Review

Philipp Koehn

30 April 2020

Exam

- Date: Thursday, May 13, all day (but should take at most 3 hours)
- Posted on Piazza, to be submitted to Gradescope
- Format
  - open book
- Grading: homework is 60%, exam is 40%

Lectures Covered By Exam

- Artificial Intelligence in Context — not covered

- Intelligent Agents, Heuristic Search, and Game Playing
  - Intelligent Agents
  - Basic Search
  - Informed Search
  - Game Playing
  - Constraint Satisfaction

- Logic and Knowledge Representation
  - Logical Agents
  - First Order Logic
  - Inference in First-Order Logic
  - Knowledge Representation
  - Planning

- Uncertainty
  - Probabilistic Reasoning
  - Bayesian Networks
  - Markov Decision Processes
  - Decision Theory

- Machine Learning
  - Statistical Learning
  - Neural Networks
  - Reinforcement Learning
  - Deep Reinforcement Learning

- Natural Language

intelligent agents

Intelligent Agents

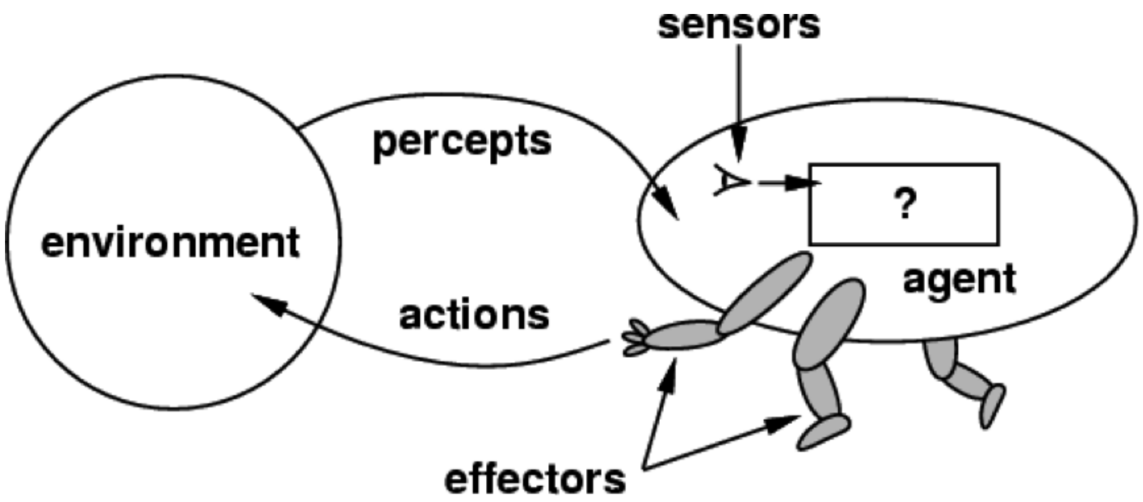

- Types of environments: (in)accessible, (non-)deterministic, (non)-episodic
- Types of agents: reflex, with memory, with goals, with learning, utility-based

Basic Search

- Problem solving agents

- Analysis
  - completeness
  - time complexity
  - space complexity
  - optimality

- Basic search algorithms
  - tree search
  - breadth / depth-first search
  - iterative deepening

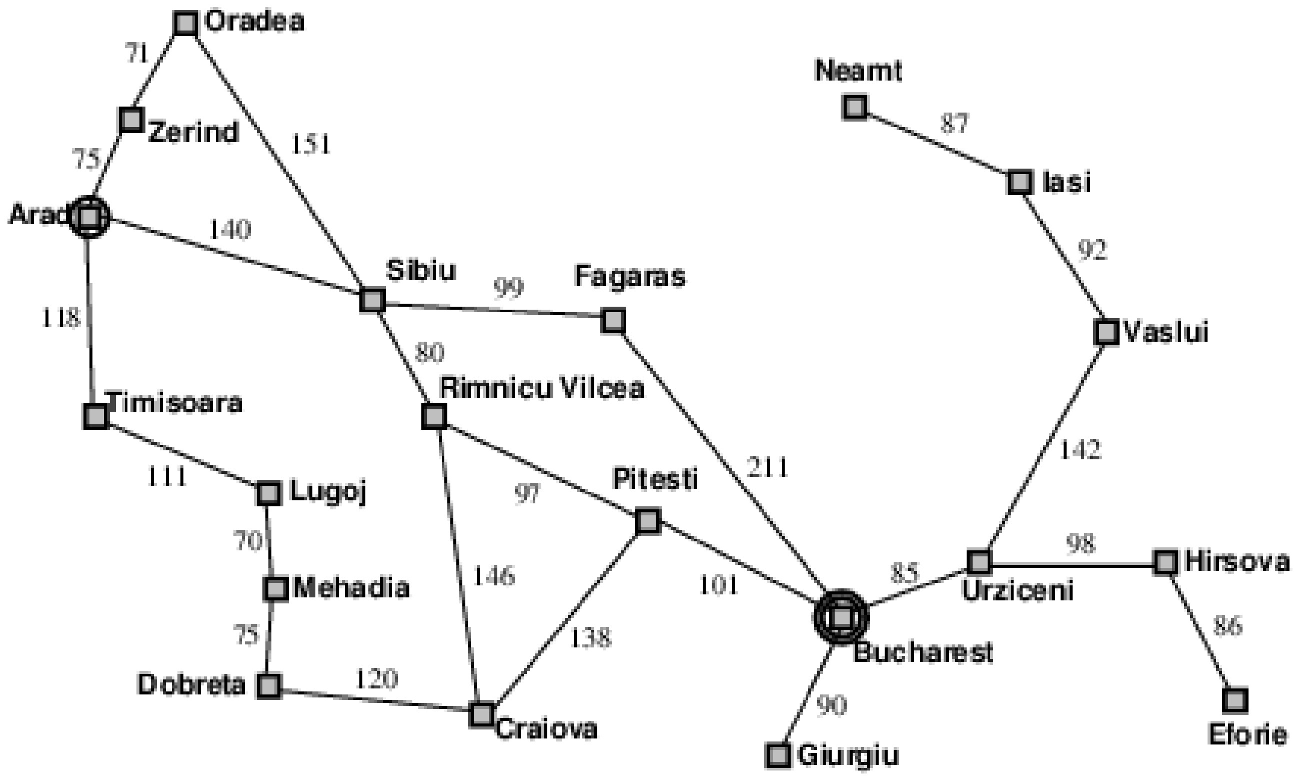

Informed Search

- Best-first search
- A^* search
- Heuristic algorithms
  – hill-climbing
  – simulated annealing

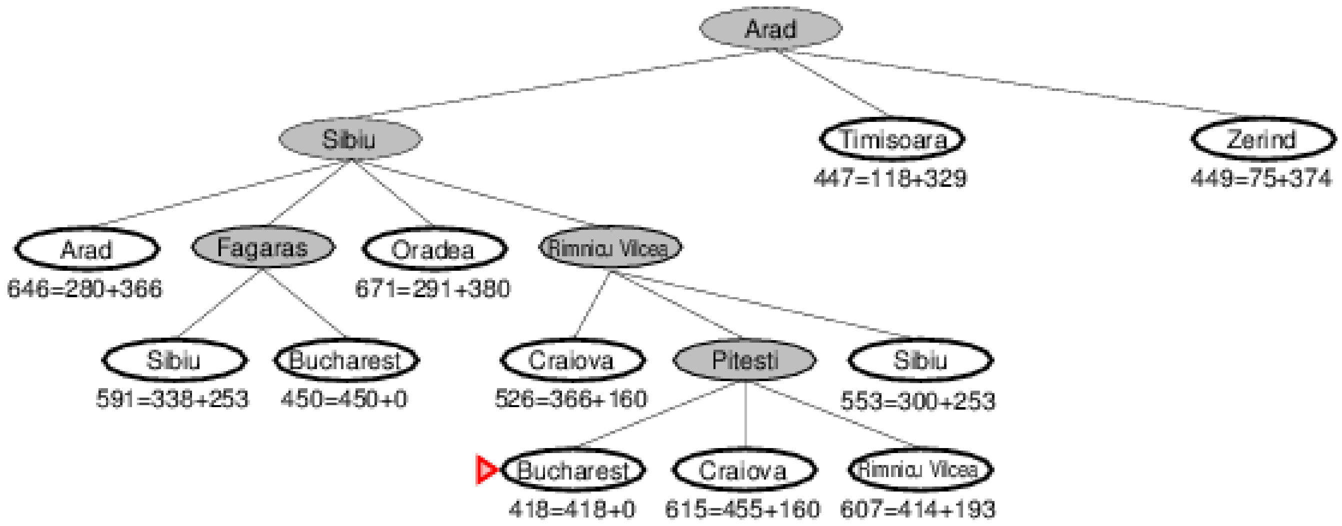

Game Playing

- Types of games
  - deterministic / probabilistic
  - (im)perfect information
- Search over game tree
  - minimax algorithm
  - α-β pruning
  - evaluation functions
- Solvable games, but typically resource limits
- Probabilistic games: pruning with bounds

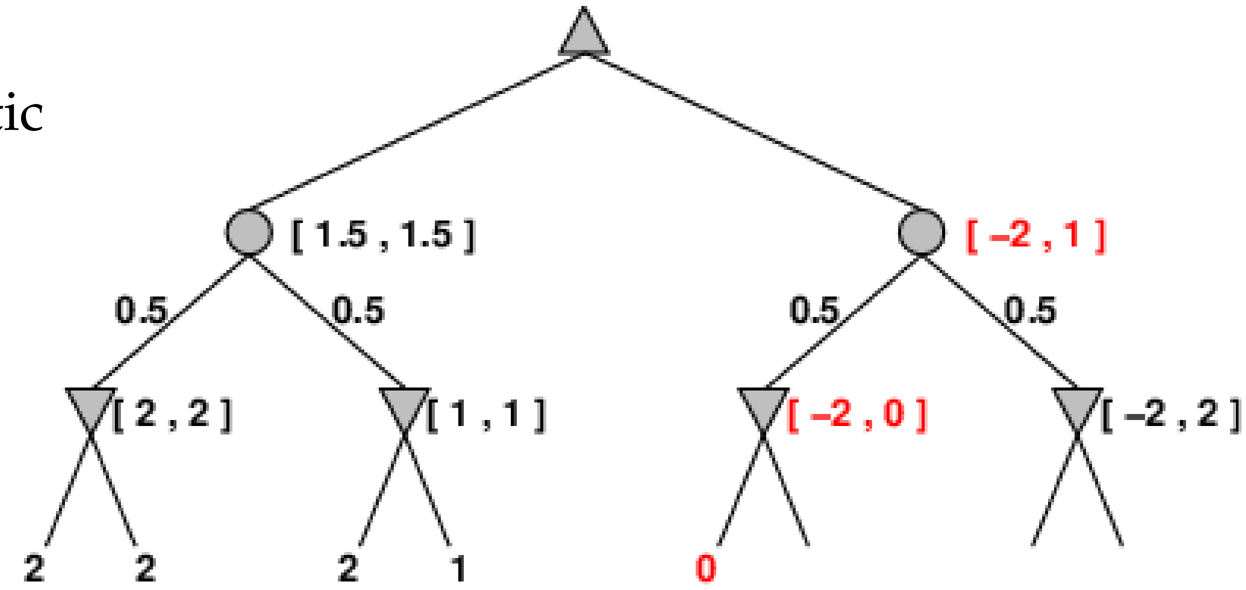

Constraint Satisfaction

- Variables, domains, constraints
- Backtracking search
- Constraint propagation
  - forward checking
  - arc consistency
- Problems structure
- Iterative algorithms

5	3			7				
6			1	9	5			
	9	8					6	
8				6				3
4			8		3			1
7				2				6
	6					2	8	
			4	1	9			5
				8			7	9

logic

Logical Agents

- Knowledge-based agents
  - internal representations
  - incorporate new percepts
  - deduce hidden properties of the world
- Logic
  - formal language (syntax)
  - truth in real world (semantics)
  - entailment and inference
- Algorithms
  - forward chaining
  - backward chaining
  - resolution

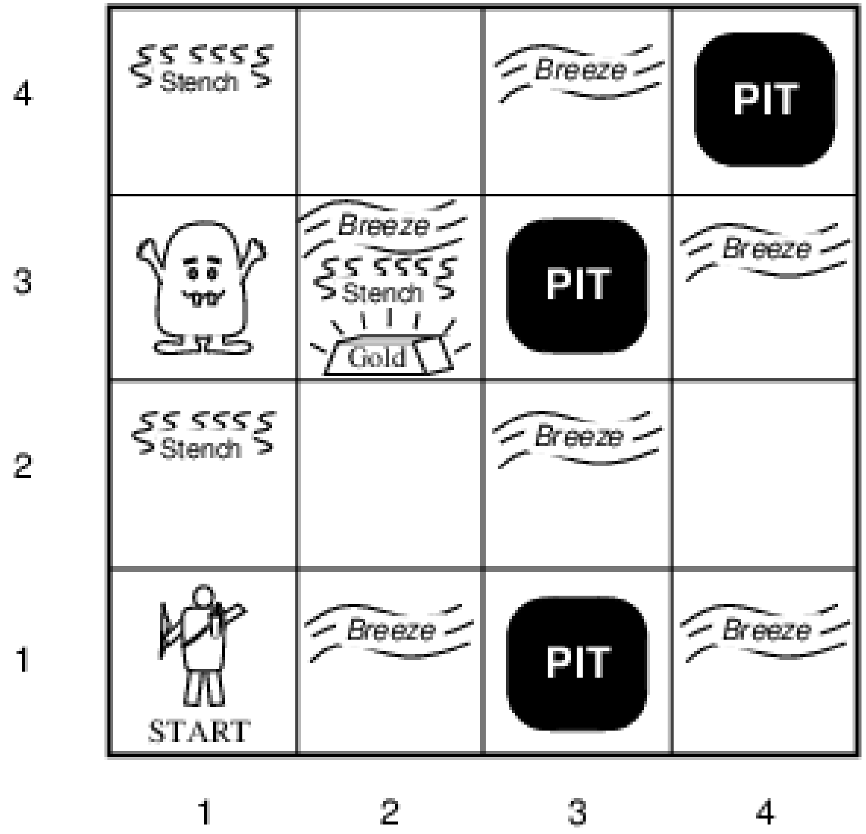

First Order Logic

- Adding
  - variables
  - relations
  - functions
  - quanitifiers

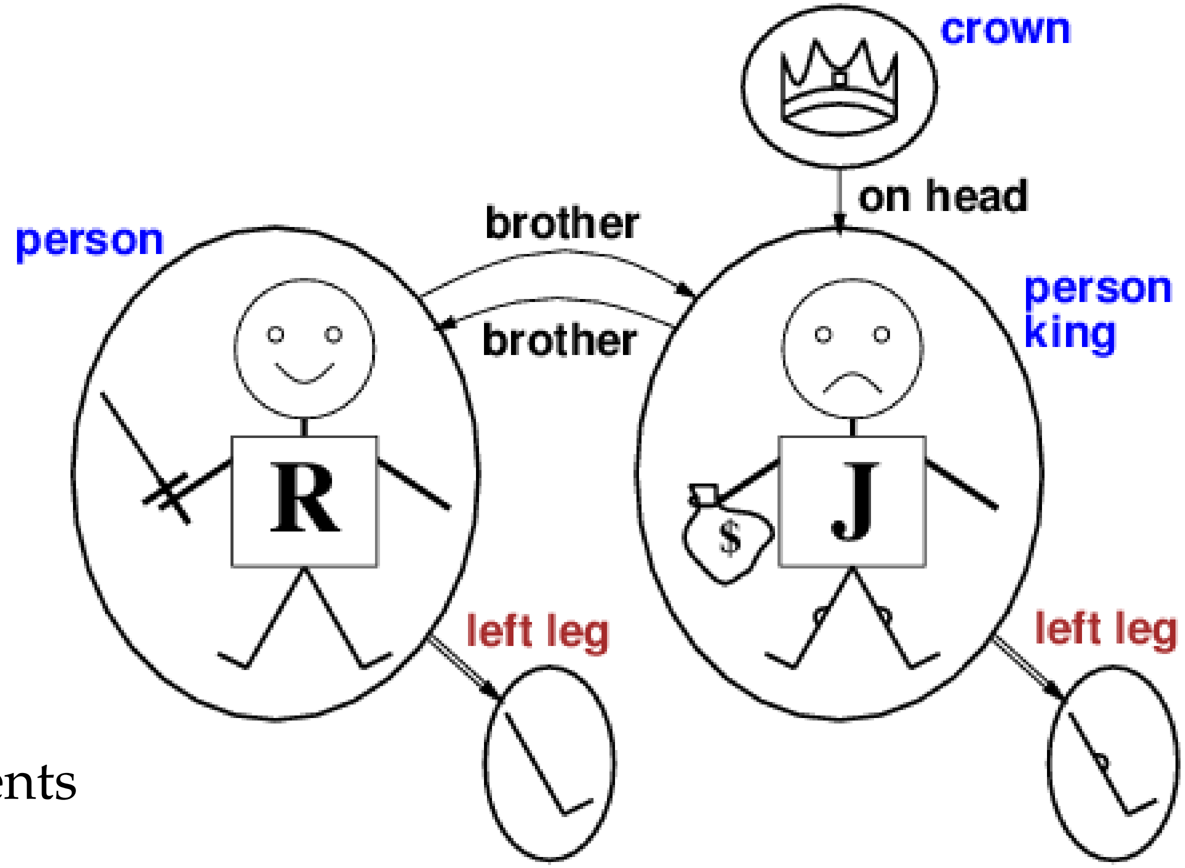

- Modeling natural language

- Dynamic world: states and fluents

- Situation calculus

Inference in First-Order Logic

- Reducing first-order inference to propositional inference
- Unification
- Generalized modus ponens
- Forward and backward chaining
- Logic programming (Prolog)
- Resolution

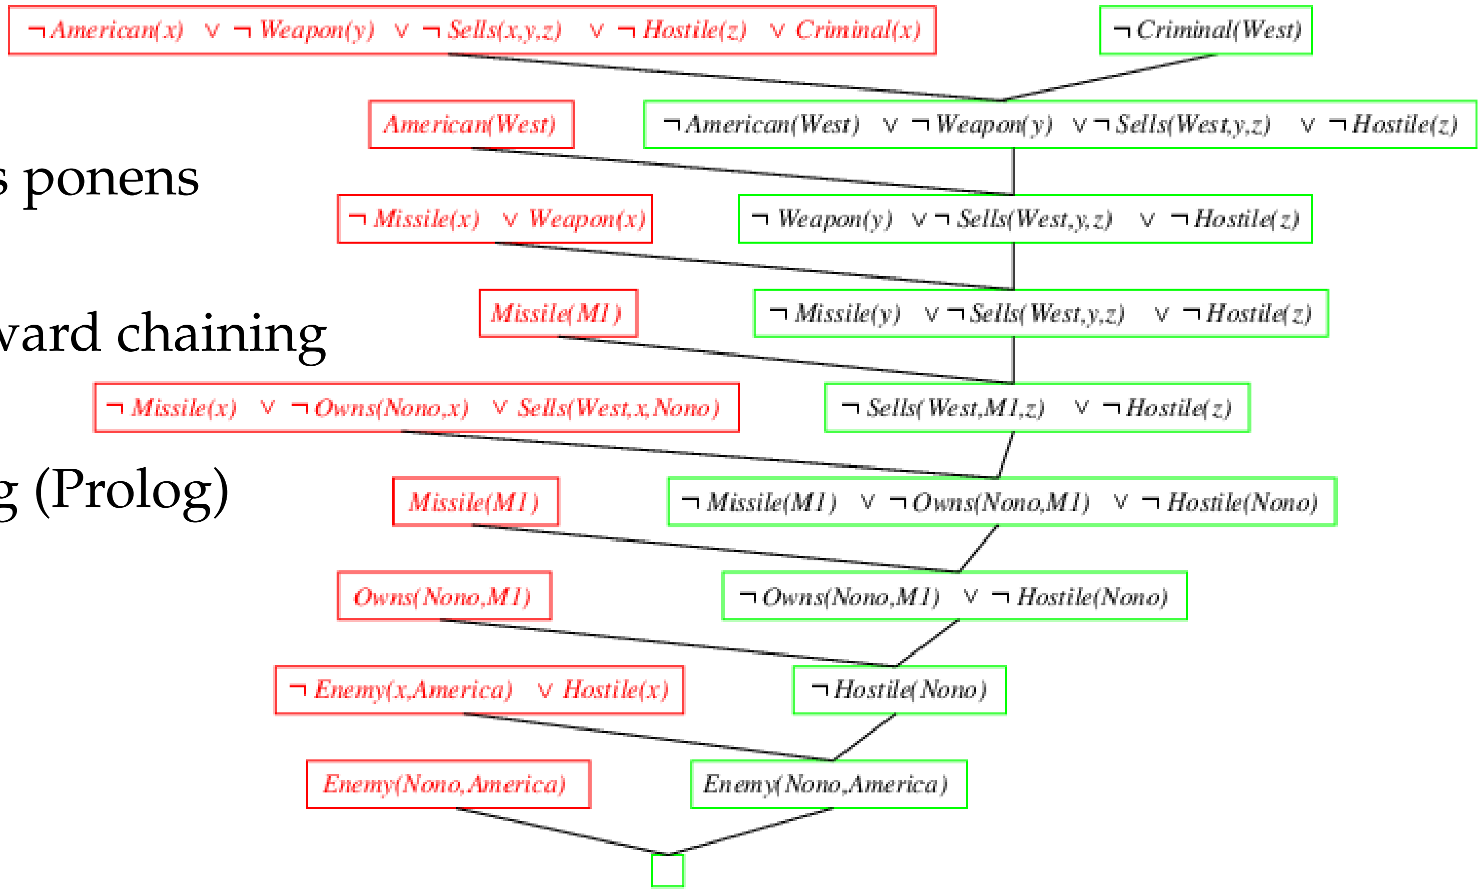

Knowledge Representation

- Representation systems
- Categories and objects
  → ontologies
- Frames
- Events and scripts
- Practical examples
  - Cyc
  - Semantic web

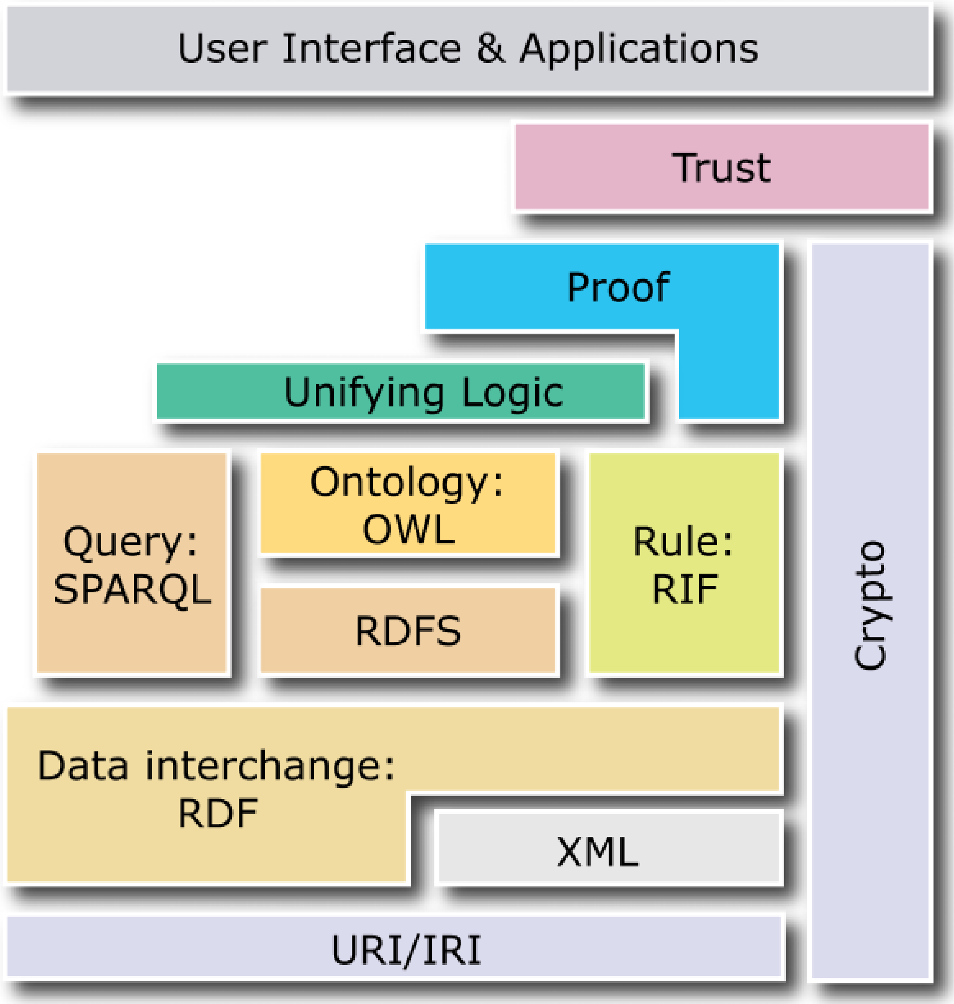

Planning

- Search vs. planning
- STRIPS operators
- Partial-order planning
- The real world
  - incomplete information
  - incorrect information
  - quantification problem
- Conditional planning
- Monitoring and replanning

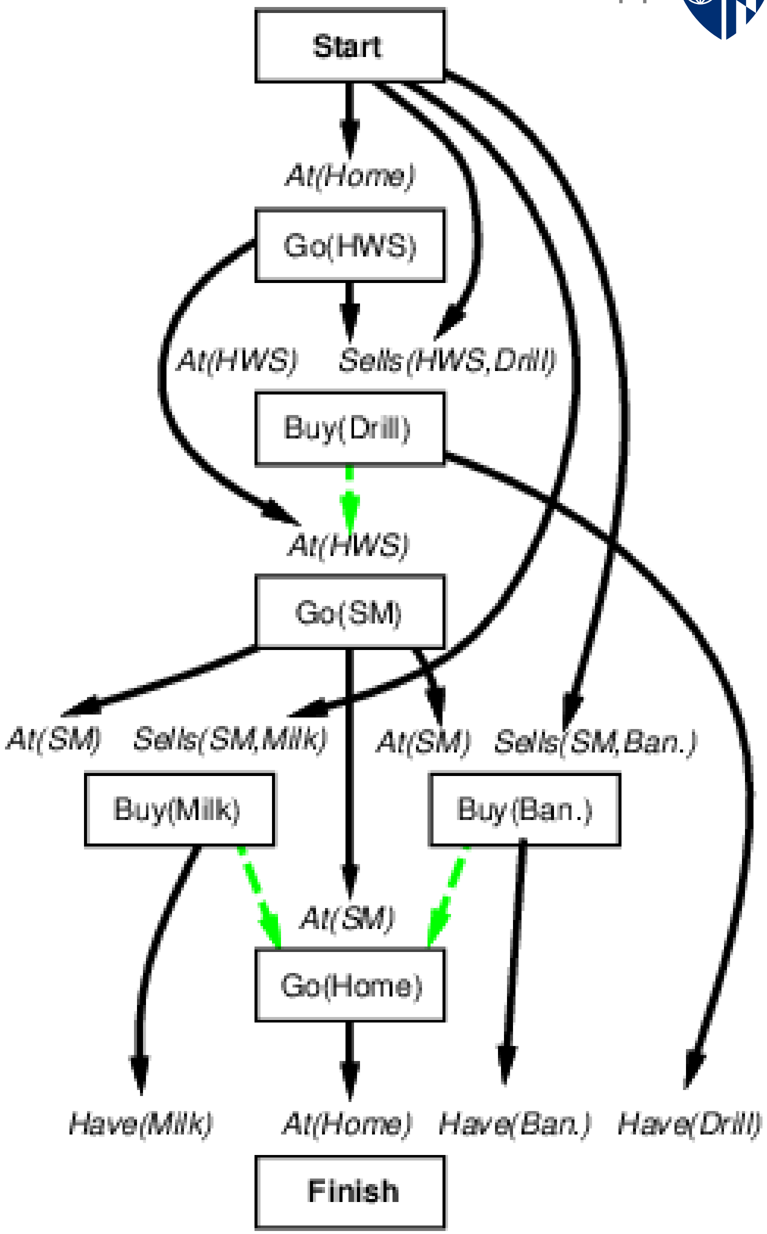

uncertainty

Probabilistic Reasoning

- Uncertainty
- Probability
  - conditional probability
  - independence
  - Bayes rule
- Inference
- Independence and Bayes' Rule

Bayesian Networks

- Bayesian Networks
- Parameterized distributions
- Exact inference
  - inference by enumeration
  - variable elimination
- Approximate inference
  - rejection sampling
  - likelihood weighting
  - Markov chain Monte Carlo

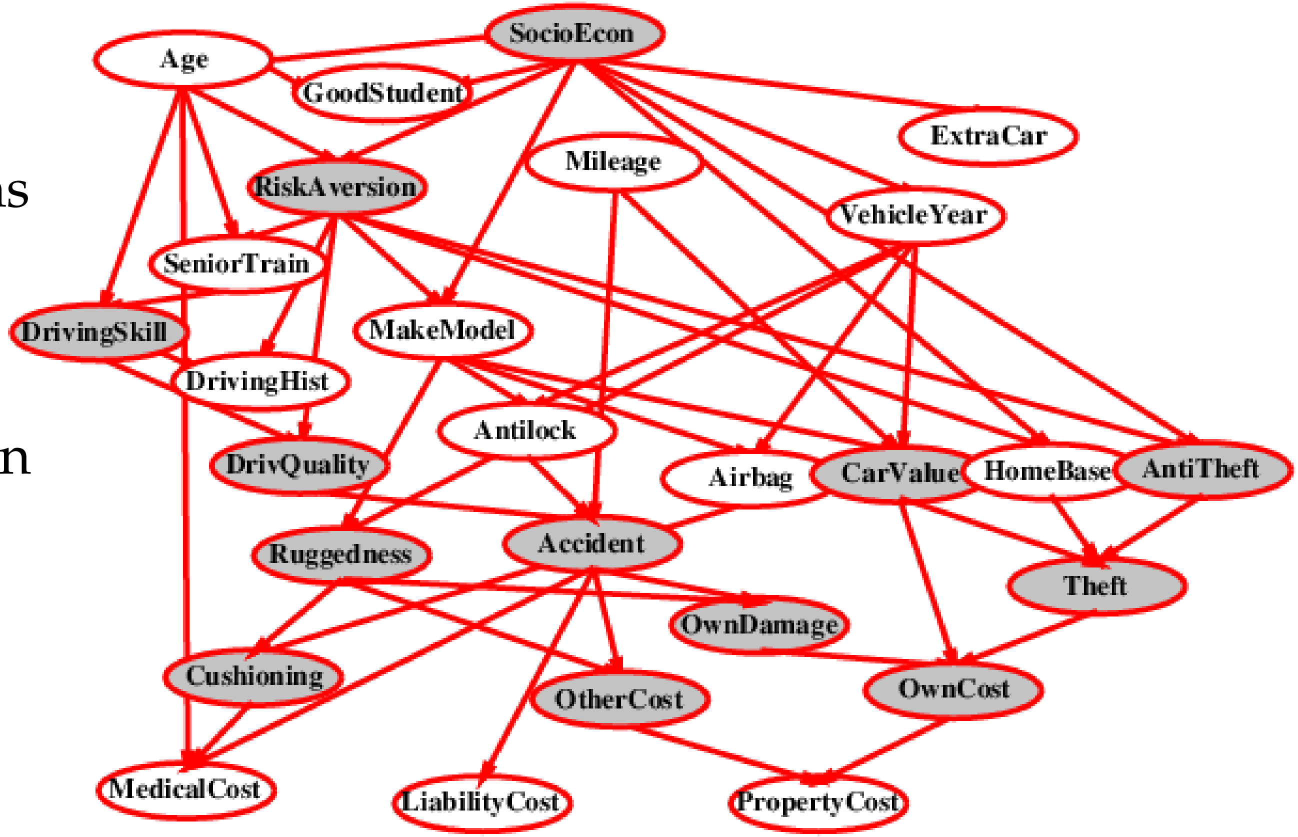

Markov Decision Processes

- Temporal processes
- Hidden Markov models
- Inference
  - filtering
  - smoothing
  - best sequence
- Kalman filters
- Dynamic Bayesian nets
- Example: speech recognition

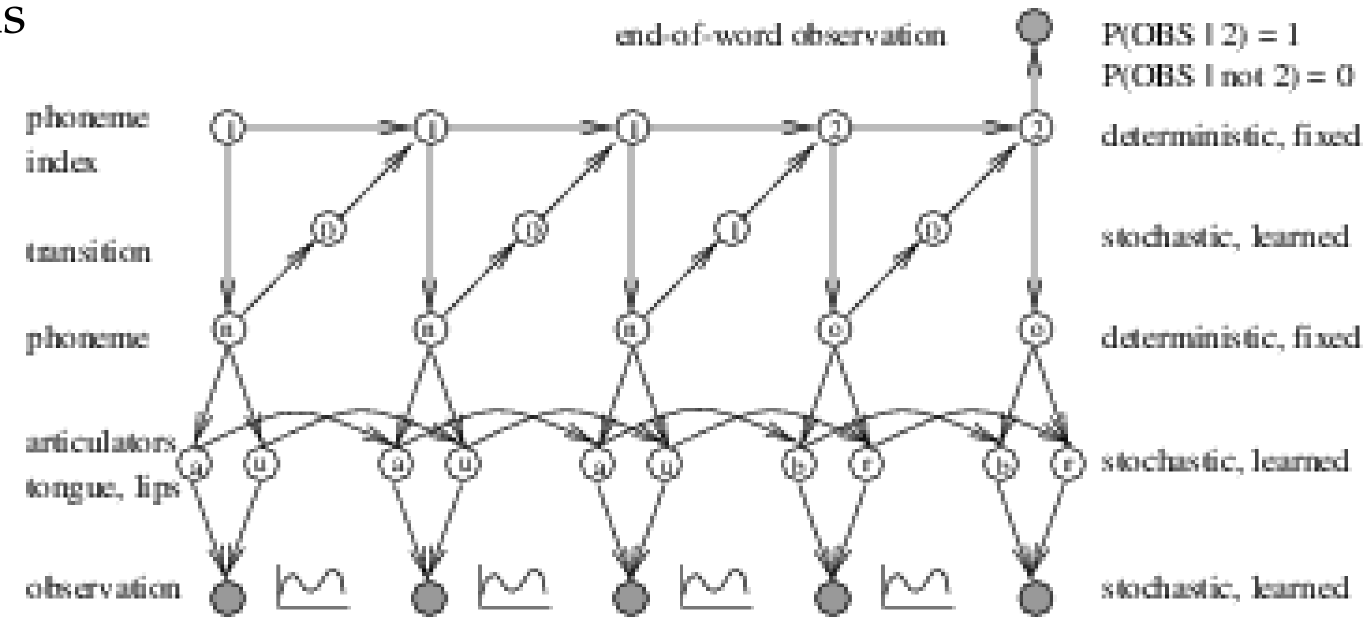

Decision Theory

- Rational preferences
- Utilities
- Decision networks
- Value of information
- Markov decision processes
- Inference algorithms
  - value iteration
  - policy iteration
- Partially observable Markov decision processes (POMDP)

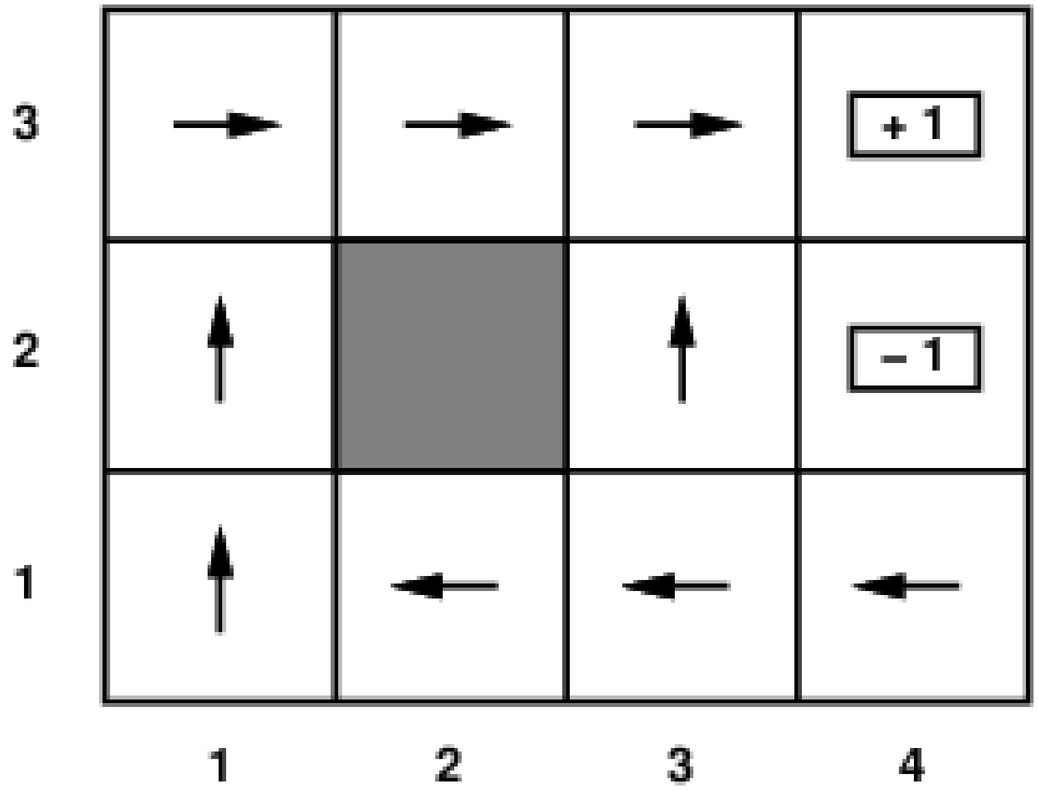

Reinforcement Learning

- Rewards, often delayed

- Passive reinforcement learning
  (compute utility of policy)
  - adaptive dynamic programming
  - temporal difference learning

- Active Reinforcement Learning
  - greedy agent
  - extended adaptive dynamic programming
  - Q-learning

- Generalizations over the state space

- Policy search

exam questions

Sample Exam

- Exam will assess
  - understanding of core concepts
  - understanding of algorithms
    → ability to carry them out by hand

- Exam will be similar to last years (posted on web site)

Sample Question

Logical knowledge representation

Which of the following are semantically and syntactically correct translations of *"Everyone's zipcode within a state has the same first digit"*?

1. $\forall$ x, s, z1 [State(s) $\wedge$ LivesIn(x, s) $\wedge$ Zip(x) = z1] $\Rightarrow$
   [$\forall$ y, z2 LivesIn(y, s) $\wedge$ Zip(y) = z2 $\Rightarrow$ Digit(1, z1) = Digit(1, z2)].

2. $\forall$ x, s [State(s) $\wedge$ LivesIn(x, s) $\wedge$ $\exists$ z1 Zip(x) = z1] $\Rightarrow$
   [$\forall$ y, z2 LivesIn(y, s) $\wedge$ Zip(y) = z2 $\wedge$ Digit(1, z1) = Digit(1, z2)].

3. $\forall$ x, y, s State(s) $\wedge$ LivesIn(x, s) $\wedge$ LivesIn(y, s) $\Rightarrow$ Digit(1, Zip(x) =Zip(y)).

4. $\forall$ x, y, s State(s) $\wedge$ LivesIn(x, s) $\wedge$ LivesIn(y, s) $\Rightarrow$ Digit(1, Zip(x)) = Digit(1, Zip(y)).

Sample Question

Game playing

Consider the game of 2×2 tictactoe.

1. Draw the full game tree down to depth 2. You need not show nodes that are rotations or reflections of siblings already shown.

2. Circle any node that would not be evaluated by alpha–beta during a left-to-right exploration of your tree.

questions?