



- Reflex Agent - pick action based on current state
- Planning Agent - use model to simulate actions
- Search Problem
 - State Space - all possible states
 - Successor Function - inputs: state, action; output: cost, successor state
 - start state
 - goal test
 - worst state - all in G about worst
 - search state - only info needed for problem
 - state space graph - each possible state & goal, directed
 - search tree - state can appear > 1 time
- Uniformed Search
 - idea: fringe, expand, goal test
 - completeness? if infinite, find it?
 - optimality? lowest cost?
 - branching factor b = # children of each node, $O(b^d)$ at depth d
 - maximum depth m
 - depth of shallowest solution s
- Depth First Search (DFS)
 - idea: select deepest node to expand
 - implementation: LIFO stack
 - completed? NO
 - optimal? NO
 - time? $O(b^m)$
 - space? $O(bm)$ // minimal edit solution
- Breadth First Search (BFS)
 - idea: shallowest node expand
 - implementation: FIFO queue
 - completed? YES
 - optimal? NO, unless cost equal
 - time? $O(b^d)$
 - space? $O(b^d)$
- Uniform Cost Search (UCS)
 - idea: lowest cost expand
 - implementation: heap PQ, use "bucket sort"
 - completed? YES
 - optimal? YES, if nonnegative costs
 - time? $O(b^{C/\epsilon})$ // C : optimal cost; ϵ : smallest behavior error
 - space? $O(b^{C/\epsilon})$
- Informed Search
 - idea: use info about goal
 - Heuristics
 - inputs: state, output: estimated cost to goal
 - used: heuristic, usually preferable to actual problem
 - ex. Manhattan distance in a maze?
 - Greedy Search
 - idea: expand lowest heuristic value
 - implementation: PQ, like UCS, but extended frontier set
 - completed? NO (highly depend, varies a lot)
 - optimal? NO
 - A^* Search
 - idea: lowest estimated total cost
 - implementation: PQ, UCS + greedy
 - completed? YES (if good heuristic)
 - optimal? YES
 - Admissibility and Consistency
 - g(n) - actual cost to goal; u(n) - cost to goal
 - h(n) - heuristic function cost; g(n) optional
 - f(n) = g(n) + h(n)
 - Admissibility
 - heuristic for optimality
 - $0 \leq h(n) \leq h^*(n)$
 - Graph Search
 - idea: from search but don't expand visited nodes
 - optimization: prevent whole computation
 - issue: make A^* nonoptimal, unclear consistent?
 - Consistency
 - h(A) - h(B) $\leq c_{AB} + h(B)$ cost(A, B)
 - make A^* graph search optimal &
 - Dominance
 - if h2(n) dominates h1(n) $\Rightarrow \forall n: h_2(n) \geq h_1(n)$
 - heuristic more of better heuristic
 - taking care of multiple heuristics usually of advantage
 - goal for making better heuristics

- Constraint Satisfaction Problems (CSP)
 - search function problem: just need nodes, no path
 - Variables - X1, X2, ... Xn
 - Domain - set of values each var can take
 - Constraints
 - NP-hard
 - lower, O(d) domain, $O(d^{nm})$ assignments
 - use partial assignments to make search problem
 - use heuristics to prune search
 - Constraint Types (Constraint Graphs)
 - unary - 1 var, primes domains
 - binary - 2 var, edges
 - higher-order - 3+ var, w/ nodes edges
 - Backtracking Search
 - idea: retrace steps
 - fix order for vars like X1, X2, ...
 - if state value feasible, if not, go back
 - if not, backtrack, pick another
 - generally for solving CSPs
 - Filtering
 - idea: prune unassigned vars domain by removing values which will backtrack
 - Forward Checking
 - without goal \Rightarrow prevent H from constraint
 - Arc Consistency
 - each unassigned edge \geq 2 distinct arcs
 - Algorithm
 - check all arcs in queue Q
 - remove A \Rightarrow B if var in A domain has value A
 - if Z is removed from A, add all X \rightarrow A
 - if Q contains all arcs, done
 - if Q empty, stop: if domain empty, backtrack
 - AC-3
 - heuristic
 - $O(d^3)$
 - E: # arcs / domain sizes
 - d: size of largest domain
 - k-consistency: any set of k nodes \Rightarrow consistent assignment to k-1 with guarantee \exists a value to consistent value
 - strong k-consistency: all k, k-1, k-2, ... consistent
 - Ordering
 - does not affect ordering
 - Minimum Remaining Values (MRV)
 - least valid values / most constrained
 - Least Constraining Values (LCV)
 - how high value
 - value that prunes least values
 - takes extra computation
 - Stubborn
 - first chosen CSP: $O(n^3)$
 - Pick next $Parent(i) \rightarrow X_i$
 - heuristics: not used
 - backtracks: prevent recursive
 - perfect forward assignment
 - cache assignment
 - reassign orig. value, make it leaf
 - $O(d^c(n-c)^2)$
 - c: count of consistent variables
 - Local Search
 - random assignment
 - pick random initial var, reassign n variables
 - Result consistency (inconsistent heuristic)
 - surprisingly good, but complex if subjective
 - almost not used for large n
 - critical value R: $\frac{H(\text{start})}{H(\text{goal})}$ and autohonest
 - Bayes net (Sampling)
 - Prior Sampling - with sample, can be visited
 - Rejection Sampling - stop if not visited
 - likelihood weighting - reassign a sample
 - Goal: e to be at
 - var X_1, \dots, X_n
 - if X_i is evidence var e_i
 - $X_i = e_i$: assign value w_i
 - $w_i = \frac{P(e_i | X_i)}{P(e_i)}$
 - else X_i sample from $P(X_i | \text{Parents}(X_i))$
 - return w_1, \dots, w_n, ω
 - if not prob, use ω as a dummy, normalizing
 - Gibbs Sampling
 - all var random first, let flow around
 - reassign var one constraint var (evidence)
 - repeat for long time

- Adversarial Search Problems (Games)
 - Looking at deterministic two player games
 - Minimax
 - opponent behaves optimally
 - terminal value
 - state value
 - U_{max} , V_{min} in $V^*(s)$
 - V_{min} , V_{max} in $V^*(s)$
 - V terminal, $V(s)$ known
 - predicate: terminal
 - slow: $O(b^m)$
 - Alpha-Beta Pruning
 - α : MAX's best option
 - β : MIN's best option
 - stop if terminal or $\alpha > \beta$, otherwise
 - utility of child / estimate
 - $def: \max V(s, \alpha, \beta)$
 - $V = -\infty$
 - for each successor of node
 - V : min value, value (cost, α, β)
 - if $Z \geq \beta$: return β / prune
 - $\alpha = \max(\alpha, V)$
 - return V
 - $def: \min V(s, \alpha, \beta)$
 - $V = \infty$
 - for each successor of s :
 - V : max value, value (cost, α, β)
 - if $V \leq \alpha$: return α / prune
 - $\beta = \min(\beta, V)$
 - return V
 - Evaluation Function
 - guess utility of state
 - Expectimax
 - choose $V(s) = \sum_{Player} P(s) V(s)$
 - need minimax
 - Mixed Layer Types
 - mix of state
 - Circular Game (not two-person)
 - multi-agent states (type of states)
 - Utilities
 - Principle of Maximum Utility
 - rational preference
 - $A > B // A$ prefer B
 - $A = B //$ don't care
 - history // possibility
 - Actions // Rationality
 - Ordinality $(A > B) \vee (B > A) \vee (A = B)$
 - Transitivity $(A > B) \wedge (B > C) \Rightarrow A > C$
 - Continuity $A > B > C \Rightarrow \exists p \in (A, B) \cup (B, C)$
 - Substitutability $A > B \wedge C \sim D \Rightarrow (A, C) \sim (B, D)$
 - Hedonicity
 - $A \geq B \Rightarrow \exists p \in (A, B) \cup (B, C) \Rightarrow C > B > A$
 - Utility Function U
 - $U(A) \geq U(B) \iff A \geq B$
 - $U(C, p, R, S) \iff U(C) \geq U(S)$
 - risk neutral / averse / seeking
 - Non-Deterministic Search (Markov Decision Processes (MDPs))
 - MDP formalism
 - S : set of states
 - A : set of actions
 - start state
 - terminal state
 - T : success factor
 - q: state/action state - like expectation change rule
 - Y : present value of reward, given state s
 - Markov decision property
 - known past not impact future independent
 - $T(s, a, s') = P(s', a | s)$
 - Bellman Equation (DP)
 - $V^*(s)$: optimal value of s // 0 for terminal state
 - $Q^*(s, a)$: optimal value of q, state
 - $V^*(s) = \max_a \sum_{s'} T(s, a, s') [R(s, a, s') + \gamma V^*(s')]$
 - $Q^*(s, a) = \sum_{s'} T(s, a, s') [R(s, a, s') + \gamma V^*(s')]$
 - $V^*(s) = \max_a Q^*(s, a)$
 - Value Iteration (DP)
 - convergence $V_{t+1}(s) = \max_a \sum_{s'} T(s, a, s') [R(s, a, s') + \gamma V_t(s)]$
 - $V_{t+1}(s) = \max_a \sum_{s'} T(s, a, s') [R(s, a, s') + \gamma V_t(s)]$
 - $V_t(s) = \max_a \sum_{s'} T(s, a, s') [R(s, a, s') + \gamma V_{t-1}(s)]$
 - Policy Iteration
 - $V_t(s) = \max_a \sum_{s'} T(s, a, s') [R(s, a, s') + \gamma V_t(s)]$
 - $V_t(s) = \max_a \sum_{s'} T(s, a, s') [R(s, a, s') + \gamma V_t(s)]$
 - $V_t(s) = \max_a \sum_{s'} T(s, a, s') [R(s, a, s') + \gamma V_t(s)]$
 - $V_t(s) = \max_a \sum_{s'} T(s, a, s') [R(s, a, s') + \gamma V_t(s)]$
 - $V_t(s) = \max_a \sum_{s'} T(s, a, s') [R(s, a, s') + \gamma V_t(s)]$
 - $V_t(s) = \max_a \sum_{s'} T(s, a, s') [R(s, a, s') + \gamma V_t(s)]$
 - $V_t(s) = \max_a \sum_{s'} T(s, a, s') [R(s, a, s') + \gamma V_t(s)]$
 - $V_t(s) = \max_a \sum_{s'} T(s, a, s') [R(s, a, s') + \gamma V_t(s)]$
 - $V_t(s) = \max_a \sum_{s'} T(s, a, s') [R(s, a, s') + \gamma V_t(s)]$
 - $V_t(s) = \max_a \sum_{s'} T(s, a, s') [R(s, a, s') + \gamma V_t(s)]$
 - Policy Iteration
 - any policy π // value $V_t(s)$
 - stop when $V^*(s) = \sum_{s'} T(s, a, s') [R(s, a, s') + \gamma V^*(s)]$
 - $\pi_{t+1}(s) = \arg \max_a \sum_{s'} T(s, a, s') [R(s, a, s') + \gamma V_t(s)]$
 - $\pi_{t+1} = \pi_t \Rightarrow$ convergence

- Reinforcement Learning (RL)
 - Online learning / exploration
 - Model Based Learning
 - $Q(s, a, s')$ heuristic back by sample
 - $Q(s, a, s')$ learn about state
 - Law of Large Numbers, make pretty well
 - use policy iteration often
 - Model Free Learning
 - Direct Evaluation
 - fixed policy π
 - eventually learn $V^*(s)$
 - Temporal Difference (TD)
 - Sample: $R(s, a, s') + \gamma V^*(s')$
 - $V(s) = (1 - \alpha) V(s) + \alpha \cdot \text{sample}$
 - α : learning rate, decrease over time
 - Q-learning
 - q-value iteration ideas like value iteration
 - Sample: $R(s, a, s') + \gamma Q(s', a)$
 - $Q(s, a) = (1 - \alpha) Q(s, a) + \alpha \cdot \text{sample}$
 - α : learning rate, decrease over time
 - Approximate Q-learning
 - S bits to use feature vector
 - linear value function
 - difference: $R(s, a, s') + \gamma \sum_{a'} Q(s, a')$
 - $W = w_1 \cdot w + \text{difference}$, $\alpha \cdot Q(s, a)$
 - $\alpha Q(s, a) = Q(s, a) - \alpha(Q(s, a) - R(s, a, s'))$
 - Gradient and Expectation
 - E: Greedy Blunt
 - Policy Gradient
 - backprop through loss, need to have non-linear
 - Expectation Function
 - don't need to know Q
 - $Q(s, a) = (1 - \alpha) Q(s, a) + \alpha [R(s, a, s') + \gamma \sum_{a'} Q(s, a')]$
 - $F(s, a)$: expectation function, prefer to know $Q(s, a)$ first, values are variable
 - $\frac{d}{dt} F(s, a) = \frac{d}{dt} Q(s, a) + \frac{d}{dt} [R(s, a, s') - Q(s, a)]$
 - Probabilistic Inference
 - J: set distribution - all prob. allow
 - Inference by enumeration (CSP)
 - $P(Q_1, \dots, Q_n | e_1, \dots, e_n)$
 - Query Var Q_i : left
 - Evidence Var e_i : right, known
 - Hidden var: no data
 - collect 1 example
 - Bayes Net (Approximate)
 - report, use local PDF, DAG
 - each node is conditionally independent
 - all ancestor nodes, given parent
 - $P(A|B) = \frac{P(B|A)P(A)}{P(B)}$
 - also $P(X | \text{Parents}(X))$
 - locally DAG a conditional probability table (CPT)
 - Bayes Net (Inference)
 - inference: call joint PDF for q, e
 - eliminate var X
 - join parents of factor of X
 - sum out X
 - Bayes Net (O-Approximation)
 - Control Chain
 - $X \rightarrow Y \rightarrow Z$
 - $X \perp\!\!\!\perp Y, Z$ if Y known
 - $X \perp\!\!\!\perp Z$ if Y unknown
 - Common Cause
 - $X \rightarrow Y, Z$
 - $X \perp\!\!\!\perp Y, Z$
 - $X \perp\!\!\!\perp Y, Z$ if Y known
 - Common Effect
 - $X, Y \rightarrow Z$
 - $X \perp\!\!\!\perp Y$ if Z known
 - O : directed acyclic graph
 - $X \perp\!\!\!\perp Y, Z$ if Y known
 - $X \perp\!\!\!\perp Y, Z$ if Y unknown



- Markov Models

- chain = memoryless property
- $P(w_1, \dots, w_n) = P(w_1) \prod_{i=2}^n P(w_i | w_{i-1})$
- assume stationary = same
- Mini-Forward Algorithm
- $P(w_{i+1}) = \sum_{w_i} P(w_{i+1} | w_i) P(w_i)$
- Stationary Distribution
- $P(t, t) = P(t)$

- Hidden Markov Model

- collect evidence that affects beliefs
- $\begin{matrix} \text{state} \\ \text{evidence} \end{matrix}$
- assume $P(F_i | w_i)$ stationary
- only know evidence
- Forward Algorithm
- $B(w_i) = P(w_i | F_1, \dots, F_i)$
- $B'(w_i) = P(w_i | F_1, \dots, F_{i-1})$
- $B'(w_{i+1}) = \sum_{w_i} P(w_{i+1} | w_i) B'(w_i)$
- $B(w_{i+1}) = P(F_{i+1} | w_{i+1}) B'(w_{i+1})$
- normalize to get new beliefs
- $B(w_{i+1}) \propto P(F_{i+1} | w_{i+1}) \sum_{w_i} P(w_{i+1} | w_i) B(w_i)$

- Viterbi Algorithm

- most likely sequence of states given evidence
- forward and backward pass

- Particle Filtering

- n particles, d states, n < d
- Time Step Update
- sample new particles from dist
- Observation Update

 1. weight each particle w/ $P(F_i | T_i)$
 2. get total weight each state
 3. if sum all weights all states = 0, resample all
 4. else resample based on weight dist

- Decision Networks

- chance, action, utility nodes
- Maximum Expected Utility (MEU)
- $EU(a|e) = \sum_{s, a} P(s, a|e) U(s, a)$
- argmax over a , basically expected B vs U
- Value of Perfect Information (VPI) - (non-evidential)
- $MEU(e) = \max_a \sum_s P(s|e) U(s, a)$
- $MEU(e, e') = \max_a \sum_s P(s|e, e') U(s, a)$
- $MEU(e, e') = \sum_s P(e'|e) MEU(e, e')$
- $VPI(e'|e) = MEU(e, e') - MEU(e)$
- Properties
- nonnegativity, $VPI \geq 0$
- non-additivity, in general
- order-independent

- Machine Learning

- default: training, validation, test
- Naive Bayes - classification independent
- Features = label, action feature, label
- Parameter Estimation
- Maximum Likelihood Estimation (MLE)
- assume identically distributed, independent
- assume all parameters unless equal probs, 0
- $L(\theta) = \prod_{i=1}^n P(y_i | x_i)$ return $\frac{\partial}{\partial \theta} L(\theta) = 0$
- Laplace Smoothing = k
- avoid overfitting
- assume k extra for each possibility
- $k \rightarrow \infty$ == all equal probs

- Perceptron

- Linear Classifier
- activation = linear comb of features (dot product)
- Binary Perceptron
- 2 classes, hyperplane
- all weights = 0
- if activation $> w^T f > 0$, else $-1 = y$
- if $y \neq y^*$, $w = w + y^* f(x)$
- if w update, exit
- Bias
- shift hyperplane from origin
- just add bias if values 1 and 0 marks
- Multiclass Perceptron
- if wrong, add to correct weight subtract from wrong weights

- Neural Networks

- Multistage
- nonlinear responses
- Idea - multistage perception
- universal function approximator
- Multilayer Feedforward Neural Networks
- don't apply diff nonlinearity
- perception - rectify by $\frac{1}{1+e^{-x}}$
- convex Belief
- Sigmoid $\frac{1}{1+e^{-x}}$

- Gradient Ascent/Descent

- $\theta_{w, l+1} = \left[\frac{\partial \ell(\theta)}{\partial w_l} \dots \right]$
- $w = w + \alpha \cdot \nabla \ell(w)$
- learning rate
- batch gradient descent - use all points (slow)
- mini-batching - takes batches
- stochastic gradient descent (SGD) - $k=1$

- Backpropagation

- back calc gradients
- forward pass - normal calc
- backward pass - get gradients