Artificial Intelligence Recap
CSE P573

Mausam

What is intelligence?

- (bounded) Rationality
  - We have a performance measure to optimize
  - Given our state of knowledge
  - Choose optimal action
  - Given limited computational resources

- Human-like intelligence/behavior

Search in Discrete State Spaces

- This is different from Web Search ☺
- Every discrete problem can be cast as a search problem.
- (states, actions, transitions, cost, goal-test)
- Types
  - uninformed systematic: often slow
    - DFS, BFS, uniform-cost, iterative deepening
  - Heuristic-guided: better
    - Greedy best first, A*
    - relaxation leads to heuristics
  - Local: fast, fewer guarantees; often local optimal
    - Hill climbing and variations
    - Simulated Annealing: global optimal
    - Genetic algorithms: somewhat non-local due to crossing over
  - (Local) Beam Search

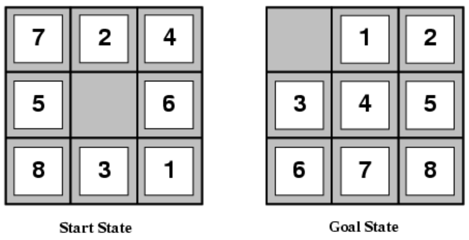

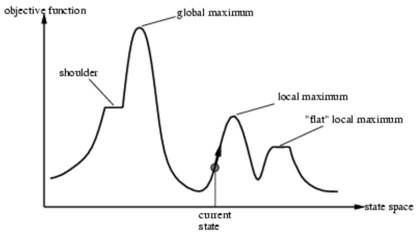

Search Example: Game Playing

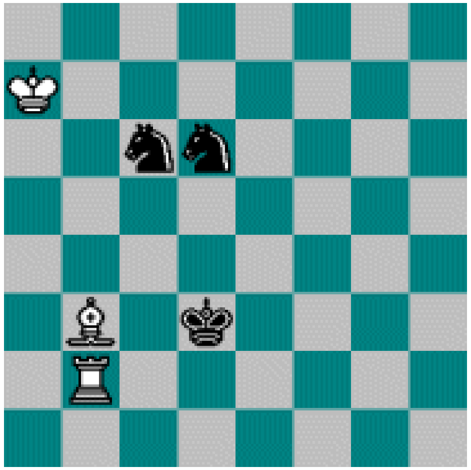

- Game Playing
  - AND/OR search space (max, min)
  - minimax objective function
  - minimax algorithm (~dfs)
    - alpha-beta pruning
  - Utility function for partial search
    - Learning utility functions by playing with itself
  - Openings/Endgame databases
    - Secondary search/Quiescence search

Knowledge Representation and Reasoning

- Representing: what I know
- Reasoning: what I can infer

- CSP
- Logic

- Bayes Nets
- Markov Decision Process

- Decision Trees
- Neural Network

KR&R Example: Propositional Logic

- Representation: Propositional Logic Formula
  – CNF, Horn Clause,…
- Reasoning: Deduction
  – Forward Chaining
  – Resolution

- Model Finding
  – Enumeration
  – SAT Solving

Search+KR&R Example: CSP

- Representation
  - Variables, Domains, Constraints
- Reasoning: Constraint Propagation
  - Node consistency, Arc Consistency, k-Consistency
- Search
  - Backtracking search: partial var assignments
    - Heuristics for choosing which var/value next
  - Local search: complete var assignments
- Tree structured CSPs: polynomial time
- Cutsets: vars assigned → converts to Tree CSP

Search+KR&R Example: SAT Solving

- Representation: CNF Formula
- Reasoning
  - pure literals; unit clauses; unit propagation
- Search
  - DPLL (~ backtracking search)
    - MOM's heuristic
  - Local: GSAT, WalkSAT
- Advances
  - Clause Learning: learning from mistakes
  - Restarts in systematic search
  - Portfolio of SAT solvers; Parameter tuning

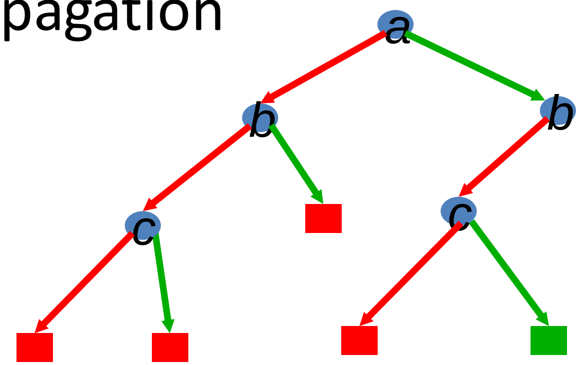

- Phase Transitions in SAT problems

Search+KR&R Example: Planning

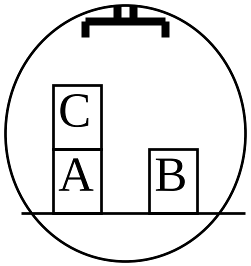

- Representation: STRIPS
- Reasoning: Planning Graph
  – Polynomial data structure
  – reasons about constraints on plans (mutual exclusion)
- Search
  – Forward: state space search
    - planning graph based heuristic
  – Backward: subgoal space search
  – Local: FF (enforced hill climbing)
- Planning as SAT: SATPlan

KR&R Part 2: Continuous Spaces

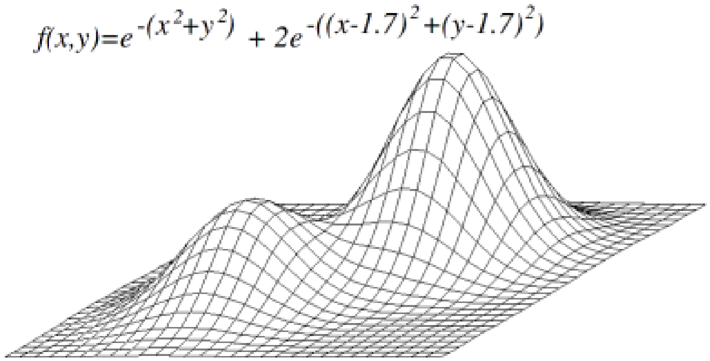

- Search
  - Gradient Descent
  - Newton Raphson
  - Optimization (convex/non-convex…)

- Constraint Optimization (we didn’t study this)
  - Linear Programming
  - Integer Linear Programming
  - Mixed Integer Linear Programming

KR&R: Probability

- Representation: Bayesian Networks
  - encode probability distributions compactly
    - by exploiting conditional independences

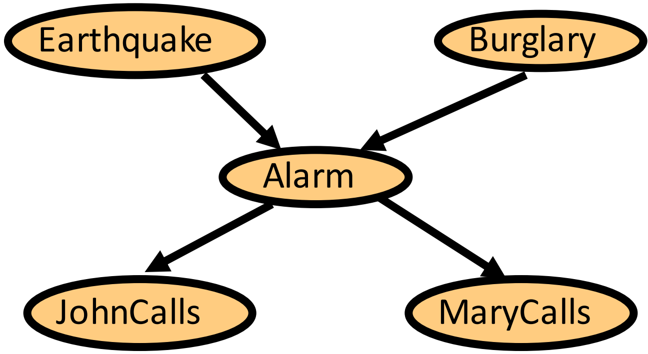

- Reasoning
  - Exact inference: var elimination
  - Approx inference: sampling based methods
    - rejection sampling, likelihood weighting, Gibbs sampling

KR&R: Hidden Markov Models

- Representation
  - Spl form of BN
  - Sequence model
  - One hidden state, one observation

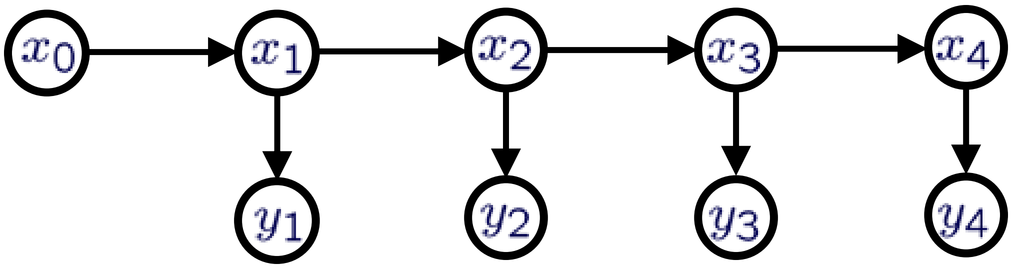

- Reasoning/Search
  - most likely state sequence: Viterbi algorithm
  - marginal prob of one state: forward-backward

KR&R: One-step Decision Theory

- Representation
  - actions, probabilistic outcomes, rewards
- Reasoning
  - expected value/regret of action
  - Expected value of perfect information

	States of Nature	
Actions	**Favorable Market**	**Unfavorable Market**
Large plant	**$200,000**	**-$180,000**
Small plant	**$100,000**	**-$20,000**
No plant	**$0**	**$0**

- Non-deterministic uncertainty
  - Maximax, maximin, eq likelihood, minimax regret..

- Utility theory: value of money…

KR&R: Markov Decision Process

- Representation
  - states, actions, probabilistic outcomes, rewards
  - ~AND/OR Graph (sum, max)
- Reasoning: V*(s)
  - Value Iteration: search thru value space
  - Policy Iteration: search thru policy space
- State space search
  - LAO* (AND/OR version of A*)

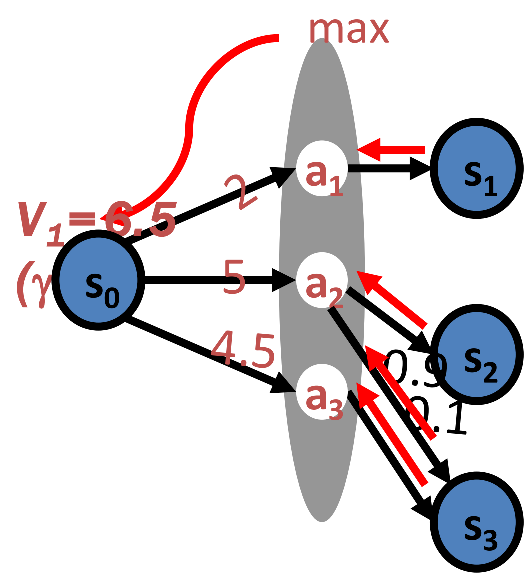

Learning: BNs/HMMs/NB

- ML estimation. max $P(D|\theta)$
  - counting; smoothing
- MAP estimation max $P(\theta|D)$..
- Hidden data
  - Expectation Maximization (EM) {local search}

- Structure learning (BN)
  - Local search thru structure space
  - Trade off structure complexity and data likelihood

- HMM: Hidden State Space
  - Baum Welch (like EM)

Learning: Decision Tree

- Representation
  – tree with one variable at each node
- Reasoning
  – just follow the appropriate path
- Learning
  – Greedy search: split one var at a time
    - post pruning/early stopping

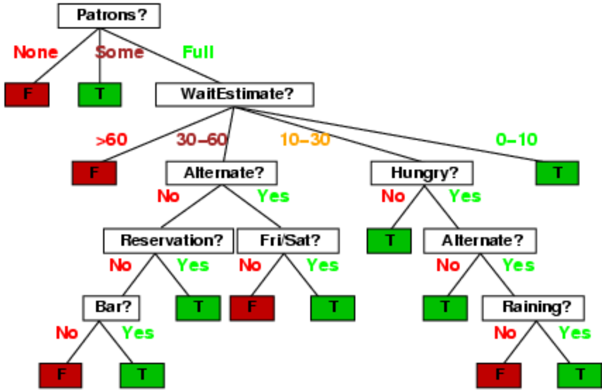

Learning: Perceptron

- Representation: perceptron

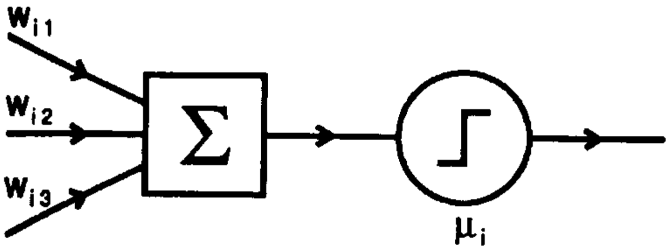

- Learning
  - local search in weight space to minimize errors
  - contrast with SVM
    - maximize margin from support vectors

- Perceptron: linear separator
- Neural network: layers of perceptrons

Learning: Nearest neighbor

- Representation: none!
- Reasoning: weighted average of k-nearest pts
- Learning: none!

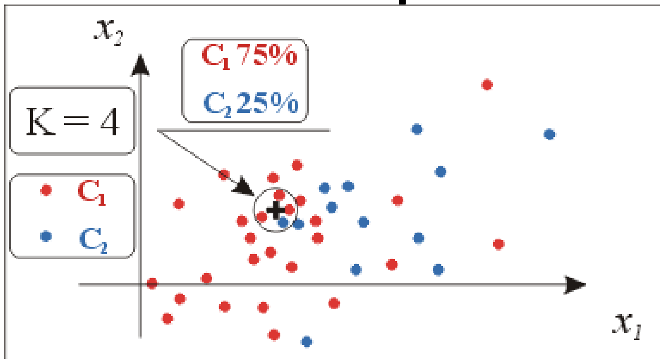

- can represent any decision boundary
  – requires huge data (needs all space to be filled)
  – makes error close to boundary

Agents

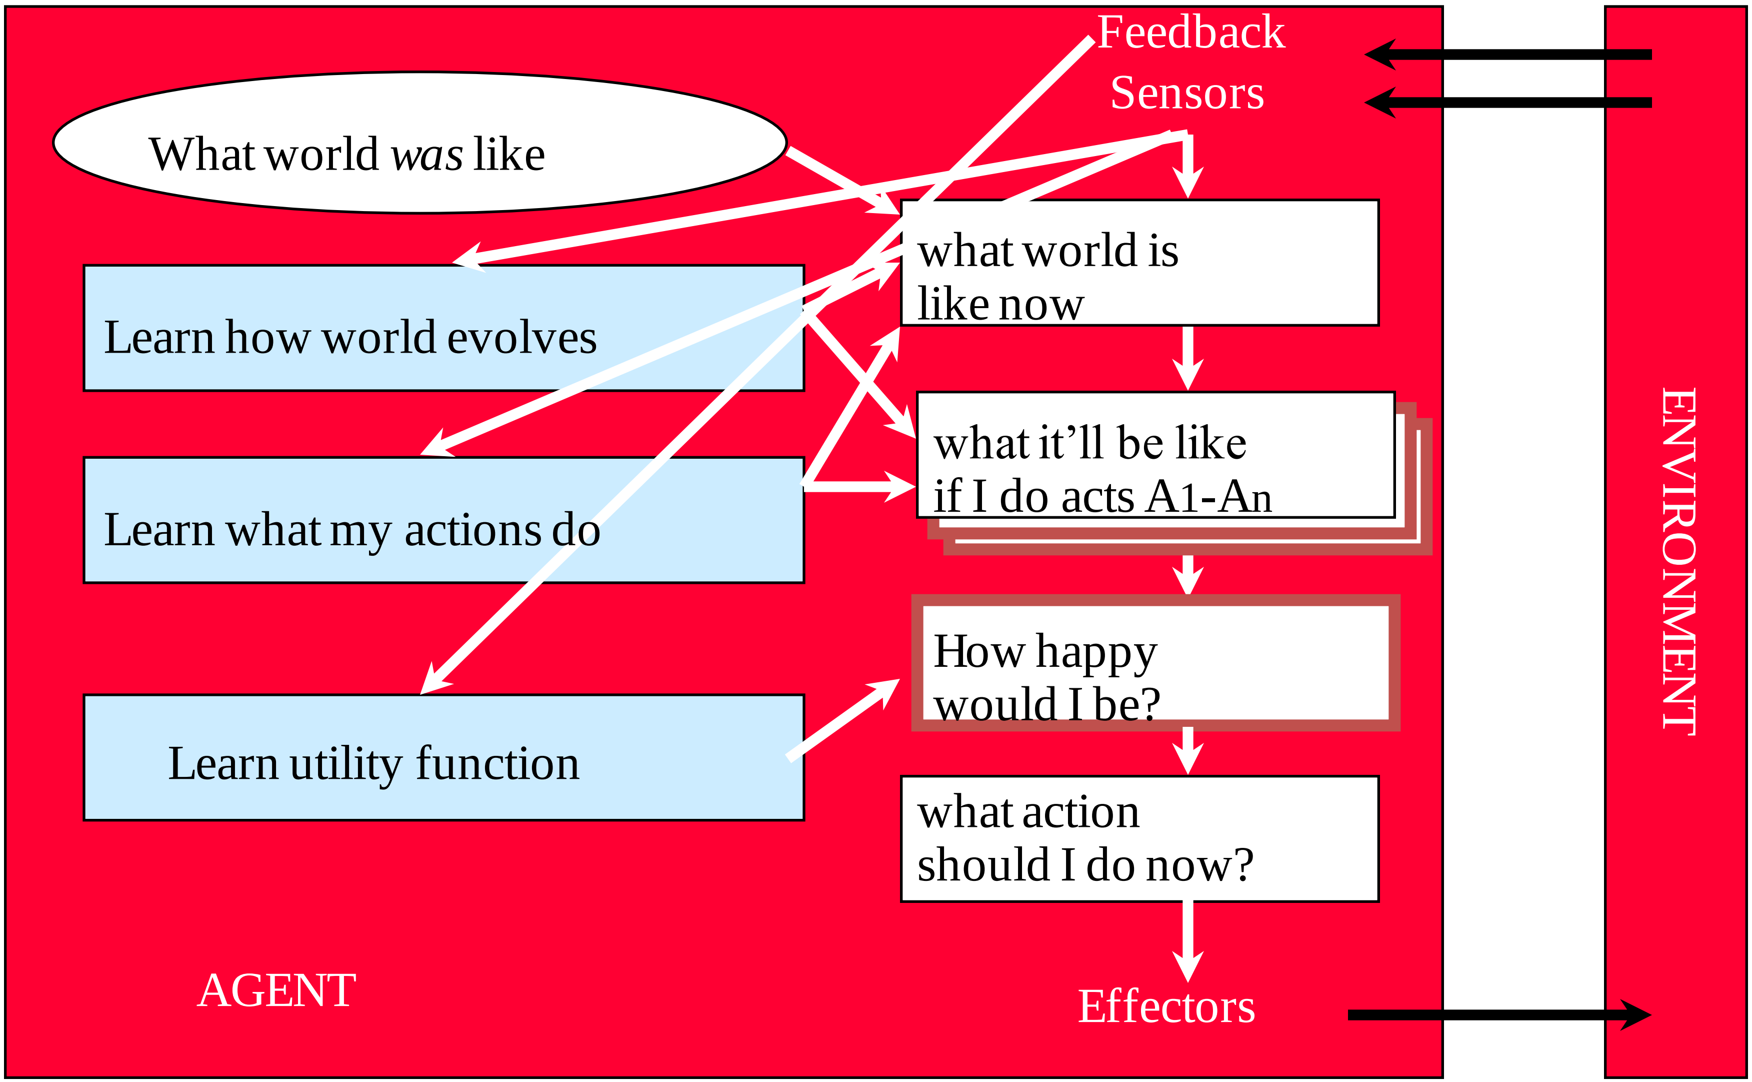
Feedback
Sensors
What world *was* like
what world is
like now
Learn how world evolves
what it'll be like
if I do acts A1-An
Learn what my actions do
How happy
would I be?
Learn utility function
what action
should I do now?
AGENT
Effectors
ENVIRONMENT

Popular Themes

- Weak AI vs. Strong AI
- Syntax vs. Semantics
- Logic vs. Probability

Weak AI vs. Strong AI

- Weak – general methods
  - primarily for problem solving
  - A*, CSP, Bayes Nets, MDPs…
- Strong -- knowledge intensive
  - – more knowledge $\Rightarrow$ less computation
  - – achieve better performance in specific tasks
  - – POS tagging, Chess, Jeopardy

Syntax vs. Semantics

- Syntax: what can I say
  - Sentence in English
  - Logic formula in Prop logic
  - CPT in BN

- Semantics: what does it mean
  - meaning that we understand
  - A ^ B: both A and B are true
  - Conditional independence …

Logic vs. Probability

- Discrete || Continuous
- Hill climbing || Gradient ascent

- SAT solving || BN inference
- Tree structured CSP || Polytree Bayes nets
- Cutset || Cutset

- Classical Planning || Factored MDP
- Bellman Ford || Value Iteration
- A* || LAO*

Advanced Ideas in AI

- Factoring state/actions…
- Hierarchical decomposition
  - Hierarchy of actions
- Approximation by sampling
  - Markov Chain Monte Carlo
  - UCT algorithm: game playing
  - Particle filters: belief tracking in robotics
- Context sensitive independence
  - Cutsets
  - Backbones in logic
- Combining probability and logic
  - Markov Logic Networks, Probabilistic Relational Models

AI we didn’t cover

- Ontologies
- Information retrieval/web search
- Robotics
- Vision
- Mechanism design
- Computational Neuroscience
- Reinforcement learning
- …

Applications of AI

- Sumit: automatic accompaniment of music
  - probabilistic reasoning, machine learning (HMMs)
- Ashish: hardware/software verification, combinatorial design, subprobs in many domains
  - SAT solving
- Joseph: fraud detection, market/risk assessment, personalization, recommender systems…
  - Machine learning
- Matthai: elderly care
  - Machine learning, probabilistic reasoning

Applications of AI

- Mars rover: planning
- Jeopardy: NLP, info retrieval, machine learning
- Puzzles: search, CSP, logic
- Chess: search
- Blackjack: MDP
- Text categorization: machine learning

- Self-driving cars: robotics, prob. reasoning, ML…

Ethics of Artificial Intelligence

- Robots
  - Robot Rights
  - Three Laws of Robotics
- AI replacing people jobs
  - Any different from industrial revolution?
- Ethical use of technology
  - Dynamite vs. Speech understanding
- Privacy concerns
  - Humans/Machines reading freely available data on Web
  - Gmail reading our news
- AI for developing countries/improving humanity

AI-Centric World ☺