistency makes f nondecreasing along a path.

Study-note section AI18

> Admissibility suffices for optimal A* tree search; graph search that never reopens CLOSED states normally requires consistency. With an admissible inconsistent heuristic, reopen states when g improves. Test a goal when removed at minimum f. Completeness also needs finite branching and positive lower-bounded costs in infinite spaces. Worst-case time and memory can be exponential; no heuristic universally guarantees speed.

Question sources

Artificial Intelligence review questions · Question 13

Question 188 | NS0014

Iterative Deepening Search (IDS) combines the advantages of which two algorithms?

A. UCS and A*

B. BFS and DFS

C. Greedy Search and BFS

D. DLS and UCS

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer B: Breadth-first solution depth with depth-first memory use.

AI16 defines IDS as DLS runs at limits 0,1,2,..., finding the shallowest solution under finite branching with O(bd) frontier space. It obtains BFS's shallow-solution behavior and DFS-style memory use. It does not combine UCS priorities, greedy ordering or A* cost thresholds; IDA* is a separate threshold-based variant.

Why each choice is right or wrong

A - Incorrect: IDS does not combine UCS costs with A* evaluation.

B - Correct: It obtains breadth-first shallow-solution behavior through repeated depth-first limited searches.

C - Incorrect: No greedy heuristic ordering is required for IDS.

D - Incorrect: DLS is used, but UCS ordering is not part of standard iterative deepening.

Primary note proof

Study-note section AI16

DLS runs DFS to depth l and reports cutoff separately from failure. It misses goals deeper than l. Iterative deepening runs DLS with limits 0,1,2,...; for finite b it finds the shallowest solution with O(b^d) time and O(bd) frontier space, optimal for equal step costs. Repeated shallow work is small compared with the deepest level.

Study-note section AI13

BFS removes the oldest frontier entry using a FIFO queue. Start with the root, remove the head, test it, append unseen successors at the tail, and repeat. It explores depths 0,1,2,... and is complete for finite b and a finite-depth goal. It minimizes number of edges, so minimum cost requires equal positive step costs.

Study-note section AI14

DFS follows the deepest pending branch using a LIFO stack or recursion. FILO means the same order as LIFO. With the same children as BFS, a left-first DFS visits S,A,C before B; push rightmost children first to pop the leftmost first.

Question sources

Artificial Intelligence review questions · Question 14

Question 189 | NS0015

A zero-sum game evaluator measures a position from the designated MAX player's perspective. What does a positive evaluation value indicate?

A. A draw or neutral position.

B. An estimated advantage for the designated MAX player.

C. An estimated advantage for the opposing MIN player.

D. An invalid game state.

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer B: An estimated advantage for the designated MAX player.

Under this convention, positive values estimate an advantage for MAX, negative values favor MIN, and zero is neutral. A positive heuristic score does not guarantee an eventual win, because it may be only an estimate at a search cutoff. The reference player must be specified before a sign can be interpreted; player names alone do not determine the convention.

Why each choice is right or wrong

A - Incorrect: A neutral result is represented by zero under the stated sign convention.

B - Correct: A positive value favors the designated MAX player under this convention.

C - Incorrect: An advantage for MIN would have a negative sign under this convention.

D - Incorrect: The sign estimates desirability, not whether the position is legally valid.

Primary note proof

Study-note section AI23

> At a depth cutoff use an evaluation function estimating position desirability. Positive favors the designated MAX player, negative MIN; the scale is a convention, with +1/0/−1 often win/draw/loss. Material counts weighted pieces; mobility counts legal options; territory scores controlled regions. These are different features, often combined. Heuristics improve limited computation, not guarantee victory.

Question sources

Artificial Intelligence review questions · Question 15

Question 190 | NS0016

Which knowledge representation technique uses interconnected nodes and arcs to show concept relationships?

A. Frames

B. Logic

C. Semantic Networks

D. Probabilistic Models

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer C: Semantic networks.

AI24 defines semantic networks using concept nodes and labeled relations such as is-a. That matches the node/arc representation in the stem. Frames organize slots, logic uses formal sentences and probabilistic models attach uncertainty; those may have related graph forms, but the specific concept-relationship network term is semantic network.

Why each choice is right or wrong

A - Incorrect: Frames represent stereotyped entities through slots and values rather than this named concept-link network.

B - Incorrect: Logic represents claims with formal expressions, not specifically the node/arc scheme described.

C - Correct: Semantic networks connect concept nodes with labeled relationship arcs.

D - Incorrect: Probabilistic models focus on uncertainty rather than the specific semantic-relation representation requested.

Primary note proof

Study-note section AI24

> A semantic network connects concept nodes by labeled relations such as is-a. Frames describe stereotyped entities through slots, values and defaults. Production systems use IF conditions THEN conclusions or actions, matching rules against working memory and selecting among applicable rules. A graph suits relationships such as frequently bought together; frames suit attributes. No representation is always best for recommendation without task requirements.

Study-note section AI24

> Knowledge representation encodes facts, objects, relations, events and rules so a machine can use them. Syntax defines legal expressions; semantics defines their meaning. A knowledge base stores assertions; an inference system derives consequences. Declarative knowledge says what is true; procedural knowledge specifies how to act.

Study-note section AI29

A Bayesian network is a directed acyclic graph with conditional probability tables; its joint distribution factors as the product of each node's probability given its parents. A Markov model assumes the next state depends on the current state rather than the entire past. These represent uncertainty; fuzzy membership represents a different issue.

Question sources

Artificial Intelligence review questions · Question 16

Question 191 | NS0017

Which listed group comprises unstructured input modalities from which deep neural models can learn hierarchical features?

A. Tabular numerical data only

B. Structured relational data

C. Unstructured data like images, sound, and text

D. Small labeled datasets

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer C: Unstructured data like images, sound, and text

Deep learning uses multiple representation layers and can learn features from images, sound and text. Those are the unstructured modalities identified here. Deep models can also process other data, so this is not a claim of universal superiority over every alternative model. Structured relational data is a different organization, and small labeled datasets describe size and supervision rather than modality.

Why each choice is right or wrong

A - Incorrect: Tabular numerical data only excludes the non-tabular modalities requested and is not an inherent deep-learning restriction.

B - Incorrect: Structured relational data is not the unstructured modality group asked for.

C - Correct: Images, sound and text are explicit modalities for learned multilayer representations.

D - Incorrect: A small labeled dataset describes sample size and feedback, not an input modality group.

Primary note proof

Study-note section AI34

Feedforward networks pass signals toward outputs without recurrent state loops; recurrent networks carry state across sequence steps. Training adjusts weights, often through gradients and backpropagation. Backpropagation is not the same as a recurrent connection. Deep learning uses multiple representation layers and can learn features from images, sound and text. CNNs

share filters across image locations. ANN is a brain-inspired abstraction, not a literal brain simulation; DL⊂ML⊂AI.

Study-note section AI35

Overfitting means fitting training-specific details so performance drops on unseen data; underfitting means the model is too limited or poorly fitted even on training data. A learning curve plots performance against training size or training progress. Complexity control, regularization, pruning and more representative data can help.

Question sources

Artificial Intelligence review questions · Question 17

Question 192 | NS0018

Which agent type improves its performance over time through experience and feedback?

A. Simple Reflex Agent

B. Model-Based Reflex Agent

C. Goal-Based Agent

D. Learning Agent

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer D: A learning agent.

AI08 defines the learning element as changing the agent to improve future performance from experience and feedback. This identifies the learning-agent category. Reflex, model-based and goal-based describe decision architectures that may be enhanced by learning, but do not by themselves require improvement over time; D explicitly names that capability.

Why each choice is right or wrong

A - Incorrect: Simple reflex rules alone do not require adaptation from experience.

B - Incorrect: Maintaining a model does not by itself require learning or improving it over time.

C - Incorrect: Planning toward a goal does not itself establish experience-based policy improvement.

D - Correct: Learning agents explicitly use experience and evaluative feedback to improve future behavior.

Primary note proof

Study-note section AI08

> The performance element selects actions. The learning element changes the agent to improve future performance. The critic evaluates behavior against a performance standard and supplies feedback. The problem generator suggests exploratory actions likely to yield useful experience.

Study-note section AI08

> Learning is improvement on task T, measured by P, through experience E. For spam filtering, T is classifying email, E is past labeled messages, and P might be recall at a required precision. Exploration trades immediate reward for information; exploitation uses the best action currently known. Learning can improve reflex, goal-based or utility-based agents.

Question sources

Artificial Intelligence review questions · Question 18

Question 193 | NS0019

What is abductive reasoning best described as?

A. Drawing certain conclusions from known facts

B. Making educated guesses from patterns

C. Finding the best explanation for a set of observations

D. Using probability to decide outcomes

Hint

Translate the statement carefully, then check the scope of its quantifiers, connectives or inference rule. Distinguish guaranteed consequences from plausible generalizations or explanations.

Answer C: Finding a plausible best explanation for observations.

AI28 illustrates abduction with wet ground suggesting rain while allowing sprinklers as another cause. It proposes an explanatory hypothesis rather than guaranteeing truth or merely generalizing a pattern. C gives that explanatory role. Probability may rank hypotheses, but numerical probability alone is not the definition of abductive reasoning.

Why each choice is right or wrong

A - Incorrect: Certain truth-preserving conclusions describe valid deduction rather than explanatory abduction.

B - Incorrect: Pattern-based generalization is closer to induction and omits the explanatory-hypothesis requirement.

C - Correct: Abduction seeks a plausible explanation for observed evidence without guaranteeing that explanation.

D - Incorrect: Probability can support a diagnosis but does not itself define the explanatory reasoning form.

Primary note proof

Study-note section AI28

> Deduction applies truth-preserving rules: all people are mortal; Ada is a person; therefore Ada is mortal. Induction generalizes observations to a broader rule and can be overturned by new examples. Abduction proposes a plausible explanation: wet ground suggests rain, although sprinklers could explain it. Analogy transfers a pattern between similar situations.

Study-note section AI29

> Probability describes uncertainty about events. $P(\neg A)=1-P(A)$; $P(A|B)=P(A\cap B)/P(B)$, for $P(B)>0$. Bayes gives $P(H|E)=P(E|H)P(H)/P(E)$: posterior equals likelihood times prior divided by evidence probability. For mutually exclusive exhaustive hypotheses, $P(E)=\Sigma_i P(E|Hi)P(Hi)$.

Question sources

Artificial Intelligence review questions · Question 19

Question 194 | NS0020

How is ordinary untimed chess classified with respect to transition outcomes, change during deliberation and state/action values? Assume both players' moves are specified when evaluating a transition.

A. Stochastic, Dynamic, Continuous

B. Deterministic, Static, Discrete

C. Stochastic, Static, Discrete

D. Deterministic, Dynamic, Continuous

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer B: Deterministic, Static, Discrete

Chess has discrete positions and legal moves. Given the selected move, the resulting position is determined by the rules rather than random outcome sampling. In untimed play, the environment does not change while the player deliberates, so the relevant classification is static. Opponent choice is strategic uncertainty rather than dice-like stochasticity. Specifying untimed play excludes the semidynamic effect of a score-affecting clock.

Why each choice is right or wrong

A - Incorrect: Chess moves are discrete, and opponent strategy is not dice-like stochasticity.

B - Correct: Untimed chess is deterministic given chosen moves, static while deliberating and discrete.

C - Incorrect: Ordinary chess transition outcomes are deterministic given chosen moves rather than stochastic.

D - Incorrect: Chess is discrete rather than continuous, and a clock gives semidynamic rather than this classification.

Primary note proof

Study-note section AI09

> Fully observable means sensors reveal relevant state; partially observable means relevant information is hidden or noisy. Deterministic means state and action determine the next state; stochastic means outcomes involve randomness. Strategic uncertainty about an opponent is distinct from dice-like randomness.

Study-note section AI09

> Episodic decisions are independent; sequential decisions affect later situations. Static worlds do not change while deliberating; dynamic ones do. A clock affecting score makes an otherwise static game semidynamic. Discrete quantities have distinct alternatives; continuous quantities vary along real-valued ranges. Known versus unknown concerns knowledge of the transition model, not observability.

Question sources

Artificial Intelligence review questions · Question 20

Question 195 | NS0021

What does the 'deep' in Deep Learning refer to?

A. The depth of the training dataset

B. The depth of layers in a neural network

C. The complexity of the problem being solved

D. The deep analysis of heuristic functions

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer B: Multiple neural representation layers.

AI34 defines deep learning through multiple representation layers in a neural model. Dataset size and problem difficulty can affect training, but neither is what the word deep denotes. Heuristic analysis belongs to search methods rather than the architecture depth of a neural network.

Why each choice is right or wrong

A - Incorrect: Dataset size or depth does not define the deep part of deep learning.

B - Correct: Deep refers to multiple layers of learned representations in the neural architecture.

C - Incorrect: Problem difficulty does not determine whether a model architecture is deep.

D - Incorrect: Heuristic analysis is not the neural-layer definition of deep learning.

Primary note proof

Study-note section AI34

> Feedforward networks pass signals toward outputs without recurrent state loops; recurrent networks carry state across sequence steps. Training adjusts weights, often through gradients and backpropagation. Backpropagation is not the same as a recurrent connection. Deep learning uses multiple representation layers and can learn features from images, sound and text. CNNs share filters across image locations. ANN is a brain-inspired abstraction, not a literal brain simulation; DL⊂ML⊂AI.

Question sources

Artificial Intelligence review questions · Question 21

Question 196 | NS0022

Which search strategy is called 'blind search' because it uses no extra problem knowledge?

A. Informed Search

B. Heuristic Search

C. Uninformed Search

D. Best-First Search

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer C: Uninformed search.

AI17 says uninformed methods use the problem definition and known costs without a problem-specific remaining-cost estimate. This is the usual blind-search meaning. Informed or heuristic search explicitly adds such estimates, and best-first is a broad ordering family that may use heuristic information rather than naming the blind category.

Why each choice is right or wrong

A - Incorrect: Informed search adds problem-specific guidance beyond the basic problem definition.

B - Incorrect: Heuristic search uses the extra estimate excluded by the stem.

C - Correct: Uninformed search is the category using only the problem definition and known costs.

D - Incorrect: Best-first specifies evaluation-based ordering and is not generally synonymous with blind search.

Primary note proof

Study-note section AI17

> A heuristic h(n) estimates remaining cost or evaluates promise using problem-specific knowledge. It can reduce explored work, but computing it also costs time. Uninformed algorithms use only the problem definition and known costs; UCS remains uninformed despite using g.

Study-note section AI15

> UCS uses a priority queue ordered by actual accumulated path cost g(n). Remove a minimum-g node; accept a goal when removed, not when first generated. Relax each edge: if a cheaper path is found, update its frontier cost and parent.

Question sources

Artificial Intelligence review questions · Question 22

Question 197 | NS0023

In Reinforcement Learning, what guides an agent to improve its behavior?

A. Labeled training examples

B. Hard-coded rules

C. Rewards and penalties from environment interactions

D. Pre-trained heuristic functions

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer C: Reward feedback from interactions.

AI33 describes an agent acting, receiving reward and a new observation, and improving its policy to maximize expected cumulative reward. Rewards may be positive or negative and delayed; they need not name the correct action. Labeled examples, fixed rules and predefined heuristics do not provide this distinctive reinforcement feedback mechanism.

Why each choice is right or wrong

A - Incorrect: Correct target labels characterize supervised learning rather than reinforcement feedback.

B - Incorrect: Fixed rules alone do not provide the reward-based adaptation defining RL.

C - Correct: Rewards and penalties evaluate outcomes and guide improvement of expected cumulative return.

D - Incorrect: A pretrained heuristic may help implementation but is not the defining RL feedback signal.

Primary note proof

Study-note section AI33

> An RL agent observes state information, takes an action, receives reward and a new observation, and improves its policy to maximize expected cumulative reward. Rewards evaluate outcomes rather than supply the correct action label. Rewards may be delayed; not every correct-looking action receives immediate reward.

Study-note section AI31

> Supervised learning fits a mapping from labeled input-target examples. Classification predicts discrete labels, such as institution generations or spam/nonspam; regression predicts numeric quantities, such as prices. Predictive analytics uses historical patterns to forecast outcomes and can use either. Specify the target before choosing a method.

Question sources

Artificial Intelligence review questions · Question 23

Question 198 | NS0024

For greedy best-first search on an implicit finite tree with branching factor b and maximum depth m, what is the standard worst-case node-count time bound?

A. O(m)

B. O(b*m)

C. O(b^m)

D. O(log m)

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer C: O(b^m)

Greedy best-first ordering can be misled by a heuristic and may need to explore an exponentially large portion of the finite search tree. With branching factor b and depth m, the standard worst-case generated-node bound is O(b^m). This is a node-count convention rather than a claim about every priority-queue implementation or all best-first algorithms. The explicit algorithm and tree model determine the bound.

Why each choice is right or wrong

A - Incorrect: A misleading heuristic can force exponentially many nodes, so a bound merely linear in maximum depth m is insufficient.

B - Incorrect: O(bm) is a familiar DFS storage form, not this generic search-time guarantee.

C - Correct: The specified implicit finite-tree worst case can require O(b^m) nodes.

D - Incorrect: A logarithmic-in-depth bound is insufficient when the search can explore an exponential number of tree nodes.

Primary note proof

Study-note section AI17

> Best-first is a family that orders OPEN by an evaluation function. Greedy best-first uses f=h and removes the smallest estimate, ignoring the cost already paid. It can choose an expensive route, loop without duplicate checks, or follow an infinite misleading branch. On a finite graph with full duplicate handling it is complete, but not generally optimal. Its implicit finite-tree worst case can be O(b^m).

Question sources

Artificial Intelligence review questions · Question 24

Question 199 | NS0025

Which knowledge representation uses templates with slots to describe stereotyped situations?

A. Semantic Networks

B. Frames

C. Logic

D. Bayesian Networks

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer B: Frames.

AI24 defines frames as stereotyped entity descriptions with slots, values and defaults. This is the template-and-slot representation in the question. Semantic networks focus on relationships among concept nodes, formal logic expresses sentences and Bayesian networks represent probabilistic dependencies; none is the specific slot-template term requested.

Why each choice is right or wrong

A - Incorrect: Semantic networks emphasize concept nodes and labeled relations rather than the slot-template representation.

B - Correct: Frames organize stereotyped descriptions into slots with values and defaults.

C - Incorrect: Logic uses formal sentences rather than the specifically described frame slots.

D - Incorrect: Bayesian networks represent conditional probabilities, not the named stereotyped slot template.

Primary note proof

Study-note section AI24

> A semantic network connects concept nodes by labeled relations such as is-a. Frames describe stereotyped entities through slots, values and defaults. Production systems use IF conditions THEN conclusions or actions, matching rules against working memory and selecting among applicable rules. A graph suits relationships such as frequently bought together; frames suit attributes. No representation is always best for recommendation without task requirements.

Study-note section AI29

> A Bayesian network is a directed acyclic graph with conditional probability tables; its joint distribution factors as the product of each node's probability given its parents. A Markov model assumes the next state depends on the current state rather than the entire past. These represent uncertainty; fuzzy membership represents a different issue.

Question sources

Artificial Intelligence review questions · Question 25

Question 200 | NS0026

What is Natural Language Processing (NLP) primarily focused on?

A. Visual image interpretation

B. Physical robot movement

C. Enabling machines to understand, interpret, and generate human language

D. Optimizing database queries

Hint

Match the input, output and purpose to the appropriate language, perception or physical-action task.

Answer C: Understanding, interpreting and generating human language.

AI37 directly defines NLP as interpreting and generating human language and lists dialogue, translation and summarization. Visual interpretation belongs to computer vision, physical movement to robotics and query optimization to database processing. NLP can interact with those systems, but its primary language-processing domain remains distinct.

Why each choice is right or wrong

A - Incorrect: Interpreting visual images is primarily computer vision rather than language processing.

B - Incorrect: Physical robot movement concerns control and robotics, not the definition of NLP.

C - Correct: Human-language interpretation and generation are the primary NLP tasks named in AI37.

D - Incorrect: Database-query optimization is not the human-language-processing focus described.

Primary note proof

Study-note section AI37

> NLP interprets and generates human language: translation, dialogue, summarization, question answering and text generation. Speech recognition converts audio to text; speech synthesis generates speech. They overlap language processing but solve different input/output tasks.

Study-note section AI39

> Computer vision interprets images or video: classification names image content, detection locates objects, segmentation labels pixels, and recognition identifies objects or people. Occlusion, lighting change, noise, unfamiliar conditions and incomplete context make open-world perception difficult even when a face detector works well.

Study-note section AI39

> Robotics combines mechanical, electrical, control and computing systems to sense, plan and act physically. Components include sensors, controller, power supply, actuators and an end effector. Locomotion may be legged, wheeled or tracked, with trade-offs in terrain, stability and

efficiency. Software can act autonomously without being a robot. Humanoid shape is not evidence of general intelligence.

Question sources

Artificial Intelligence review questions · Question 26

Question 201 | NS0027

Which listed term is a reasoning method rather than one of the common learning regimes classified by availability of target labels?

A. Supervised Learning

B. Unsupervised Learning

C. Deductive reasoning

D. Semi-supervised Learning

Hint

Translate the statement carefully, then check the scope of its quantifiers, connectives or inference rule. Distinguish guaranteed consequences from plausible generalizations or explanations.

Answer C: Deductive reasoning

Supervised, unsupervised and semi-supervised learning differ in the target labels available during training. Deductive reasoning instead derives consequences from premises using truth-preserving rules. Deduction can be used inside a learning system, so the distinction is not that the two can never coexist. The question separates a reasoning method from the conventional label-availability taxonomy.

Why each choice is right or wrong

A - Incorrect: Supervised learning uses labeled input-target pairs, so it is a label-availability learning regime.

B - Incorrect: Unsupervised learning has no supplied targets and discovers structure, so it is a label-availability regime.

C - Correct: Deduction derives truth-preserving consequences from premises rather than naming a label-availability regime.

D - Incorrect: Semi-supervised learning combines labeled and unlabeled data, so it is a label-availability regime.

Primary note proof

Study-note section AI31

Supervised learning fits a mapping from labeled input-target examples. Classification predicts discrete labels, such as institution generations or spam/nonspam; regression predicts numeric quantities, such as prices. Predictive analytics uses historical patterns to forecast outcomes and can use either. Specify the target before choosing a method.

Study-note section AI32

Unsupervised learning seeks structure without target labels: clustering groups similar examples; association finds co-occurrence; dimensionality reduction compresses features. Good clusters have high within-cluster similarity and low between-cluster similarity under an appropriate distance. K-means alternates assignment to nearest centroids and centroid updates.

Study-note section AI32

Semi-supervised learning combines some labels with additional unlabeled data. Active learning chooses informative cases for an oracle to label; it still uses labels. “Exploratory learning” is not an exclusive standard synonym: it can mean unsupervised exploration or active data acquisition, so define the feedback. A task named summarization or time-series analysis is not automatically supervised; its training targets determine that.

Study-note section AI28

Deduction applies truth-preserving rules: all people are mortal; Ada is a person; therefore Ada is mortal. Induction generalizes observations to a broader rule and can be overturned by new examples. Abduction proposes a plausible explanation: wet ground suggests rain, although sprinklers could explain it. Analogy transfers a pattern between similar situations.

Question sources

Artificial Intelligence review questions · Question 27

Question 202 | NS0028

In adversarial search, what makes it different from traditional search?

A. It only uses one agent

B. Multiple agents with conflicting goals compete in the same search space

C. It guarantees the optimal solution always

D. It does not use evaluation functions

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer B: Competing agents with conflicting objectives.

AI22 defines adversarial search by agents with competing objectives and gives minimax for opposition. That distinguishes it from a single-agent path search. It can use evaluation functions at cutoffs, and its best-outcome guarantee is conditional on the modeled game and opponent, not an unconditional guarantee that every implementation finds a globally optimal solution.

Why each choice is right or wrong

A - Incorrect: A single decision-making agent does not capture the competing-objective feature of adversarial search.

B - Correct: Competing agents and conflicting objectives are its defining distinction.

C - Incorrect: Guarantees depend on the algorithm, model and search depth, not merely the adversarial label.

D - Incorrect: Depth-limited adversarial search commonly uses heuristic evaluation functions.

Primary note proof

Study-note section AI22

> Adversarial search models agents with competing objectives. A game specifies states, player to move, legal actions, successor function, terminal test and utility. Perfect information means previous moves and current state are visible; it does not reveal future choices or imply absence of chance. Zero-sum utility makes one player's gain the other's loss.

Study-note section AI23

> At a depth cutoff use an evaluation function estimating position desirability. Positive favors the designated MAX player, negative MIN; the scale is a convention, with +1/0/−1 often win/draw/loss. Material counts weighted pieces; mobility counts legal options; territory scores controlled regions. These are different features, often combined. Heuristics improve limited computation, not guarantee victory.

Question sources

Artificial Intelligence review questions · Question 28

Question 203 | NS0029

A Depth-Limited Search (DLS) is NOT complete in itself because:

A. It uses too much memory

B. Solutions beyond the preset depth limit will be missed

C. It explores nodes in random order

D. It requires a heuristic function

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer B: A solution beyond the fixed depth limit is missed.

AI16 defines DLS as DFS stopped at depth l and explicitly says it misses goals deeper than that. A fixed limit therefore does not establish completeness for arbitrary problem depths. This is independent of randomized ordering, heuristics or memory use; increasing limits through iterative deepening addresses the depth restriction under its stated assumptions.

Why each choice is right or wrong

A - Incorrect: Depth-limited DFS is memory-efficient; excessive memory is not the reason goals are missed.

B - Correct: A preset cutoff prevents exploring a solution whose depth exceeds that limit.

C - Incorrect: DLS need not use random order; the depth cutoff alone causes the limitation.

D - Incorrect: No heuristic is required for depth-limited DFS.

Primary note proof

Study-note section AI16

> DLS runs DFS to depth l and reports cutoff separately from failure. It misses goals deeper than l. Iterative deepening runs DLS with limits 0,1,2,...; for finite b it finds the shallowest solution with O(b^d) time and O(bd) frontier space, optimal for equal step costs. Repeated shallow work is small compared with the deepest level.

Study-note section AI14

> DFS follows the deepest pending branch using a LIFO stack or recursion. FILO means the same order as LIFO. With the same children as BFS, a left-first DFS visits S,A,C before B; push rightmost children first to pop the leftmost first.

Question sources

Artificial Intelligence review questions · Question 29

Question 204 | NS0030

Which probabilistic model represents knowledge in terms of conditional probabilities among variables?

A. Semantic Networks

B. Frames

C. Bayesian Networks

D. Ontologies

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer C: Bayesian networks.

AI29 defines a Bayesian network as a DAG with conditional probability tables and joint factorization by parents. That is the probabilistic knowledge representation in the stem. Semantic networks and frames encode relations or slots without inherently supplying conditional distributions, and ontologies specify conceptual structure rather than this particular probabilistic factorization.

Why each choice is right or wrong

A - Incorrect: Semantic networks encode labeled concept relations rather than inherently conditional probability distributions.

B - Incorrect: Frames encode slots, values and defaults rather than the named probabilistic dependency structure.

C - Correct: Bayesian networks explicitly store conditional probability relationships among variables.

D - Incorrect: An ontology's conceptual organization is not itself the conditional-probability model specified.

Primary note proof

Study-note section AI29

> A Bayesian network is a directed acyclic graph with conditional probability tables; its joint distribution factors as the product of each node's probability given its parents. A Markov model assumes the next state depends on the current state rather than the entire past. These represent uncertainty; fuzzy membership represents a different issue.

Study-note section AI24

> A semantic network connects concept nodes by labeled relations such as is-a. Frames describe stereotyped entities through slots, values and defaults. Production systems use IF conditions THEN conclusions or actions, matching rules against working memory and selecting among applicable rules. A graph suits relationships such as frequently bought together; frames suit attributes. No representation is always best for recommendation without task requirements.

Study-note section AI-S06

An ontology describes entities or individuals, classes and properties and may state logical relationships and restrictions. It organizes explicit conceptual knowledge. Conditional probability tables and a probabilistic factorization are additional modeling commitments, not inherent in the word ontology.

Question sources

Artificial Intelligence review questions · Question 30

Question 205 | NS0031

What is an 'expanded node' in a search algorithm?

A. A node that has been visited but not explored further

B. A node selected for further exploration whose neighboring nodes have been generated

C. A node removed from the open list

D. A node that has reached the goal state

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer B: A node whose successors have been generated.

AI12 defines expansion as generating successors. B states that action, unlike mere visitation, removal from OPEN or reaching a goal. A node can be removed and goal-tested without expanding it, and discovered frontier nodes await expansion. The criterion is what was generated, not whether the node was seen or selected.

Why each choice is right or wrong

A - Incorrect: Visiting without generating successors does not meet the definition of expansion.

B - Correct: Generating the selected node’s neighboring successor nodes is expansion.

C - Incorrect: Removal alone may precede a goal test without generating any successors.

D - Incorrect: Goal status is not the definition of expansion; a goal may be accepted without expanding it.

Primary note proof

Study-note section AI12

> A search node records a state plus parent, action, depth and path cost. Expanding a node generates its successors. OPEN or frontier stores discovered candidates awaiting expansion; CLOSED records expanded states. Reaching the same state through different paths does not imply the same node or cost.

Study-note section AI13

> BFS removes the oldest frontier entry using a FIFO queue. Start with the root, remove the head, test it, append unseen successors at the tail, and repeat. It explores depths 0,1,2,... and is complete for finite b and a finite-depth goal. It minimizes number of edges, so minimum cost requires equal positive step costs.

Question sources

Artificial Intelligence review questions · Question 31

Question 206 | NS0032

For which type of problem would Supervised Learning with CNNs be most appropriate?

A. Customer segmentation from unlabeled data

B. Classifying a large labeled image dataset

C. Learning chess actions from rewards obtained through trial-and-error play.

D. Finding hidden patterns in transaction data

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer B: Classifying a large labeled image dataset

The labeled images provide input-target pairs for supervised classification, and CNN filters exploit spatial image structure. Unlabeled customer segmentation and transaction-pattern discovery seek structure without class targets. Learning chess actions from rewards is a reinforcement objective.

Why each choice is right or wrong

A - Incorrect: Unlabeled segmentation seeks groups rather than supervised image classes.

B - Correct: Labeled images supply class targets, and CNN filters learn spatial image features.

C - Incorrect: Trial-and-error chess with rewards defines a reinforcement objective rather than supplied image-class targets.

D - Incorrect: Hidden transaction-pattern discovery seeks unlabeled structure rather than labeled image classification.

Primary note proof

Study-note section AI31

> Supervised learning fits a mapping from labeled input-target examples. Classification predicts discrete labels, such as institution generations or spam/nonspam; regression predicts numeric quantities, such as prices. Predictive analytics uses historical patterns to forecast outcomes and can use either. Specify the target before choosing a method.

Study-note section AI34

> Feedforward networks pass signals toward outputs without recurrent state loops; recurrent networks carry state across sequence steps. Training adjusts weights, often through gradients and backpropagation. Backpropagation is not the same as a recurrent connection. Deep learning uses multiple representation layers and can learn features from images, sound and text. CNNs share filters across image locations. ANN is a brain-inspired abstraction, not a literal brain simulation; DL⊂ML⊂AI.

Study-note section AI32

Unsupervised learning seeks structure without target labels: clustering groups similar examples; association finds co-occurrence; dimensionality reduction compresses features. Good clusters have high within-cluster similarity and low between-cluster similarity under an appropriate distance. K-means alternates assignment to nearest centroids and centroid updates.

Study-note section AI33

An RL agent observes state information, takes an action, receives reward and a new observation, and improves its policy to maximize expected cumulative reward. Rewards evaluate outcomes rather than supply the correct action label. Rewards may be delayed; not every correct-looking action receives immediate reward.

Question sources

Artificial Intelligence review questions · Question 32

Question 207 | NS0033

What does 'perfect information' mean in the context of game theory?

A. The agent knows the future

B. Both players have complete and accurate knowledge of the entire game state

C. The game has no randomness

D. Each player can see only their own pieces

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer B: The current game state and previous moves are visible to the players.

AI22 defines perfect information as visibility of previous moves and the current state. This makes B the correct state-knowledge description. It does not reveal future choices and does not by itself remove chance from a game. Seeing only one's own pieces instead describes hidden relevant state, inconsistent with perfect information.

Why each choice is right or wrong

A - Incorrect: Future choices are not known merely because the present game state is visible.

B - Correct: Complete visibility of current state and relevant move history is the intended information condition.

C - Incorrect: A game can include chance while still revealing the state and prior moves.

D - Incorrect: Seeing only one's own pieces leaves relevant state hidden and is not perfect information.

Primary note proof

Study-note section AI22

> Adversarial search models agents with competing objectives. A game specifies states, player to move, legal actions, successor function, terminal test and utility. Perfect information means previous moves and current state are visible; it does not reveal future choices or imply absence of chance. Zero-sum utility makes one player's gain the other's loss.

Study-note section AI04

> Omniscience would mean knowing actual action outcomes in advance. Rationality does not require omniscience or guaranteed success. Even complete outcome knowledge cannot create a successful action when every available action fails. Information gathering and exploration may therefore be rational actions.

Question sources

Artificial Intelligence review questions · Question 33

Question 208 | NS0034

Semi-supervised Learning is typically used when:

A. All data is labeled

B. No data is labeled

C. Obtaining a fully labeled dataset is time-consuming or costly

D. The problem requires reinforcement feedback

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer C: When labels are costly and some labeled data can be supplemented by unlabeled data.

AI32 defines semi-supervised learning as combining some labels with additional unlabeled data. Costly full labeling provides the motivation described in C, provided at least some labeled targets exist. Fully labeled data is ordinary supervised learning; wholly unlabeled structure fits unsupervised learning, and environmental rewards specify reinforcement feedback rather than this mixed dataset regime.

Why each choice is right or wrong

A - Incorrect: A fully labeled dataset does not give the defining labeled-plus-unlabeled mixture.

B - Incorrect: With no labels at all, the described regime is unsupervised rather than semi-supervised.

C - Correct: Expensive labeling motivates using a limited labeled subset together with plentiful unlabeled examples.

D - Incorrect: Reward feedback defines reinforcement learning rather than a mixture of labeled and unlabeled examples.

Primary note proof

Study-note section AI32

> Semi-supervised learning combines some labels with additional unlabeled data. Active learning chooses informative cases for an oracle to label; it still uses labels. “Exploratory learning” is not an exclusive standard synonym: it can mean unsupervised exploration or active data acquisition, so define the feedback. A task named summarization or time-series analysis is not automatically supervised; its training targets determine that.

Study-note section AI31

> Supervised learning fits a mapping from labeled input-target examples. Classification predicts discrete labels, such as institution generations or spam/nonspam; regression predicts numeric quantities, such as prices. Predictive analytics uses historical patterns to forecast outcomes and can use either. Specify the target before choosing a method.

Study-note section AI33

> An RL agent observes state information, takes an action, receives reward and a new observation, and improves its policy to maximize expected cumulative reward. Rewards evaluate outcomes rather than supply the correct action label. Rewards may be delayed; not every correct-looking action receives immediate reward.

Question sources

Artificial Intelligence review questions · Question 34

Question 209 | NS0035

In the context of AI agents, what are 'effectors'?

A. Sensors that gather environmental data

B. Components that allow an agent to take actions in its environment

C. Internal memory modules

D. Heuristic evaluation functions

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer B: Components through which an agent acts on its environment.

AI03 uses actuators or effectors for acting components and sensors for perception. Therefore B gives the correct role. Memory stores internal information and heuristic functions estimate promise; neither is the acting interface. An effector may be physical hardware or a software action mechanism depending on the agent’s architecture.

Why each choice is right or wrong

A - Incorrect: Sensors gather percepts; effectors implement actions rather than observations.

B - Correct: Effectors are the components through which selected actions affect the environment.

C - Incorrect: Internal memory maintains information but is not the general acting interface.

D - Incorrect: Heuristic functions evaluate promise and do not themselves execute environmental actions.

Primary note proof

Study-note section AI03

> An agent perceives an environment through sensors and acts through actuators or effectors. Cameras sense; motors move. Software sensors may be incoming messages and software actions may be API requests. Humans are agents but are not artificial agents.

Study-note section AI06

A model-based reflex agent maintains internal state using the previous state, previous action, current percept, and models of world change and action effects. It then applies rules to that state. Both kinds can use condition-action rules; memory and a world model distinguish them.

Study-note section AI17

A heuristic h(n) estimates remaining cost or evaluates promise using problem-specific knowledge. It can reduce explored work, but computing it also costs time. Uninformed algorithms use only the problem definition and known costs; UCS remains uninformed despite using g.

Question sources

Artificial Intelligence review questions · Question 35

Question 210 | NS0036

Which of the following best describes a Utility-Based Agent?

A. It only reacts to current perceptions using condition-action rules

B. It selects actions that maximize expected utility, handling conflicting goals through trade-offs

C. It plans sequences of actions toward a single fixed goal

D. It learns entirely from environmental feedback

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer B: Choose actions with maximum expected utility, balancing outcome trade-offs.

AI07 defines utility-based agents by comparing desirability and trade-offs and maximizing expected utility over uncertain outcomes. This is broader than merely satisfying a fixed goal. Current-percept rule matching describes simple reflex behavior, and experience-based improvement describes learning, which may coexist with utility choice but does not define it.

Why each choice is right or wrong

A - Incorrect: Current-percept condition-action matching describes a simple reflex architecture.

B - Correct: Expected-utility maximization explicitly compares trade-offs among uncertain outcomes.

C - Incorrect: Goal-reaching planning alone does not rank degrees of desirability among outcomes.

D - Incorrect: Learning describes adaptation from feedback, not the utility-based decision criterion itself.

Primary note proof

Study-note section AI07

A goal-based agent considers future action sequences that reach a desired state. A utility-based agent compares degrees of desirability and trade-offs, choosing maximum expected utility: $EU(a)=\Sigma s\ P(s|a,evidence)U(s)$. Reaching the destination is a goal; balancing arrival time, safety and cost requires preferences among outcomes.

Study-note section AI06

A simple reflex agent matches condition-action rules against the current percept only: if dirty, clean. It has no persistent world-state memory. This is adequate when the current observation contains everything needed for its rule, but partial observability can make the same percept require different actions.

Study-note section AI08

The performance element selects actions. The learning element changes the agent to improve future performance. The critic evaluates behavior against a performance standard and supplies feedback. The problem generator suggests exploratory actions likely to yield useful experience.

Question sources

Artificial Intelligence review questions · Question 36

Question 211 | NS0037

What is Computer Vision in AI primarily concerned with?

A. Processing and generating spoken language

B. Enabling machines to understand and interpret visual information

C. Designing physical robotic systems

D. Managing database indexing

Hint

Match the input, output and purpose to the appropriate language, perception or physical-action task.

Answer B: Interpreting visual information from images or video.

AI39 defines computer vision as interpreting images or video and includes classification, detection, segmentation and recognition. Language generation, robot mechanics and database indexing are separate task areas, although vision may be one component of a robot. B states the visual information-processing purpose directly.

Why each choice is right or wrong

A - Incorrect: Speech and language generation belong to language/speech processing rather than visual interpretation.

B - Correct: Computer vision extracts and interprets information from visual inputs such as images and video.

C - Incorrect: Robotic system design may use vision but is broader than the visual-processing field.

D - Incorrect: Database indexing is not the interpretation of visual scenes or objects.

Primary note proof

Study-note section AI39

> Computer vision interprets images or video: classification names image content, detection locates objects, segmentation labels pixels, and recognition identifies objects or people. Occlusion, lighting change, noise, unfamiliar conditions and incomplete context make open-world perception difficult even when a face detector works well.

Study-note section AI37

> NLP interprets and generates human language: translation, dialogue, summarization, question answering and text generation. Speech recognition converts audio to text; speech synthesis generates speech. They overlap language processing but solve different input/output tasks.

Study-note section AI39

> Robotics combines mechanical, electrical, control and computing systems to sense, plan and act physically. Components include sensors, controller, power supply, actuators and an end effector. Locomotion may be legged, wheeled or tracked, with trade-offs in terrain, stability and efficiency. Software can act autonomously without being a robot. Humanoid shape is not evidence of general intelligence.

Question sources

Artificial Intelligence review questions · Question 37

Question 212 | NS0038

In a zero-sum Tic-Tac-Toe terminal utility convention, MAX wins score +1 and MIN wins score −1. Which listed value represents a draw giving neither player an advantage?

A. +1

B. -1

C. +Infinity

D. 0

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer D: 0

The convention assigns equal and opposite utilities to the two decisive outcomes. A draw gives neither side an advantage, so its neutral utility is zero under that symmetric scale. Positive one favors MAX and negative one favors MIN; positive infinity would strongly favor MAX rather than indicate neutrality. This is an explicit utility convention, not an intrinsic numerical property of the game outcome.

Why each choice is right or wrong

A - Incorrect: Under the common stated scale, +1 is a win rather than a draw.

B - Incorrect: Under the common stated scale, −1 is a loss rather than a draw.

C - Incorrect: Positive infinity is not the draw value in the note's common scale.

D - Correct: Zero is the neutral utility between +1 and −1 under this convention.

Primary note proof

Study-note section AI23

> At a depth cutoff use an evaluation function estimating position desirability. Positive favors the designated MAX player, negative MIN; the scale is a convention, with +1/0/−1 often win/draw/loss. Material counts weighted pieces; mobility counts legal options; territory scores controlled regions. These are different features, often combined. Heuristics improve limited computation, not guarantee victory.

Question sources

Artificial Intelligence review questions · Question 38

Question 213 | NS0039

Treat a whole independently reset chess match as one episode with no effects on later matches; treat the other options as ongoing tasks at the action level described. Which task is episodic under those units?

A. A self-driving car navigating city traffic

B. A chess game (each individual match)

C. A robot cleaning a multi-room house continuously

D. A stock trading algorithm operating over months

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer B: A chess game (each individual match)

An episodic task has independent decision units. Under the stated convention, each whole chess match is reset and cannot affect later matches, so matches can be treated as episodes. Individual moves within a match remain sequential. Driving, continuous cleaning and continuing stock trading all have actions that affect later situations within the ongoing task. The classification therefore follows the explicit unit and independence assumptions.

Why each choice is right or wrong

A - Incorrect: Driving decisions affect later positions and traffic interactions, making the task sequential.

B - Correct: Whole independently reset matches are independent episodes under the stated convention.

C - Incorrect: Cleaning actions change later room states and decisions in the continuing task.

D - Incorrect: Trading decisions affect later holdings and opportunities, so the continuing task is sequential.

Primary note proof

Study-note section AI09

> Episodic decisions are independent; sequential decisions affect later situations. Static worlds do not change while deliberating; dynamic ones do. A clock affecting score makes an otherwise static game semidynamic. Discrete quantities have distinct alternatives; continuous quantities vary along real-valued ranges. Known versus unknown concerns knowledge of the transition model, not observability.

Question sources

Artificial Intelligence review questions · Question 39

Question 214 | NS0040

The heuristic function h(n) in A* represents:

A. The exact cost from start to node n

B. The total path cost already traveled

C. The estimated cost from node n to the goal state

D. The branching factor at node n

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer C: Estimated remaining cost from n to a goal.

AI17 defines h as an estimate of remaining cost or promise, and AI18 adds it to accumulated g in A*. Thus C is the heuristic component. Exact start-to-node cost belongs to g, and branching factor b is a search-size parameter rather than a cost-to-go estimate.

Why each choice is right or wrong

A - Incorrect: Exact cost from the start belongs to g(n), not the heuristic h(n).

B - Incorrect: Already-traveled path cost is g(n), not estimated future cost.

C - Correct: h estimates the remaining cost from the current node to a goal.

D - Incorrect: Branching factor counts successor alternatives and is not the heuristic cost estimate.

Primary note proof

Study-note section AI17

> A heuristic h(n) estimates remaining cost or evaluates promise using problem-specific knowledge. It can reduce explored work, but computing it also costs time. Uninformed algorithms use only the problem definition and known costs; UCS remains uninformed despite using g.

Study-note section AI18

> A* orders by f(n)=g(n)+h(n). Admissibility means $0 \le h(n) \le h^*(n)$, the true cheapest remaining cost. Consistency means $h(n) \le c(n,n') + h(n')$ on every edge, with h(goal)=0; it implies admissibility on states with goal paths. Consistency makes f nondecreasing along a path.

Study-note section AI12

> Completeness means finding a solution if one exists under stated assumptions; optimality means returning the best solution. Time counts work, often generated nodes; space counts simultaneous storage. Let b be branching factor, d shallowest goal depth, m maximum search depth, and l a depth limit. Graph-search bounds in vertices V and edges E differ from implicit-tree bounds.

Question sources

Artificial Intelligence review questions · Question 40

Question 215 | NS0041

When a doctor proposes a disease as a plausible cause explaining observed symptoms, which term names this explanatory direction of reasoning, rather than a method for numerically ranking hypotheses?

A. Deductive Reasoning

B. Inductive Reasoning

C. Abductive Reasoning

D. Probabilistic Reasoning

Hint

Translate the statement carefully, then check the scope of its quantifiers, connectives or inference rule. Distinguish guaranteed consequences from plausible generalizations or explanations.

Answer C: Abductive Reasoning

Abduction proposes an explanation for observed evidence, such as a disease that could account for symptoms. Other causes may also fit, so the conclusion is not logically guaranteed. Probability can help rank the possible explanations, but that numerical method is distinct from the explanatory direction. Deduction preserves truth from premises, and induction generalizes patterns across cases.

Why each choice is right or wrong

A - Incorrect: Deduction follows truth-preserving implications from premises, not a merely plausible cause inferred from its effects.

B - Incorrect: Induction generalizes across observations rather than specifically proposing a cause that explains an observation.

C - Correct: Abduction proposes a plausible explanation for observed evidence.

D - Incorrect: Probability can rank disease hypotheses, but does not itself name the observation-to-explanation direction.

Primary note proof

Study-note section AI28

> Deduction applies truth-preserving rules: all people are mortal; Ada is a person; therefore Ada is mortal. Induction generalizes observations to a broader rule and can be overturned by new examples. Abduction proposes a plausible explanation: wet ground suggests rain, although sprinklers could explain it. Analogy transfers a pattern between similar situations.

Study-note section AI29

> Probability describes uncertainty about events. $P(\neg A)=1-P(A)$; $P(A|B)=P(A\cap B)/P(B)$, for $P(B)>0$. Bayes gives $P(H|E)=P(E|H)P(H)/P(E)$: posterior equals likelihood times prior divided by evidence probability. For mutually exclusive exhaustive hypotheses, $P(E)=\Sigma_i P(E|Hi)P(Hi)$.

Question sources

Artificial Intelligence review questions · Question 41

Question 216 | NS0042

What major ethical concern does AI raise related to decision-making systems?

A. AI systems use too little data

B. Bias in algorithms leading to unfair or discriminatory outcomes

C. AI always makes perfectly rational decisions

D. AI cannot process unstructured data

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer B: Biased decisions causing unfair outcomes.

AI40 says bias in data or evaluation can create unfair outcomes and calls for checking errors across groups, robustness and oversight. B states the ethical concern directly. Data scarcity may be a technical issue, but little data is not the stated general ethical principle; perfect rationality and inability to process unstructured inputs are false generalizations.

Why each choice is right or wrong

A - Incorrect: Using too little data is not the specific fairness concern described here.

B - Correct: Biased data or evaluation can produce unfair or discriminatory decisions.

C - Incorrect: AI is not guaranteed perfect rational decisions or error-free outcomes.

D - Incorrect: The note explicitly includes models that process images, sound and text.

Primary note proof

Study-note section AI40

> Rules offer inspectable domain reasoning but can be brittle and costly to maintain. Learned models handle data-rich patterns but depend on representative data and validation. Search suits explicit action alternatives; probabilistic models suit uncertain evidence. Bias in data or evaluation can create unfair outcomes. Evaluate errors across relevant groups, robustness, privacy and human oversight rather than assuming automation is always beneficial.

Study-note section AI34

> Feedforward networks pass signals toward outputs without recurrent state loops; recurrent networks carry state across sequence steps. Training adjusts weights, often through gradients and backpropagation. Backpropagation is not the same as a recurrent connection. Deep learning uses multiple representation layers and can learn features from images, sound and text. CNNs share filters across image locations. ANN is a brain-inspired abstraction, not a literal brain simulation; DL⊂ML⊂AI.

Study-note section AI04

A performance measure defines success externally: safety, useful deliveries, elapsed time, or cost. Rationality depends on that measure, prior knowledge, percept history, and available actions. A rational agent selects the action with the best expected performance given its information.

Question sources

Artificial Intelligence review questions · Question 42

Question 217 | NS0043

In a continuous environment like self-driving cars, what challenge does 'continuous' specifically pose?

A. A finite set of states makes decisions trivial

B. Infinite or large state/action spaces require more complex computations

C. The environment never changes

D. All outcomes are deterministic

Hint

Read the relevant agent component, action-selection criterion or environment dimension. Use only properties established by the description.

Answer B: Real-valued state/action variables create large or infinite spaces to represent and search.

AI09 defines continuous quantities as varying along real-valued ranges, unlike distinct discrete alternatives. Such ranges can create unboundedly many state/action possibilities, motivating approximation or richer computation rather than exhaustive finite enumeration. Continuity is separate from whether a world changes over time or whether outcomes are stochastic.

Why each choice is right or wrong

A - Incorrect: Continuous ranges do not reduce the task to a trivial finite set of alternatives.

B - Correct: Real-valued ranges can produce very large or infinite possibilities needing suitable representation and computation.

C - Incorrect: No world change defines static rather than continuous.

D - Incorrect: Determinism is a separate transition-property dimension and does not follow from continuity.

Primary note proof

Study-note section AI09

> Episodic decisions are independent; sequential decisions affect later situations. Static worlds do not change while deliberating; dynamic ones do. A clock affecting score makes an otherwise static game semidynamic. Discrete quantities have distinct alternatives; continuous quantities vary along real-valued ranges. Known versus unknown concerns knowledge of the transition model, not observability.

Question sources

Artificial Intelligence review questions · Question 43

Question 218 | NS0044

The successor function in game search is responsible for:

A. Evaluating terminal game states

B. Generating new states by applying legal moves to the current game state

C. Pruning unpromising branches

D. Calculating the heuristic value of each state

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer B: Generate successor states by applying legal moves.

AI22 lists legal actions and a successor function as game components; AI11 identifies successors as states reached by actions. In game search, applying each legal move yields its successor position. Terminal utilities and heuristic scores evaluate positions, while alpha-beta pruning avoids exploring branches; neither is the transition-generation function.

Why each choice is right or wrong

A - Incorrect: Terminal utility evaluation scores a completed outcome rather than generating the next position.

B - Correct: The successor function maps legal moves at the current position to resulting game states.

C - Incorrect: Pruning skips exploration under a bound; it is not state-transition generation.

D - Incorrect: A heuristic evaluation scores a state's promise rather than producing legal successor states.

Primary note proof

Study-note section AI22

> Adversarial search models agents with competing objectives. A game specifies states, player to move, legal actions, successor function, terminal test and utility. Perfect information means previous moves and current state are visible; it does not reveal future choices or imply absence of chance. Zero-sum utility makes one player's gain the other's loss.

Study-note section AI11

> First formulate the goal; then formulate the problem by specifying initial state, legal actions, transition model, goal test, and action costs. Search selects a sequence; execution applies it. A state abstracts away irrelevant world details.

Study-note section AI23

> At a depth cutoff use an evaluation function estimating position desirability. Positive favors the designated MAX player, negative MIN; the scale is a convention, with +1/0/−1 often win/draw/loss. Material counts weighted pieces; mobility counts legal options; territory scores

controlled regions. These are different features, often combined. Heuristics improve limited computation, not guarantee victory.

Question sources

Artificial Intelligence review questions · Question 44

Question 219 | NS0045

For implicit-tree search with branching factor b and shallowest goal depth d, which is a key memory advantage of iterative deepening over BFS when each depth-limited pass stores pending siblings?

A. It finds deeper solutions faster

B. It requires storing all nodes at each level in memory

C. It uses O(bd) frontier space rather than storing an entire BFS level.

D. It does not guarantee completeness

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer C: It uses O(bd) frontier space rather than storing an entire BFS level.

Iterative deepening repeats depth-limited DFS while increasing the limit. A pass retains a depth-first path and pending siblings, giving O(bd) frontier space rather than BFS's potentially exponential level frontier. It remains complete for finite branching and a finite-depth solution, although it repeats some shallow work. The option states the correct memory advantage without claiming that every implementation stores only the current path.

Why each choice is right or wrong

A - Incorrect: IDS repeats shallow work and does not universally find deeper solutions faster than BFS.

B - Incorrect: Storing all nodes at a level is BFS-style memory behavior, not the IDS advantage.

C - Correct: The stated implementation keeps only a depth-first frontier, requiring O(bd) rather than the BFS level frontier.

D - Incorrect: IDS is complete under finite branching and a finite-depth solution, so loss of completeness is not its benefit.

Primary note proof

Study-note section AI16

DLS runs DFS to depth l and reports cutoff separately from failure. It misses goals deeper than l. Iterative deepening runs DLS with limits 0,1,2,...; for finite b it finds the shallowest solution with O(b^d) time and O(bd) frontier space, optimal for equal step costs. Repeated shallow work is small compared with the deepest level.

Study-note section AI13

The tree contains 1+b+...+b^d nodes through depth d, giving O(b^d) under generation-time goal testing. Expanding depth-d nodes before testing can generate depth d+1 and yields O(b^(d+1)); specify the convention. Explicit finite graph traversal with visited-state tracking takes O(V+E) time and O(V) space.

Study-note section AI14

DFS follows the deepest pending branch using a LIFO stack or recursion. FILO means the same order as LIFO. With the same children as BFS, a left-first DFS visits S,A,C before B; push rightmost children first to pop the leftmost first.

Question sources

Artificial Intelligence review questions · Question 45

Question 220 | NS0046

Representation learning in Deep Learning refers to:

A. Manually engineering input features

B. The ability to automatically discover and learn relevant features from raw data

C. Storing knowledge in frames and ontologies

D. Using labeled data with a teacher signal

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer B: Automatically learning useful features from data.

AI34 says deep models learn features from images, sound and text through multiple representation layers. That is representation learning rather than manually specifying every input feature. A label signal may guide training but is not required for every deep objective, and frames/ontologies are separate explicit knowledge structures rather than the learned representations meant here.

Why each choice is right or wrong

A - Incorrect: Manual feature engineering is contrasted with learning the representations automatically.

B - Correct: The model learns task-useful features or representations from the supplied data.

C - Incorrect: Frames and ontologies are explicit knowledge representations, not the deep feature-learning mechanism described.

D - Incorrect: Teacher labels define supervision; representation learning can also use unsupervised or reinforcement objectives.

Primary note proof

Study-note section AI34

Feedforward networks pass signals toward outputs without recurrent state loops; recurrent networks carry state across sequence steps. Training adjusts weights, often through gradients and backpropagation. Backpropagation is not the same as a recurrent connection. Deep learning uses multiple representation layers and can learn features from images, sound and text. CNNs share filters across image locations. ANN is a brain-inspired abstraction, not a literal brain simulation; DL⊂ML⊂AI.

Study-note section AI43

ANNs, genetic search and reinforcement learning all improve candidate behavior or performance through experience in different ways. Deep learning is an architecture family usable with supervised, unsupervised or reinforcement objectives. Decision trees and deep models may both be supervised. To distinguish learning regimes, specify labels, reward or unlabeled structure.

Study-note section AI24

A semantic network connects concept nodes by labeled relations such as is-a. Frames describe stereotyped entities through slots, values and defaults. Production systems use IF conditions THEN conclusions or actions, matching rules against working memory and selecting among applicable rules. A graph suits relationships such as frequently bought together; frames suit attributes. No representation is always best for recommendation without task requirements.

Study-note section AI-S06

An ontology describes entities or individuals, classes and properties and may state logical relationships and restrictions. It organizes explicit conceptual knowledge. Conditional probability tables and a probabilistic factorization are additional modeling commitments, not inherent in the word ontology.

Question sources

Artificial Intelligence review questions · Question 46

Question 221 | NS0047

What makes Best-First (Greedy) Search non-optimal?

A. It always explores all possible paths

B. It relies entirely on heuristics and can overlook cheaper paths, making suboptimal choices

C. It uses actual path costs rather than estimates

D. It explores nodes in breadth-first order

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer B: It uses only h and can ignore a cheaper route with a larger estimate.

AI17 defines greedy best-first as f=h, ignoring cost already paid. AI19 demonstrates a route of cost 11 selected greedily while another route costs 4. The heuristic ranking can therefore prefer a worse total solution. Greedy is neither exhaustive nor breadth-first and does not order primarily by actual path cost.

Why each choice is right or wrong

A - Incorrect: Greedy search may stop at an early found goal rather than exploring every path.

B - Correct: Ignoring accumulated cost can favor a seemingly close but expensive route over a cheaper solution.

C - Incorrect: Using actual g alone describes UCS, whereas greedy uses h.

D - Incorrect: Breadth-first depth order is not the greedy minimum-heuristic policy.

Primary note proof

Study-note section AI17

> Best-first is a family that orders OPEN by an evaluation function. Greedy best-first uses f=h and removes the smallest estimate, ignoring the cost already paid. It can choose an expensive route, loop without duplicate checks, or follow an infinite misleading branch. On a finite graph with full duplicate handling it is complete, but not generally optimal. Its implicit finite-tree worst case can be O(b^m).

Study-note section AI19

> For S→A cost 2, S→B cost 1, A→G cost 2, B→G cost 10, set h(A)=2, h(B)=1, h(G)=0. After expanding S, greedy chooses B because 1<2 and can then choose G, returning cost 11. A* compares A:f=2+2=4 and B:f=1+1=2; after B, its G entry costs 11. It next expands A and replaces G with cost 4, which is removed and accepted.

Question sources

Artificial Intelligence review questions · Question 47

Question 222 | NS0048

In AI, what does 'Acting Humanly' (Turing Test concept) primarily evaluate?

A. Whether a machine can solve math problems

B. Whether a machine's behavior is indistinguishable from a human's in conversation

C. Whether a machine can perceive visual stimuli

D. Whether a machine follows logical rules perfectly

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer B: Humanlike conversational behavior judged by indistinguishability.

AI02 asks whether a judge can reliably distinguish a machine from a human through dialogue. This operational test concerns observed conversation, not a proof of consciousness or perfect logical reasoning. Mathematical success and perception may contribute capabilities but are not the primary distinguishing conversational criterion. B states that criterion directly.

Why each choice is right or wrong

A - Incorrect: Mathematical ability alone does not supply the test's conversational indistinguishability criterion.

B - Correct: The test evaluates whether conversational behavior can be reliably distinguished from a human's.

C - Incorrect: Visual perception alone is not the conversational test described in AI02.

D - Incorrect: Perfect logic is a rationality criterion, not the test's humanlike observed-behavior requirement.

Primary note proof

Study-note section AI02

> Thinking humanly models human mental processes; thinking rationally derives conclusions by valid logic. Acting humanly seeks humanlike behavior; acting rationally selects actions expected to achieve objectives. Thought processes and observed behavior are different evaluation targets.

Study-note section AI02

> The conversational Turing test asks whether a judge can reliably distinguish a machine from a human through their dialogue. It requires language use, knowledge, reasoning and often learning. It tests an operational behavior criterion, not proof of consciousness. Cognitive models can be tested against response times, errors and psychological experiments.

Question sources

Artificial Intelligence review questions · Question 48

Question 223 | NS0049

Which of the following is the correct relationship between AI, ML, and DL?

A. DL ⊂ AI ⊂ ML

B. AI ⊂ ML ⊂ DL

C. ML ⊂ DL ⊂ AI

D. DL ⊂ ML ⊂ AI

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer D: DL⊂ML⊂AI.

AI34 explicitly gives deep learning inside machine learning inside AI. Deep neural architectures are a subset of learning methods, and learning is one component of the broader intelligence field, which also includes planning and reasoning. Every other ordering reverses at least one of these stated inclusions.

Why each choice is right or wrong

A - Incorrect: This wrongly puts AI inside machine learning rather than the reverse.

B - Incorrect: Both inclusions are reversed relative to the note's hierarchy.

C - Incorrect: Machine learning contains deep learning, not the reverse.

D - Correct: Deep learning is contained in machine learning, which is contained in AI.

Primary note proof

Study-note section AI34

> Feedforward networks pass signals toward outputs without recurrent state loops; recurrent networks carry state across sequence steps. Training adjusts weights, often through gradients and backpropagation. Backpropagation is not the same as a recurrent connection. Deep learning uses multiple representation layers and can learn features from images, sound and text. CNNs share filters across image locations. ANN is a brain-inspired abstraction, not a literal brain simulation; DL⊂ML⊂AI.

Study-note section AI33

> An RL agent observes state information, takes an action, receives reward and a new observation, and improves its policy to maximize expected cumulative reward. Rewards evaluate outcomes rather than supply the correct action label. Rewards may be delayed; not every correct-looking action receives immediate reward.

Question sources

Artificial Intelligence review questions · Question 49

Question 224 | NS0050

A recommendation system must directly represent products and labeled “frequently bought together” relationships, then traverse those relationships to find connected products. Which listed representation most directly fits that requirement?

A. Rule-based IF-THEN logic

B. Frames describing product attributes

C. Graph-based model with products as nodes and relationships as edges

D. Probabilistic Bayesian network of user demographics

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer C: Graph-based model with products as nodes and relationships as edges

A semantic graph represents entities as nodes and their labeled relationships as edges, making connected products directly traversable. Frames primarily organize attributes, rules express conditions and conclusions, and a Bayesian demographic model represents uncertain dependencies rather than the requested product-edge structure. The answer follows the explicit relational representation requirement; it does not claim that graphs are universally best for every recommendation system.

Why each choice is right or wrong

A - Incorrect: Rules can encode recommendations, but do not directly provide the requested product-node relationship structure.

B - Incorrect: Frames describe product attributes rather than making product-to-product links the primary structure.

C - Correct: Products as nodes and labeled relationships as edges directly support the required traversal.

D - Incorrect: A demographic Bayesian network models uncertain dependencies rather than the specified product-edge structure.

Primary note proof

Study-note section AI24

> A semantic network connects concept nodes by labeled relations such as is-a. Frames describe stereotyped entities through slots, values and defaults. Production systems use IF conditions THEN conclusions or actions, matching rules against working memory and selecting among applicable rules. A graph suits relationships such as frequently bought together; frames suit attributes. No representation is always best for recommendation without task requirements.

Study-note section AI29

> A Bayesian network is a directed acyclic graph with conditional probability tables; its joint distribution factors as the product of each node’s probability given its parents. A Markov model

assumes the next state depends on the current state rather than the entire past. These represent uncertainty; fuzzy membership represents a different issue.

Question sources

Artificial Intelligence review questions · Question 50

Question 225 | NS0156

What is the difference between supervised and unsupervised learning?

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer: Supervised learning fits labeled input-target examples; unsupervised learning discovers structure without target labels.

AI31 defines supervised learning through labels, with classification for discrete targets and regression for numeric targets. AI32 defines unsupervised discovery, including clustering, association and dimensionality reduction. The distinction is the supplied training feedback, not whether the data are images, text or numbers; the same modality can support either regime.

Primary note proof

Study-note section AI31

Supervised learning fits a mapping from labeled input-target examples. Classification predicts discrete labels, such as institution generations or spam/nonspam; regression predicts numeric quantities, such as prices. Predictive analytics uses historical patterns to forecast outcomes and can use either. Specify the target before choosing a method.

Study-note section AI32

Unsupervised learning seeks structure without target labels: clustering groups similar examples; association finds co-occurrence; dimensionality reduction compresses features. Good clusters have high within-cluster similarity and low between-cluster similarity under an appropriate distance. K-means alternates assignment to nearest centroids and centroid updates.

Question sources

Artificial Intelligence review questions · Question 1

Question 226 | NS0157

Under what search and heuristic conditions does A* find a least-cost path, and why?

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer: A* is optimal under its heuristic and search assumptions: admissible tree search, or consistent graph search without reopening; otherwise reopen improved states.

With admissible h, A* tree search uses f=g+h as a lower bound on solution cost through each frontier node. Removing a minimum-f goal cannot skip a cheaper represented solution. Graph search that never reopens closed states normally requires consistency; with an admissible inconsistent heuristic, improved states must be reopened. Goal testing is on removal. Finite branching and positive lower-bounded step costs supply the usual infinite-space completeness conditions.

Primary note proof

Study-note section AI18

> A* orders by f(n)=g(n)+h(n). Admissibility means 0≤h(n)≤h*(n), the true cheapest remaining cost. Consistency means h(n)≤c(n,n′)+h(n′) on every edge, with h(goal)=0; it implies admissibility on states with goal paths. Consistency makes f nondecreasing along a path.

Study-note section AI18

> Admissibility suffices for optimal A* tree search; graph search that never reopens CLOSED states normally requires consistency. With an admissible inconsistent heuristic, reopen states when g improves. Test a goal when removed at minimum f. Completeness also needs finite branching and positive lower-bounded costs in infinite spaces. Worst-case time and memory can be exponential; no heuristic universally guarantees speed.

Question sources

Artificial Intelligence review questions · Question 2

Question 227 | NS0158

What is the Turing Test and what does it measure?

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer: The conversational Turing test asks a judge to distinguish machine from human through dialogue; it measures humanlike behavior, not consciousness.

AI02 sets the observable criterion: whether a judge can reliably tell a machine's conversation from a human's. Language use, knowledge, reasoning and often learning support the performance. Passing would meet this operational behavioral standard, but it does not prove consciousness, actual understanding or that the machine internally thinks in the same way as a human.

Primary note proof

Study-note section AI02

> The conversational Turing test asks whether a judge can reliably distinguish a machine from a human through their dialogue. It requires language use, knowledge, reasoning and often learning. It tests an operational behavior criterion, not proof of consciousness. Cognitive models can be tested against response times, errors and psychological experiments.

Question sources

Artificial Intelligence review questions · Question 3

Question 228 | NS0159

What is the difference between precision and recall?

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer: Precision=TP/(TP+FP): how many predicted positives are correct. Recall=TP/(TP+FN): how many actual positives are found.

AI36 distinguishes precision's predicted-positive denominator from recall's actual-positive denominator. False positives reduce precision; false negatives reduce recall. With TP=30, FP=10 and FN=20, precision is 30/40=75% and recall is 30/50=60%. The positive class and zero-denominator convention should be fixed before comparing scores.

Primary note proof

Study-note section AI36

> With a fixed positive class, TP means true positives, FP false positive predictions, FN missed positives, TN true negatives. Accuracy=(TP+TN)/N; precision=TP/(TP+FP); recall=TP/(TP+FN); F1=2PR/(P+R). State the averaging and zero-denominator convention when relevant.

Question sources

Artificial Intelligence review questions · Question 4

Question 229 | NS0160

Why is tokenization important in natural language processing?

Hint

Match the input, output and purpose to the appropriate language, perception or physical-action task.

Answer: Tokenization divides text into processing units such as words, punctuation or subwords so later NLP stages can operate on a consistent input sequence.

AI37 defines tokenization as dividing text into processing units and lists several permitted granularities. Establishing that sequence lets later tagging, recognition and parsing refer to units consistently. Tokens need not equal whitespace-separated words, and tokenization by itself does not determine grammatical structure or resolve word meaning; those belong to later analyses.

Primary note proof

Study-note section AI37

> NLP interprets and generates human language: translation, dialogue, summarization, question answering and text generation. Speech recognition converts audio to text; speech synthesis generates speech. They overlap language processing but solve different input/output tasks.

Study-note section AI37

> Tokenization divides text into processing units; tokens may be words, punctuation or subwords. POS tagging assigns grammatical roles; named-entity recognition identifies people, organizations and locations. Parsing identifies grammatical structure; semantics concerns meaning; discourse links sentences; pragmatics uses context and speaker intent. Sentiment analysis predicts evaluative tone. Ambiguity can occur at every level: “bank” has multiple senses, and “I saw the man with a telescope” has multiple attachments.

Question sources

Artificial Intelligence review questions · Question 5

Question 230 | NS0285

Which AI approach most directly supports understanding biological intelligence behavior as a scientific goal?

A. Weak AI

B. Applied AI

C. Cognitive AI

D. Strong AI

Hint

Read the definition and the classification axis requested; distinguish the defining purpose from supporting techniques or implementation resources.

Answer C: Cognitive AI.

AI01 explicitly describes cognitive AI as using computational models to investigate biological intelligence, making it the direct scientific-goal match. Applied AI emphasizes useful solutions. Narrow AI describes capability scope, while strong AI has a philosophical understanding sense; neither is the named research approach for testing models against biological behavior.

Why each choice is right or wrong

A - Incorrect: Weak AI denotes bounded capability rather than the specific scientific modeling goal.

B - Incorrect: Applied AI builds useful solutions rather than primarily explaining biological intelligence.

C - Correct: Cognitive AI uses computational models to investigate biological intelligence and behavior.

D - Incorrect: Strong AI concerns broad or actual-understanding claims, not this specific empirical modeling approach.

Primary note proof

Study-note section AI01

> Artificial intelligence studies systems that perform tasks requiring intelligence: reasoning, learning, planning, perception and language. Applied AI builds useful solutions; cognitive AI uses computational models to investigate biological intelligence. A model should be compared with human behavior or experimental observations before claiming it explains human thought.

Study-note section AI01

> Narrow or weak AI addresses bounded tasks, such as face recognition, spam filtering or a voice assistant. General AI denotes broad competence transferable across domains; strong AI also has a philosophical meaning involving actual understanding. Do not equate a humanoid body, autonomous vehicle, or many features with general intelligence.

Question sources

Computer Science Exit Exam · 2018 · Question 2

Question 231 | NS0291

Why is Greedy Best-First Search not guaranteed to be optimal or complete?

A. It ignores path cost g(n) and relies only on h(n)

B. It always chooses the shallowest node

C. It expands nodes in FIFO order

D. It requires full knowledge of the goal state in advance

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer A: It uses h alone, ignoring accumulated g.

AI17 says greedy best-first chooses f=h and can take an expensive route or follow an infinite misleading branch. Ignoring g explains nonoptimality, while cycles or infinite paths without appropriate handling explain failure of completeness. On a finite graph with full duplicate handling it is complete, so the lack of a general guarantee does not mean every implementation fails.

Why each choice is right or wrong

A - Correct: Using only h ignores total path cost and can follow misleading or unbounded branches.

B - Incorrect: Shallowest-first describes breadth-first order rather than greedy heuristic order.

C - Incorrect: FIFO is breadth-first queue behavior, not minimum-h priority selection.

D - Incorrect: Heuristics estimate promise; exact full goal-state knowledge is not required by the algorithm.

Primary note proof

Study-note section AI17

> Best-first is a family that orders OPEN by an evaluation function. Greedy best-first uses f=h and removes the smallest estimate, ignoring the cost already paid. It can choose an expensive route, loop without duplicate checks, or follow an infinite misleading branch. On a finite graph with full duplicate handling it is complete, but not generally optimal. Its implicit finite-tree worst case can be O(b^m).

Study-note section AI19

> For S→A cost 2, S→B cost 1, A→G cost 2, B→G cost 10, set h(A)=2, h(B)=1, h(G)=0. After expanding S, greedy chooses B because 1<2 and can then choose G, returning cost 11. A* compares A:f=2+2=4 and B:f=1+1=2; after B, its G entry costs 11. It next expands A and replaces G with cost 4, which is removed and accepted.

Question sources

Computer Science Exit Exam · 2018 · Question 8

Question 232 | NS0305

Which search strategy tries to expand the node that is closest to the goal, on the grounds that this is likely to lead to a solution quickly; Thus, it evaluates nodes by using just the heuristic function: f (n) = h (n).

A. Uniform-cost search

B. Greedy best-first search

C. Depth-first search

D. A* search

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer B: Greedy best-first search.

AI17 explicitly defines greedy best-first by removing the smallest remaining-cost estimate h while ignoring paid cost g. This is exactly the stem's equation. UCS uses g, A* combines g and h, and DFS uses depth/stack discipline rather than a minimum heuristic evaluation.

Why each choice is right or wrong

A - Incorrect: UCS orders by accumulated g rather than h alone.

B - Correct: Greedy best-first selects the node with the smallest heuristic estimate h.

C - Incorrect: DFS follows a LIFO/deepest-branch order rather than minimum h.

D - Incorrect: A* uses g+h, while the stem explicitly uses only h.

Primary note proof

Study-note section AI17

> Best-first is a family that orders OPEN by an evaluation function. Greedy best-first uses f=h and removes the smallest estimate, ignoring the cost already paid. It can choose an expensive route, loop without duplicate checks, or follow an infinite misleading branch. On a finite graph with full duplicate handling it is complete, but not generally optimal. Its implicit finite-tree worst case can be O(b^m).

Study-note section AI15

> UCS uses a priority queue ordered by actual accumulated path cost g(n). Remove a minimum-g node; accept a goal when removed, not when first generated. Relax each edge: if a cheaper path is found, update its frontier cost and parent.

Study-note section AI18

> A* orders by f(n)=g(n)+h(n). Admissibility means 0≤h(n)≤h*(n), the true cheapest remaining cost. Consistency means h(n)≤c(n,n′)+h(n′) on every edge, with h(goal)=0; it implies admissibility on states with goal paths. Consistency makes f nondecreasing along a path.

Study-note section AI14

DFS follows the deepest pending branch using a LIFO stack or recursion. FILO means the same order as LIFO. With the same children as BFS, a left-first DFS visits S,A,C before B; push rightmost children first to pop the leftmost first.

Question sources

Computer Science Exit Exam · 2018 · Question 23

Question 233 | NS0312

Which of the following best defines a heuristic in the context of AI search?

A. A function that estimates how close a state is to the goal

B. A guaranteed optimal path to the goal

C. A method that always avoids revisiting states

D. A brute-force enumeration of all possible states

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer A: A problem-specific estimate of remaining cost or promise.

AI17 defines a heuristic h(n) as estimating remaining cost or evaluating promise. A's closeness wording captures that estimate rather than an exact path guarantee. A heuristic does not inherently prevent repeated states and is not exhaustive enumeration; those are separate search-policy or bookkeeping choices. Its quality and computation cost affect practical search.

Why each choice is right or wrong

A - Correct: Estimating closeness or remaining cost is the heuristic role defined in AI17.

B - Incorrect: An estimate is not itself a guaranteed optimal path.

C - Incorrect: Repeated-state avoidance requires duplicate handling rather than following from heuristic estimation.

D - Incorrect: Exhaustive enumeration is a search method, not the definition of a heuristic estimate.

Primary note proof

Study-note section AI17

> A heuristic h(n) estimates remaining cost or evaluates promise using problem-specific knowledge. It can reduce explored work, but computing it also costs time. Uninformed algorithms use only the problem definition and known costs; UCS remains uninformed despite using g.

Study-note section AI12

> A search node records a state plus parent, action, depth and path cost. Expanding a node generates its successors. OPEN or frontier stores discovered candidates awaiting expansion; CLOSED records expanded states. Reaching the same state through different paths does not imply the same node or cost.

Question sources

Computer Science Exit Exam · 2018 · Question 35

Question 234 | NS0326

Which type of heuristic search strategy that evaluates nodes by combining g (n), the cost to reach the node, and h (n.), the cost to get from the node to the goal: f (n) = g (n) + h (n) Since g(n) gives the path cost from the start node to node n, and h(n) is the estimated cost of the cheapest path from n to the goal, we have f (n) = estimated cost of the cheapest solution through n ?

A. A* search

B. Greedy Best-First Search

C. Uniform-cost search

D. Depth-first search

Hint

Check how the method selects its next candidate, what information it uses, and which cost, depth or heuristic assumptions the question supplies.

Answer A: A* search.

AI18 defines A* by f(n)=g(n)+h(n), exactly as the long stem specifies. The sum estimates the cheapest solution through the node when h estimates remaining cost. Greedy ignores g, UCS ignores h, and DFS does not prioritize this sum, so none of those matches the stated evaluation.

Why each choice is right or wrong

A - Correct: A* adds actual cost-to-node g and estimated remaining cost h.

B - Incorrect: Greedy best-first uses only h and omits accumulated path cost.

C - Incorrect: UCS uses only g and does not add a remaining-cost heuristic.

D - Incorrect: DFS uses stack/depth ordering rather than minimum g+h.

Primary note proof

Study-note section AI18

> A* orders by f(n)=g(n)+h(n). Admissibility means 0≤h(n)≤h*(n), the true cheapest remaining cost. Consistency means h(n)≤c(n,n′)+h(n′) on every edge, with h(goal)=0; it implies admissibility on states with goal paths. Consistency makes f nondecreasing along a path.

Study-note section AI17

> Best-first is a family that orders OPEN by an evaluation function. Greedy best-first uses f=h and removes the smallest estimate, ignoring the cost already paid. It can choose an expensive route, loop without duplicate checks, or follow an infinite misleading branch. On a finite graph with full duplicate handling it is complete, but not generally optimal. Its implicit finite-tree worst case can be O(b^m).

Study-note section AI15

UCS uses a priority queue ordered by actual accumulated path cost g(n). Remove a minimum-g node; accept a goal when removed, not when first generated. Relax each edge: if a cheaper path is found, update its frontier cost and parent.

Study-note section AI14

DFS follows the deepest pending branch using a LIFO stack or recursion. FILO means the same order as LIFO. With the same children as BFS, a left-first DFS visits S,A,C before B; push rightmost children first to pop the leftmost first.

Question sources

Computer Science Exit Exam · 2018 · Question 52

Question 235 | NS0349

In which of the following Learning methods an agent receives feedback in the form of rewards or penalties to maximize cumulative reward?

A. Supervised learning

B. Reinforcement learning

C. Semi-supervised learning

D. Unsupervised learning

Hint

Identify the training feedback, prediction target or measured error. Keep learning objectives separate from model architecture and data modality.

Answer B: Reinforcement learning.

AI33 defines RL through interaction, rewards and improvement of a policy to maximize expected cumulative reward. The reward may be delayed and does not provide a correct-action label. Supervised and semi-supervised learning instead concern target labels, while unsupervised discovery lacks the defining action-reward interaction described here.

Why each choice is right or wrong

A - Incorrect: Supervised learning fits labeled targets rather than defining behavior through cumulative interaction rewards.

B - Correct: Reward or penalty feedback guiding cumulative-return improvement is the defining RL mechanism.

C - Incorrect: Semi-supervised learning mixes labeled and unlabeled examples, not necessarily action rewards.

D - Incorrect: Unsupervised learning discovers unlabeled structure rather than optimizing an interaction reward policy.

Primary note proof

Study-note section AI33

An RL agent observes state information, takes an action, receives reward and a new observation, and improves its policy to maximize expected cumulative reward. Rewards evaluate outcomes rather than supply the correct action label. Rewards may be delayed; not every correct-looking action receives immediate reward.

Study-note section AI31

Supervised learning fits a mapping from labeled input-target examples. Classification predicts discrete labels, such as institution generations or spam/nonspam; regression predicts numeric quantities, such as prices. Predictive analytics uses historical patterns to forecast outcomes and can use either. Specify the target before choosing a method.

Study-note section AI32

Semi-supervised learning combines some labels with additional unlabeled data. Active learning chooses informative cases for an oracle to label; it still uses labels. “Exploratory learning” is not an exclusive standard synonym: it can mean unsupervised exploration or active data acquisition, so define the feedback. A task named summarization or time-series analysis is not automatically supervised; its training targets determine that.

Study-note section AI32

Unsupervised learning seeks structure without target labels: clustering groups similar examples; association finds co-occurrence; dimensionality reduction compresses features. Good clusters have high within-cluster similarity and low between-cluster similarity under an appropriate distance. K-means alternates assignment to nearest centroids and centroid updates.

Question sources

Computer Science Exit Exam · 2018 · Question 83