CS Exit Ethiopia

Revision Outlines

Computer Science Exit Examination

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

Introduction to Artificial Intelligence

Revision Outlines

AI capabilities and scope | Exam focus

- Artificial intelligence scope
  - Reasoning; learning; planning; perception; language
  - Applied AI: useful solutions
  - Cognitive AI: biological intelligence models
  - Human behavior and experimental validation
- Narrow, general and strong AI
  - Weak AI: bounded tasks
  - Face recognition; spam; voice assistants
  - General AI: transferable domain competence
  - Strong AI: actual understanding
  - Humanoid appearance ≠ general intelligence
  - Autonomous vehicles ≠ general intelligence
  - Feature count ≠ general intelligence
- Exam trap Data analysis-AI overlap
  - Learning or inference criterion
  - Categorical exclusion trap

Four approaches and the Turing test | Exam focus

- Four AI approaches
  - Thinking humanly: mental processes
  - Thinking rationally: valid logic
  - Acting humanly: humanlike behavior
  - Acting rationally: expected objective achievement
  - Thought versus behavior evaluation
- Conversational Turing test
  - Judge: human-machine dialogue distinction
  - Language; knowledge; reasoning; learning
  - Operational behavior criterion

- Consciousness-proof trap
- Cognitive models: response times and errors
- Psychological experiments

• Worked example Approach identification
- Best action sequence: acting rationally
- Human-error modeling: thinking humanly

Agents, percepts and action | Exam focus

• Agent interaction
- Sensors perceive; actuators/effectors act
- Cameras versus motors
- Software messages versus API requests
- Human agents versus artificial agents

• Agent function and architecture
- Percept: single observation
- Percept sequence: entire observation history
- f:P*→A
- Agent = architecture + program
- Program implements function
- Environment beyond search space
- Autonomy: experience over designer instructions

Rationality and performance | Exam focus

• Rationality and success measures
- External performance: safety; deliveries; time; cost
- Prior knowledge and percept history
- Available actions
- Best expected performance given information

• Rationality versus omniscience
- Omniscience: advance outcome knowledge
- Rationality without guaranteed success
- All actions fail boundary case
- Rational information gathering and exploration

- Worked example Expected versus realized route time
  - Rational expectation: 20 minutes
  - Unforeseen closure: 40 minutes

Specify the task with PEAS | Exam focus

- PEAS task specification
  - Performance measure; Environment; Actuators; Sensors
  - Specification before agent design
  - Taxi performance: safety and travel time
  - Taxi environment: roads; traffic; passengers
  - Taxi actuators: steering; brakes; displays
  - Taxi sensors: cameras; positioning; speed
- PEAS category traps
  - Sensor influence ≠ actuation
  - Performance measure: success criterion
  - Percept sequence: observation evidence
  - Agent observations versus designer knowledge

Simple and model-based reflex agents

- Simple reflex agent
  - Current-percept condition-action rules
  - Dirty → clean
  - No persistent world-state memory
  - Sufficient current information assumption
  - Partial observation: identical percept, different actions
- Model-based reflex agent
  - Previous state and previous action
  - Current percept
  - World-change and action-effect models
  - Internal-state condition-action rules
  - Shared rule use
  - Distinguishing memory and world model
- Worked example Hidden cleaned-room memory

- Internal state beyond sensor range
- Current dirt detection: no memory requirement

Goals, utility and model updates

- Goals versus utility
  - Goal: desired-state action sequences
  - Utility: desirability and trade-offs
  - Maximum expected utility
  - EU(a)=Σs P(s|a,evidence)U(s)
  - Destination goal
  - Arrival time; safety; cost preferences
- Internal model versus external action
  - Post-action state prediction
  - State update ≠ environment modification
  - Unused goal variable trap
  - BEST_ACTION dependence on goal
  - Coexisting model, goal, utility, learning capabilities

Learning-agent components | Exam focus

- Learning-agent components
  - Performance element: action selection
  - Learning element: future improvement
  - Critic: standard-based evaluation and feedback
  - Problem generator: useful exploratory experience
- Learning: task T, performance P, experience E
  - Spam T: email classification
  - Spam E: labeled messages
  - Spam P: recall at required precision
  - Exploration: information versus immediate reward
  - Exploitation: currently best-known action
  - Learning across reflex, goal, utility agents

Environment dimensions | Exam focus

- Observation and transition dimensions
  - Fully observable: relevant state revealed
  - Partially observable: hidden or noisy information
  - Deterministic: state/action determine successor
  - Stochastic: random outcomes
  - Strategic opponent uncertainty versus randomness
- Environment dimension contrasts
  - Episodic independence versus sequential effects
  - Static deliberation versus dynamic change
  - Semidynamic: clock changes score
  - Discrete alternatives versus continuous ranges
  - Known/unknown: transition-model knowledge
  - Model knowledge ≠ observability
- Worked example Environment classification examples
  - Chess: sequential; discrete; deterministic given moves
  - Untimed chess: static
  - Timed chess: semidynamic
  - Fixed crossword: static
  - Driving: partial; sequential; dynamic; partly continuous

Single-state, multiple-state and contingency problems

- State uncertainty representations
  - Single-state: known deterministic, fully observed
  - Sensorless/multiple-state: no observation
  - Known deterministic transitions
  - Belief state: possible-state set
- Contingency versus exploration
  - Nondeterminism or partial observation
  - Percept-dependent plan branches
  - Unknown state space or transition model
  - Exploration/online search
  - Observation alone insufficient

- Overlapping classifications
- Partial deterministic description: ambiguous category

Formulate before searching | Exam focus

- Goal then problem formulation
  - Initial state and legal actions
  - Transition model
  - Goal test and action costs
  - Search: sequence selection
  - Execution: applying actions
  - State abstraction: irrelevant details removed
- Search problem examples
  - Routes: cities; roads; neighboring successors
  - Accumulated distance or time
  - Solution: legal goal-ending path
  - Optimal: minimum path cost
  - Automatic programming/program synthesis
  - Program satisfying specification
  - Conditional plans and loops: richer representations

Frontiers and evaluation criteria | Exam focus

- Search nodes and frontier
  - State; parent; action; depth; path cost
  - Expansion generates successors
  - OPEN/frontier: pending candidates
  - CLOSED: expanded states
  - Same state ≠ same node/cost
- Search evaluation criteria
  - Completeness under stated assumptions
  - Optimality: best solution
  - Time: generated work
  - Space: simultaneous storage
  - b: branching; d: shallowest goal

- m: maximum depth; l: limit
- Graph V,E versus implicit-tree bounds

Breadth-first search | Exam focus

- BFS: FIFO oldest-first
  - Root → head removal → goal test
  - Unseen successors appended at tail
  - Depth order: 0,1,2,…
  - Complete: finite b, finite goal depth
  - Minimum number of edges
  - Minimum cost: equal positive step costs
- BFS complexity conventions
  - 1+b+…+b^d
  - Generation-time testing: O(b^d)
  - Depth-d expansion: O(b^(d+1))
  - Explicit visited graph: O(V+E) time
  - Explicit visited graph: O(V) space
- Worked example BFS queue trace
  - Children S:[A,B], A:[C,D]
  - Removal: S,A,B,C,D
  - Child-order assumption

Depth-first search | Exam focus

- DFS: deepest pending branch
  - LIFO stack or recursion
  - FILO = LIFO
  - Left-first trace: S,A,C before B
  - Push rightmost first
- DFS guarantees and bounds
  - Implicit-tree time: O(b^m)
  - Pending-sibling space: O(bm)
  - Path-only recursive space: O(m)
  - Generally nonoptimal

- Infinite-path and cycle failure
- Finite visited graph: complete
- Visited-state storage: O(V)
- Bank space O(b^m): time confusion

Uniform-cost search and Dijkstra | Exam focus

- UCS: minimum actual g(n)
  - Priority queue
  - Goal acceptance on removal
  - First-generation goal trap
  - Edge relaxation: cheaper path
  - Updated frontier cost and parent
- UCS and Dijkstra conditions
  - Finite branching; step cost ≥ε>0
  - UCS completeness and optimality
  - Finite graph: nonnegative edges, duplicate handling
  - Dijkstra: reachable minimum distances
  - UCS: goal stop, on-demand states
  - Queue initialization: implementation choice
  - Negative edges: finalized-distance failure

Depth Limits and Bidirectional Search | Exam focus

- Depth-limited and iterative deepening
  - DLS: DFS to l
  - Cutoff versus failure
  - Goals beyond l missed
  - Limits: 0,1,2,…
  - Finite b: shallowest solution
  - Time O(b^d); frontier space O(bd)
  - Optimal: equal step costs
  - Deepest level dominates repeated work
- Bidirectional search and IDA*
  - Start and goal frontiers

- Reverse actions and meeting tests
- Equal-cost BFS: O(b^(d/2)) time/space
- Weighted search: sound stopping bound
- First intersection ≠ cheapest
- IDA*: f=g+h threshold deepening

Heuristics and greedy best-first search | Exam focus

• Heuristic versus uninformed search
- h(n): remaining cost or promise
- Problem-specific knowledge
- Reduced exploration versus heuristic computation cost
- Uninformed: definition and known costs
- UCS uses g, remains uninformed

• Greedy best-first search
- Best-first family: evaluation-ordered OPEN
- Greedy f=h; minimum estimate
- Past cost ignored
- Expensive route trap
- Loops without duplicate checks
- Infinite misleading branches
- Finite graph, full duplicate handling: complete
- Generally nonoptimal
- Finite-tree worst case: O(b^m)

Admissibility, consistency and A* | Exam focus

• A*, admissibility and consistency
- f(n)=g(n)+h(n)
- Admissible: 0≤h(n)≤h*(n)
- h*: true cheapest remaining cost
- Consistent: h(n)≤c(n,n′)+h(n′)
- h(goal)=0
- Consistency ⇒ admissibility on goal-reaching states
- Consistency: nondecreasing path f

- A* guarantees and traps
  - Tree optimality: admissible h
  - No-reopen graph search: consistent h
  - Admissible inconsistent h: reopen improved g
  - Goal test on minimum-f removal
  - Infinite-space completeness: finite branching
  - Positive lower-bounded step costs
  - Exponential worst-case time and memory
  - No universal heuristic speed guarantee

Work a heuristic-search example | Exam focus

- Greedy versus A* trace
  - Edges: SA=2, SB=1, AG=2, BG=10
  - h(A)=2; h(B)=1; h(G)=0
  - Greedy B→G: cost 11
  - A*: f(A)=4, f(B)=2
  - Expand B; G cost 11
  - Expand A; replace G cost 4
  - Remove goal: optimal cost 4
- Heuristic lower-bound examples
  - Unit sliding-puzzle moves
  - Misplaced tiles excluding blank
  - Summed Manhattan distances excluding blank
  - Straight-line ≤ road distance
  - Travel-time conversion: valid speed bound

Local optimization and evolutionary search | Exam focus

- Local search and hill climbing
  - Complete candidate; objective improvement
  - Path may be irrelevant
  - First-improvement: first better neighbor
  - Steepest-ascent: all neighbors
  - Stochastic choice among improvements

- Local maxima; plateaus; ridges
- Random restarts; limited sideways steps
- No guaranteed success

• Genetic and means-ends search
- Population and fitness evaluation
- Parent selection
- Crossover; mutation; replacement
- No global optimality guarantee
- Means-ends: current-goal difference
- Difference-reducing operator
- Recursive precondition solving
- Backward A* misconception

Constraint satisfaction | Exam focus

• Constraint satisfaction problem
- Variables X; domains D; constraints C
- Solution: every variable, every constraint
- Complete: all variables assigned
- Consistent: no violated constraints
- Complete assignment may be inconsistent

• CSP search and propagation
- Backtracking: assign, try, undo failure
- Fewest remaining values
- Degree tie-breaker
- Least-constraining value
- Forward checking: incompatible neighbor removal
- Arc consistency: unsupported binary values removed
- Problem framework, not fourth uninformed traversal

Games and minimax | Exam focus

• Adversarial game formulation
- Competing objectives
- States and player to move

- Legal actions and successor function
- Terminal test and utility
- Perfect information: visible state/history
- Future choices still unknown
- Perfect information may include chance
- Zero-sum: opposite gains/losses

• Minimax and chance nodes
- Terminal utilities
- MIN: minimum; MAX: maximum
- Maximum-valued root action
- Best guarantee against optimal opposition
- Modeled-tree limitation
- Chance nodes: probability-weighted outcomes

• Worked example Minimax backup example
- MIN[3,5]=3
- MIN[2,9]=2
- MAX chooses first branch: 3

Alpha-beta and game evaluation | Exam focus

• Alpha-beta pruning
- Alpha: MAX lower bound
- Beta: MIN upper bound
- Bounds passed downward
- MAX updates alpha; MIN updates beta
- Prune: alpha≥beta
- Same-tree minimax value preserved
- Good move ordering
- Worst O(b^m); ideal O(b^(m/2))

• Cutoff evaluation functions
- Estimated position desirability
- Positive MAX; negative MIN
- +1/0/−1: win/draw/loss convention
- Material: weighted pieces

- Mobility: legal options
- Territory: controlled regions
- Combined features
- Limited computation, no victory guarantee

Represent knowledge for inference | Exam focus

- Knowledge representation and inference
  - Facts; objects; relations; events; rules
  - Syntax: legal expressions
  - Semantics: meaning
  - Knowledge base: assertions
  - Inference system: derived consequences
  - Declarative: what true
  - Procedural: how to act
- Knowledge representation structures
  - Semantic networks: labeled concept relations
  - is-a links
  - Frames: slots; values; defaults
  - Production systems: IF-THEN rules
  - Working-memory matching and rule selection
  - Graph: frequently-bought-together relationships
  - Frames: attributes
  - Recommendation requirements determine suitability

Propositional logic and models | Exam focus

- Propositional connectives
  - Atoms: whole truth-valued claims
  - ¬P: negation
  - P∧Q: both; P∨Q: at least one
  - P→Q false: P true, Q false
  - P↔Q: matching truth values
  - P→Q ≡ ¬P∨Q
  - De Morgan: ¬(P∧Q) ≡ ¬P∨¬Q

- Models, validity and entailment
  - Model: satisfying interpretation
  - Valid: every interpretation
  - Satisfiable: at least one
  - Unsatisfiable: no interpretations
  - Invalid ≠ unsatisfiable
  - P∨¬P: valid
  - P: satisfiable, invalid
  - P∧¬P: unsatisfiable
  - KB⊨α: all KB models satisfy α

First-order logic and quantifier scope | Exam focus

- First-order logic vocabulary
  - Objects; predicates; functions; quantified variables
  - Human(a): property
  - Likes(a,b): relation
  - father(a): function term
  - Propositional logic: no object quantifiers
  - Assigned atom values determine truth
- Quantifier translation and scope
  - All humans: ∀x(Human(x)→Mortal(x))
  - Tall student: ∃x(Student(x)∧Tall(x))
  - Translation patterns, not connective prohibitions
  - ∀x∃y Likes(x,y): potentially different y
  - ∃y∀x Likes(x,y): shared y
  - ¬∀x P(x) ≡ ∃x¬P(x)
  - Negation swaps quantifiers
  - Bound occurrences within correct scope
- Worked example Existential conjunction trap
  - ∃x(StandingAt(x,door)∧Tall(x))
  - Same individual, both predicates
  - Implication: unrelated vacuous witness

Inference, unification and chaining | Exam focus

- Inference and unification
  - Modus ponens: $P, P \rightarrow Q \vdash Q$
  - Modus tollens: $P \rightarrow Q, \neg Q \vdash \neg P$
  - Affirming consequent trap
  - Universal instantiation: object substitution
  - Consistent unifying substitution
  - Knows(x,Ada),Knows(Ben,y): x=Ben,y=Ada
  - Distinct constants cannot unify
- Resolution and chaining
  - $P \vee Q, \neg P \vee R \vdash Q \vee R$
  - Refutation: KB plus $\neg \alpha$
  - Clause form and contradiction
  - Forward chaining: facts → rules
  - Backward chaining: query → required premises
  - Soundness: truth preservation
  - Completeness: all entailed conclusions
  - Specified-setting qualification

Deduction, induction and abduction | Exam focus

- Reasoning forms
  - Deduction: truth-preserving rules
  - Mortal people + Ada → mortal Ada
  - Induction: observations → generalization
  - New examples may overturn induction
  - Abduction: plausible explanation
  - Wet ground: rain or sprinklers
  - Analogy: similar-situation pattern transfer
- Monotonic and default reasoning
  - Classical entailment: consequences persist
  - Added premises preserve deductions
  - Nonmonotonic defaults: retractable conclusions
  - Bird usually flies; penguin exception

- Incomplete knowledge; no truth guarantee
- Blind search ≠ reasoning definition

Probability and uncertain models | Exam focus

- Probability and Bayes
  - $P(\neg A)=1-P(A)$
  - $P(A|B)=P(A\cap B)/P(B)$, $P(B)>0$
  - $P(H|E)=P(E|H)P(H)/P(E)$
  - Posterior; likelihood; prior; evidence
  - Exclusive exhaustive hypotheses
  - $P(E)=\Sigma_i P(E|Hi)P(Hi)$
- Uncertain graphical models
  - Bayesian network: directed acyclic graph
  - Conditional probability tables
  - Joint: product of parent-conditioned probabilities
  - Markov: next depends on current state
  - Entire-history dependence excluded by assumption
  - Uncertainty versus fuzzy membership
- Worked example Bayes numerical method
  - Prior 0.1; likelihood 0.8
  - Alternative likelihood 0.2
  - Evidence: $0.08+0.18=0.26$
  - Posterior: $0.08/0.26\approx 0.308$
  - Likelihood ≠ posterior

Fuzzy logic and vague categories | Exam focus

- Fuzzy membership versus probability
  - Graded membership: $0\leq\mu(x)\leq 1$
  - Warm/tall vague concepts
  - Warm degree 0.7 ≠ probability 70%
  - Boolean: true/false
  - Probability: uncertain events
- Fuzzy controller pipeline

- Fuzzify measurements
- Graded IF-THEN rules
- Combine outputs
- Defuzzify numerical action
- Fuzzy control: linguistic degrees
- RL: reward-based adaptation
- Complex/uncertain environment: insufficient method criterion

Supervised learning and prediction | Exam focus

- Supervised prediction
  - Labeled input-target mapping
  - Classification: discrete labels
  - Institution generations; spam/nonspam
  - Regression: numeric targets, prices
  - Predictive analytics: historical-pattern forecasts
  - Classification or regression
  - Target before method
- Decision trees and impurity
  - Internal attribute tests; branches; leaf predictions
  - ID3: high information gain
  - $H(S)=-\Sigma p_i \log_2 p_i$
  - Gain: parent minus weighted child entropy
  - $Gini=1-\Sigma p_i^2$
  - Pure class: Gini zero
  - Pruning against overfitting
  - Supervised neural networks, SVMs, trees

Unsupervised, Semi-supervised and Active Learning | Exam focus

- Unsupervised structure discovery
  - No target labels
  - Clustering: similar-example groups
  - Association: co-occurrence
  - Dimensionality reduction: feature compression

- High within-cluster similarity
- Low between-cluster similarity
- Appropriate distance
- K-means: nearest assignment ↔ centroid updates

• Semi-supervised and active learning
- Semi-supervised: labeled plus unlabeled
- Active: informative oracle-label queries
- Active learning still uses labels
- Exploratory learning: ambiguous terminology
- Unsupervised exploration versus active acquisition
- Feedback definition
- Summarization/time-series labels depend on training targets

Reinforcement learning | Exam focus

• Reinforcement learning interaction
- Observe → act → reward/new observation
- Policy improvement
- Expected cumulative reward objective
- Outcome reward ≠ correct-action label
- Delayed rewards
- Correct-looking action without immediate reward

• RL return and policy
- $Gt=rt+1+\gamma rt+2+\gamma^2 rt+3+\ldots$
- $0\leq\gamma<1$
- Policy: states/histories → action choices
- Exploration: information
- Exploitation: currently estimated reward
- RL⊂ML⊂AI
- No labels: RL or unsupervised
- Distinguishing reward-driven interaction

Neural networks and deep learning | Exam focus

• Neuron and activation functions

- z=Σiwixi+b; a=φ(z)
- x=(2,1), w=(0.5,−1), b=0.2
- Weighted sum: z=0.2
- Linear: z
- Sigmoid: 1/(1+e^(−z))
- tanh: bipolar
- ReLU: max(0,z)
- Nonlinear multilayer relationships

• Network architectures and training
- Feedforward: output-directed, no recurrent loops
- Recurrent: state across sequence steps
- Weight adjustment through gradients/backpropagation
- Backpropagation ≠ recurrent connection
- Deep learning: multiple representation layers
- Learned image, sound, text features
- CNN: shared filters across image locations
- ANN: brain-inspired abstraction
- Literal brain-simulation trap
- DL⊂ML⊂AI

Generalization and validation | Exam focus

• Generalization failures and controls
- Overfitting: training details, poor unseen performance
- Underfitting: limited or poorly fitted model
- Learning curve: performance versus training size/progress
- Complexity control
- Regularization; pruning
- More representative data

• Training, validation and test splits
- Training: parameter fitting
- Validation: model/settings selection
- Held-out test: final estimation
- k-fold: k models, k−1 training folds

- Remaining-fold evaluation; average scores
- Stratification: approximate label proportions
- Preprocessing fitted inside training folds
- Time-respecting splits
- Repeated test tuning: information leakage

Measure the right errors | Exam focus

• Classification confusion matrix
- Fixed positive class
- TP: true positives; TN: true negatives
- FP: false alarms; FN: missed positives
- Accuracy=(TP+TN)/N
- Precision=TP/(TP+FP)
- Recall=TP/(TP+FN)
- F1=2PR/(P+R)
- Averaging and zero-denominator conventions

• Metric calculation and selection
- TP=30, FP=10, FN=20, TN=40
- Accuracy=70%; precision=75%; recall=60%
- F1≈66.7%
- Rare-class failure hidden by accuracy
- Missed positives: recall
- Costly false alarms: precision
- MAE=mean$|y-\hat{y}|$
- MSE=mean$(y-\hat{y})^2$
- Squared errors emphasize large misses

Language processing tasks | Exam focus

• Natural language processing tasks
- Language interpretation and generation
- Translation; dialogue; summarization
- Question answering; text generation
- Speech recognition: audio → text

- Speech synthesis: generated speech
- Overlapping fields, distinct input/output tasks

• Language processing levels
- Tokenization: words; punctuation; subwords
- POS tagging: grammatical roles
- NER: people; organizations; locations
- Parsing: grammatical structure
- Semantics: meaning
- Discourse: sentence links
- Pragmatics: context and speaker intent
- Sentiment: evaluative tone
- Bank: lexical ambiguity
- Telescope: attachment ambiguity

Morphology, sounds and prosody

• Morphology and speech sounds
- Meaningful word parts
- un-happi-ness: prefix; stem; suffix
- cats: cat + plural -s
- Tokenization: processing boundaries
- Phonetics: physical speech sounds
- Phonology: sound patterns and contrasts

• Prosody and contextual meaning
- Stress; rhythm; timing
- Prominence and intonation
- Syllables and phrases
- Rising pitch: possible question
- Stress shifts: emphasis
- Dialogue and speech-processing cues
- Pragmatics: intended contextual meaning
- Pragmatics ≠ word-internal structure

Perception, vision and robotics | Exam focus

- Computer vision tasks and limits
  - Image/video interpretation
  - Classification: content names
  - Detection: object locations
  - Segmentation: pixel labels
  - Recognition: object/person identity
  - Occlusion; lighting changes; noise
  - Unfamiliar conditions; incomplete context
  - Face detection versus open-world perception
- Robotic sensing, planning and action
  - Mechanical; electrical; control; computing systems
  - Sensors; controller; power supply
  - Actuators and end effector
  - Legged; wheeled; tracked locomotion
  - Terrain; stability; efficiency trade-offs
  - Autonomous software ≠ robot
  - Humanoid shape ≠ general intelligence

Expert systems and choosing methods | Exam focus

- Expert system components
  - Specialized domain knowledge; equipment diagnosis
  - Domain expert: expertise
  - Knowledge engineer: elicitation and encoding
  - Inference engine: knowledge base/working facts
  - User interface: interaction
  - Explanation facility: why/how conclusions
- Method selection and responsible evaluation
  - Rules: inspectable, brittle, maintenance cost
  - Learned patterns: representative data and validation
  - Search: explicit action alternatives
  - Probability: uncertain evidence
  - Data/evaluation bias and unfairness

- Group-specific errors
- Robustness; privacy; human oversight
- Automation benefit assumption trap

Clarify search and logic questions

- Search-question clarifications
  - A*: forward best-first, g+h
  - No inherent backward frontier
  - A*/IDA*: conditional optimality
  - A*/IDA*: exponential worst cases
  - Minimum heuristic = estimated closest
  - LIFO = FILO
- Ambiguous search and validity stems
  - Current-state information ≠ heuristic proof
  - Estimated remaining cost criterion
  - False: invalid iff unsatisfiable
  - False: satisfiable implies valid
  - Multiple false options
  - Repair: all false or specified claim

Distinguish agent categories and components

- Agent-class naming conventions
  - Five including learning
  - Four decision architectures before learning
  - Count requires convention
  - Target-based may mean goal-based
  - Unity-based ≠ utility-based
- Agent classification evidence
  - Generic rule: memory unspecified
  - Simple versus model-based reflex memory
  - Unused goal ≠ goal-based policy
  - Sensors/effectors: interaction only
  - Interaction ≠ omniscience/intelligence

- Effector/actuator options: nonunique without distinction

Clarify overlapping learning and AI labels

- Overlapping learning labels
  - ANN; genetic search; RL improvement
  - Experience-based performance changes
  - Deep learning: architecture family
  - Supervised; unsupervised; reinforcement objectives
  - Trees and deep models: potentially supervised
  - Regime evidence: labels, rewards, unlabeled structure
- AI-label question traps
  - Face recognition: applied and narrow
  - Emulating intelligence and building learning systems
  - Humanlike problem-solving: insufficient expert-system criterion
  - Python; C++; Java implementations
  - Common language: nonunique answer
  - McCarthy: naming, not sole invention

Source Damage and Assumption Repairs

- Quantifier negation
  - $\forall x \exists y \exists t \neg F(x,y,t)$
  - Each person: failed fooling person/time
  - Equivalent: $\neg \exists x \forall y \forall t\ F(x,y,t)$
  - Nobody fools everybody all the time
  - Correct D; bold C wrong
  - C: stronger prohibition on every success
  - Read the formula carefully
- Implementation and task assumptions
  - Dijkstra/UCS: stated implementations
  - Queue initialization: nonintrinsic difference
  - Chess moves: sequential
  - Independently reset matches: episodic
  - Recommendation suitability requires requirements

- Graph relationships: support, not exclusivity