

Probability Basics

Uncertainty may come from partial observability, noisy sensors, complexity of modeling and lack of ^{real} world dynamics
 Probability summarizes the effects of laziness (don't list all) and ignorance
 Bayesian probability relates probability to one's own state of knowledge

The set Ω ("sample space") defines worlds
 Events are subsets of Ω , probability is the sum over all possible worlds: $\sum_{w \in A} P(w)$
 $0 \leq P(w) \leq 1$ $\sum_{w \in \Omega} P(w) = 1$

Given joint distributions, one can marginalize "sum out" over variables partial assignment
 $P(a|b) = \frac{P(a,b)}{P(b)}$ conditional probabilities allow us to use evidence

Inference by Enumeration

Evidence variables: $E_1, \dots, E_n = e_1, \dots, e_n$
 Query var: Q Hidden var: H_1, \dots, H_r
 Want: $P(Q|e_1, \dots, e_n)$
 1) Select entries consistent with evidence
 2) Sum out H to get joint Q and E
 3) Normalize $\rightarrow P(Q|e_1, \dots, e_n) = \frac{1}{Z} P(Q, e_1, \dots, e_n)$
 $Z = \sum_Q P(Q, e_1, \dots, e_n)$
 Worst case time/Space: $O(d^n)$

Product Rule: $P(X)P(X|Y) = P(X, Y)$
 Chain Rule: $P(x_1, x_2, \dots, x_n) = \prod P(x_i | x_1, \dots, x_{i-1})$
 Bayes' Rule: $P(X|Y) = \frac{P(Y|X)P(X)}{P(Y)}$

Bayes Nets and Bayesian Inference

Models describe how world works, accord d^n
 blowup of joint distributions by independence

Independence
 Absolute if $\forall x, y P(x, y) = P(x)P(y)$
 Conditional if $\forall x, y, z P(x, y | z) = P(x|z)P(y|z)$

Conditional independence represented by Bayes Nets
 Bayesian Networks

Helps factor joint distribution into product of conditional
 Reduces table size from $O(d^n)$ to $O(n \cdot d^k)$
 Syntax: set of variable nodes X_i , DAG, conditional probability table for each node given parents

Encode joint distribution as $P(x_1, \dots, x_n) = \prod P(x_i | \text{Parents}(x_i))$
 Variables independent of non-descendants given parents
 Markov blanket includes parents, children and children's parents and all variables independent of others given Markov blanket

Inference by Enumeration
 "Summing out" hidden variables from joint table \rightarrow exponentially many products

Variable elimination
 Move summation expressions inwards as possible and perform calculation from inside out using factors
 1) Compute the pointwise product of factors
 2) Sum out variables from the factor

Pick a query variable Y of evidence, start Y initial factors, while H left (push H , join, sum) and join + normalize
 Time and space complexity determined by largest factor
 Ordering can affect size of largest factor, not always

Polytrees - DAG with no undirected cycles
 Variable elimination linear $O(n)$ if start VE from leaves

Human Utilities, Money and the Optimal's Curse

Given $EV(U) = px + (1-p)y$, $U(L) = pU(x) + (1-p)U(y)$
 Certainty equivalent $CE(L) =$ cash amount such that $U(CE) \sim U(L)$
 Insurance premium is $EV(U) - CE(L)$
 Given a set of estimates, we choose one value and with high probability the actual value is much less than V^*
 Assess utility by lottery "best possible price" vs "worst catastrophe" to price A and adjust $p, 1-p$

Bayes Nets and Approximate Inference

Sampling allows us to get approximate answer fast with simple model, little memory
 Prior Sampling
 Generate samples from root nodes and relevant

Rejection Sampling
 Given a query $P(a|e, \dots)$, reject sample if not consistent otherwise no different from prior sampling

Likelihood weighting
 Fixes problem of unlikely evidence and rejection sampling
 Instead, fix evidence, sample rest and weight by $P(e|Parents(e))$
 $S_w(z_1, \dots, z_n) = \prod_{i=1}^n P(z_i | \text{Parents}(z_i)) \prod_{j=1}^m P(e_j | \text{Parents}(e_j))$
 Upstream vals not influenced by downstream evidence

Gibbs Sampling
 Markov chain Monte Carlo, random walk w/ sampling and averaging
 Fix evidence variables, pick variable each time and sample for a value conditioned on all others
 Eventually you will sample from true posterior

Markov Models

Introduce retention of time by connecting state nodes
 Transition model $P(X_t | X_{t-1})$ specifies evolution
 Stationarity assumption - transition P does not change
 Markov assumption: X_t indep. of X_0, \dots, X_{t-2} given X_{t-1} is 1st order, k th order have k previous

Joint distribution $\rightarrow P(x_0, \dots, x_T) = P(x_0) \prod_{t=1}^T P(x_t | x_{t-1})$
 Forward Algorithm
 At time t , $P(x_t) = \sum_{x_{t-1}} P(x_t | x_{t-1}) P(x_{t-1})$
 Assuming stationarity, can also multiply by trails^k from $t-1$ to t
 Stationary distribution can be reached, $P_{\infty} = P_{\infty+1} = P_{\infty}$

Hidden Markov Models $\rightarrow \text{state} \rightarrow \text{obs} \rightarrow \text{state} \rightarrow \text{obs} \rightarrow \dots$
 Markov chain over states X , observe E at each time-step
 Define by initial dist $P(x_0)$, t s model $P(x_t | x_{t-1})$, sensor model $P(e_t | x_t)$
 Joint distribution given by $P(x_0, \dots, x_T, E) = P(x_0) \prod_{t=1}^T P(x_t | x_{t-1}) \prod_{t=1}^T P(e_t | x_t)$

Inference Tasks
 Filtering: $P(x_t | e_{1:t})$
 Start with uniform π_0 and use a predict/update cycle to keep track of current belief state

Filtering algorithm
 We want a recursive alg of form $P(x_{t+1} | e_{1:t+1}) = g(e_{t+1}, P(x_t | e_{1:t}))$
 $P(x_{t+1} | e_{1:t+1}) = \alpha P(e_{t+1} | x_{t+1}) \sum_{x_t} P(x_t | e_{1:t}) P(x_{t+1} | x_t)$
 Lost per time step: $O(|X|^2)$, time and space $O(1)$

Linear algebra: $F_{t+1} = A O_{t+1} T F_{t+1}$, O has E_t of obs and diagonal
 Most Likely Explanation: $\arg \max_{x_{1:t}} P(x_{1:t} | e_{1:t})$
 State trellis: states + transitions
 Each arc is $x_{t-1} \rightarrow x_t$ with weight $P(x_t | x_{t-1}) P(e_t | x_t)$
 Viterbi algorithm

Find max probability of any path to a state t
 $m_{t+1} = \text{Viterbi}(m_{1:t+1}) = P(x_{t+1} | x_t) \max_{x_t} P(x_{t+1} | x_t) m_{t+1}$
 Time: $O(|X|^2 T)$, Space: $O(|X|T)$, # paths: $O(|X|T^2)$
 Dynamic Bayes Nets

HMM with multiple vars
 Sparse dependencies exp fewer parameters
 Exact Inference
 Offline: Variable elimination $P(x_{1:t} | e_{1:t})$, unroll network for T steps
 Online: Eliminate variable from previous time, store factors for current but largest factor contains all variables for current time

Particle Filtering
 Likelihood weighting needs exponentially many samples w/ T , so initialize particles and resample the fittest
 Start w/ particles, $N \ll |X|$, $P(x)$ approximated by # particles
 Propagate forward: $x_{t+1} \sim P(x_{t+1} | x_t)$, sample from prev t , m
 Resample: $w_t P(x_t | x_t)$, weight samples based on evidence, re-sample

Decision Theory

A rational agent maximizes expected utility given knowledge
 worst-case minimax does not care about exact terminal value scale, optimal decision invariant under monotone trans.
 Magnitudes are important for average-case expected max
 Utilities are functions from real outcomes to real numbers

Axioms of Rationality
 Orderability: $(A \succ B) \vee (B \succ A) \vee (A \sim B)$
 Transitivity: $(A \succ B) \wedge (B \succ C) \Rightarrow (A \succ C)$
 Continuity: $(A \succ B) \Rightarrow \exists p \in [0, 1] pA + (1-p)C \sim B$
 Substitutability: $A \sim B \Rightarrow [pA + (1-p)C] \sim [pB + (1-p)C]$
 Monotonicity: $A \succ B \Rightarrow [pA + (1-p)C] \succ [qA + (1-q)C]$

Decision Networks

Bayes Net + \square action nodes + \diamond utility nodes
 Fix evidence E , for each action a : fix action node to a , compute posterior $P(W|E, a)$ for $W \in U$, compute $\sum_{w \in U} P(w|E, a) U(w)$, return highest utility

Value of Information
 Expected improvement in decision quality from observation of var
 $VI(E; I) = \sum_e P(e|I) \arg \max_a EU(a|e, I) - \arg \max_a EU(a|I)$
 Non-negative: $VI(E; I) \geq 0$, not usually additive: $VI(E_1; E_2, I) \neq VI(E_1; I) + VI(E_2; I)$
 Order-independent: $VI(E_1, E_2; I) = VI(E_2, E_1; I)$

Markov Decision Process

Definition: set of states $s \in S$, set of actions $a \in A$
 transition model $T(s, a, s')$, reward func $R(s, a, s')$ or just $R(s)$
 (s, s') start state, terminal state with zero rewards
 Fully observable but probabilistic search problem
 Policy π^* : $S \rightarrow A$ maximizes expected utility for all states
 Utility $U(s, \pi) = r_0 + \gamma r_1 + \gamma^2 r_2 + \dots$ where reward bounded by $R_{max}/(1-\gamma)$
 Allow infinite streams or if terminal state reached stop (geom. series)
 $\pi^*(s)$ - optimal action, states $V^*(s) = V^{\pi^*}(s)$ - value of state (optimality from s)
 $Q^*(s, a)$ - expected util of (s, a) optimal where $V^*(s) = \max_a Q^*(s, a)$
 Bellman Update Equations: $V_{t+1}(s) \leftarrow \max_a \sum_{s'} P(s'|s, a) [R(s, a, s') + \gamma V_t(s')]$
 Repeat until converge or ϵ , B is a contraction by γ \rightarrow exp fast conv.
 Policy extractor: $\pi^*(s) = \arg \max_a \sum_{s'} P(s'|s, a) [R(s, a, s') + \gamma V^*(s')]$
 Policy iteration: $\pi_{t+1}(s) = \arg \max_a \sum_{s'} P(s'|s, a) [R(s, a, s') + \gamma V_t^{\pi_t}(s')]$

Reinforcement Learning
 Maximize expected rewards based on observed sample transitions
 Model-based Learning
 Learn empirical MDP by counting ϕ for s, a , number $P(s'|s, a)$, $R(s, a, s')$
 2) solve learned MDP by value/policy iteration
 Multi good use of experience but not scalable to large S
 Direct-Utility Estimation
 Use sum of discounted rewards to S and average over trials
 No PR required, optimal, but requires state connections \rightarrow 10^6 states

Temporal Difference Learning
 TD rule: $V^{\pi}(s) \leftarrow V^{\pi}(s) + \alpha [R(s, \pi(s), s') + \gamma V^{\pi}(s') - V^{\pi}(s)]$
 Approximate Bellman update, can't use value/imp policy + trans. model

Q-Learning
 Q-learning gives us both optimal policy/value
 $Q(s, a) \leftarrow (1-\alpha) Q(s, a) + \alpha [R(s, a, s') + \gamma \max_{a'} Q(s', a')]$
 Converges as long as we explore enough and decay α appropriately: $\sum_t \alpha_t = \infty$ and $\sum_t \alpha_t^2 < \infty$ like $\alpha_t = O(1/t)$

Exploration vs. Exploitation
 Pure exploitation often gets stuck in rut \rightarrow optimal policy
 ϵ -greedy: flip coin to try random action, make sure to learn ϵ
 Exploration function: e.g. $f(x, t) = w + \text{noise}$ encourage exploration at 1st trial
 $Q \leftarrow Q(s, a) + \alpha [R(s, a, s') + \gamma \max_{a'} Q(s', a') - Q(s, a)]$
 Regret measures total cost of youthful errors, minimizing also encourages optimal learning being optimal

Approximate Q-Learning
 Describe state using vector of features for generalization
 $V(s) = w_0 + w_1 \phi_1(s) + \dots + w_n \phi_n(s)$, $Q_w(s, a) = F_w(s, a) + w_n \phi_n(s, a)$
 update: $w \leftarrow w + \alpha [R(s, a, s') + \gamma \max_{a'} Q_w(s', a') - Q_w(s, a)] \phi(s, a)$
 Linear feature approx \rightarrow TD V converge, Q diverge but non-linear \rightarrow will better

using neural network instead of linear function
 Machine Learning
 Supervised Learning (correct answers)
 Learn unknown target function f given training set of labeled (x_i, y_i)
 and outputs h (classification regression \rightarrow discrete, real-valued)
 Tradeoff between degree of fit and complexity

Decision Trees
 Partition input space, assign label to each group
 Increasing hypothesis language expressiveness \rightarrow more complex, hyp, bad at N
 Choose attributes by f gain/entropy for info about classification
 $B(p) = -H(p) = -\sum_i p_i \log p_i$: # bits to classify new example
 Information gain: $B(p)/p(n) = \sum_i p_i \log(p_i/p(n))$
 Generate hyp on training set, test by validation, accuracy on test set
 Accurate: fraction correct, Learning curve: accuracy vs. training size

Neural Networks
 Input: $\sum_i w_i x_i + a$ (bias) output: $a_i = g(\text{in}_i) = g(\sum_j w_{ij} x_j + a_i)$
 Activation functions g : threshold $(0, 1)$, sigmoid $(0, 1)$
 Threshold perception as linear separator, $g(x) = 1$ if $x > 0$
 Perceptron learning rule: $w \leftarrow w + \alpha (y - h(x)) x$
 Linearly separable if \exists hyperplane, converge to min-error if \exists well
 Multilayer perceptron - feedforward neural network ≥ 2 hidden, sigmoid
 Loss: $\sum_i (y_i - h(x_i))^2$, $w \leftarrow w - \alpha \text{Loss}$ by backpropagation output hidden
 Using enough hidden nodes can represent any function

Statistical Learning
 Choices: hypothesis space H , degree of fit, degree of model complexity,
 finding good h by knowing if h will predict well
 Classical stat/ML uses linear feature comb, loss function, regularization
 and complexity penalties, optimization (discrete search, besting)

Maximum Likelihood Estimation
 $P(x_1, \dots, x_n) = \prod_{i=1}^n P(x_i | x_{1:i-1})$, $\arg \max_{\theta} P(\theta | \{x_i\}_{i=1}^n)$
 optimization/discrete search, empirical process theory - iid
 Evidence: $x = x_1, \dots, x_n$ Likelihood $P(x_1, \dots, x_n | \theta)$, Max likelihood
 by solving for $\theta = \arg \max_{\theta} P(\theta | x)$
 Laplace smoothing with α : $G_{\theta} = (n + \alpha) \theta$
 Naive Bayes
 polytree for $P(x|X) = \prod_{i=1}^n P(x_i | x_{1:i-1})$, $P(x_i | x_{1:i-1}) = \prod_{j \in \text{pa}(x_i)} P(x_i | x_j)$
 Linear separator
 Estimate $P(x)$ by ML, $P(x|C)$ by some Laplace smoothing

Agents and Rationality

Agent acts, operates autonomously, perceives environment, ^{perceives} adapts
 Rational agent acts to increase (expected) outcomes
 Agents perceive environment by sensors and acts by actuators
 Agent function is abstract mathematical concept, requires implement.
 Performance measure determines rational behavior, should design with desired environment in mind
 Agent should learn, autonomy measures, reliance on ^{designer} prior knowledge

Environments

Task environment: performance measure, environment, actuators, sensors
 Fully/partially observable - sensor give full/partial knowledge of env
 Single/multi-agent - other agents may be adversarial \rightarrow against objects
 Deterministic/stochastic - next state probabilities may (not) be important, distinguish between part obs, deterministic vs stochastic
 Episodic/sequential - current decisions affect later decisions
 Static/dynamic - env changing while agent deliberates, semi-persistent
 Discrete/continuous - time handling, chess vs. real-time taxi
 Known/unknown - know if outcome/prob known, else unknown

Types of Agents

Simple Reflex Agents
 Base action on current percept (no history). \forall table (percept, action)
 In finite loops, as possible, can randomize actions
 Model-based Agents
 Maintain internal state based on percept history, world state
 Use model, update state from percept/action to "best guess"
 Goal-based Agents
 Make decision based on state as well as goal
 More flexible than reflex-based, goal-oriented decisions
 Utility-based Agents
 Generalization of goal-based agents for optimality
 Rational can make tradeoffs to maximize expected utility

Search Problems

Environment is observable, discrete, known and deterministic
 "Family tree, search, execute" - offline search \forall done in execution state
 Formulation: 1 Initial state 2 Actions 3 Transition model (R(s,a): s' space)
 Can form graph \forall states connected by links = actions
 4 Goal Test (check s: goal state) 5 Path cost function $c(s,a,s')$
 Optimal solution entail lowest path cost, tree-search + expand \rightarrow graph search
 In search node in search tree, remember n: state, n: parent, n: action, n: path cost
 Completeness measures find solution if exists

Uniformed Search Strategies

Breadth-first search (queue as fringe)
 Goal test applied when node generated, finds shallowest node d
 Time/space $O(b^d)$, dominated by function size
 Uniform-cost search (priority queue as fringe)
 Expands node \forall lowest path cost $g(n)$, optimal path when expanded
 Complete if every step γ small, $\forall \epsilon$ Time/space $O(b^{U+\epsilon/\gamma})$
 Depth-first search (stack as fringe)
 Graph search in finite state space complete while tree search not
 Time/space $O(b^m)$, $O(b^m)$ \rightarrow good for memory DFS-like
 Depth Limited Search - fix depth limit, iterative deepening - heap expand
 Iterative lengthening of path costs, Bidirectional Search (BFS/UCS) $\rightarrow b^{d/2} + b^{d/2} \ll b^d$

Informed (Heuristic) Search

Note use of combination of path cost $g(n)$ and estimate cost $h(n)$
 Greedy-Best-First Search
 $f(n) = h(n)$, tree search infinite loop, graph search is complete
 A* Search
 $f(n) = g(n) + h(n)$, complete and optimal if $h(n)$ admissible
 for tree search and consistent for graph search
 Node expanded \Rightarrow optimal path found, consistent \Rightarrow $h(n) \leq \text{true distance}$ $+ h(n)$
 Expands nodes along goal contour
 Iterative deepening using $f(n)$, Reuse best-first behavior in multi-state
 Heuristic Functions

Local Search

Evaluate and modify current state; path cost is not relevant
 Use small amounts of memory, can usually find reasonable sol in large/minib spaces
 Complete - always finds solution, optimal - always finds global min/max
 Hill-Climbing Search
 Steepest-ascent/greedy local search; complete state formulation
 local maxima - stuck, ridges - seq of local maxima, plateau - flat area, shoulder
 Stochastic - random turn uphill moves with weighted probability
 Simulated Annealing
 Retire state by picking random move, and take it better
 otherwise accept \forall probability $\propto T$ which you gradually decrease
 to shake the system out of local min/max
 Local-beam search
 Start \forall k random states and generate all successors, keep top k
 Parallel search threads can pass info around
 Stochastic beam search - choose k successors at random & how "good"
 Genetic Algorithm
 Similar to stochastic beam search but generate successors by combining
 two parent states
 Use objective function as fitness, cross over point and mutation

Agent Program Representations

Atomic "Block box"
 Each state of world indivisible and has no internal structure
 Factored Representation
 Split state into variables and attributes, each has a value
 Can use booleans and continuous to represent uncertainty
 Structured Representation
 Model how parts of the world are connected to each other
 Learning Agents
 Learning element makes improvements \rightarrow part element
 Performance element makes decisions, learning uses critic feedback
 Problem generator suggests to performance element \rightarrow informative exp

Adversarial Search

Search problem: S - Initial state, player(s) - which player has move
 action(s) - set of legal moves, terminal test, Utility (s, p):
 objective/payload function \rightarrow zero-sum - payoff same in all games
 Minimax Search
 Alternate min and max plies, accomplished by DFS to terminal
 Multiplayer games may use vector (color) by player
 Alpha-Beta Pruning
 As we search, keep values of best max, β (best min value)
 and cut off exploration if value worse than β
 May use transposition table of known states to cut down
 cut-off search \forall evaluation to make imperfect walking decision
 Can modify α - β to use eval beyond certain point
 Good for quietest positions; horizon is trickier
 Searching with Nondeterministic Actions
 Percepts help in partially observable or nondeterministic environment
 Contingency planning - by all possible seq of percepts
 Solutions look like trees (if-then-else), rather than sequences
 AND-OR search trees alternate OR (you) and ALL (env) options
 and solution must be able to handle \forall OR and all AND
 some nondeterministic environments may require optimal sol

Searching with Partial Observations

Belief state is set of all possible physical states you may be in
 No observation
 Belief states: U to 2^N unique sets of belief states = belief state space
 Initial state: set of all states in P initially
 Actions: If illegal actions have no effect, then
 $U \rightarrow \text{Actions}(s)$, else $\text{Actions}(s, a) \Rightarrow \text{Safe}$
 Transition Model: deterministic \Rightarrow $\beta = \text{Result}(b, a) = \{s' : s' = \text{Result}(s, a) \text{ and } s \in b\}$
 Goal Test: Check that all belief states satisfy else \cup Results $\{s' : a\}$
 Path cost may be different between states
 With observations
 Prediction: $\hat{b} = \text{Predict}(b, a)$
 Observation prediction: determine set of possible percepts observable
 Update: Partitions and returns set of belief states that
 correspond to intake percept $o \Rightarrow b_o = \text{update}(b, o)$
 Nondeterminism just enlarges the number of belief states

Constraint Satisfaction Problems

Utilize a factored representation of variable assignment to
 simultaneously satisfy all constraints
 X - set of variables $\{X_1, \dots, X_n\}$, D - set of domains, c - each var
 C - set of constraints \Rightarrow allowable combinations of values
 constraint graphs - to visualize systems of binary constraints
 global constraint - arbitrary # of variables
 node consistency - enforce unary, arc enforce binary
 path consistency - tuples $\{X_i, X_j\} \forall X_m$ if $\{X_i, X_m\}, \{X_m, X_j\}$
 Solutions are commutative
 Backtracking Search (DFS one var at a time, assume backtrack if fail)
 MRV - minimum remaining value
 Degree - prioritize items \forall higher degree
 LCV - least constraining value
 Forward Checking - choose variable, establish arc consistency
 from neighbors and continue, backtrack if no domain
 MAC (Maintaining Arc Consistency) - call AC3 instead of assuming
 from neighbors
 Chronological backtracking - backjumping
 min conflicts - pick random variables and min conflicts, repeat
 Cutset Conditioning
 Choose subset C of graph to form a tree
 Keep through cutset and run tree CSP solver $O(d^{2n})$
 which ensures backtrack consistency then runs forward checking

Logical Agents

Knowledge base \equiv set of sentences in knowledge representation
 language that describes the world
 TELL and ASK operations both involve inference to derive
 new sentences from old

Logic

Entailment implies β follows logically from α , all worlds
 where α are true require β to be true
 sound - can derive entailed sentences
 complete - can derive any sentence that is entailed
 Propositional Logic
 $\neg, \wedge, \vee, \Rightarrow, \Leftrightarrow$
 Truth tables allow you to express truthfulness of sentences
 Model-checking, Theorem-Proving
 Valid - true in all models, Satisfiable - true in some model
 Prove β from α by unsatisfiability of $\alpha \wedge \neg \beta$
 (Contradiction)
 Initial state: Knowledge Base
 Actions: All inference rules applied to sentences
 making up top half of inference rule
 Result: Add sentence in bottom half
 Goal State we are trying to prove
 Conjunctive normal form - \wedge 's of \vee 's
 Horn form - at most one positive literal in \vee
 PPL - pure symbol (same sign), unit clause \rightarrow
 early termination