1 These are all the glosary terms from my AI class Csce876 from UNL

A^* Search	The most widely-known form of the best-first search, A^* evaluates nodes using the evaluation function: $$f(n) = g(n) + h(n)$$ A^* has several nice properties. Using the Tree-Search Algorithm and an *admissible* heuristic function we have: • Optimality • Complete • Optimally efficient If the heuristic function is consistent, then we have: • Optimality • Complete
C^*	This is the actual cost of the optimal solution path.
$f(n)$ (Evaluation Function)	The evaluation function is "the" problem-specific knowledge beyond the definition of the problem itself that is used to determine which node is chosen for expansion in the search process.
$g(n)$ (Cost to node function)	This function gives the actual cost to reach node n on the current path.
$h(n)$ (Heuristic function)	The estimated cost of the cheape4st path from n to a (the) goal.
$k-$consistency	A CSP is $k-$consistent if, for any set of $k-1$ variables and for any consistent assignment to those variables, a consistent value can always be assigned to any k^{th} variable. CITATION(aitextbook) (p 147)
n-queens problem	The the goal of the $n-queens$ problem is to place n queens on an $n \times n$ chess board where no queen is in a position to attack another queen.
Abstraction (definition, goal, example)	The process of removing detail from a representation is called abstraction. We say the abstraction is valid if we can expand any abstract solution into a solution in the more detailed world.
Admissible Heuristic	$h(n)$ is an admissible heuristic provided that $h(n)$ never overestimates the cost to reach the goal. The direct consequence of this is that $f(n)$ for A^* never overestimates the true cost of a solution through n.
Agent	An agent is just something that acts CITATION(aitextbook)
Alldiff constraint	A constraint that indicate that all involved variables must be assigned a different value. CITATION(aitextbook) (p 140)
Alliances	Alliances occur in many games involving more than two players. Alliances involve a group of players that are cooperating in a competitive environment. CITATION(aitextbook) (p 166)
Alpha-beta pruning	This procedure is based on a DFS. Consider a node n somewhere in the tree such that Player has a choice of moving to that node. If Player has a better choice m either at the parent node of n or at any choice point further up, then *n will never be reached in actual play.* So once we have found out enough about n (by examining its descendants in a DFS) to reach this conclusion, we can prune it. Alpha-beta pruning gets its name from the two parameters: $\alpha =$ The value of the best choice we have found so far at any choice point along the path for max. $\beta =$ The value of the best choice we have found so far at any choice point along the path for min. CITATION(aitextbook) (p 168-169)
And-elimination	$\frac{A \wedge B}{A}$ CITATION(aitextbook) (p 211)
Anytime algorithm	An anytime algorithm is an algorithm whose output quality improves gradually over time, so that it has a reasonable decision ready whenever it is interrupted. CITATION(aitextbook) (p 971).

Arc consistency	An arc from X to Y is consistent if, for every value x of X there is some value y of Y that is consistent with x. CITATION(aitextbook) (p 145)
Arity (of a relation or a function)	Arity is the number of arguments CITATION(aitextbook) (p 246)
Assertion	A FOL logic sentence added to the KB are called an assertion. CITATION(aitextbook) (p 253)
Atmost constraint	The atmost constraint has a number P and variables $P_1, ..., P_k$ that are assigned numbers such that $\sum_{i=1}^{k} P_i \leq P$. CITATION(aitextbook) (p 148)
Atomic sentences	Consists of a single proposition symbol. CITATION(aitextbook) (p 204)
Autonomy (autonomous agent)	Not relying on prior knowledge of its designer, but rather on its own precepts. CITATION(aitextbook)
Axiom	Axioms provide the basic factual information form which conclusions can be derived. CITATION(aitextbook) (p 255)
Background knowledge	This is the sentences of information that a knowledge base is initialized with. CITATION(aitextbook) (p 196)
Backjumping	The backjumping method backtracks to the most recent variable in the conflict set. CITATION(aitextbook) (p 149)
Backtracking Search	A variant of the depth-first search it uses still less memory. In backtracking, only on successor is generated at a time rather than all successors; each partially expanded node remembers which successor to generate next. Thus the memory requirement is only $O(m)$ instead of $O(bm)$.
Backtracking search	The backtracking search is a depth-first search that chooses values for one variable at a time and backtracks when a variable has no legal values left to assign. CITATION(aitextbook) (p 141)
Backward chaining	Backward chaining starts with the query for the fact q and finds those implications in the knowledge base that conclude q. If all of the premises of one of these implications can be proved true (through backward chaining), then q is true. CITATION(aitextbook) (p 220)
Belief State	When in a sensorless problem, the agent may only know what states it can possibly start in. By applying the successor function to each of these states it can determine a set of states that could be reached by a specific action. The set of states that the agent could be in at any one time is the *belief state*.
Best-first Search	Best-first search is an instance of the general Tree-Search or Graph-Search algorithm in which a node is selected for expansion based on an evaluation function, $f(n)$.
Biconditional	$A \Leftrightarrow B$ is a biconditional that states that A is true if and only if B is true or alternately if $A \Rightarrow B$ and $B \Rightarrow A$. CITATION(aitextbook) (p 205)
Bidirectional search	The strategy is to run two simultaneous searches, one from the initial state forward and one from the goal state backwards. This is not always possible, and requires that the successor function be invertible. This should cut the search depth by a factor of two so that the storage requirements are $O(b^{d/2}) << b^d$.
Binary constraint	A binary constraint relates two variables. CITATION(aitextbook) (p 140)
Binding list	When posing a query $\exists x\ Predicate(x)$, answering yes is not very helpful. The standard form of an answer to such a query includes a set of values that satisfy the query. This set is called a substitution or binding list. CITATION(aitextbook) (p 254)
Blind Search	See Uninformed search
Body (of a Horn clause)	The negative literals of a horn clause form the body. CITATION(aitextbook) (p 218)
Boolean CSPs	Boolean CSPs are CSPs where the domains of the variables are either true or false. CITATION(aitextbook) (p 139)
Bounded differences (constraint of)	A bounded difference constraint has the form of $const_1 \quad \theta \quad x - y \quad \theta \quad const_2$ where $\theta \in <, >, \leq, \geq$. CITATION(ainotes)
Bounds Propagation	We say that a CSP is bounds-consistent if for every variable X, and for both the lower bound and upper bound values of X, there exists some value of Y that satisfies the constraint between X and Y, for every variable Y. This kind of constraint propagation is called *bounds propagation*. CITATION(aitextbook) (p 148)
Branching Factor	The *branching factor* is the maximum number of successors of any node.
Breadth-first Search	The BFS expands nodes from a FIFO queue. That is, as each node is discovered it is put into a FIFO queue. This results in all nodes being expanded at each level before a new level is expanded. The storage requirement is $O(b^{d+1})$. This requirement is prohibitive.

Chinese room experiment	The system consists of a human, who understands only English (CPU), equipped with a rule book (PROGRAM), written in English, and various stacks of paper, some blank, some with indecipherable inscriptions (memory). The system is inside a room with a small opening to the outside. Through the opening appear slips of paper with indecipherable symbols. The human finds matching symbols in the rule book, and follows the instructions. The instructions may include writing symbols on new slips of paper, finding symbols in the stacks, rearranging the stacks and so on. Eventually, the instructions will cause one or more symbols to be transcribed onto a piece of paper that is passed back to the outside world with the correct response. Searle argues against the Turing test in that just running the right program does not generate understanding. CITATION(aitextbook)
Chronological backtracking	The process of a backtrack search where we backup to the preceding variable and try a different value is called chronological backtracking, because the most recent decision point is revisited. CITATION(aitextbook) (p 148)
Coercion	Coercion is the ability of an agent to reach a goal state even in a sensorless problem. The agent knows that it starts in any of the states but can reach a goal by reasoning about sets of states. See the Vacuum cleaner agent on p84 for a good example.
Cognitive science	The interdisciplinary field of cognitive science brings together computer models from AI and experimental techniques from psychology to try to construct precise and testable theories of the workings of the human mind. CITATION(aitextbook)
Competitive Agent	In a multiagent system, when one's performance depends upon the performance of another in an inverse relationship, then the two agents are in competition and may be called competitive agents.
Complementary literals	Literals such that one is a negation of the other. CITATION(aitextbook) (p 214)
Complete-state formulation	A *complete-state formulation* starts with a complete proposed solution and then changes to it to find a solution that satisfies the goal test.
Completeness	If the exists a solution, then the algorithm is guaranteed to find a solution.
Completeness (of inference)	An inference i is complete if whenever $\text{KB} \models \alpha$, it is also true that $\text{KB} \vdash_i \alpha$. That is, the procedure will answer any question whose answer follows from what is know by the KB. CITATION(ainotes) (10.23)
Compositional (languages)	In a compositional language, the meaning of a sentence is a function of the meaning of its parts. CITATION(aitextbook) (p 241)
Compositionality	In a compostional language, the meaning of a sentence is a function of the meaning of its parts. CITATION(aitextbook) (p 241)
Condition-action rule	A condition action rule is a tuple (Condition, Action). When the condition is matched by the sensor input, the Simple Reflex agent will perform the "Action".
Conflict set	A more intelligent approach to backtracking than chronological backtracking is to go all the way back to one of the set of variables that *caused the failure*. This set is called the **conflict set**. CITATION(aitextbook) (p 148)
Conflict-directed backjumping	The backtracking method that backs up to the source of a conflict that prohibits the finding of a consistent assignment to the remaining variables. CITATION(ainotes,aitextbook) (p 149)
Conjunctive normal form	A sentence expressed as a conjunction of disjunctions of literals is said to be in conjunctive normal form. CITATION(aitextbook) (p 215)
Connectionism	Connectionism models mental or behavioral phenomena as the emergent processes of interconnected networks of simple units. There are many different forms of connectionism, but the most common forms utilize neural network models. CITATION(Wikipedia:Connectionism)
Consistency (as a property of the h function)	A heuristic $h(n)$ is consistent if, for every node n and every successor n' of n generated by any action a, the estimated cost of reaching the goal from n is no greater than the step cost of getting to n' plust the estimated cost or reaching the goal from n':
Consistent assignment	An assignment that does not violate any constraints is called a consistent or legal assignment. CITATION(aitextbook) (p 137)
Constant symbol	A symbol in FOL that stand for objects. CITATION(aitextbook) (p 246)
Constraint graph	A graph where nodes correspond to the variables of the problem and arcs correspond to the constraints. CITATION(aitextbook) (p 138)
Constraint hypergraph	A constraint hypergraph contains two types of nodes: variables and constraints. arcs from constraint nodes to variable nodes indicate that the variable is in-

Constraint propagation	This is the general term for propagation the implications of a constraint on one variable onto other variables. CITATION(aitextbook) (p 145)
Constraint scope	The scope of a constraint is the set of variables that are involved in the constraint CITATION(ainotes).
Constructive Search	This is the standard search process were we assign a value to a variable at each node starting from the first variable to the last. CITATION(ainotes).
Contingency problem	If the environment is partially observable or if actions are uncertain, then the agent's percepts provide new information after each action. Each possible percept defines a *contingency* that must be planned for. A contingency problem is *adversarial* if the uncertainty is caused by the actions of other agents.
Continuous domains	Where the variables require continuous values such as the real numbers we say we are using continuous domains. Problems that require precise values often require continuous domains. CITATION(aitextbook) (p 139)
Contours	The fact that f-costs are nondecreasing along any path also means that we can draw contours in the state space, just like the contours in a topographic map.
Crossover	In genetic algorithms when you cross two states you choose a point to split the states and "cross" the two states. State 1's first half with state 2's second half, and similarly for the second state. The split point is called the "crossover" point.
Current State	Local search algorithms operate using a single "current state" and generally move only to neighboring states. The current state is then a complete description of a solution that may not meet all of the constraints or may not be optimal. (Of course it could be optimal too, but then it would be a solution)
Cutset conditioning	Cutset conditioning is the process of finding the smallest cycle cutset. CITATION(aitextbook) (p 154)
Cycle cutset	A subset S called the cycle cutset is a set of variables of a CSP such that the constraint graph becomes a tree after the removal of S CITATION(aitextbook) (p 153).
Data driven	Data driven reasoning is reasoning in which the focus of attention starts with the known data. CITATION(aitextbook) (p 219)
Data Mining	Data mining, also known as knowledge-discovery in databases (KDD), is the practice of automatically searching large stores of data for patterns. CITATION(Wikipedia:Datamining)
Decision problem	Is any problem that must be answered yes or no.
Deduction	Logical inference. That is to use the rules of logic to infer a conclusion. We say a conclusion is reached by deduction if it is logically inferred by the premise(s).
Deduction theorem	For any sentences α and β, $\alpha \models \beta$ if and only if the sentence $(\alpha \Rightarrow \beta)$ is valid. CITATION(aitextbook) (p 210)
Definite clause	Horn clauses with exactly one positive literal are called definite clauses. CITATION(aitextbook) (p 218)
Degree (ordering heuristic)	Attempts to reduce the branching factor by selecting the variable that is involved in the largest number of constraints. CITATION(aitextbook) (p 144)
Depth-First Search	This strategy expands nodes from a LIFO queue, also known as a stack. This results in nodes being expanded as deep as they can before exploring other nodes at the same level. Thus it is called the DFS. The space requirements for DFS are $O(bm)$. There is no guarantee that this will ever find a goal state because it may get caught in a loop.
Depth-Limited Search	This is just a DFS search where we limit the depth that the search may descend to. It is also not guaranteed to find a solution in general, but we may guarantee a goal state is found if we know a goal exists within a certain depth. This has the same requirements as DFS
Diameter (of a graph)	Given a graph $G = (V, E)$, the distance between two vertices u and v in G, denoted $d(u, v)$, is the length of the shortest $u - v$ path in G. The *diameter* of G, denoted diam(G), is the maximum distance between any two vertices in G. CITATION(GraphTheoryThulasiramanSwamy)
Domain (of a model)	The domain of a model is the set of objects that it contains. CITATION(ainotes) (p 245)
Domain (of a variable)	Given a CSP the domain of a variable are the values that the variable may be assigned CITATION(ainotes,aitextbook) (p 137)
Domain/degree heuristics (not in textbook)	The degree heuristic chooses the variable to assign that is involved in the largest number of constraints on other unassigned variables. CITATION(aitextbook) (p 144 - it is in the book)

Domination (of two functions)	We say a heuristic function $h_1(n)$ dominates another heuristic function $h_2(n)$ if $$h_1(n) \geq h_2(n)$$
Effective branching factor CITATION(ainotes)	One way to characterize the quality of a heuristic is using the *effective branching factor*. In words this is the the branching factor that a uniform tree of depth d would have in order to contain $N+1$ nodes. Where d is the depth of the search tree and N is the number of nodes generated. That is at each depth starting at 0 we have: $$N = 1 + b^* + (b^*)^2 + ... + (b^*)^d = \frac{(b^*)^{d+1} - 1}{b^* - 1}$$
Entailment	Let KB be a knowledge base of sentences and let α be a sentence. α is entailed by a KB if the models of the KB are *all* models of α. CITATION(ainotes) (11.5) For my own notes: Entailment is similar to implication except the domain of the variables are sentences (or propositions). KB $\models \alpha$ iff all models of KB are models of α. (*i.e.*, $M(KB) \subseteq M(\alpha)$) CITATION(ainotes) (11.5)
Environment: Deterministic vs. Stochastic	If the next state of the environment is completely determined by the current state and the action executed by the agent, then we say the environment is deterministic, otherwise we say it is stochastic.
Environment: Discrete vs. Continuous	The discrete/continuous distinction can be applied (1) to the state of the environment - is it constantly changing?, (2) to the way *time* is handled - is it continuous in the application or is it handled as discrete points?, and (3) to the *percepts* and *actions* of the agent - How can the agent perceive the environment, and how does it act (continuously or discretely)?.
Environment: Episodic vs. Sequential	If the agents experience is broken up into atomic episodes that do not affect actions in previous episodes, we call it episodic. If on the other hand current actions depend upon the sequence of actions taken previously, then we call the task environment sequential.
Environment: Observable: Fully vs Partially	A task environment is effectively fully observable if the sensors detect all aspects that are *relevant* to the choice of action. Otherwise we say the task environment is partially observable.
Environment: Single agent vs. Multiagent	The key distinction is whether an agent's behavior is best described as maximizing a performance measure whose value depends on other agents' behavior.
Environment: Static vs. Dynamic	If the environment can change while the agent is deliberating, we say the environment is **dynamic**. If the environment does not change while the agent is deliberating, we call the environment **static**. If the score changes (decreases) as the agent deliberates in a static environment, we call the environment **semi-dynamic**.
Environment: Strategic	If the environment is deterministic except for the actions of other agents, we say that the environment is strategic. (This was under Deterministic vs Stochastic in the book.)
Epistemological level	Abstract description of what the agent knows about the world. CITATION(ainotes) (10.4)
Exploration Problem	When the states and actions of the environment are unknown (e.g. you are dropped somewhere where you don't speak the language) the agent must act to discover them. This type of problem can be viewed as an extreme case of the contingency problem.
Extended interpretation	An extended interpretation specifies a domain element in addition to the mapping given by the interpretation. CITATION(aitextbook) (p 249).
Extensively defined constraint (not in textbook)	Extensionally defined constraints: all allowed tuples are listed. This is practical for defining arbitrary constraints. CITATION(ainotes)
Factoring	Factoring is the last step in resolution and involves removing multiple copies of literals from the expression. CITATION(aitextbook) (p 214)
Finite domains	Variables with finite (or discrete) domains are the simplest kind of CSPs. CITATION(aitextbook) (p 139)

First-Order Logic	First-Order logic adds to propositional • symbols: Objects, properties, functions and relations • connectives: quantifiers $\exists$ and $\forall$ CITATION(ainotes) (12.4,5)
Fitness function	In genetic algorithms each state produced in the next generation is evaluated or rated using a fitness function. This determines how ”good” the produced state is.
Fixed point	The point at which applying inference rules produces no new facts. CITATION(aitextbook) (p 218)
Forward chaining	If all the premises of an implication are known, then its conclusion is added to the list of known facts. The process continues until the desired fact is known. This process of chaining horn clauses together is called forward chaining. CITATION(aitextbook) (p 218)
Forward Checking	Whenever a variable X is assigned, the forward checking process looks at each unassigned variable Y that is connected to X by a constraint and deletes from Y's domain any value that is inconsistent with the value chosen for X. CITATION(aitextbook) (p 144)
Fringe	The collection of nodes that have been generated but not yet expanded is called the fringe (p 70).
Function	In terms of FOL, a function is a relation in which there is only one ”value” for a given ”input.” CITATION(aitextbook) (p 242)
Function (mathematical definition)	A relation $f : A \rightarrow B$ is said to be a function if and only if f maps each $a \in A$ to only one value in B. (Dr. Anderson)
Genetic algorithm	A genetic algorithm is a variant of stochastic beam search in which successor states are generated by combining two parent states.
Global constraint	Involves all the variables in a CSP. CITATION(ainotes)
Global optimum (maximum or minimum)	If elevation in a state space landscape corresponds to cost, then the aim is to find the lowest valley (minimum). If the landscape corresponds to an objective function, then the aim is to find the highest peak.
Goal Test	The goal test determines whether a given state is a goal state.
Goal-directed reasoning	Goal directed reasoning is useful for answering specific questions such as ”What shall I do now?” and ”Where are my keys?” Backward chaining is an example of goal directed reasoning. CITATION(aitextbook) (p 220)
Gradient ascent	The gradient defines the direction of the steepest ascent (in this case) or descent. (From calculus)
Greedy Algorithm	A greedy algorithm always makes the choice that looks best at the moment. CITATION(IntroToAlgorithms) The Best-first search, also called the ”greedy search” or ”greedy best-first search”, tries to expand the node that is closest to the goal, on the grounds that this is likely to lead to a solution quickly. Thus it evaluates nodes by using just the heuristic function $$f(n) = h(n)$$
Greedy local search	Hill climbing is sometimes called a greedy local search because it grabs a good neighbor state without thinking ahead about where to go next.
Ground term	A term with no variables is called a ground term. CITATION(aitextbook) (p 249)
Grounding	Grounding is the process of determining if the logical reasoning processes and the real environment in which the agent exists are consistent. The question is, ”how do we know that KB is true in the real world?” Of course if you worried about this, you would never have time to worry about anything else. CITATION(aitextbook) (p 204)
Hanhatten distance	This is the sum distances to the object in each dimension. Sometimes called the city block distance because you can't cut through the buildings.
Head (of a Horn clause)	The positive literal of a horn clause is called the head. CITATION(aitextbook) (p 205)
Heuristic Function	See $h(n)$.
Heuristic Search	See ”Informed search”
Higher-Order Logic	Higher-Order logic views the relations and functions in first-order logic as objects in themselves. CITATION(aitextbook) (p 244)

Hill climbing: random restart	Because hill climbing will find it harder and harder to find a better state as time goes on, we may stop the search process and start it from a new random location. This is called random restart.
Hill climbing: stochastic	Stochastic hill climbing chooses at random from among the uphill moves where the probability of selection can vary with the steepness of the uphill move.
Horn clauses	A horn clause is a disjunction of literals of which at most one is positive. CITATION(aitextbook) (p 204)
Incremental formulation	An incremental formulation involves operations that augment or *build* the state description. The goal is to build a state that satisfies the goal test and at any point in the process, if we find we can not satisfy the goal state using the last augmentation, we may back track.
Induction	General rules aquired by exposure to repeated associations between their elements. CITATION(aitextbook)
Inference	Deriving new sentences from old. In the case of logical agents, inference must obey the fundamental requirement that when one asks a question of the knowledge base, the answer should follow from what has been told the knowledge base previously. CITATION(aitextbook) (p 195)
Infinite domains	Domains where the size of the set is infinite. This can include discrete domains such as the integers. CITATION(aitextbook) (p 139)
Informed Search	Strategies that know whether one non-goal state is more promising than another non-goal state are called informed or heuristic searches.
Initial State	The state that the agent starts in.
Instantiated variable	A variable is instantiated if it has been assigned a value from its domain. CITATION(ainotes)
Intended interpretation	The intended interpretation is the one possible interpretation that we specify. CITATION(aitextbook) (p 246)
Intensively defined constraint (not in textbook)	Intensionally defined constraints are used when it is not practical or even possible to list all tuples. This is the form we normally see using variables and relationship predicates. CITATION(ainotes)
Interpretation	The semantics of FOL must relate sentences to models in order to determine truth. An interpretation specifies exactly which objects, relations and functions are referred to by the constant, predicate and function symbols.CITATION(aitextbook) (p 246) An *interpretation* for a set X of formulas **is a domain** D together with a rule that • assigns to each n-place predicate symbol (that occurs in a formula) of X an n-place predicate in D; • assigns to each n-place operation symbol of X an n-place operation in D; • assigns to each constant symbol of X an element of D; • assigns to $=$ the identity predicate $=$ in D, defined by: $a = b$ is true if and only if a and b are the same. CITATION(FOML)
Intractability	We say an algorithms is intractable if the time to run the algorithm increases exponentially as a function of the size of the problem. (From 310 or 823?)
Intractability	A problem is called intractable if the time required to solve instance of the problme grows exponentially with the size of the instances. (p 8)
Iterative-Deepening Search	Also called the *Iterative deepening depth-first search*, this is just a DFS performed to a depth d where d iteratively increases from 0 to some limiting value l until a goal is found. this will occur when d reaches the shallowest goal state. Space complexity will be the same as the DFS $O(bd)$. Unlike DFS it is complete and finds the optimal goal when the path cost is a non-decreasing function of the depth of the node. *In general, iterative deepening is the preferred uninformed search method when there is a large search space and the depth of the solution is not known.*
Iterative-Lengthening Search	Identical to the Iterative-Deepening Search except that the search uses an increasing path-cost limit instead of an increasing depth limit. This guarantees optimality while avoiding expensive memory requirements.
k-CNF	A sentence expressed as a conjunction of disjunctions of literals is said to be in conjunctive normal form. A sentence is $k-$CNF if it has exactly (or at most) k literals per clause. CITATION(aitextbook) (p 215)

Knowledge acquisition	Knowledge acquisition is the process of acquiring knowledge. This can be done by exploration or by (p261) applying to an outside expert to retrieve the knowledge required.
Knowledge base	A set (conjunction) of sentences that represent facts about the world. CITATION(aitextbook) (p 195)
Knowledge representation	One of the six disciplines of AI dedicated to the "representation" of knowledge that an entity knows or senses. CITATION(aitextbook)
Knowledge representation language	A formal language used to represent sentences of information in the knowledge base. CITATION(aitextbook) (p 195)
Leaf Node	A leaf node is an element of the fringe, that is, a node with no successors in the tree.
Limited rationality	It is not always possible to do precisely the right thing in complicated environments do to high computational demands. CITATION(aitextbook)
Linear constraints	Constraints in which each variable appears only in linear form. CITATION(ainotes)
Linear Programming	Linear programming involves CSPs where constraints must be linear inequalities forming a convex region. CITATION(aitextbook) p 140
Literal	A literal is single symbol or its negation. CITATION(aitextbook) (p 204)
Local Optimum	This corresponds to a peak or a valley in the state space landscape that is not the best solution. This is were a search may get stuck and need a random restart.
Local Search	Local searches do not care about the path too the solution, and so they don't keep track of where they have been. The use very little memory and they can often find solutions in large or infinite (continuous) state spaces for which systematic algorithms are unavailable.
Logical connectives	Not, conjunction, disjunction implication, biconditional which are respectively: $\neg \wedge, \vee, \Rightarrow, \Leftrightarrow$. CITATION(aitextbook) (p 204-205)
Logical equivalence	Two sentences A and B are logically equivalent if they are true in the same set of models. This is written: $A \equiv B$. CITATION(aitextbook) (p 210)
Min-conflict heuristic	The min-conflicts heuristic chooses a value for a variable that causes the fewest conflicts with other variables. CITATION(ainotes)
Minimax algorithm	The minimax algorithm computes the minimax decision form the current state. It uses a simple recursive computation of the minimax values of each successor state, directly implementing the defining equations. The recursion proceeds all the way down to the leaves of the tree, and then the minimax values are backed up through the tree as the recursions unwinds. CITATION(aitextbook) (p 165)
Minimax decision	When the Minimax value has been computed for each child node in the tree for the current node, the Minimax decision is the largest value of the child nodes. That is we choose the action that is the optimal choice (max value) because it guarantees the highest minimum value. CITATION(aitextbook) (p 165)
Minimax value	$Minimax\text{-}Value(n) = \begin{cases} Utility(n) & \text{if } n \text{ is a terminal state} \\ \max_{s \in Succ(n)} Minimax\text{-}Val(s) & \text{if } n \text{ is a max node} \\ \min_{s \in Succ(n)} Minimax\text{-}Val(s) & \text{if } n \text{ is a min node.} \end{cases} \quad (1)$
Minimum remaining values (a.k.a. least domain variable)	A heuristic that chooses the variable with the fewest legal values.CITATION(aitextbook) (p 143)
Missionaries and cannibals problem	Three missionaries and three cannibals are on one side of a river, along with a both that can hold one or two people.. Find a way to get everyone to the other size without ever leaving a group of missionaries in one place outnumbered by the cannibals in that place (p 90).
Model (in general)	A model is a world in which a sentence is true under a particular interpretation. CITATION(ainotes) (11.4)
Model checking	Given a knowledge base KB and a sentence α, model checking enumerates all possible models to check that a sentence α is true in all models in which KB is true. OR, model checking ensures that α is true in each model contained in KB. CITATION(aitextbook) (p 202)
Model in propositional logic	A model is a mapping from proposition symbols directly to truth or falsehood. The models of a sentence are the mappings that make the sentence true. CITATION(ainotes)
Modus ponens	$\frac{\alpha \Rightarrow \beta, \quad \alpha}{\beta}$ CITATION(aitextbook) (p 211)

Monotonicity (as a property of the h function	A heuristic function $h(n)$ is said to be monotone if it satisfies: $$h(n) \leq c(n, n') + h(n') \forall n, n' \| n' \in SCS(n)$$ where $SCS(n)$ gives the successors of n.
Monotonicity (of inference)	Monotonicity says that the set of entailed sentences can only increase as information is added to the knowledge base. CITATION(aitextbook) (p 213)
Mutation	In genetic algorithms after a crossover between two states has been completed we will sometimes randomly change the state by randomly changing one element of the state. This is called a mutation.
Natural Language Processing	The ability of the computer to process English or some other natural language to communicate successfully. CITATION(aitextbook)
Node consistency	A term meaning that a CSP is 1-consistent, or that any node by itself is consistent. CITATION(aitextbook) (p 147)
Node expansion (in search)	After we examine a state (node) x and determine that it is not a goal state, we determine the successor states (nodes) using the successor function. These set of states (nodes) are the expansion of x (p 69).
NP-completeness	The theory of NP-completeness, provides a method to reconize an intractable problem. Any problem class to which the class of NP-complete problems can be reduced is likely to be intractable. Cook and Karp showed the existence of large classes of canonical combinatorial search and resoning problems that are NP-complete. CITATION(aitextbook)
Objective function	Objective functions are part of the description of an optimization problem that measures states to determine the best state. (from Glossary 6)
Objective function (in optimization problem)	Objective functions are part of the description of an optimization problem that measures states to determine the best state.
Occam's (or Ockham's) razor	is a principle attributed to the 14th-century English logician and Franciscan friar, William of Ockham. It forms the basis of methodological reductionism, also called the principle of parsimony or law of economy. In its simplest form, Occam's Razor states that one should make no more assumptions than needed. Put into everyday language, it says: Given two equally predictive theories, choose the simpler. CITATION(Wikipedia:OccamsRazor)
Omniscience	Literally "all knowing". In the AI sense it means knowing all the information necessary to make the best choice that optimizes the performance measure.
Open-loop	When given a **static**, **observable**, **discrete**, **deterministic** environment, solutions can be executed without pay attention to the percepts. Thus an agent carries out its plan with its eyes closed. Control theorists call this an **open-loop** because it breaks the loop between agent and environment.
Operator	Successor functions gives a description of the possible actions available to the agent. An alternate formulation uses a set of operators that can be applied to a state to generate successors. (p 62)
Optimal solution	According to AIAMA: The optimal solution has the lowest path cost among all solutions.
Optimally efficient	If an algorithm is optimally efficient it means that no other optimal algorithm is guaranteed to expand fewer nodes except possibly through tie-breaking among nodes with $f(n) = C^*$. *We note that A^* is optimally efficient* **and** *that it is also exponential in practice.*
Optimization Problem	Optimization problem define as a goal to find the best state according to an objective function.
Path	A *path* in the state space is a sequence of states connected by a sequence of actions.
Path consistency	Any pair of adjacent variables can always be extended to a third neighboring variable - also called $3-$consistency. CITATION(aitextbook) (p 147)
Path Cost	A *path cost* function assigns a numeric cost to each path from the root node to any node in the tree.
Pathmax Equation CITATION(astar)	The PathMax function is an optimization routine that ensures monotonicity: $$\hat{f}(m) = \max\{f(n), f(m)\}$$ for all $m \in \{succ(n)\}$.

Payoff function	Also called the **objective function** or **utility function**, it gives the numeric value for the terminal states. (Objective and utility functions generally give a value to each state). CITATION(aitextbook) (p 163)
Percept Sequence	An agent's percept sequence is the complete history of everything the agent has ever perceived.
Performance Measure	A performance measure embodies the criterion for success of an agent's behavior.
Planning	Planning is the process of finding a sequence that achieves a desired effect given a suitable constructive inference algorithm (p 330).
Plateau	An area in the state space landscape where the evaluation function is flat. I could be a flat local maximum or a shoulder.
Ply	This is essentially a turn in a turn taking game. One move is made up of two different players in a two player game each taking a turn. Each turn is called a half move or ply. CITATION(aitextbook) (p 163)
Predecessors	Predecessors of a state x are all those states that have x as a successor.
Predicate symbol	Stand for relations in FOL. CITATION(aitextbook) (p 246)
Premise	The premise or antecedent is the left hand side of an implication ($\Rightarrow$). CITATION(aitextbook) (p 205)
Problem Formulation	Problem formulation is the process of deciding what actions and states to consider, given a goal. A problem can be defined formally by four components: 1. initial state 2. successor function 3. goal test 4. path cost
Proof	A sequence of applications of inference rules is called a proof. CITATION(aitextbook) (p 212)
Property	Unary relations CITATION(aitextbook) (p 242)
Proposition symbol	A symbol that can be assigned true or false which represents a proposition (or statement). CITATION(aitextbook) (p 204)
Propositional logic	Propositional logic is a very simple language consisting of proposition symbols and logical connectives. It can handle propositions that are known true, known false, or completely unknown. CITATION(aitextbook) (p 233) Propositional logic is made up of a set of symbols and boolean connectives $\wedge, \vee, \neg, \Rightarrow, \Leftrightarrow$ formed into sentences using the formal grammar form: Backus-Naur Form. The semantics are defined by the operation of the connectives as either True or False. CITATION(ainotes) (11.3, 11.7)
Pruning	When a node n' is explored and placed in the fringe as part of a node expansion of n and $h(n') > h(n)$ then A^* will not never expand this node, because $f(n') > C^*$. When this happens we say that n' is *pruned*.
Pruning (in general)	Pruning is the process of eliminating possibilities from consideration without having to examine them. In a tree structure we often say that we prune a subtree or branch by eliminating it from consideration. CITATION(aitextbook) (p 100)
Rationality	Rationality (restricted rationality) is the ideal concept of intelligence (restrcited rationality is listed above) It simply means to do the right thingCITATION(aitextbook)
Reduction ad absurdum	Proving β from α by checking the unsatisfiability of $(\alpha \wedge \neg\beta)$ corresponds exactly to the standard mathematical proof technique of *reductio ad aburdum*. This is also called refutation: $\alpha \models \beta$ if and only if the sentence $(\alpha \wedge \neg\beta)$ is unsatisifiable. CITATION(aitextbook) (p 211)
Refutation completeness	Any complete search algorithm, applying only the resolution rule, can derive any conclusion entailed by any knowledge base in propositional logic. There is a caveat: resolution is complete in a specialized sense. Given that A is true, we can not use the resolution to automatically generate the consequence $A \vee B$ is true. This is called **refutation completeness**, meaning that resolution can always be used to either confirm or refute a sentence, but it cannot be used to enumerate true sentences. CITATION(aitextbook) (p 214)
Relation	A relation between A and B is defined as $R_{A,B} \subseteq A \times B$ or $R_{A,B} \subseteq \{(x,y) \| x \in A, y \in B\}$ CITATION(ainotes)
Relation (mathematical definition)	A relation between A and B is defined as $R_{A,B} \subseteq A \times B$ or $R_{A,B} \subseteq \{(x,y) \| x \in A, y \in B\}$ CITATION(ainotes)

Relaxed problem	A problem with fewer restrictions on the actions is called a *relaxed problem.* The cost of an optimal solution to a relaxed problem is an admissible heuristic for the original problem.
Resolution	If l_i and m_j are complementary literals, then we have resolution defined as: $$\frac{l_1 \vee ... \vee l_k, \quad m_1 \vee ... \vee m_n}{l_1 \vee ... \vee l_{i-1} \vee l_{i+1} \vee ... \vee l_k \vee m_1 \vee ... \vee m_{j-1} \vee m_{j+1} \vee ... \vee m_n} \quad (2)$$ CITATION(aitextbook) (p 214)
Resolution closure	Given a set of clauses S, resolution closure is the set of all derivable clauses. CITATION(aitextbook) (p 217)
Route-finding problem	The Route-finding problem is defined in terms of locations and transitions along links between them.
Satisfiability	A sentence is satisfiable if it is true in some model. CITATION(aitextbook) (p 211)
Search	The search process consists of finding a path from the initial state to a goal state (p 60). The essence of a search is to check a state (node) to determine if it satisfies the goal state and expanding it to find other options to examine later if it does not. (p 69)
Search cost (vs. total cost)	Search cost typically depends on the time complexity , but can also include a term for memory usage. Total cost combines the search cost and the path cost of the solution found (p 72).
Search node	The search node corresponds to the node representing the initial state. (p69)
Search Strategy	The *search strategy* determines which which state to expand next.
Search Tree	A tree structure generated from the initial state and children found using the successor function defines the search tree.
Semantics (of a logic)	The semantics of the language defines the truth of each sentence with respect to each possible world. CITATION(aitextbook) (p 207)
Sensorless problem	(Also called conformant problems) If an agent has no sensors at all, then it could be in one of several possible initial states, and each action might lead to one of several possible successor states. To know that a goal state has been reached, all states in the successor set must be goal states.
Sentence	A sentence is a representation of a fact in the world in a *formal language.* CITATION(ainotes) (10.3)
Sideway move	When reaching a shoulder, we can not find a state to move to that is better than the state we are in. To fix this we allow moving to states that are equal in hopes of finding a way off a shoulder. This is called a sideways move.
Simulated annealing	This combines hill climbing with random walk in such a way that the amount of randomness slowly decreases (cools) over time.
Size of a CSP	The size of a CSP is characterized by the number of constraints. CITATION(ainotes)
Solution quality	Solution quality is measured by the path cost function.
Soundness (of inference)	An inference algorithm that derives only entailed sentences is called sound. CITATION(aitextbook) (p 203)
State space	Together, the initial state and successor function implicitly define the state space which is the set of all states reachable from the initial state.
Step Cost	The step cost is the cost of taking an action a to go from state x to state y is denoted by $c(x, a, y)$.
Straight-line distance	The shortest distance between two points. This is an especially interesting heuristic, because it will always be optimistic no mater how we define points and measures we use for distance in our problem since the shortest distance between two points is a straight line.
Strong $k-$consistency	A graph is strongly $k-$consistent if it is $i-$consistent for $i \in 1, .., k$. CITATION(aitextbook) (p 147)
Substitution	When posing a query $\exists x\ Predicate(x)$, answering yes is not very helpful. The standard form of an answer to such a query includes a set of values that satisfy the query. This set is called a substitution or binding list. CITATION(aitextbook) (p 254)
Successor Function (in search)	The successor function is a function $S : x \in s_i \rightarrow S(s_i) : \{< A_j, S^j_{i+1} >^*\}$ where s_i is a state, A_j are actions and S^j_{i+1} is a state. [notes]
Syllogism	A Syllogism is an inference in which one proposition (conclusion) follows necessarily from two other premises. It is irrefutable reasoning processes (or patterns) that consequently always yield correct conclusions. CITATION(aitextbook)

Syntax (of a logic)	Defines the sentences in a language (grammar) CITATION(ainotes) (10.16)
Task Environment*	PEAS: Performance, Environment, Actuators, Sensors.
Tautology	A sentence that is valid. CITATION(aitextbook) (p 210)
Term	A term is a logical expression that referes to an object. CITATION(aitextbook) (p 248)
Terminal states	States where the game has ended are called terminal states. CITATION(aitextbook) (p 163)
Terminal test	A terminal test determines when the game is over. CITATION(aitextbook) (p 163)
Ternary constraint	A constraint involving 3 variables. CITATION(ainotes)
Theorem	Theorems are sentences entailed by axioms. CITATION(aitextbook) (p 255)
Total Turing test	A human judge engages in a natural language conversation with two other parties, one a human and the other a machine; if the judge cannot reliably tell which is which, then the machine is said to pass the test. To this "Turing test" we add a video signal and the opportunity to pass physical object "through the hatch" to get a total Turing test. [**?**, **?**]
Touring Problem	Visit every city (vertice) given at least once starting and ending at a specific city. Or more precisely given a graph $G = (V, E)$ and an initial vertex v_0 give a sequence of edges such that each vertex is visited at least once (p 67).
Toy Problems	A toy problem is intended to illustrate or exercise various problem-solving methods. It can be given a concise, exact, description and is easily used by different researchers to compare the performance of algorithms.
Transpositions	In games, repeated states occur frequently because of transpositions which are different permutations of the move sequence that end up in t6he same position (or state). CITATION(aitextbook) (p 170)
Traveling Salesperson Problem (TSP - give a *formal* definition) CITATION(NPCompleteness)	INSTANCE: A finite set $C = c_1, c_2, ..., c_m$ of cities, a distance $d(c_i, c)j) \in Z^+$ for each pair of cities $c_i, c_j \in C$, a bound $B \in Z^+$. QUESTION: Is there a "tour" of all the cities in C having total length no more than B, that is, an ordering $< c_{\pi(1)}, c\|\pi(2), ..., c_{\pi(m)} >$ of C such that $$\left(\sum_{i=1}^{m-1} d(c_{(i)}, c_{(i+1)})\right) + d(c_{\pi(m)}, c_{\pi(1)} = B? \quad (3)$$
Tree decomposition (of a CSP)	The process of decomposing a constraint graph into a tree of connected components. A tree decomposition must satisfy the following three requirements: (1) Every variable in the original problem appears in at least one of the subproblems. (2) If two variables are connected by a constraint in the original problem, they must appear together (along with the constraint) in at least one of the subproblems. (3) If a variable appears in two subproblems in the tree, it must appear in every subproblem along the path connecting the those subproblems. CITATION(aitextbook) (p 154) **Definition 1** *A* tree-decomposition *of a graph G is a tree T whose nodes, called bags, are subsets of $V(G)$ such that:* *1.* $\cup_{x \in V(T)} X = V(G)$; *2.* $\forall \{u, v\} \in E(G)$, $\exists X \in V(T)$ *such that* $u, v \in X$; *and* *3.* $\forall X, Y, Z \in V(T)$, *if* Y *is on the path from* X *to* Z *in* T *then* $X \cap Z \subseteq Y$. *From Dourisboure and Gavoille CITATION(doorisboure03)*
Tree width of a graph (check a book on graph theory)	Let T be a tree-decomposition fo a graph G. The width of T is $$width(T) = \max_{X \in T}\{\|X\| - 1\} \quad (4)$$ Let D be the set of all decompositions of the tree T. Then the tree width is $\min_{T \in D} width(T)$. CITATION(doorisboure03)
Triangle inequality	In AI we say $$h(n) \leq c(n, a, n') + h(n')$$ is the triangle inequality for paths from one node to another in the search tree.
Truth table	A truth table specifies the truth value of a complex sentence for each possible assignment of truth values. CITATION(aitextbook) (p 206)

Turing machine	A Turing machine is a 7-tuple, $(Q, \Sigma, \Gamma, \delta, q_0, q_{accept}, q_{reject})$, where Q is the set of start states, Σ is the input alphabet not containing the blank symbol, Γ is the tape alphabet, where $\_ \in \Gamma$ and $\Sigma \subseteq \Gamma$, $\delta : Q \times \Gamma \to Q \times \Gamma \times \{L, R\}$ is the transition function, $q_0 \in Q$ is the start state, $q_{accept} \in Q$ is the accept state, and $q_{reject} \in Q$ is the reject state, where $q_{accept} \neq q_{reject}$. CITATION(IntroToTheTheoryOfComputation2)
Turing test	A human judge engages in a natural language conversation with two other parties, one a human and the other a machine; if the judge cannot reliably tell which is which, then the machine is said to pass the test. CITATION(Wikipedia:TuringTest)
Unary constraint	The simplest kind of constraint which restricts a single variable. CITATION(aitextbook) (p 140)
Uniform-cost search	This strategy expands the node n with the lowest path cost. Thus this will use a $min-priority-queue$ to determine which nodes to explore and expand next. If all step costs are equal this is identical to the BFS. Let C^* be the cost of the optimal solution, and assume that every action costs at least ϵ. Then the algorithms's worst-case time and space complexity is $O(b^{1+\lfloor C^*/\epsilon \rfloor})$ which can be much greater than b^d.
Uninformed search	These types of searches have no information beyond what is given in the problem definition (this is also called the blind search). All they can do is generate successors and distinguish a goals state from a non-goal state. There are five given in Section 3.4: BFS, Uniform-Cost Search, DFS, Depth-limited Search, Iterative deepening depth-first search, and bidirectional search.
Unit clause	A unit clause is a cloause with just one literal. CITATION(aitextbook) (p 221)
Universal constraint	A statement that the two variable are not involved in a constraint together. CITATION(ainotes)
Utility Function	A utility function maps a state (or sequence of states) onto a real number, which describes the associated degree of happiness. This allows rational decisions to be made when goals are inadequate. (1) When goals conflict and (2) When uncertainty is attached to several goals the utility function can give a way to determine which is goal is most likely achievable.
Validity	A sentence is valid if it is true in all models. CITATION(aitextbook) (p 210)
Validity (of a logical sentence)	A sentence is valid if it is true in all models. CITATION(ainotes)
World Value ordering heuristic	Determines the order in which you choose values to assign to a variable from the domain. CITATION(aitextbook) (p 143)
Variable (of a CSP)	Variables are part of a CSP that can take values defined by their domains. CITATION(aitextbook) (p 137)
Variable ordering heuristic	Determines the order in which you assign variables. CITATION(aitextbook) (p 143)
Zero-sum games	Means environments in which the utility values at the end of the game are always equal and opposite. CITATION(aitextbook) (p 161)

All definitions without a citation come from CITATION(aitextbook)