Study questions on artificial intelligence

1. (a) ; (b) ; (c) ; (d) ; (e) none of these
2. (a) ; (b) ; (c) ; (d) ; (e) none of these
3. (a) ; (b) ; (c) ; (d) ; (e) none of these
4. (a) ; (b) ; (c) ; (d) ; (e) none of these
5. (a) ; (b) ; (c) ; (d) ; (e) none of these

Study questions on Introduction

1. What this course is about

1. Intelligence has been credibly described as (a) very fast information processing; (b) the firing of neurons; (c) remembering correct answers; (d) storing information; (e) making wise choices about what to do next
2. Real-world problems are distinguished from ____ ones (a) imaginary; (b) algorithmic; (c) interactive; (d) toy; (e) none of these
3. The early history of AI is best associated with (a) real brains; (b) the physical symbol system hypothesis; (c) expert systems; (d) rational agents; (e) none of these
4. AI in the late 20^{th} century is best associated with (a) real brains; (b) the physical symbol system hypothesis; (c) expert systems; (d) rational agents; (e) none of these
5. AI in the 21^{st} century is best associated with (a) real brains; (b) the physical symbol system hypothesis; (c) expert systems; (d) rational agents; (e) none of these
6. The instructor associates AI most closely with (a) logic; (b) feelings; (c) hard computational problems; (d) knowledge; (e) algorithms
7. The instructor associates AI most closely with (a) logic; (b) feelings; (c) rational adaptive computational behavior; (d) knowledge; (e) algorithms
8. "Hard computational problems" are identified mathematically with (a) undecidability; (b) algorithms; (c) functions; (d) NP-hardness and exponential time; (e) none of these
9. This course will most emphasize (a) easy problems; (b) expert systems; (c) robotics; (d) interaction; (e) none of these
10. Rationality is associated most closely here with (a) humanness; (b) maximizing reward; (c) proofs; (d) creativity; (e) none of these
11. The behavior that interests us most is (a) algorithmic; (b) fast; (c) based on facts; (d) adaptive; (e) human
12. Behavior is (a) human action; (b) any action; (c) any output; (d) goal-driven action; (e) collaboration
13. Computation always involves (a) silicon; (b) deterministic algorithms; (c) processing of symbols; (d) numbers; (e) none of these
14. We point to a distinction between (a) good and bad AI; (b) elegant and simplistic algorithms; (c) toy and real-world problems; (d) logic and inference; (e) none of these
15. Early AI research explored (a) artificial neurons; (b) simulation of an entire brain; (c) autonomous agents; (d) reinforcement learning; (e) multi-agent systems
16. AI in two decades after 1969 explored (a) expert systems; (b) simulation of an entire brain; (c) autonomous agents; (d) reinforcement learning; (e) multi-agent systems

2. Course background

1. Hard computational problems are defined in the theory of computational complexity as ones (a) for which no heuristics exist; (b) for which no algorithms exist; (c) that have no solutions; (d) that humans don't try solving; (e) that are believed to require exponential time

Data structures

1. Arrays are structures that are (a) linked; (b) branching; (c) linear; (d) dynamically allocated; (e) none of these
2. A tree is a kind of (a) list; (b) array; (c) graph; (d) all of these; (e) none of these
3. (T-F) A stack is a last-in, first out structure.
4. The height of a binary tree is (a) the number of nodes it contains; (b) the maximum path length between two leaf nodes; (c) the number of leaf nodes; (d) the maximum path length from the root to a leaf node; (e) infinite
5. A graph is (a) a set of integers; (b) a set of vertices; (c) a set of vertices and a set of edges; (d) a set of edges; (e) a set of paths
6. A tree is a graph that is (a) connected and cyclic; (b) connected and acyclic; (c) unconnected and cyclic; (d) unconnected and acyclic; (e) none of these
7. A graph is defined in part by (a) exactly one ordered pair of vertices; (b) a relation; (c) a cycle; (d) one path joining each pair of vertices; (e) none of these.

Mathematical foundations

1. A language is a (a) string; (b) number; (c) set of numbers; (d) sequence of strings; (e) set of strings
2. When A and B are sets, $(A \times B)$ is (a) a set of ordered pairs; (b) an arithmetic expression; (c) a sequence of values; (d) all of these; (e) none of these
3. For array A, $|A|$ is (a) the absolute value of the sum of A's elements; (b) the absolute value of A; (c) the smallest element of A; (d) the number of elements in A; (e) none of these
4. $\subseteq$ denotes (a) set membership; (b) union; (c) conjunction; (d) a relation between sets; (e) negation
5. $\wedge$ denotes (a) set membership; (b) union; (c) AND; (d) a relation between sets; (e) negation
6. $\neg$ denotes (a) set membership; (b) union; (c) AND; (d) a relation between sets; (e) logical negation
7. $\cup$ denotes (a) set membership; (b) union; (c) AND; (d) a set; (e) negation
8. $\emptyset$ denotes (a) set membership; (b) union; (c) AND; (d) a set; (e) negation
9. $\in$ denotes (a) set membership; (b) union; (c) AND; (d) a relation between sets; (e) negation
10. A string is a (a) collection; (b) set; (c) tree; (d) sequence; (e) list
11. Where B and C are sets, $(B \times C)$ is (a) a pair of sets; (b) a set of pairs; (c) an arithmetic product; (d) a sequence; (e) a concatenation
12. $\vee$ denotes (a) set membership; (b) union; (c) AND; (d) OR; (e) implication

13. $\Rightarrow$ denotes (a) set membership; (b) union; (c) AND; (d) OR; (e) implication
14. {1,2,3} $\cup$ {2,4,5} = (a) {}; (b) {1,2}; (c) 2; (d) {2}; (e) {1,2,3,4,5}
15. A *relation* on set *A* is (a) an element of *A*; (b) a subset of *A*; (c) an element of $A \times A$; (d) a subset of $A \times A$; (e) none of these
16. A function *f*: {1,2,3} $\rightarrow$ {0,1} is a set of (a) integers; (b) ordered pairs; (c) sets; (d) relations; (e) none of these

Longer answer

1. Coming into this course, what do you think are essential features of intelligence?
2. Coming into this course, do you have a definition of intelligence or AI?
3. Consider these axioms.
   a. a is true
   b. f is false
   c. $a \rightarrow b$
   d. $b \rightarrow \neg d$
   e. $b \rightarrow c$
   f. $\neg a \rightarrow e$
   g. $\neg f \rightarrow \neg d$
   h. $e \rightarrow \neg f$
   i. $g \rightarrow f$
   j. $h \rightarrow \neg c$
   k. *false* $\rightarrow q$ for any q (contradiction)
   l. $((p \rightarrow q) \wedge p) \rightarrow q$ for any p, q (modus ponens)
   m. $((p \rightarrow q) \wedge \neg q) \rightarrow \neg p$ for any p, q (modus tollens)

3. Prove one of the following, corresponding to your classroom number, referring to the axiom(s) used:
   0. c
   1. $b \vee \neg e$
   2. $c \vee d$
   3. $c \wedge \neg d$
   4. $e \rightarrow \neg d$
   5. $\neg b \vee \neg f$
   6. $\neg d$
   7. $a \vee e$
   8. $\neg g$
   9. $\neg h$
   10. $f \rightarrow \neg a$
4. Write an algorithm, in pseudocode, that has an array of numbers as a parameter and that computes
   0. Disjunction (OR), assuming all eleements are 0 or 1 (false or true)
   1. Conjunction (AND), assuming all eleements are 0 or 1 (false or true)
   2. Sum of all elements
   3. Number of elements, starting with first, that are all the same
   4. Number of zeroes
   5. Smallest element
   6. Subscript of the leftmost 1
   7. Subscript of largest element
   8. *True* if all values are the same, otherwise *false*
   9. *True* if all values are in ascending order, otherwise *false*
   10. Length of longest ascending sequence that starts with first element

Study questions on Topic 1: Cognition and computation

1. Cognition

1. Cognitive science is (a) an independent discipline; (b) interdisciplinary; (c) a subdiscipline of AI; (d) a subdiscipline of psychology; (e) none of these
2. The theory of cognitive science presented asserts that thinking can best be understood in terms of (a) data processing; (b) knowledge representation; (c) inference; (d) stimulus-response; (e) representational structures and computational procedures
3. The ___-based approach to cognitive science is concerned with the form of reasoning (a) logic; (b) rule; (c) concept; (d) analogy; (e) connections
4. The ___-based approach to cognitive science is concerned with parallel distributed processing as in the brain (a) logic; (b) rule; (c) concept; (d) analogy; (e) connections
5. The ___-based approach to cognitive science is concerned with condition-action pairs (a) logic; (b) rule; (c) concept; (d) analogy; (e) connections
6. The ___-based approach to cognitive science is concerned with typical entities or situations (a) logic; (b) rule; (c) concept; (d) analogy; (e) connections
7. The ___-based approach to cognitive science is concerned with guidance derived from past situations (a) logic; (b) rule; (c) concept; (d) analogy; (e) connections
8. In rule-based systems, rules represent (a) the cortex; (b) sensations; (c) actions; (d) short-term memory; (e) long-term memory
9. Dualism holds that (a) mind is not a different substance from body; (b) mind and body are separate; (c) what is mental dominates what is material; (d) every question has two valid answers; (e) life’s value is in owning things
10. Materialism holds that (a) mind is not a different substance from body; (b) mind and body are separate; (c) what is mental dominates what is material; (d) every question has two valid answers; (e) life’s value is in owning things
11. Idealism holds that (a) mind is not a different substance from body; (b) mind and body are separate; (c) what is mental dominates what is material; (d) every question has two valid answers; (e) life’s value is in owning things
12. In computational systems, physical states are (a) representations; (b) determined by mental states; (c) random; (d) static; (e) none of these
13. A representation is a structure that (a) has self-evident meaning; (b) has no meaning; (c) stands for something else; (d) acts on other structures; (e) none of these
14. Representations of typical entities or situations are (a) processes; (b) programs; (c) concepts; (d) analogies; (e) inferences
15. Analogic reasoning (a) makes deductive inferences; (b) is rule-based; (c) adapts thinking about familiar situations to new ones; (d) applies definitions; (e) none of these
16. In the brain, a concept is (a) a logical formula; (b) a reflex; (c) a synapse; (d) a neuron; (e) a pattern of neuron activations
17. Emotions enable (a) sensory input; (b) inference; (c) rule application; (d) focus and action; (e) none of these
18. Scientific evidence of consciousness is associated with (a) sensory input; (b) reflex actions; (c) spiritual development; (d) emotions; (e) brain processes
19. The study of thinking via cooperation is (a) distributed cognition; (b) introspection; (c) psychology; (d) ethics; (e) none of these
20. Interfaces between neurons are called (a) axons; (b) dendrites; (c) synapses; (d) potentials; (e) receptors
21. A neuron fires when (a) it receives an impulse; (b) it is not inhibited; (c) it receives impulses beyond a certain threshold; (d) a receptor molecule binds to a transmitter; (e) none of these
22. Memory consists of (a) creation of neurons; (b) creation of synapses; (c) creation of axons; (d) changes in weights of synapses; (e) storage of electrical potential
23. The cerebral cortex supports (a) vision; (b) reflex; (c) executive function; (d) fight-flight responses; (e) none of these
24. The amygdala supports (a) vision; (b) reflex; (c) executive function; (d) fight-flight responses; (e) none of these
25. Personality is in the (a) amygdala; (b) cerebellum; (c) cortex; (d) spinal cord; (e) none of these
26. Brain processing uses (a) chemicals; (b) electricity; (c) neurons; (d) neurons, glial cells, and chemicals; (e) modus ponens

2. The rational-agent approach to AI

1. Rationality is associated with (a) interaction; (b) common sense; (c) humans; (d) inference and expected reward; (e) none of these
2. An agent receives ___ from the environment (a) knowledge; (b) actions; (c) percepts; (d) instructions; (e) none of these
3. An agent performs (a) knowledge; (b) actions; (c) percepts; (d) instructions; (e) none of these
4. The easiest environment below is (a) stochastic, dynamic, fully observable; (b) deterministic, static, fully observable; (c) stochastic, static, partially observable; (d) stochastic, dynamic, partially observable; (e) none of these

5. The most difficult environment below is (a) stochastic, dynamic, fully observable; (b) deterministic, static, fully observable; (c) stochastic, static, partially observable; (d) stochastic, dynamic, partially observable; (e) none of these
6. A reflex agent (a) learns from its environment; (b) reasons based on past percepts; (c) acts only on current percept; (d) maintains a model of its environment; (e) none of these
7. Rationality maximizes (a) correctness of inference; (b) immediate reward; (c) information; (d) actual long-term reward; (e) expected long-term reward
8. A rational agent (a) makes a correct deduction; (b) gathers maximum information; (c) acts optimally; (d) acts as well as possible; (e) none of these
9. The most difficult environment below is (a) deterministic and fully observable; (b) episodic and static; (c) dynamic and partially observable; (d) discrete and single-agent; (e) fully observable and stochastic
10. Utility-based agents seek mainly (a) reward; (b) truth; (c) points; (d) to be helpful; (e) knowledge
11. Good Old Fashioned AI sees intelligence as related to (a) interaction; (b) agent behavior; (c) emergent; (d) a set of logical propositions about the world; (e) data retrieval
12. An *agent* is (a) autonomous and situated in an environment; (b) given instructions by a master; (c) a physical object; (d) an algorithm; (e) none of these
13. Newell and Simon hypothesized that a necessary and sufficient condition for intelligence is (a) emotion; (b) rationality; (c) adaptation; (d) symbol manipulation; (e) embodiment
14. The Turing Test measures intelligence as (a) a machine's ability to detect a human's presence; (b) a human's ability to detect a machine's presence; (c) a human's inability to detect a machine's presence; (d) use of a Turing machine; (e) an IQ test for machines
15. A well-known way to define machine intelligence is (a) using computability theory; (b) using predicate logic; (c) the Turing Test; (d) using complexity theory; (e) as processing speed comparable to that of the brain

3. Models of computation

1. The Turing machine model is said to capture (a) regular languages; (b) interaction; (c) efficient computation; (d) algorithmic computation; (e) all of these
2. A Turing machine (a) lacks an alphabet; (b) has tape instead of states; (c) can compute any mathematical function; (d) stores data on a tape; (e) none of these
3. The Church-Turing thesis associates the Turing machine with (a) regular languages; (b) parsing; (c) lexical analysis; (d) algorithms; (e) interaction
4. The Church-Turing Thesis refers to (a) a formalization of an intuitive notion; (b) a theorem; (c) provable; (d) disprovable; (e) a paper written at Harvard
5. A transition system is defined by (a) a set of states and a relation on them; (b) a set of points and a mapping among them; (c) a set of symbols and rules for sequencing them; (d) a set of strings; (e) none of these
6. The ____ is a widely-used model of computation (a) PC; (b) Macintosh; (c) operating system; (d) transition system; (e) principle of mathematical induction
7. A pushdown automaton may have (a) more states than a finite automaton; (b) random-access memory; (c) an infinite alphabet; (d) a stack; (e) faster transitions
8. A pushdown automaton has (a) a finite number of configurations; (b) fewer configurations than a DFA; (c) an infinite number of configurations; (d) a queue; (e) a tape
9. The Turing machine model is said to capture (a) regular languages; (b) interaction; (c) efficient computation; (d) algorithmic computation; (e) all of these
10. A Turing machine is distinguished from finite and pushdown automata by (a) a set of states; (b) a stack; (c) RAM; (d) a tape; (e) a simpler transition function
11. A Turing machine has ____ memory (a) random-access; (b) limited; (c) unbounded; (d) stack; (e) queue
12. A Turing machine's special feature that distinguishes it from finite automata and pushdown automata is (a) a set of states; (b) a stack; (c) RAM; (d) a tape; (e) a simpler transition function
13. A Turing machine (a) has no alphabet; (b) replaces states with a tape; (c) uses tape as storage; (d) has finite memory; (e) can compute any mathematical function
14. An state-transition system with tape is a (a) finite transducer; (b) DFA; (c) NFA; (d) PDA; (e) Turing machine
15. Turing machine-computable functions are the same as those computable on a (a) DFA; (b) PDA; (c) random-access machine; (d) control device; (e) none of these
16. Unlike Turing machines, random-access machines have (a) tape; (b) stack; (c) queue; (d) addressable storage; (e) hard disk
17. The Church-Turing thesis associates the Turing machine with (a) regular languages; (b) parsing; (c) lexical analysis; (d) algorithms; (e) interaction
18. The Church-Turing Thesis refers to (a) a formalization of an intuitive notion; (b) a theorem; (c) provable; (d) disprovable; (e) a paper written at Harvard
19. Connectionism is (a) silicon inspired; (b) parallel distributed processing; (c) symbolic; (d) inference based; (e) abductive
20. Connectionist models of computation are based on (a) the bit; (b) the neuron; (c) transition systems; (d) pseudocode; (e) the Ethernet protocol
21. Parallel computing is a kind of (a) concurrency; (b) Von Neumann architecture; (c) serial paradigm; (d) single-core computing; (e) none of these
22. The Parallel Random Access Machine model of computation assumes that memory is (a) all distributed; (b) shared; (c) infinite; (d) magnetic; (e) none of these

23. Algorithms (a) compute functions; (b) provide services; (c) accomplish missions in multi-agent systems; (d) may execute indefinitely; (e) none of these
24. A feature of algorithmic computation is (a) alternation of input and output; (b) processing before input; (c) output before processing; (d) input, then processing, then output; (e) none of these
25. A feature of interactive computation is (a) alternation of input and output; (b) processing before input; (c) output before processing; (d) input, then processing, then output; (e) none of these
26. Reactive systems (a) compute functions; (b) provide services; (c) accomplish multi-agent missions; (d) execute only finitely; (e) none of these
27. I/O in reactive systems is (a) static; (b) dynamic; (c) finite; (d) constrained; (e) none of these
28. Interaction is distinguished from algorithmic computation by the presence of (a) finite input; (b) persistent state; (c) input; (d) processing; (e) none of these
29. A mutual causal effect between two agents occurs in all (a) interaction; (b) algorithms; (c) communication; (d) computing; (e) none of these
30. Synchrony entails (a) communication; (b) taking turns; (c) input; (d) autonomy; (e) none of these
31. Stream I/O characterizes (a) interaction; (b) algorithms; (c) functions; (d) O(1) processes; (e) none of these
32. Persistent state stores (a) parameters; (b) loop invariants; (c) memory of past interactions; (d) algorithmic computation; (e) none of these
33. The form of input/output in interactive computation is (a) static; (b) finite; (c) streams; (d) strings; (e) none of these
34. A *service* is characteristic of (a) an algorithm; (b) an interactive process; (c) a multi-agent system; (d) a parallel system; (e) none of these

4. Intractability

1. Problems for which no polynomial-time solutions are known are called (a) undecidable; (b) intractable; (c) NP; (d) optimization; (e) none of these
2. The set of intractable problems is associated with (a) P; (b) divide-and-conquer algorithms; (c) greedy algorithms; (d) NP; (e) NP-completeness and exponential time
3. *P* is the set of (a) algorithms that execute in O(n) time; (b) problems decidable in O(n^k) time for some constant k; (c) problems *not* decidable in O(n^k) time; (d) intractable problems; (e) exponential-time problems
4. Intractable problems (a) are undecidable; (b) lack acceptable approximate versions; (c) are decidable but take an unacceptably long time; (d) lack solutions; (e) none of these
5. Analysis (a) computes a function; (b) separates something into parts; (c) puts components together; (d) writes a program; (e) is the entire problem-solving process
6. Best case for an algorithm (a) takes the same time for all data; (b) assumes the data that the algorithm handles in the greatest time; (c) assumes the data that the algorithm handles in the least time; (d) is the expected time considering all possible input data; (e) none of these
7. Worst case for an algorithm (a) takes the same time for all data; (b) assumes the data that the algorithm handles in the greatest time; (c) assumes the data that the algorithm handles in the least time; (d) is the expected time considering all possible input data; (e) none of these
8. Average case for an algorithm (a) takes the same time for all data; (b) assumes the data that the algorithm handles in the greatest time; (c) assumes the data that the algorithm handles in the least time; (d) is the expected time considering all possible input data; (e) none of these
9. Bubble sort's worst-case running time function is determined by (a) a single loop; (b) nested loops; (c) a series of loops; (d) the input data; (e) none of these
10. A loop nested to two levels, each with roughly n iterations, has running time (a) O(1); (b) O(n); (c) O(n^2); (d) O(n lg n); (e) O(2^n)
11. A loop nested to n levels has running time (a) O(1); (b) O(n); (c) O(n^2); (d) O(n lg n); (e) O(2^n)
12. The running time function of an algorithm is determined by (a) the number of operations in a sequence structure; (b) the number of branches in a selection structure; (c) the time of the slowest of a series of loops; (d) the data; (e) none of these
13. An approach to algorithm design often used to address intractable problems is (a) divide and conquer; (b) greedy; (c) brute force; (d) dynamic programming; (e) probabilistic
14. One way to find an *adequate though inexact* solution to an intractable optimization problem may be (a) brute force; (b) approximation; (c) divide and conquer; (d) greedy algorithm; (e) none of these
15. An *adequate though inexact* solution to an intractable optimization problem may be (a) brute force; (b) probabilistic; (c) divide and conquer; (d) O(2^n); (e) none of these
16. For _______ problems, sometimes global maxima/minima differ from local ones (a) optimization; (b) O(n); (c) BST search; (d) sorting; (e) none of these

Short and longer answer

1. Relate cognition to computation.
2. What is the computational-representational understanding of mind? (1a)
3. In the mind, what is representation and what is computation? (1a)
4. In a computer and in a brain, what is representation and what is computation? (1a)
5. What features of software could satisfy the definition of *mind* under the computational-representational understanding of mind? (1a)
6. Distinguish toy problems from real-world ones.
7. Describe some dimensions of classification of environments. (1b)
8. Distinguish fully observable from partially observable environments, saying which which is harder. Give examples. (1b)
9. Distinguish deterministic from stochastsic environments, saying which which is harder. (1b)
10. Distinguish episodic from sequential environments, saying which which is harder. (1b)
11. Distinguish static from dynamic environments, saying which which is harder. (1b)
12. Distinguish discrete from continuous environments, saying which which is harder. (1b)
13. Describe some classes of environment that are more difficult than static, fully observable, episodic, deterministic ones, and give features agents would need for more difficult environments. (1b)
14. How does the brain support intelligence? (1c)
15. Describe some significant distinctions between a brain and a PC that are noted in AI research. (1c)
16. Describe the connectionist model of computation, with examples. (1c)
17. Describe the transition-system model of computation, with examples. (1c)
18. What is an algorithm? (1c)
19. Distinguish algorithms from interaction. (1c)
20. Give two algorithmic models of computation. (1c)
21. What is a special features of recursively-definable functions? (1c)
22. Give reasons why intelligence has a social character; reasons why it uses non-interactive processing.
23. Describe a *reflex agent*. (1d)
24. What is the rational-agent approach to AI, and what is one kind of rational agent? (1d)
25. In what environments can a reflex agent operate effectively? (1d)
26. What kind of agent can operate in a fully observable, deterministic, static, discrete environment, and why? (1d)
27. What is a goal-based agent?
28. How does the brain support intelligence?
29. Is CRUM a good model for intelligence? Why?
30. Name five significant distinctions between a brain and a PC.
31. What would a computer program require to have consciousness?
32. Describe the computational-representational understanding of mind and any defects or especially useful features you see in it.
33. Which of the six major approaches to cognitive science seems most useful to you? Why?
34. Show whether the following formula
    $p \wedge (\neg q \vee r) \wedge (p \vee q)$
    in propositional logic is
    (a) valid
    (b) satisfiable
35. Is a car a rational agent, and why/why not? What would a rational agent require to get from your home to FSC? Use the definition provided.
36. Relate intelligence to interaction.
37. (For Theory of Computing students) Can a (DFA, PDA, TM, random-access machine, recursively-definable function) be intelligent?

For Analysis of Algorithms students:

38. Relate intelligence to intractable problems.
39. Is a brain subject to the same limitations as a computer in relation to the notions of tractability and decidability?

Study questions on Topic 2: State-space search

1. State-space search

1. Games and puzzles are simple examples of (a) embodied intelligence; (b) state-space search; (c) inference; (d) agent interaction; (e) adaptation
2. The breadth-first search (a) uses a queue; (b) uses a stack; (c) searches an array; (d) searches a tree; (e) none of these
3. The depth-first search (a) uses a queue; (b) uses a stack; (c) searches an array; (d) searches a tree; (e) none of these
4. Goal-driven state-space search arrives at (a) goals from facts; (b) facts from goals; (c) rules from facts; (d) search strategies; (e) heuristics
5. Data-driven state-space search arrives at (a) goals from facts; (b) facts from goals; (c) rules from facts; (d) search strategies; (e) heuristics
6. Exploration may be useful for environments that are (a) fully observable; (b) partially observable; (c) episodic; (d) one-state; (e) not observable
7. A goal is a(n) (a) path; (b) action; (c) percept; (d) set of states; (e) number
8. A state space is a set of (a) three-dimensional coordinates; (b) locations in the physical universe; (c) governmental entities; (d) actual arrangements of values; (e) possible arrangements of values
9. A set of possible arrangements of values is a(n) (a) state space; (b) path; (c) combination; (d) random variable; (e) none of these
10. Optimizing search compares (a) costs of paths; (b) costs of information; (c) reward values; (d) costs of algorithm design; (e) none of these
11. One well-known strategy for state-space search is called (a) measure and estimate; (b) generate and test; (c) try and abandon; (d) forward and back; (e) design and revise

2. Hard problems in constraint and optimization

1. AI problems tend to involve (a) large numbers; (b) combinatorial explosion of running time; (c) easy choices once understood; (d) straightforward inference; (e) none of these
2. Constraint-satisfaction problems aim at (a) constraints that rule out all but a few cases; (b) constraints that rule out only a few cases; (c) constraints that involve several variables; (d) optimization of values; (e) none of these
3. Satisfiability is a(n) (a) optimization problem; (b) constraint satisfaction problem; (c) algorithm; (d) heuristic; (e) interactive problem
4. Local search (a) solves any problem; (b) is always effective; (c) is never effective; (d) reduces difficulty of some constraint satisfaction problems; (e) is more thorough than global search
5. Combinatorial explosion is (a) sudden increase in difficulty of an environment; (b) a cause of collapse of an agent's effectiveness; (c) exponential size of state space; (d) failure of search; (e) none of these
6. A problem of finding values of several variables such that a certain condition holds is called (a) graph search; (b) tree traversal; (c) constraint satisfaction; (d) sorting; (e) optimization
7. Constraint satisfaction is a problem of (a) finding values of a set of variables such that a certain condition holds is called; (b) SAT; (c) finding a maximal or minimal value; (d) optimizing a path; (e) none of these
8. Bounded rationality is associated with (a) optimality; (b) constraint satisfaction; (c) well ordering; (d) satisficing; (e) tractability
9. An optimization problem finds a maximum or minimum value that satisfies a certain (a) formula in predicate logic; (b) constraint; (c) time specification; (d) user; (e) protocol
10. A problem of finding a set of values that yields the highest or lowest return value when used as parameters to a function is (a) constraint satisfaction; (b) optimization; (c) maximization; (d) minimization; (e) central tendency
11. Exponential time is closely associated with (a) tractability; (b) combinatorial explosion; (c) constraint problems; (d) sorting problem; (e) interaction

3. Heuristics

1. Heuristics are (a) axioms; (b) inference rules; (c) rules that guide state-space search; (d) results of inference; (e) none of these
2. Heuristics must often be used in (a) logical inference; (b) state-space search; (c) backtracking; (d) robotic sensing; (e) abstraction
3. A rule of thumb that guides state-space search is a(n) (a) axiom; (b) inference rule; (c) heuristic; (d) theorem; (e) adaptation
4. Best-first search uses a(n) (a) inference rule; (b) heuristic; (c) form of knowledge representation ; (d) protocol; (e) none of these
5. Minimax is a(n) (a) inference rule; (b) heuristic; (c) form of knowledge representation ; (d) protocol; (e) none of these
6. Hill climbing is a(n) (a) problem; (b) heuristic strategy; (c) best-first search; (d) expression of consciousness; (e) form of representation
7. A drawback of hill climbing is (a) very long running time; (b) undecidability; (c) tendency to become stuck at local maxima; (d) the absence of goals; (e) none of these
8. Admissibility, informedness, and monotonicity are features of all (a) algorithms; (b) heuristics; (c) formulas in predicate logic; (d) problems; (e) robots
9. Minimax is a (a) problem; (b) algorithm; (c) game; (d) neural-network design; (e) form of consciousness

10. The assumption that a game opponent will make the best possible move is made in (a) depth-first search; (b) breadth-first search; (c) all two-player games; (d) the minimax algorithm; (e) none of these

Short and longer answer

1. What does a state-space search do? (2a)
2. What are the rules of thumb that can reduce the number of steps in a state-space search? (2a)
3. What is the state space to search in
   a. tic tac toe
   b. a maze
   c. the traveling-salesperson problem
   d. the game of chess
   e. choosing courses to register for in the next semester
   (2a)
4. Distinguish data-driven from goal-driven state-space search. (2a)
5. What are two ways to sequence the exhaustive search of a state space? Explain. (2a)
6. In a goal-driven AI approach to state-space search, what does the goal consist of ? Give an example. (2a)
7. Describe a kind of agent that may require exploration of a state space. (2a)
8. Describe a kind of agent that must execute more than just a reflex. (2a)
9. What is the name for a set of desirable states, and how does an agent operate in an environment defined in this way? (2a)

2b. Game trees

Consider the following game trees. Player *X* is at the game state denoted by the root vertex and may choose between move *a* and move *b*, denoted by edges, leading to a game state that is a vertex adjacent to the root. Then player *Y* will move, followed by player *X*. After *X*'s second move, that player will immediately be in a game state that is a win (W), tie (T), or loss (L). Decide whether player *X*'s better move is *a* or *b*, and explain. What is the maximum size of the state space of paths for a game tree of depth *n*, where players have a maximum of *k* choices in making a move? (Adapted from Brookshear.)

1.

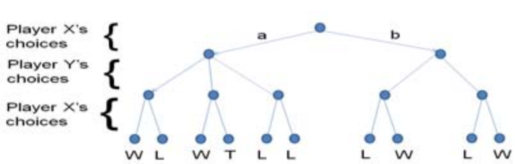

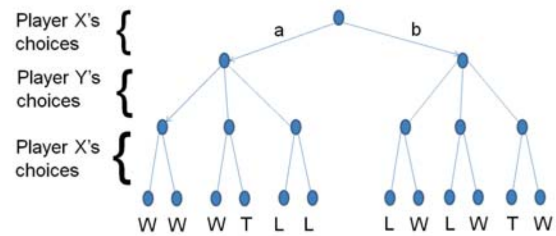

2.

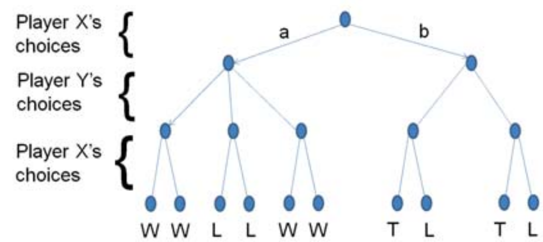

3.

Consider the mazes below, where the player starts at *S* and tries to reach goal *G*.
(a) Convert to graph form
(b) Perform partial depth-first and breadth-first searches by giving the order in which the graph vertices are to be visited
(c) If *n* two-way branches are encountered in a maze, then how large is the state space of paths?

4.

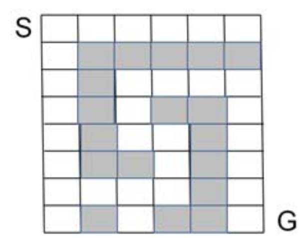

5.

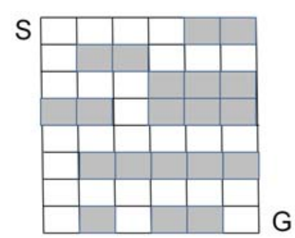

6.

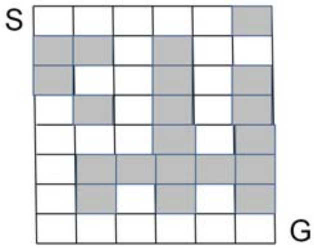

Consider the following tic-tac-toe positions.

(a) Draw a game tree, using the notation

1 2 3
4 5 6
7 8 9

to denote moves

(b) Use game tree to choose X's next move and predict game outcome

7.

X	O	X
	O	
X		O

8.

X	O	X
	O	O
	X	

9.

X	X	O
	X	
	O	O

10.

	X	O
O	X	X
	O	

11.

X		
X	O	O
O		X

2c. Exhaustive search and heuristics

1. Explain the hill-climbing heuristic. (2c)
2. Why are there often too many states to exhaustively search in the state space? What are the rules of thumb that can reduce the number of states visited? Give an example. (2c)
3. Which heuristic for game play assumes that the opponent will play optimally? Explain. (2c)
4. Describe the minimax heuristic. (2c)
5. What is a heuristic, and what is its use? Give an example. (2c)
6. What does "hard computational problems" mean, formally and in practice? (2c)
7. What sorts of running times is *intractability* associated with? (2c)
8. Describe in mathematical terms the kinds of computational problems addressed by AI. (2c)
9. For what sorts of problems are approximation algorithms, randomization, and heuristics used? What are heuristics and when are they used? (2c)
10. Luger05, pp. 121-122, #2-3

Study questions on Topic 3: Knowledge representation and rule-based inference

1. Knowledge and beliefs

1. A key component of an agent may be a (a) compiler; (b) processor; (c) state set; (d) database; (e) knowledge base
2. Knowledge may be commonly represented in the language of (a) mathematics; (b) logic; (c) a country; (d) computers; (e) databases
3. One significant formalism for knowledge representation in AI is (a) the ASCII code; (b) arithmetic; (c) predicate logic; (d) Java; (e) none of these
4. Rules for model-based reasoning include (a) network protocols; (b) communication primitives; (c) rules of engagement; (d) diagnostic and causal rules; (e) grammar rules
5. A situation is a (a) search path; (b) set of all possible states; (c) state of the environment; (d) belief state; (e) knowledge base
6. A state of the environment that results from a previous state and an action is a(n) (a) model; (b) space; (c) location; (d) situation; (e) inference
7. Planning may start with (a) original state only; (b) goal state; (c) either original or goal state; (d) intermediate or goal state; (e) original or intermediate state
8. Planning is part of ____ seeking (a) knowledge; (b) goal; (c) algorithm; (d) belief; (e) victory
9. 0Belief is a(n) (a) inference; (b) axiom; (c) propositional attitude; (d) grammar rule; (e) instance of knowledge
10. Knowledge is (a) belief; (b) proof; (c) inference; (d) justified true belief; (e) theorems
11. Justified true belief is (a) logic; (b) wisdom; (c) inference; (d) knowledge; (e) none of these
12. Logic-based agents (a) answer queries; (b) adapt to their environments; (c) operate in multi-agent systems; (d) seek to prove assertions; (e) use training to improve ability to act
13. Knowledge-based agents (a) answer queries; (b) adapt to their environments; (c) operate in multi-agent systems; (d) seek to prove assertions; (e) use training to improve ability to act
14. A knowledge base is expressed in a ______ language (a) knowledge-representation; (b) regular; (c) natural; (d) programming; (e) query
15. Information expressed in a knowledge representation language is used by _____ agents (a) reflex; (b) goal-based; (c) knowledge-based; (d) stochastic; (e) reinforcement-learning
16. Application of an inference rule is (a) a greater-knowledge to less-knowledge state transition; (b) a state transition from less knowledge to greater knowledge; (c) a full state-space search; (d) a transition to an initial state; (e) none of these
17. Semantic networks represent an approach to (a) inference; (b) adaptation; (c) knowledge representation; (d) predicate calculus; (e) logic programming
18. Capturing deep semantic aspects of language is the objective of (a) LISP; (b) Prolog; (c) case frames; (d) heuristics; (e) none of these
19. Creation of a formal representation of a knowledge base of domain-knowledge rules is (a) inference; (b) database design; (c) policy search; (d) knowledge engineering; (e) none of these
20. Knowledge-based agents use percepts and _____ knowledge (a) domain; (b) universal; (c) first-order logic; (d) robotic; (e) scientific

2. Concepts and instances

1. Semantic networks include (a) nodes and edges that are concepts and objects; (b) servers and clients; (c) protocols; (d) inferences; (e) none of these
2. The closed-world assumption is (a) part of first-order logic; (b) a form of default reasoning; (c) a theorem; (d) provable; (e) none of these
3. Case frames capture (a) syntax; (b) inference; (c) first-order logic; (d) semantics; (e) real-time data
4. To represent concepts, an alternative to predicate calculus is (a) first-order logic; (b) semantic networks; (c) context-free grammars; (d) heuristics; (e) none of these
5. Case-based reasoning is (a) natural entailment; (b) arithmetic; (c) analogic; (d) abductive; (e) none of these
6. Containment is a ___ relationship (a) kind-of; (b) part-of; (c) inheritance; (d) mutual; (e) exclusive
7. Inheritance is a ___ relationship (a) kind-of; (b) part-of; (c) multi-way; (d) mutual; (e) exclusive
8. An *ontology* is a (a) logic; (b) set of percepts; (c) set of facts about the world chosen to be in a knowledge base; (d) database; (e) none of these
9. Abstract knowledge about concepts in the world, including things, actions, and relationships, comprise a(n) (a) database; (b) ontology; (c) query; (d) knowledge base; (e) none of these
10. A *frame* is a (a) fact about an instance of a concept; (b) knowledge base; (c) relationship; (d) scheme to express connections among concepts; (e) none of these
11. Fuzzy logic is related to a theory of (a) computation; (b) possibility; (c) probability; (d) knowledge; (e) none of these
12. Truth is quantified as a real number between 0 and 1 by ____ logic (a) predicate; (b) propositional; (c) fuzzy; (d) modal; (e) temporal
13. *Tall-persons* is a ____ set (a) fuzzy; (b) null; (c) large; (d) infinite; (e) sub-

3. Logical inference

1. Predicate logic is a(n) (a) algorithm; (b) language of assertions; (c) language of arithmetic expressions; (d) set of symbols; (e) set of operations
2. An assertion's value is (a) true; (b) s symbol; (c) a number; (d) true or false; (e) none of these
3. $(\forall x)\ x = x + 1$ is (a) a numeric expression; (b) false; (c) true; (d) an assignment; (e) none of these
4. $(\exists x)\ x = x + 1$ is (a) a numeric expression; (b) false; (c) true; (d) an assignment; (e) none of these
5. An interpretation is (a) an assignment of truth values to variables; (b) the value of an assertion; (c) the meaning of a program; (d) a formula; (e) none of these
6. A formula is satisfiable if it has a(n) ____ under which it is true (a) operation; (b) algorithm; (c) number; (d) interpretation; (e) none of these
7. Quantifiers ____ variables for meaningful use (a) give values to; (b) take values from; (c) bind; (d) assign; (e) declare
8. A sentence that is not true under any interpretation is (a) complete; (b) incomplete; (c) consistent; (d) a contradiction; (e) valid
9. A sentence that is true under any interpretation is (a) complete; (b) incomplete; (c) a a contradiction; (d) inconsistent; (e) valid
10. Inference rules maintain (a) completeness; (b) consistency; (c) validity; (d) satisfiability; (e) falsehood
11. An inference rule that never produces contradictions is (a) complete; (b) incomplete; (c) inconsistent; (d) sound; (e) useless
12. $(p \wedge (p \rightarrow q)) \rightarrow q$ is (a) false; (b) Modus Ponens; (c) inconsistent; (d) not always true; (e) none of these
13. An interpretation of a set of formulas in predicate logic is (a) a logical inference; (b) a heuristic; (c) an assignment of truth values to symbols; (d) a theorem; (e) a truth value
14. The sentence, $\alpha \models \beta$ (in every interpretation where α is true, β is true), is an instance of (a) entailment; (b) negation; (c) validity; (d) satisfiability; (e) falsehood
15. Predicate calculus extends propositional logic with (a) inference; (b) negation; (c) implication; (d) variables; (e) quantifiers
16. Predicate calculus extends propositional logic with (a) inference; (b) negation; (c) implication; (d) variables; (e) functions
17. A formula in predicate logic is valid if (a) it is true for some interpretation; (b) it is true for all interpretations; (c) it is true for no interpretation; (d) it is an axiom; (e) it is not disproven
18. A formula in predicate logic is satisfiable if (a) it is true for some interpretation; (b) it is true for all interpretations; (c) it is true for no interpretation; (d) it is an axiom; (e) it is not disproven
19. A formula in predicate logic is a contradiction if (a) it is true for some interpretation; (b) it is true for all interpretations; (c) it is true for no interpretation; (d) it is an axiom; (e) it is not disproven
20. A validity-maintaining procedure for deriving sentences in predicate logic from other sentences is a(n) (a) proof; (b) theorem; (c) algorithm; (d) inference rule; (e) inference chain
21. Satisfiability is ___ validity (a) weaker than; (b) equivalent to; (c) stronger than; (d) a subset of; (e) none of these
22. An algorithm that determines what substitutions are needed to make two sentences match is (a) resolution; (b) inference; (c) unification; (d) contradiction; (e) nonexistent
23. Unification is (a) an algorithm for making substitutions so that two sentences match; (b) a proof method; (c) an inference rule; (d) a theorem; (e) a knowledge-representation scheme
24. Inference rules enable derivation of (a) axioms; (b) other inference rules; (c) new knowledge; (d) percepts; (e) none of these
25. The problem of evaluating a formula in propositional logic is (a) intractable; (b) undecidable; (c) tractable; (d) $\Omega(2^n)$; (e) polymorphic
26. Deciding whether a formula in propositional logic is satisfiable is considered (a) intractable; (b) undecidable; (c) tractable; (d) decidable; (e) polymorphic
27. SAT is the problem of deciding whether a formula in propositional logic (a) holds; (b) has a set of variable assignments that make it true; (c) is not a contradiction; (d) is syntactically correct; (e) is probably true
28. *Modus ponens* asserts that (a) $p \rightarrow q$; (b) $\neg p \wedge \neg q$; (c) $p \rightarrow ((p \rightarrow q) \rightarrow q)$; (d) $(p \wedge (p \rightarrow q)) \rightarrow q$; (e) $(p \rightarrow (p \wedge q)) \rightarrow q$
29. Backward chaining is ____ driven (a) AND; (b) data; (c) goal; (d) logic; (e) inference
30. Forward chaining is ____ driven (a) AND; (b) data; (c) goal; (d) logic; (e) inference

4. Expert systems and resolution proof

1. Expert systems (a) find all inferences; (b) try to unify goals with facts; (c) try to summarize facts; (d) try to find contradictions; (e) none of these
2. A way to add to a knowledge base monotonically is (a) backward chaining; (b) inference; (c) querying; (d) arithmetic; (e) AND
3. Separating knowledge from control is a feature of (a) heuristics; (b) expert systems; (c) first-order logic; (d) predicate calculus; (e) reinforcement learning
4. Expert systems are ____ based (a) data; (b) consciousness; (c) rule; (d) proof; (e) none of these
5. Resolution proof uses (a) forward chaining; (b) contradiction; (c) abduction; (d) unification; (e) statistics

6. Prolog uses ____ proof (a) resolution; (b) unification; (c) inductive; (d) constructive; (e) none of these
7. Prolog searches for (a) data; (b) high-utility states; (c) proof of goal clauses; (d) refutations; (e) none of these
8. Prolog uses the ____ assumption (a) natural; (b) responsibility; (c) closed-world; (d) best-world; (e) optimal-utility
9. Expert systems store knowledge as (a) numbers; (b) database records; (c) inference rules; (d) proofs; (e) none of these
10. Expert systems separate (a) facts from opinions; (b) knowledge from control; (c) code from design; (d) inference from querying; (e) none of these

Short and longer answer

1. What is an algorithm for determining what substitutions are needed to make two sentences match? Why is this useful? (3a)
2. What is a *satisfiable* formula, and in what language is it usually expressed? (3a)
3. [3a: questions on inference in propositional and predicate logic.]
4. What is the relationship between belief and knowledge?
5. Distinguish causal rules from diagnostic rules.
6. What does first-order logic (predicate logic) consist of?
7. In what language is knowledge represented, traditionally, in AI? Why?
8. In what AI application is resolution proof used?
9. What do case frames represent? How? (3b)
10. What do semantic networks represent? How? (3b)
11. How is knowledge given to and received from expert systems? (3b)
12. How does inference contribute to a knowledge base? (3b)
13. Name relationships that *categories* may have. How may several category relationships be organized? (3b)
14. What is a validity-maintaining procedure for deriving sentences from other sentences in first-order logic? (3b)
15. What is unification used for? (3b)
16. What are two ways to add knowledge to a knowledge base? (3c)
17. How does a knowledge-based agent operate? (3c)
18. What application of AI stores knowledge as inference rules and how does it use the knowledge? What is a common language used for this application? (3c)
19. For what AI application is the Prolog language best known, and how does it work? (3c)
20. How is expertise represented in an expert system? (3c)
21. Distinguish knowledge-based from goal-driven agents. (3c)
22. What is a fuzzy set?
23. Luger05, pp. 77-78, #1-9. (#7, with coding is extra-credit.)

Study questions on Topic 4: Uncertainty and probabilistic reasoning

1. Acting under uncertainty

1. Planning in a partially observable environment faces challenges due to (a) adversity; (b) certainty; (c) uncertainty; (d) impossibility of prediction; (e) none of these
2. Beliefs (a) are held permanently; (b) are proven; (c) are axioms; (d) may be revised; (e) are avoided
3. Truth maintenance may alter (a) inference rules; (b) algorithms; (c) beliefs; (d) facts; (e) none of these
4. A belief state is a (a) set of transitions; (b) set of states; (c) state of the environment; (d) transition; (e) body of knowledge
5. Readiness of the environment for the next action may be verified by (a) belief state; (b) actions; (c) uncertainty; (d) action monitoring; (e) none of these
6. Readiness of the environment for the next action may be verified by (a) belief state; (b) actions; (c) uncertainty; (d) plan monitoring; (e) none of these
7. Knowledge only provides a(n) ____ for diagnosis (a) plan; (b) state; (c) degree of belief; (d) item of data; (e) none of these
8. Probabilities are employed in ____ methods (a) stochastic; (b) logical; (c) adversarial; (d) heuristic; (e) none of these
9. Stochastic methods are often used in (a) theorem proving; (b) knowledge retrieval; (c) logical inference; (d) planning under uncertainty; (e) none of these
10. Logic is ____ in that adding new facts always expands a knowledge base (a) modal; (b) propositional; (c) deontic; (d) monotonic; (e) nonmonotonic
11. Truth maintenance may require (a) inconsistency; (b) soundness; (c) percepts; (d) belief retraction; (e) none of these
12. Truth maintenance systems work with ____ logic (a) propositional; (b) predicate; (c) modal; (d) fuzzy; (e) none of these

2. Probability theory and belief

1. A random variable is a(n) (a) truth value; (b) set; (c) function; (d) relation; (e) number
2. A degree of belief in the absence of helpful information is a(n) (a) prior probability; (b) conditional probability; (c) random variable; (d) axiom; (e) event
3. A degree of belief given some helpful information is a(n) (a) prior probability; (b) conditional probability; (c) random variable; (d) axiom; (e) event
4. Any probability value is (a) 0 or 1; (b) in the range of 0 to 1; (c) some positive real number; (d) some positive or negative real number; (e) an integer
5. The possible orderings of elements of a set are (a) truth values; (b) numbers; (c) sets; (d) combinations; (e) permutations
6. The possible unordered selections from a set are (a) truth values; (b) numbers; (c) sets; (d) combinations; (e) permutations
7. A sample space is (a) a random variable; (b) a sequence; (c) a number; (d) a set of all possible outcomes; (e) an event
8. Conditional probability may apply if events have a(n) _____ relationship (a) empty; (b) noncausal; (c) independent; (d) dependent; (e) identity
9. Prior probability is (a) belief; (b) certainty; (c) conditional probability; (d) unconditional probability; (e) none of these
10. Probabilities of different event outcomes are a(n) (a) event; (b) probability distribution; (c) expected value; (d) sample space; (e) compound event
11. A set of possible outcomes is a(n) (a) random variable; (b) probaiblity distribution; (c) compound event; (d) sample space; (e) permutation
12. An outcome that is from a set of uncertain possibilities characterizes a (a) random process; (b) sample space; (c) event; (d) sequence; (e) permutation
13. The average of values for equally likely outcomes is a(n) (a) probability; (b) random variable; (c) expected value; (d) combination; (e) permutation
14. Expected value of a die throw is (a) 0; (b) 1; (c) 3.5; (d) 4; (e) 6
15. Expected value of a coin toss is (a) 0; (b) 0.25; (c) 0.5; (d) 1; (e) 2
16. For sample space S, Kolmogorov's axiom asserts that $P(S)$ = (a) 0; (b) 0.5; (c) 1; (d) 2; (e) indeterminate
17. For sample space S, Kolmogorov's axiom asserts that $P(\emptyset)$ = (a) 0; (b) 0.5; (c) 1; (d) 2; (e) indeterminate
18. Kolmogorov's axioms are considered useful for rational-agent AI because (a) they predict outcomes in many domains; (b) beliefs that violate the axioms result in poor bets; (c) they help the agent prove theorems; (d) they help the agent make inferences; (e) they are used in expert systems

3. Bayes' Theorem

1. Bayes' Theorem enables computation of probabilities of causes, given probabilities of (a) effects; (b) other causes; (c) prior world knowledge; (d) inference rules; (e) none of these
2. Evidence, in using Bayes' Theorem, consists of (a) causes; (b) effects; (c) prior world knowledge; (d) inference rules; (e) none of these
3. Bayes' Theorem is used in constructing (a) automata; (b) belief networks; (c) semantic networks; (d) knowledge bases; (e) none of these

4. ___ enables finding probabilities of causes, given effects (a) Minimax; (b) Bayes' Theorem; (c) Gödel's Theorem; (d) fuzzy logic; (e) Prolog
5. Belief networks use (a) Minimax; (b) Bayes' Theorem; (c) Godel's Theorem; (d) fuzzy logic; (e) Prolog

4. Markov models

1. A state-transition system with probabilistic transitions is a(n) (a) semantic net; (b) Bayesian net; (c) finite automaton; (d) Turing machine; (e) Markov chain
2. Observations that are probability functions of a current state characterize (a) all Markov models; (b) Bayesian networks; (c) schemas; (d) hidden Markov Models; (e) none of these
3. The assumption that the current state of a system depends only on a finite list of previous states is called (a) Bayesian; (b) Markov; (c) closed-world; (d) a predicate-logic axiom; (e) incorrect
4. Hidden Markov Models and Bayesian inference are used in (a) subsumption architecture; (b) speech recognition; (c) theorem proving; (d) expert systems; (e) none of these
5. Bayes' Theorem is about (a) entailment; (b) unconditional probability; (c) combinations; (d) inverse probability; (e) permutations
6. The Markov assumption is that (a) all outcomes have the same probability; (b) the past does not affect the future; (c) current state depends only on recent states; (d) probability is unconditional; (e) system is deterministic
7. A solution to a Markov process is a(n) (a) path; (b) number in [0..1]; (c) truth value; (d) big-O expression; (e) policy

Short and longer answer questions for topic 4

1. Why is belief maintenance nonmonotonic? (4a)
2. What does a rational agent seek to maximize? (4a)
3. What decisions are rational? (4a)
4. Distinguish belief maintenance from deduction. (4a)
5. Use standard notation to express the probability of rain, given clouds. In what cases is the probabaility of *A*, given *B*, zero? (4b)
6. If P(*A*) = 0.5 and P(*B*) = 0.3, then under what condition is P($A \wedge B$) = 0.15? (4b)
7. If P(*A*) = 0.5 and P(*B*) = 0.3, then if *A* and *B* are independent, what is P($A \vee B$)? (4b)
8. What is the expected value of the roll of two dice, and why? Three dice? Four? (4b)
9. What is the expected number of heads in two coin tosses? Three? Four? Five? (4b)
10. What is Bayesian inference used for and how? (4c)
11. What are Bayesian belief networks used for? (4c)
12. In what sort of environment is Bayesian reasoning used, and why? (4c)
13. What is the Markov assumption? (4c)
14. What is a Markov process? (4c)
15. What do Bayesian belief networks and hidden Markov models have in common? (4c)
16. Name and describe the following, including the significance of possible labels on transitions. (4c)

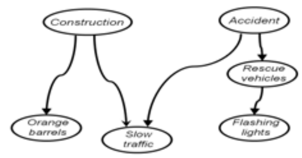

17. Distinguish “Tall(Joan) = 0.5” from “P(Joan is tall) = 0.5”.
18. Luger05, pp. 381-382, #2: problem computing confidence in a conclusion, given observations with confidence values.
19. P. 382, #16: Convert Dempster-Shafer tables to BBNs.
20. P. 383, #18: Use OMM to predict weather probabilities.

Study questions on Topic 5: Supervised learning

1. All learning (a) creates a knowledge base; (b) improves performance measure; (c) enables prediction of the future; (d) enhances utility; (e) speeds inference
2. Generalization is (a) deduction; (b) learning; (c) invalid; (d) the application of knowledge-base rules; (e) none of these
3. Concept learning is (a) random search; (b) knowledge-base querying; (c) generalization; (d) specialization; (e) deduction

1. Supervised learning

1. Supervised-learning agents (a) answer queries; (b) adapt to their environments; (c) operate in multi-agent systems; (d) seek to prove assertions; (e) use training to improve ability to act
2. Inductive inference is (a) supervised learning; (b) reinforcement learning; (c) invalid; (d) Bayesian; (e) none of these
3. Ockham's Razor recommends choosing (a) first hypothesis; (b) state-space search; (c) heuristics; (d) simplest valid hypothesis; (e) most completely-matched hypothesis
4. Abduction is (a) random search; (b) knowledge-base querying; (c) generalization; (d) specialization; (e) deduction
5. Backchaining from observations to hypotheses is (a) deduction; (b) abduction; (c) specialization; (d) knowledge-base querying; (e) random search
6. PAC learnability is an attribute of (a) first-order logic clauses; (b) knowledge bases; (c) concepts; (d) dynamic environments; (e) social environments
7. $((p \rightarrow q) \wedge q) \rightarrow p$ is (a) Modus Tollens; (b) Modus Ponens; (c) Bayes' Theorem; (d) valid; (e) abductive inference

2. Connectionist learning

1. Connectionism is (a) silicon inspired; (b) parallel distributed processing; (c) symbolic; (d) inference based; (e) abductive
2. Parallel distributed processing is (a) silicon inspired; (b) connectionism; (c) symbolic; (d) inference based; (e) abductive
3. Neural networks are (a) silicon inspired; (b) parallel distributed processing; (c) symbolic; (d) inference based; (e) abductive
4. Perceptrons are (a) first-order logic sentences; (b) neural nets; (c) social agents; (d) processors; (e) knowledge bases
5. Neurons fire according to (a) satisfaction of a first-order logic clause; (b) an algorithm; (c) hardware conditions; (d) a threshold function; (e) a *not* assertion
6. A nonsymbolic system is (a) neural nets; (b) first-order logic; (c) knowledge bases; (d) explanation-based learning; (e) context-free grammars
7. A downside of neural nets is (a) difficulty of human understanding of representation; (b) poor learning performance; (c) poor response; (d) low storage capacity; (e) none of these
8. Activity flows in loops in (a) backpropagation nets; (b) perceptrons; (c) recurrent neural nets; (d) all neural nets; (e) none of these
9. Backpropagation is used with (a) expert systems; (b) theorem proving; (c) single-layer neural nets; (d) multi-layer neural nets; (e) none of these
10. Representation of input data in neural nets is stored as (a) the quantity of output units; (b) the quantity of input units; (c) the weights of output-unit connections; (d) the weights of hidden-unit connections; (e) the number of connections
11. In training a neural net, weights of connections are changed in response to (a) agent judgment; (b) Bayesian formulas; (c) predicate-logic expression values; (d) errors detected in output units; (e) correct outputs
12. Neural nets have
    i. input units
    ii. output units
    iii. processing units
    iv. hidden units
    (a) i and ii; (b) i, ii, and iii; (c) i, iii, and iv; (d) i, ii, and iv; (e) none of these
13. Connectionist learning is (a) inference based; (b) symbol based; (c) non symbolic; (d) Bayesian; (e) none of these
14. Neural networks are associated with (a) Bayesian reasoning; (b) Ockham's Razor; (c) connectionist learning; (d) symbol-based learning; (e) none of these
15. Neurons in hidden layers are those (a) protected from firing; (b) with external inputs and outputs; (c) with external inputs but no external outputs; (d) with external outputs but no external inputs; (e) without external inputs or outputs
16. Neural nets learn by (a) abduction; (b) symbolic methods; (c) Bayesian inference; (d) adjusting weights of synapses; (e) computing rewards

3. Evolutionary computation

1. Evolutionary computation uses the technique of maximizing (a) fitness; (b) reward; (c) performance; (d) quantity of output; (e) none of these
2. Evolutionary computation (a) is deterministic; (b) seeks optimal solutions; (c) was developed in the 19th century; (d) is probabilistic; (e) none of these
3. Evolutionary computation is modeled on (a) brute force; (b) divide and conquer; (c) greediness; (d) natural selection; (e) fractals
4. Function optimization searches for (a) a function; (b) parameter values; (c) a return value; (d) an algorithm; (e) a time analysis

5. Fitness measures are (a) parameters to functions; (b) functions to be optimized; (c) return values; (d) algorithms; (e) time functions
6. Genetic algorithms are (a) greedy; (b) brute-force; (c) a way to compute fitness; (d) a form of evolutionary computation; (e) used in the human genome project
7. Ant computing is (a) greedy; (b) brute-force; (c) a way to compute fitness; (d) a form of evolutionary computation; (e) used in the human genome project
8. Evolutionary computation is (a) a brute-force method; (b) state-space search one state at a time; (c) path optimization; (d) population based; (e) DNA computing
9. The No Free Lunch Theorem states that function-optimization problems (a) are unsolvable; (b) are intractable; (c) have no solution that performs better than another, averaged over all environments; (d) can be solved over lunch; (e) cannot be solved over lunch

4. Natural language

1. Parsing comes after (a) contextual interpretation; (b) lexical analysis; (c) semantic interpretation; (d) all of these; (e) none of these
2. Semantic interpretation comes before (a) contextual interpretation; (b) lexical analysis; (c) parsing; (d) phonetic analysis; (e) none of these
3. A grammar specifies (a) semantics; (b) pragmatics; (c) phonetics; (d) syntax; (e) world knowledge
4. Semantics refers to (a) meaning; (b) grammar; (c) phonetics; (d) effects on listener; (e) content
5. Production rules occur in (a) a lexicon; (b) grammars; (c) semantic specs; (d) proofs; (e) queries
6. Parsing uses (a) a lexicon; (b) grammars; (c) semantic nets; (d) a knowledge base; (e) first-order logic proof rules
7. A language is a (a) knowledge base; (b) lexicon; (c) grammar; (d) set of strings; (e) none of these
8. The most difficult step in extracting a first-order logic sentence from a natural-language sentence may be (a) parsing; (b) lexical analysis; (c) semantic interpretation; (d) inference; (e) search
9. Disambiguation of sentences may require (a) applying grammar rules; (b) applying lexical rules; (c) querying a knowledge base; (d) use of stochastic methods; (e) none of these
10. ____ grammars are ambiguous (a) lexical; (b) some context-free; (c) all context-free; (d) stochastic; (e) fuzzy
11. Creating database entries from text is (a) inference; (b) querying; (c) learning; (d) information extraction; (e) none of these
12. Hidden state in natural-language processing includes (a) the intention of the speaker; (b) the knowledge of the listener; (c) the grammar of the language; (d) the language's semantics; (e) the state of the world

Short and longer answer

See also problems on the assignment sheet.

1. Distinguish abduction from deduction. (5a)
2. What is inductive inference? Give an example. (5a)
3. Describe some forms of supervision in supervised learning. (5a)
4. Describe two tools or techniques of concept learning. (5a)
5. What is the goal of learning? Describe two ways to attain this. (5a)
6. What is this? Describe an automated way to construct it. (5a)

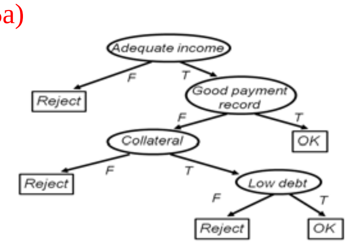

7. What is a neural net used for, and how is it prepared for operation? (5b)
8. What is a perceptron, and how does it work? (5b)
9. Describe perceptron learning. (5b)
10. What is the role of *weight adjustment* in learning? (5b)
11. Describe associative memory. (5b)
12. Describe backpropagation. (5b)
13. What is the main tool in connectionist learning and how does it work? (5b)
14. Describe neurons and how they work together. (5b)
15. What is the structure of a language called, and how is it defined? (5c)
16. What is a language, and how are some languages specified? (5c)
17. Distinguish semantics, syntax, and pragmatics. (5c)
18. What are the terms for language structure, meaning in language, and the effects of utterances? (5c)
19. What are the main aspects of natural language processing? (5c)
20. Construct and test a perceptron for some problem sets (see case, Luger05, pp. 461-462): >, addition, ($|a - b| < c$), subtraction, OR, AND
21. Explain the linear-separability constraint with empirical data
22. What is evolutionary computation? (5d)
23. How does evolutionary computation address the function-optimization problem? (5d)
24. In what kind of AI are *fitness functions* used, and how? (5d)
25. In what kind of AI are genetic algorithms used? How? (5d)
26. Solve by genetic algorithm (a) CNF-SAT; (b) TSP (Luger05, pp. 511-513)

Study questions on Topic 6: Reinforcement learning and adaptation

1. Interaction and intelligent behavior

1. Online search is necessary for (a) inference; (b) learning; (c) exploration of environment; (d) decisions; (e) none of these
2. Adaptation is required in ____ environments (a) static; (b) episodic; (c) dynamic persistent; (d) dynamic episodic; (e) none of these
3. Solutions to problems in interactive environments are often (a) inferential; (b) algorithmic; (c) adaptive; (d) entirely model-driven; (e) none of these
4. What kind of learning is required in dynamic and persistent environments? (a) supervised; (b) reinforcement; (c) connectionist; (d) training-intensive; (e) none of these
5. Dominant strategies and zero-sum are associated with ___ theory (a) game; (b) decision; (c) utility; (d) probability; (e) none of these
6. Policy search is used with which type of problem? (a) deterministic fully observable; (b) stochastic fully observable; (c) stochastic partially observable; (d) all problems; (e) no problems
7. In accessible environments, current state is identified by (a) actions; (b) percepts; (c) deduction; (d) chance; (e) none of these
8. Markov decision problems are associated with ____ problems (a) toy; (b) hard; (c) deterministic; (d) stochastic; (e) none of these
9. In game theory, a dominant strategy (a) always wins; (b) is better than others regardless of opponent strategy; (c) is a Nash equilibrium; (d) is better than random guessing; (e) none of these
10. Zero-sum games (a) are unwinnable; (b) are for one player; (c) are only win-lose; (d) may involve scores of zero; (e) none of these
11. Partial-information games are solved using (a) deduction; (b) belief states; (c) precise utility values; (d) proof; (e) none of these

2. Decision theory and expected utility

1. Expected utility is important for ____ agents (a) reflex; (b) goal-directed; (c) knowledge-based; (d) social; (e) none of these
2. Rational agents are likely to use (a) propositional logic; (b) probability theory; (c) decision theory; (d) calculus; (e) organization
3. Rational agents are concerned immediately with (a) past utility; (b) expected utility; (c) reward; (d) information; (e) past action
4. Rational agents are characterized by acting on (a) input/output rules; (b) expert knowledge; (c) a belief state; (d) a world state; (e) inference directly from facts
5. Value of information matters in (a) inference; (b) exploration; (c) belief; (d) knowledge; (e) expected reward
6. Value iteration updates (a) estimated utilities of states; (b) goals; (c) a knowledge base; (d) state-action mapping; (e) none of these
7. Policy iteration updates (a) estimated utilities of states; (b) goals; (c) a knowledge base; (d) state-action mappings; (e) none of these

3. Reinforcement learning and POMDPs

1. Policy search occurs in (a) a state space; (b) a range of possible belief values; (c) a belief state space; (d) an environment; (e) none of these
2. The Markov property holds for a system if a system will go into a given state (a) deterministically; (b) with a probability that depends on all past history; (c) with a probability that depends on recent history; (d) under all conditions; (e) none of these
3. A policy is (a) a mapping from states to actions; (b) a planned action sequence; (c) an algorithm; (d) a mapping of percepts to actions; (e) none of these
4. In interactive environments, a agent requires (a) a reflex mapping; (b) an action sequence; (c) planning under uncertainty; (d) a policy; (e) none of these
5. Utility of a state is (a) reward obtained in that state; (b) expected long-term reward; (c) unmeasurable; (d) independent of reward; (e) none of these
6. Reward is (a) observable in advance; (b) a guide to utility; (c) always obtained in a delayed way; (d) determined by querying a knowledge base; (e) none of these
7. Value of information (a) is reward; (b) is utility; (c) determines a state's utility; (d) is part of the utility of a state-action pair; (e) none of these
8. A good policy maximizes (a) information obtained by an action; (b) reward obtained by an action; (c) utility of an action; (d) knowledge; (e) none of these
9. A reward function maps (a) states to action sequences; (b) states to actions; (c) states to reals; (d) (state,action) pairs to reals; (e) none of these
10. A policy maps (a) states to action sequences; (b) states to actions; (c) states to reals; (d) states to actions; (e) none of these
11. A value function maps (a) states to action sequences; (b) states to actions; (c) states to reals; (d) (state,action) pairs to reals; (e) none of these
12. A policy is a mapping of (a) natural numbers to words; (b) income to premium; (c) states to actions; (d) percepts to outputs; (e) problems to algorithms

4. POMDPs and reinforcement learning

1. Reinforcement learning is distinguished from supervised learning in that (a) RL has a teacher; (b) RL's environment is interactive; (c) training is increased; (d) it is connectionist; (e) none of these
2. A set of hard problems in interactive environments are (a) state-space search; (b) one-player game; (c) planning action sequences; (d) POMDPs; (e) none of these
3. Belief state in solving POMDPs is (a) a set of states; (b) a fuzzy assertion; (c) a probability estimate that a state holds; (d) a probability distribution over all states; (e) none of these
4. A POMDP is a (a) stochastic differentiation problem; (b) stochastic decision problem; (c) deterministic process; (d) delayed payoff; (e) none of these
5. Reinforcement-learning agents (a) answer queries; (b) adapt to their environments; (c) operate in multi-agent systems; (d) seek to prove assertions; (e) use training to improve ability to act
6. Reinforcement learning is (a) model free; (b) model driven; (c) goal driven; (d) data driven; (e) none of these
7. Reinforcement is (a) immediate; (b) calculated; (c) utility; (d) often delayed; (e) none of these
8. Q learning learns (a) reward states; (b) utilities of states; (c) values of (state, action) pairs; (d) concepts; (e) none of these
9. Reinforcement learning searches (a) a knowledge base; (b) for a concept; (c) for a policy; (d) for an action sequence; (e) none of these
10. Reinforcement learning is appropriate with (a) concept training; (b) knowledge engineering; (c) inference; (d) partially observable interactive environments; (e) none of these

5. Robotics and embodied intelligence

1. The subsumption architecture (a) challenges the notion of explicitly centralized representation; (b) uses first-order logic; (c) supports supervised learning; (d) provides a heuristic for state-space search; (e) none of these
2. Subsumption architecture sees intelligent behavior as (a) following inference; (b) relying on case frames; (c) emerging from interaction with the environment; (d) dependent on scripts; (e) none of these
3. Robots in a multi-agent system are likely to be (a) situated and autonomous; (b) virtual and under central control; (c) situated and under central control; (d) virtual and autonomous; (e) none of these
4. In what branch of computing are agents often *situated* and is intelligence *embodied* ? (a) batch processing; (b) user interfaces; (c) robotics; (d) natural-language processing; (e) none of these
5. In intelligent robotics, unlike other areas of AI, intelligence may be (a) rational; (b) Bayesian; (c) complex; (d) situated; (e) none of these
6. In intelligent robotics, unlike other areas of AI, intelligence may be (a) rational; (b) Bayesian; (c) complex; (d) embodied; (e) none of these
7. The environment of a mobile robot may be (a) knowledge based; (b) multi agent; (c) deterministic; (d) fully observable; (e) none of these
8. Percepts of a robot are obtained via (a) wireless internet; (b) queries; (c) sensors; (d) effectors; (e) none of these
9. Actions of a robot are taken via (a) wireless internet; (b) belief maintenance; (c) message passing; (d) effectors; (e) none of these
10. Image processing includes (a) tactile sensing; (b) belief networks; (c) knowledge-base updates; (d) object recognition; (e) none of these
11. Converting sensor input to an internal representation of the environment is (a) perception; (b) information gathering; (c) knowledge base querying; (d) navigation; (e) none of these
12. An alternative to modeling the environment for robots is (a) heuristics; (b) reactive control; (c) fuzzy logic; (d) belief maintenance; (e) none of these
13. A subsumption architecture for robots relies chiefly on (a) inference; (b) probabilistic reasoning; (c) interactions among layers of a system; (d) creating a representation of the world; (e) none of these
14. An architecture that uses interacting layers of a system to organize response is (a) Markov; (b) reflex; (c) goal driven; (d) subsumption; (e) none of these

Short and longer answer

1. What special behavior is required of an intelligent agent in a dynamic environment? (6a)
2. In what sorts of environments is adaptation required? Why? (6a)
3. What are some characteristics of solutions to problems in interactive environments? (6a)
4. Relate intelligence to interaction. (6a)
5. In interactive computation, what is the arrangement of percepts and actions in time? (6a)
6. What is a *rational agent*? (6b)
7. What is a *utility-based* agent? (6b)
8. Relate utility to decision theory. (6b)
9. What has utility? What is utility? (6b)
10. Distinguish *reward* from *utility*. (6b)
11. Distinguish value of information from value of reward. (6b)
12. What is a *policy* and what artifact in AI has one? (6b)
13. Dominant strategies, zero-sum, and the prisoners' dilemma are associated with which kind of theory? Define one of the above terms. (6b)
14. What is the term for a mapping from perceived states to actions, what environments require it, and how may it be constructed? (6b)
15. What does a good policy maximize? Name a way to do this. (6b)
16. Distinguish exploration from exploitation. In which kind of AI does this matter? (6b)
17. What kind of learning is required in dynamic and persistent environments? Why? (6c)
18. Policy search is used with which type of problem? (6c)
19. What is a POMDP, and what AI approaches are used with POMDPs? (6c)
20. Describe the Markov property class of environments that have it. (6c)
21. Distinguish reinforcement learning from supervised learning. (6c)
22. What sort of learning is Q learning? What does it learn? (6c)
23. Describe policy search. (6c)
24. In what branch of computing are agents *situated* and is intelligence *embodied* ? Describe the interaction in this branch of AI. (6d)
25. Describe the environment of a robot and describe its ways of operating. (6d)
26. In what kind of AI are sensors and effectors found? Describe some. (6d)
27. In what ways may robotic systems be designed for robustness? (6d)
28. Explain how a robot may solve a Markov problem. (6d)
29. Describe features of a robotic vision system. (6d)
30. Describe the subsumption architecture for robots. (6d)

Study questions on Topic 7: Multi-stream adaptive interaction

1. Distributed AI

1. Social agents are characterized by (a) answering queries; (b) adapting to their environments; (c) operating in multi-agent systems; (d) seeking to prove assertions; (e) using training to improve ability to act
2. Distributed AI often consists of (a) coordination; (b) inference; (c) belief; (d) planning; (e) knowledge
3. Coordination is used in (a) inference; (b) state-space search; (c) probabilistic reasoning; (d) coordination; (e) none of these
4. Distributed AI is closely associated with (a) Bayesian inference; (b) reinforcement learning; (c) multi-agent systems; (d) natural language processing; (e) none of these
5. Sociogenetic adaptation is closely associated with (a) Bayesian inference; (b) reinforcement learning; (c) multi-agent systems; (d) natural language processing; (e) none of these
6. Learning in multi-agent systems is (a) supervised; (b) reinforcement; (c) sociogenetic; (d) connectionist; (e) none of these
7. Multi-agent systems enable (a) knowledge representation; (b) inference; (c) planning; (d) distributed AI; (e) none of these
8. Behavior is (a) planning; (b) inference; (c) action to obtain information; (d) action to change the environment; (e) none of these
9. Emergent intelligence is (a) processing oriented; (b) representation oriented; (c) without representation; (d) single-agent; (e) none of these
10. The most difficult environments are (a) persistent, dynamic, and virtual; (b) episodic, dynamic and physical; (c) episodic, static, and virtual; (d) persistent, dynamic, and physical; (e) none of these

2. Multi-agent systems

1. Situated agents are often found in (a) expert systems; (b) state-space search; (c) multi-agent systems; (d) supervised learning systems; (e) none of these
2. Autonomous agents are often found in (a) expert systems; (b) state-space search; (c) multi-agent systems; (d) supervised learning systems; (e) none of these
3. In which branch of AI is a concurrent action list used? (a) expert systems; (b) state-space search; (c) multi-agent systems; (d) supervised learning systems; (e) none of these
4. A *mission* is characteristic of (a) an algorithm; (b) an interactive process; (c) a multi-agent system; (d) a parallel system; (e) none of these
5. Indirect interaction requires (a) mutual causality between entities that do not exchange messages; (b) message passing; (c) synchrony; (d) static I/O; (e) none of these
6. The problem solved by a multi-agent system is called a(n) (a) algorithm; (b) function; (c) service; (d) mission; (e) process
7. Sequential-interactive agents offer a(n) (a) algorithm; (b) function; (c) service; (d) mission; (e) process
8. Indirect interaction is characteristic of systems featuring (a) mobility of agents and locality of interaction; (b) global interaction and mobility; (c) algorithmic problems; (d) knowledge-based inference; (e) none of these
9. The most difficult environments are (a) dynamic and virtual; (b) persistent and physical; (c) static and physical; (d) dynamic and episodic; (e) none of these
10. Systems featuring mobility of agents and locality of interaction often also typically feature (a) direct interaction; (b) no interaction; (c) indirect interaction; (d) semantic networks; (e) none of these

3. Stigmergy and self-organizing systems

1. Self-organization is (a) algorithmic; (b) sequential-interactive; (c) decentralized; (d) centralized; (e) none of these
2. Self-organization is associated with ____ behavior (a) deductive; (b) reactive; (c) planning; (d) emergent; (e) none of these
3. Emergent behavior is associated with (a) self-organization; (b) planning; (c) reflex agents; (d) stochastic reasoning; (e) none of these
4. Decentralized intelligence is associated with (a) early AI research; (b) reinforcement learning; (c) distributed AI; (d) expert systems; (e) none of these
5. Stigmergy makes use of (a) the environment; (b) first-order logic; (c) adaptive learning; (d) Bayesian reasoning; (e) none of these
6. Coordination via the environment is (a) minimax; (b) stigmergy; (c) centralized; (d) a heuristic; (e) none of these
7. Distributed AI uses (a) natural language; (b) coordination; (c) heuristics; (d) policy search; (e) none of these
8. Indirect interaction is in contrast to (a) deduction; (b) distributed AI; (c) stigmergy; (d) anonymous coordination; (e) none of these
9. Stigmergy uses (a) indirect interaction; (b) language; (c) reasoning; (d) a knowledge base; (e) none of these
10. Self-organization is observed in (a) English grammar; (b) garden design; (c) social insects; (d) proofs; (e) the minimax algorithm
11. Indirect interaction features (a) anonymity; (b) synchronization; (c) use of predicate logic; (d) fuzzy reasoning; (e) none of these
12. Indirect interaction features (a) space decoupling; (b) synchronization; (c) use of predicate logic; (d) fuzzy reasoning; (e) none of these
13. Ant foraging by use of chemical trails is an example of (a) inference; (b) linguistic processing; (c) stigmergy; (d) evolution; (e) none of these

Short and longer answer

1. In what branch of AI is coordination used? What is coordinated and how? (7a)
2. Describe distributed AI. (7a)
3. What is sociogenetic adaptation? Contrast it to other kinds of adaptation. (7a)
4. Describe an example of distributed AI or multi-agent systems. (7a)
5. What is a multi-agent system? (7a)
6. What are features of many multi-agent systems? (7a)
7. In which branch of AI is a concurrent action list used? How? (7a)
8. What is self-organization? (7b)
9. Is a single neuron intelligent? Explain how the brain implements distributed intelligence. (7b)
10. What is emergent behavior? (7b)
11. What is decentralized intelligence? (7b)
12. Describe an application of multi-agent systems. (7a)
13. In what kind of environment to agents in a multi-agent system operate? Explain. (7a)
14. In what branch of AI does multi-stream interaction occur? Explain. (7a)
15. In what branch of AI does indirect interaction occur? Explain. (7b)
16. What is stigmergy? Relate it to intelligent systems. (7b)
17. Contrast indirect interaction to message passing. (7b)

Study questions on Topic 8: Future prospects and philosophical considerations

1. Theories of mind

1. Ontology is the study of (a) the relationship of thought to matter; (b) what is; (c) how we know things; (d) experience; (e) none of these
2. Epistemology is the study of (a) the relationship of thought to matter; (b) what is; (c) how we know things; (d) experience; (e) none of these
3. Theory of mind is the study of (a) the relationship of thought to matter; (b) what is; (c) how we know things; (d) experience; (e) none of these
4. Phenomenology is the study of (a) the relationship of thought to matter; (b) what is; (c) how we know things; (d) experience; (e) none of these
5. Rationalism is a theory of (a) logic; (b) mind; (c) mathematics; (d) ethics; (e) ontology
6. Rationalism is a theory of (a) logic; (b) epistemology; (c) mathematics; (d) ethics; (e) ontology
7. Empiricism is a theory of (a) logic; (b) mind; (c) mathematics; (d) ethics; (e) ontology
8. Empiricism is a theory of (a) logic; (b) epistemology; (c) mathematics; (d) ethics; (e) ontology
9. _____ pictures knowledge as obtained from mind (a) empiricism; (b) rationalism; (c) relativism; (d) individualism; (e) communitarianism
10. _____ pictures knowledge as obtained from the senses (a) empiricism; (b) rationalism; (c) relativism; (d) individualism; (e) communitarianism

2. Objections to weak and strong AI

1. Weak AI is (a) pre-1980 research; (b) the claim that machines can play games; (c) the claim that machines can simulate intelligence; (d) the claim that machines can be intelligent; (e) none of these
2. Strong AI is (a) post-2000 research; (b) the claim that machines can play games; (c) the claim that machines can simulate intelligence; (d) the claim that machines can be intelligent; (e) none of these
3. The Chinese Room argument asserted that (a) a system of symbols written on paper can't have understanding; (b) a machine can't have the experience of thinking; (c) Gödel's theorem proves machines have limited capacity; (d) machines don't reference actual things in the world; (e) none of these
4. The argument against strong AI based on *phenomenology* asserted that (a) a system of symbols written on paper can't have understanding; (b) a machine can't have the experience of thinking; (c) Gödel's theorem proves machines have limited capacity; (d) machines don't reference actual things in the world; (e) none of these
5. The argument against strong AI based on *intentionality* asserted that (a) a system of symbols written on paper can't have understanding; (b) a machine can't have the experience of thinking; (c) Gödel's theorem proves machines have limited capacity; (d) machines don't reference actual things in the world; (e) none of these
6. The argument against strong AI that states that machines can't have the experience of thinking is based on (a) utilitarianism; (b) ontology; (c) ethics; (d) phenomenology; (e) epistemology
7. The instructor objects to some arguments against AI as too (a) ontological; (b) abstract; (c) rigid; (d) human-centered; (e) unmathematical

3. Future prospects

1. The best program achievable to solve a problem of adaptation to an environment has (a) rationality; (b) intelligence; (c) tractability; (d) bounded optimality; (e) none of these
2. Decision-theoretic metareasoning uses (a) utility theory; (b) theory of the value of information; (c) probability theory; (d) complexity theory; (e) none of these
3. Bounded optimality is a feature of (a) proofs; (b) AI programs; (c) problems; (d) learning; (e) none of these
4. A reflective agent architecture reasons about (a) the environment; (b) the state of the agent; (c) the states of other agents; (d) abstract environments; (e) probabilities
5. Technological singularity occurs when intelligent machines (a) understand humans; (b) rebel against humans; (c) invent intelligent machines; (d) have civil rights; (e) are simpler and simpler
6. Perfect rationality is (a) error-free deduction; (b) ability to prove any true assertion; (c) unconditional maximum utility; (d) maximum utility given computing resources; (e) none of these
7. Bounded rationality is (a) error-free deduction; (b) ability to prove any true assertion; (c) unconditional maximum utility; (d) maximum utility given computing resources; (e) none of these

Short and longer answer

1. Distinguish ontology, epistemology, theory of mind, and phenomenology. (8a)
2. Describe your theory of mind and compare it to other theories. (8a)
3. Contrast rationalism and empiricism and explain what approaches to AI would seem most consistent with each. (8a)
4. What are some ethical issues related to AI? What are your views on them? (8a)
5. Distinguish between weak and strong AI. (8b)
6. Describe arguments for or against Strong AI. (8b)
7. Describe arguments for or against Weak AI. (8b)
8. Argue that intelligence is associated with (a) humanness; (b) computation or symbol manipulation (8b)
9. Argue that human intelligence is artificial; not artificial. (8b)
10. Does AI deny the uniqueness of human thought? Defend or refute. (8b)
11. What are some future prospects of AI that have been discussed in the literature? (8b)
12. How have human intelligence and AI been compared? (8b)
13. Can AI exceed human intelligence? Explain. (8b)
14. Is AI tending to catch up with human intelligence? Explain. (8b)
15. Distinguish bounded rationality from bounded optimality. (8c)
16. Define and justify bounded optimality. (8c)
17. What are the main obstacles to a perfect intelligence that instantly knows exactly what to do next? How are the obstacles overcome, according to some AI researchers? (8c)

Study questions on course summary (multiple topics)

1. Describe some stages in the development of AI discussed in this course. (9a)
2. How do the different approaches discussed in this course implement intelligent behavior? (9a)
3. In the terms discussed in this class, contrast the notion of *reasoning and inference* with the notion of *adaptation*.
4. What are some promising areas of application of AI concepts? Give reasons and specifics.
5. What is artificial intelligence?
6. What are misconceptions about artificial intelligence you are aware of, and where do they come from?
7. Does artificial intelligence exist? If not, could it? Explain.
8. Contrast *state-space search* and *policy search*.
9. Describe different forms of *rationality*.
10. Describe a range of environments and the ways that intelligent agents behave in them.
11. Describe two or three ways in which artificial-intelligence research has been inspired by natural phenomena.
12. Describe ways to store, acquire, and maintain knowledge or belief.
13. Contrast different kinds of *learning*.