Outline for 4/2

- Introduction & Logistics
- Notion of a Problem Space
- Search Techniques
- Video: Kasparov *vs.* Deep Blue
- Constraint Satisfaction Techniques

Logistics:

- Dan Weld weld@cs.washington.edu
- Jared Saia saia@cs.washington.edu
- www.cs.washington.edu/education/courses/cse592/98sp
- Reading
  - Required: Packet of papers
  - Recommended: Russell & Norvig "AIMA"
- Grading:
  - Problem Sets
  - Class
  - Midterm, Final
  - Projects

Dimensions of the AI Definition

human-like *vs.* rational

thought *vs.* behavior		
	Systems that think like humans	Systems that think rationally
	Systems that act like humans	Systems that act rationally

AI as Science

- Origin & Laws of the Physical Universe
- Origin & Laws of Biological Life
- ***Nature of Intelligent Thought***

AI as Engineering

- Softbots & Intelligent User Interfaces
- Mobile Robots … Immobots
- Machine Learning Algorithms; Data Mining
- Medical Expert Systems...

AI Theory

- For example, *Machine Learning*
  - The concept ***orgespat***?
- Given tree structured instance space with n attributes, you need
  - $4 \log(2/\delta) + 16n \log(13/\varepsilon))/\varepsilon$ examples

Search

- Brute force
  - DFS, BFS, iterative deepening, iterative broadening
- Heuristic
  - Best first, beam, hill climbing, simulated annealing, limited discrepancy
- Optimizing
  - Branch & bound, A*, IDA*, SMA*
- Adversary Search
  - Minimax, alpha-beta, conspiracy search
- Constraint Satisfaction
  - As search, preprocessing, backjumping, forward checking, dynamic variable ordering

Knowledge Repr & Reasoning

- Propositional Logic
  - Syntax, inference: modus ponens, resolution, DPLL, stochastic methods; truth maintenance, hard problems, phase transitions
- Predicate Calculus
  - Syntax, inference: unification, resolution, compilation to propositional form
- Modal Logic
  - Representing action and change, diagnosis, planning
- Bayesian Belief Networks
  - Joint prob distrib, indepenence, Bayes rule, inference and diagnosis: clustering methods, cutset conditioning, stochastic simulation; utility functions, decision networks

Learning

- Supervised Learning
  - Methodology: N-fold cross-validation; induction of decision trees, entropy, information gain, ID3, C4.5, overfitting, ensembles of classifiers, bagging, boosting, scaleup, complete classification vs nuggets, feature selection: LOOCV, racing, schemas; inductive logic programming, FOIL, minimum description length, Grendel
- Reinforcement Learning
  - Cummulative discounted reward, dynamic programming, policy iteration, value iteration, temporal-difference learning
- Bayesian Learning

Course Topics by Week

- Search & Constraint Satisfaction
- KR 1: Propositional Logic
- Autonomous Spacecraft 1: Configuation Mgmt
- Autonomous Spacecraft 2: Reactive Planning
- KR2: Modeling Internet Information Sources
- Information Integration 2: Planning & Execution
- Supervised Learning & Datamining
- Reinforcement Learning
- Bayes Nets: Inference & Learning
- Review & Future Forecast

Search & Knowledge Representation

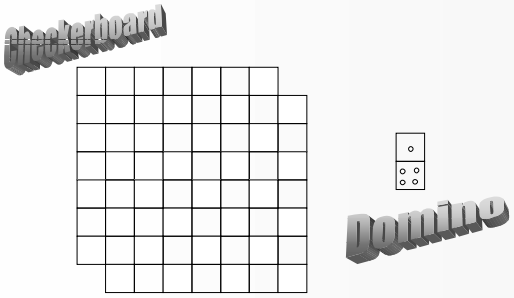

Artificial Intelligence

- Planning **(To get a good job I should get an MS/CSE…)**
- Diagnosis **(My car won't start, but I know it has gas…)**
- Design **(Space Shuttle, Boeing 777, Pentium)**
- Communicate **(Language + Social reasoning)**
- Theorem Proving **(Fermat's Last Theorem)**
- Perceive/Act **(Vision, Manipulation)**
- Game Playing **(Kasparov *vs.* Deep Blue)**
- Create Art **(MacBeth, Mona Lisa, Taj Mahal)**
- Learn to ...

11

Unifying View

- Knowledge Representation
  - Expressiveness
  - Reasoning (Tractability)
- Search
  - Space being searched
  - Algorithms & performance

Outline

- Defining a Search Space
- Types of Search
  - Blind
  - Heuristic
  - Optimization
  - Adversary Search
  - Constraint Satisfaction
- Analysis
  - Completeness
  - Time & Space Complexity

Specifying a search problem?

- What are states (nodes in graph)?
- What are the operators (arcs between nodes)?
- Initial state?
- Goal test?
- [Cost?, Heuristics?, Constraints?]

E.g., Eight Puzzle

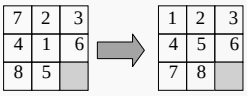

Missionaries and Cannibals

- What are states (nodes in graph)?
- What are the operators (arcs between nodes)?
- Initial state?
- Goal test?

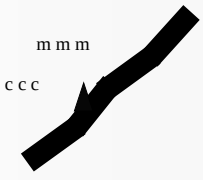

15

Supervised Concept Learning

- What are states (nodes in graph)?
- What are the operators (arcs between nodes)?
- Initial state?
- Goal test?

16

Search Strategies

- Blind Search
  - Depth first search
  - Breadth first search
  - Iterative deepening search
  - Iterative broadening search
- Heuristic Search
- Optimizing Search
- Constraint Satisfaction

17

Depth First Search

- Maintain stack of nodes to visit
- Evaluation
  - Complete?
    Not for infinite spaces
  - Time Complexity?
    O(b^d)
  - Space Complexity?
    O(d)

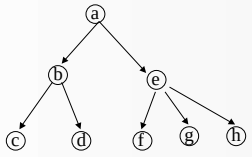

Breadth First Search

- Maintain queue of nodes to visit
- Evaluation
  - Complete?
    Yes
  - Time Complexity?
    O(b^d)
  - Space Complexity?
    O(b^d)

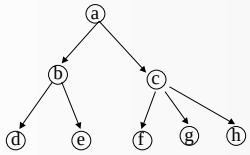

Iterative Deepening Search

- DFS with limit; incrementally grow limit
- Evaluation
  - Complete?
    Yes
  - Time Complexity?
    O(b^d)
  - Space Complexity?
    O(d)

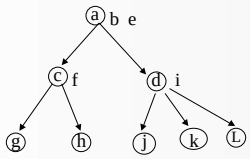

Iterative Broadening Search

- What if know solutions lay at depth N?
- No sense in doing interative ***deepening***

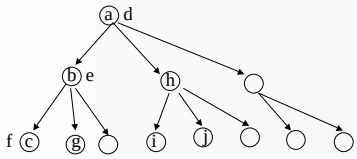

Search Strategies

- Blind Search
- Heuristic Search
  - Best-first
  - Beam
  - Hill climbing
  - Limited discrepancy
- Optimizing Search
- Constraint Satisfaction

22

Heuristic Search

- A heuristic is:
  - Function from a state to a real number
    - Low number means state is close to goal
    - High number means state is far from the goal

Designing a good heuristic is very important!
(And hard)
More on this in a bit...

Best First Search

- Idea
  - Breadth first but use priority queue instead of a queue
- Evaluation
  - Complete?
    Yes
  - Time Complexity?
    O(b^d)
  - Space Complexity?
    O(b^d)

Beam Search

- Idea
  - Best first but only keep N best items on priority queue
- Evaluation
  - Complete?
    No
  - Time Complexity?
    O(b^d)
  - Space Complexity?
    O(b + N)

Hill Climbing

- Idea
  - Always choose best child; no backtracking
- Evaluation
  - Complete?
    No - suffers from plateau, local maxima, ridges
  - Time Complexity?
    O(b^d) **but only in pathological cases**
  - Space Complexity?
    O(b)

Simulated Annealing

- Objective: avoid local minima
- Technique:
  - For the most part use hill climbing
  - Occasionally take non-optimal step
  - Reduce probability(non-optimal) over time
- Comparison to Hill Climbing
  - Completeness?
  - Speed?
  - Space Complexity?

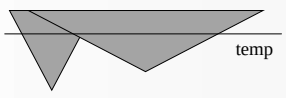

Limited Discrepancy Search

- Discrepancy bound indicates how often to violate heuristic
- Iteratively increase...

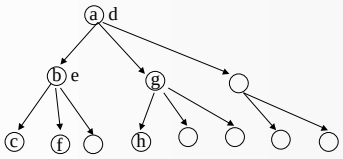

Assume that heuristic says go ***left***

Search Strategies

- Blind Search
- Heuristic Search
- Optimizing Search
  - A*
  - IDA*
  - SMA*
- Constraint Satisfaction

Objective is to find the very ***best*** solution.

29

A* Search

Underestimates cost of any solution which can reached from node

- Idea
  - Best first search with **admissible heuristic**
  - Plus keep checking until all possibilities look worse
- Evaluation
  - Finds optimal solution?
    Yes
  - Time Complexity?
    O(b^d)
  - Space Complexity?
    O(b^d)

Admissable Heuristics

- f(x) = g(x) + h(x)
- g: cost so far
- h: underestimate of remaining costs

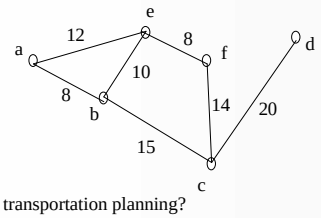

For transportation planning?

For eight puzzle?

7	2	3
4	1	6
8	5	

Importance of Heuristics

7	2	3
4	1	6
8	5	

- h1 = number of tiles in the wrong place
- h2 = sum of distances of tiles from correct loc

D	IDS	A*(h1)	A*(h2)
2	10	6	6
4	112	13	12
6	680	20	18
8	6384	39	25
10	47127	93	39
12	364404	227	73
14	3473941	539	113
18		3056	363
24		39135	1641

Iterative Deepening A*

- Like iterative deepening depth first, but...
- Depth bound modified to be an f-cost limit
- Contour lines bounding search

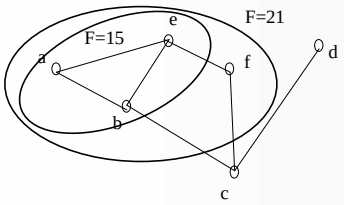

SMA*

- Problem is f-cost bound increases slowly
- Must do iterative search again and again
- Storing little state between each iteration
  - Just one number: next highest contour level
- SMA*
  - Uses all available memory to store state
  - Duplicates minimal work
  - Optimal in a number of nice ways