Agents & Environments

Chapter 2

Mausam

(Based on slides of Dan Weld, Dieter Fox, Stuart Russell)

Outline

- Agents and environments
- Rationality
- PEAS specification
- Environment types
- Agent types

Agents

- An agent is anything that can be viewed as perceiving its environment through sensors and acting upon that environment through actuators

- Human agent:
  - eyes, ears, and other organs for sensors
  - hands,legs, mouth, and other body parts for actuators

- Robotic agent:
  - cameras and laser range finders for sensors
  - various motors for actuators

Examples of Agents

  - Intelligent buildings
  - Autonomous spacecraft

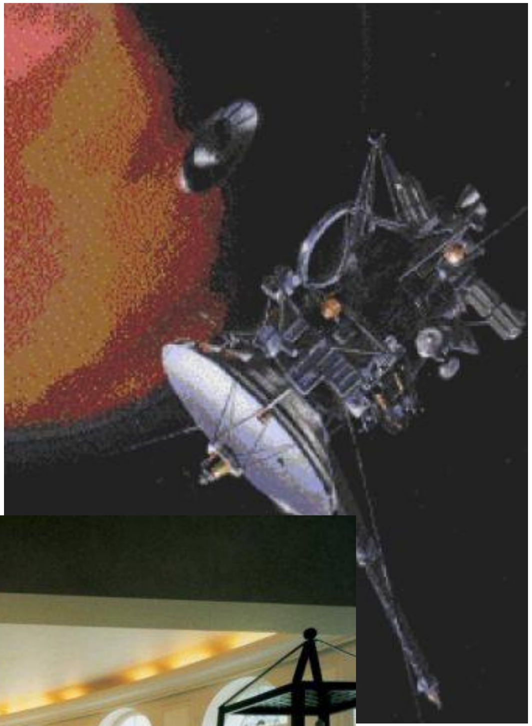

- Softbots
  - Askjeeves.com
  - Expert Systems

Intelligent Agents

- Have sensors, effectors
- Implement mapping from percept sequence to actions

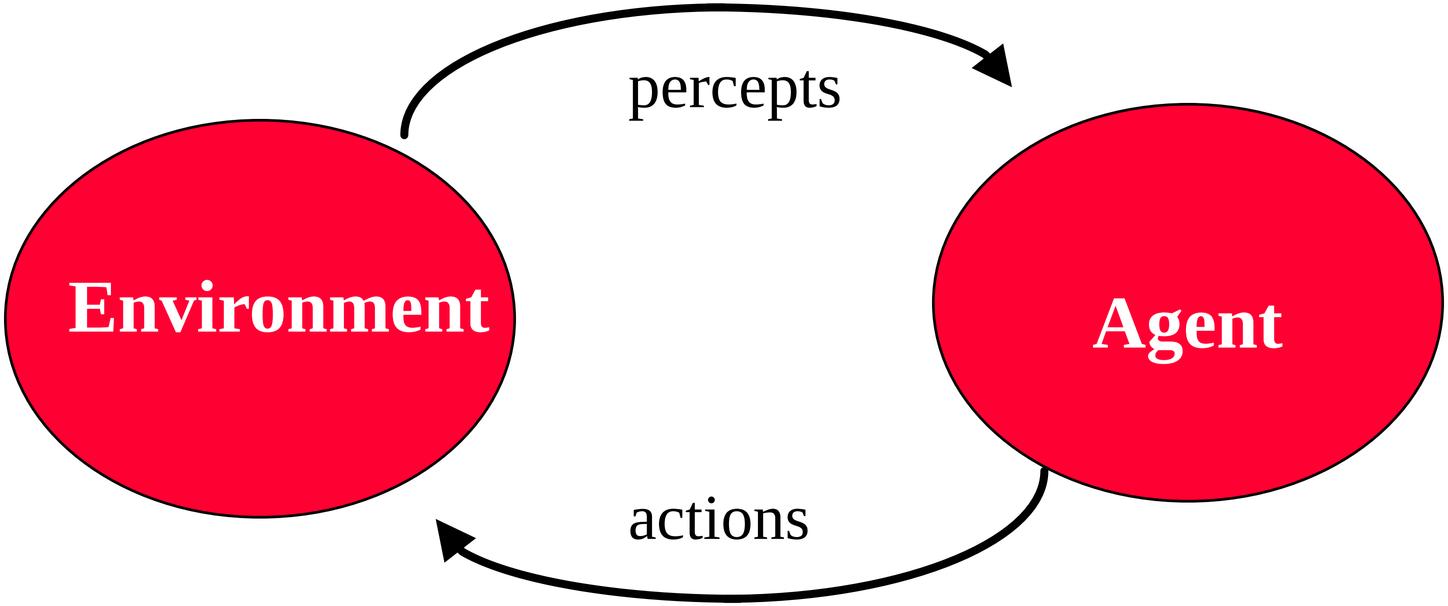

- Performance Measure

Rational Agents

- An agent should strive to do the right thing, based on what it can perceive and the actions it can perform. The right action is the one that will cause the agent to be most successful

- Performance measure: An objective criterion for success of an agent's behavior

- E.g., performance measure of a vacuum-cleaner agent could be amount of dirt cleaned up, amount of time taken, amount of electricity consumed, amount of noise generated, etc.

Ideal Rational Agent

“For each possible percept sequence, does whatever action is expected to maximize its performance measure on the basis of evidence perceived so far and built-in knowledge."

- Rationality vs omniscience?
- Acting in order to obtain valuable information

Bounded Rationality

- We have a performance measure to optimize
- Given our state of knowledge
- Choose optimal action
- ***Given limited computational resources***

PEAS: Specifying Task Environments

- PEAS: Performance measure, Environment, Actuators, Sensors

- Must first specify the setting for intelligent agent design

- Example: the task of designing an automated taxi driver:
  - Performance measure
  - Environment
  - Actuators
  - Sensors

PEAS

- Agent: Automated taxi driver

- Performance measure:
  - Safe, fast, legal, comfortable trip, maximize profits

- Environment:
  - Roads, other traffic, pedestrians, customers

- Actuators:
  - Steering wheel, accelerator, brake, signal, horn

- Sensors:
  - Cameras, sonar, speedometer, GPS, odometer, engine sensors, keyboard

PEAS

- Agent: Medical diagnosis system
- Performance measure:
  - Healthy patient, minimize costs, lawsuits
- Environment:
  - Patient, hospital, staff
- Actuators:
  - Screen display (questions, tests, diagnoses, treatments, referrals)
- Sensors:
  - (entry of symptoms, findings, patient's answers)

Properties of Environments

- Observability: full *vs.* partial *vs. non*
- Deterministic *vs.* stochastic
- Episodic *vs.* sequential
- Static *vs. Semi-dynamic vs.* dynamic
- Discrete *vs.* continuous
- Single Agent *vs.* Multi Agent (Cooperative, Competitive, Self-Interested)

RoboCup vs. Chess

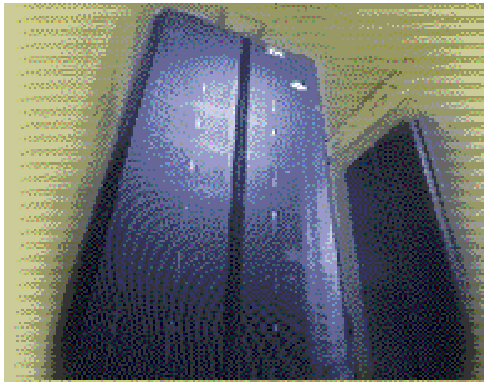

Deep Blue

Robot

- Static/Semi-dynamic
- Deterministic
- Observable
- Discrete
- Sequential
- Multi-Agent

- Dynamic
- Stochastic
- Partially observable
- Continuous
- Sequential
- Multi-Agent

More Examples

- Classical Planning
  - Static – Deterministic – Fully Obs – Discrete – Seq – Single
- Poker
  - Static – Stochastic – Partially Obs – Discrete – Seq – Multi-agent
- Medical Diagnosis
  - Dynamic – Stochastic – Partially Obs – Continuous – Seq – Single
- Taxi Driving
  - Dynamic – Stochastic – Partially Obs – Continuous – Seq – Multi-agent

Simple reflex agents

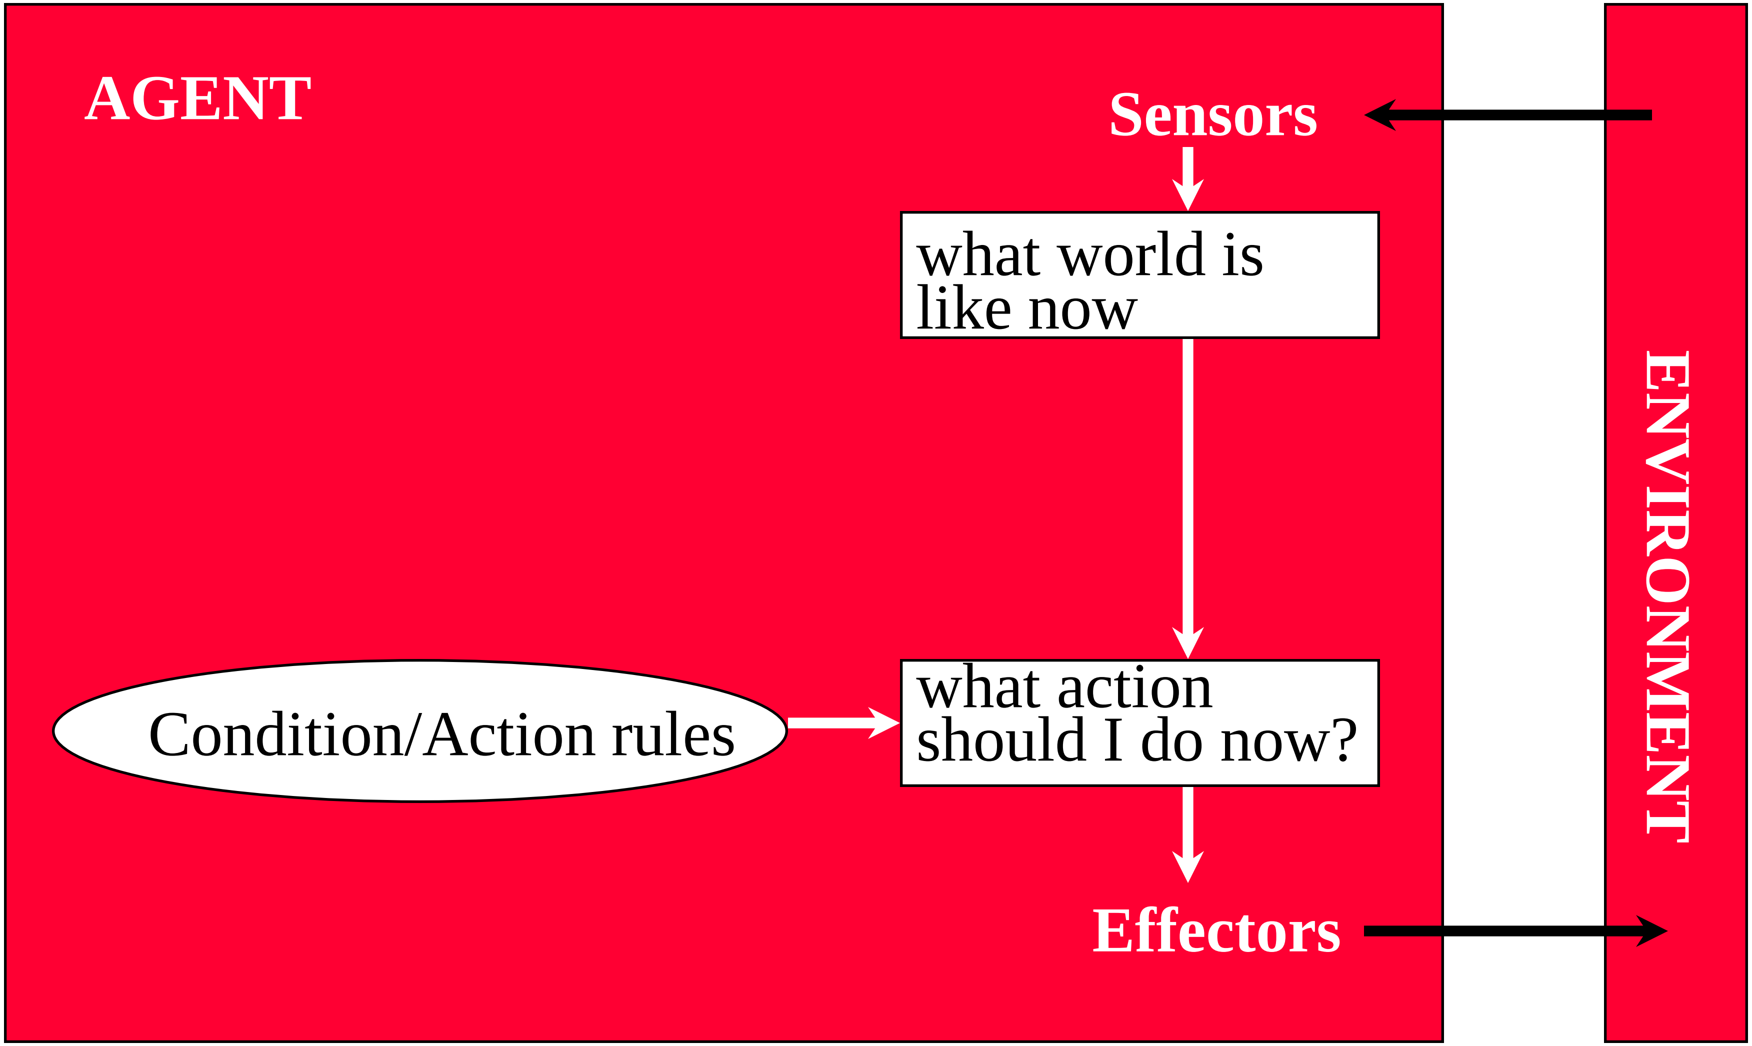

Reflex agent with internal state

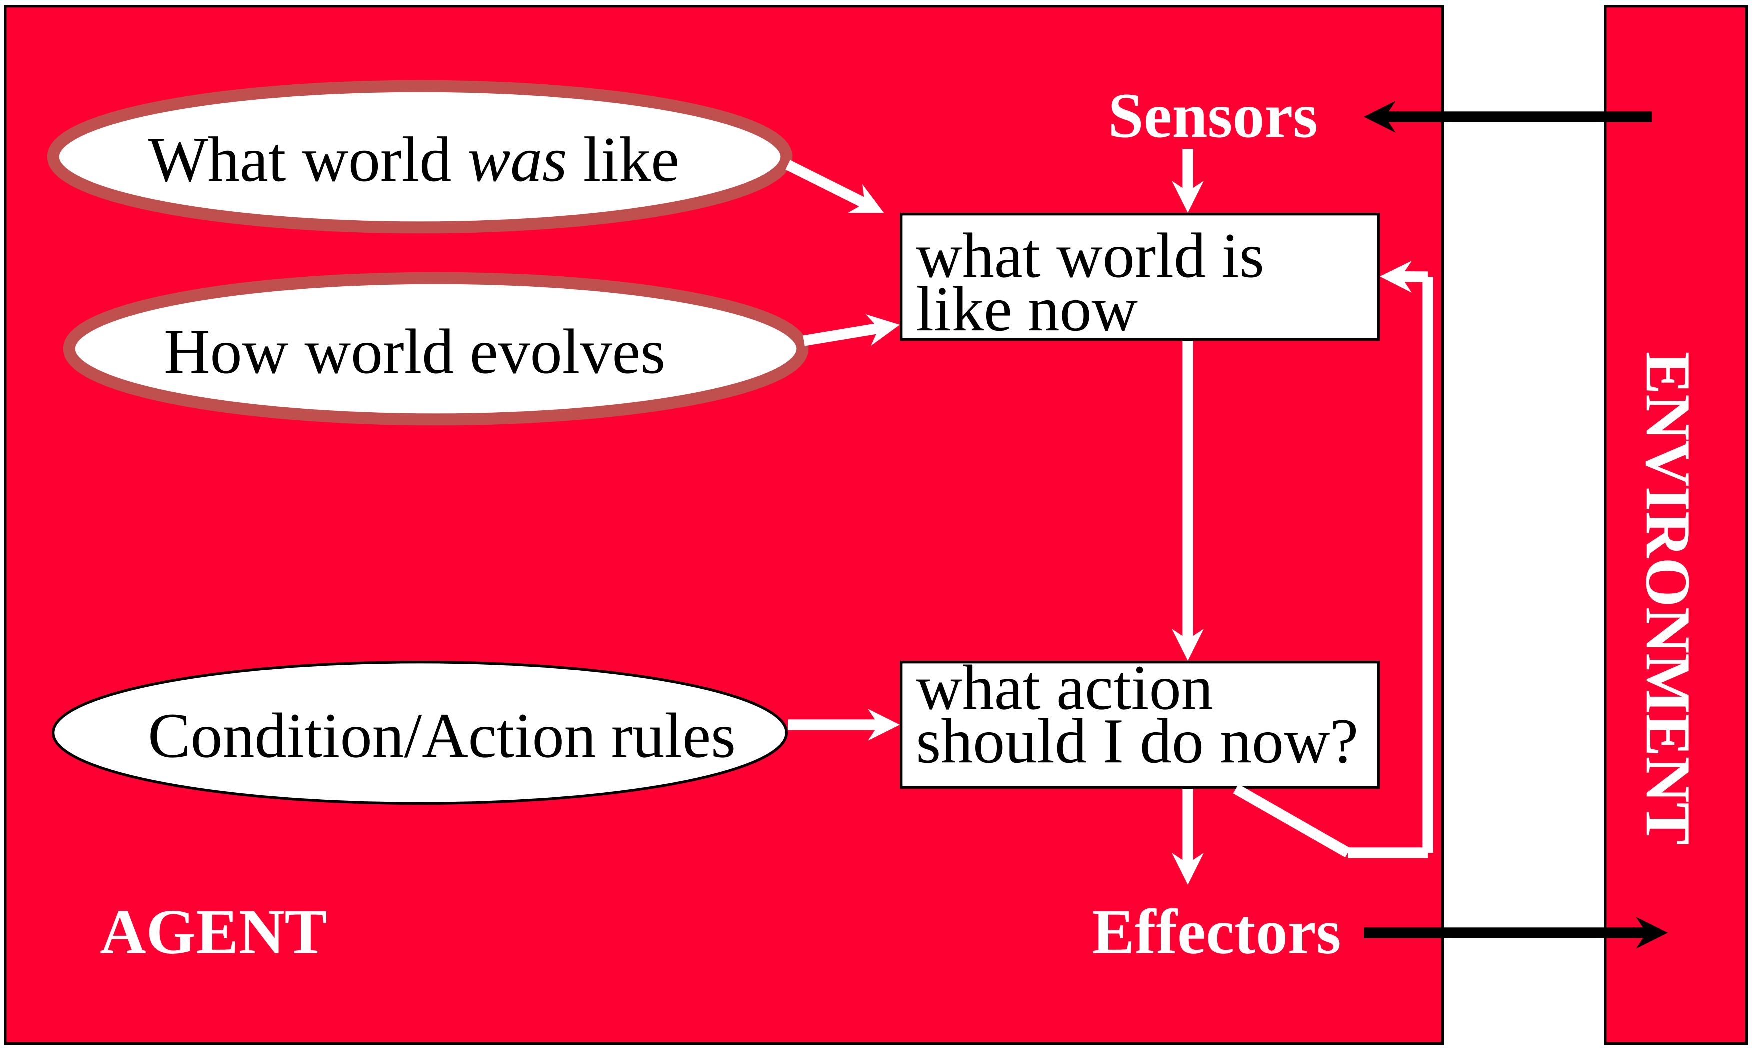

Goal-based agents

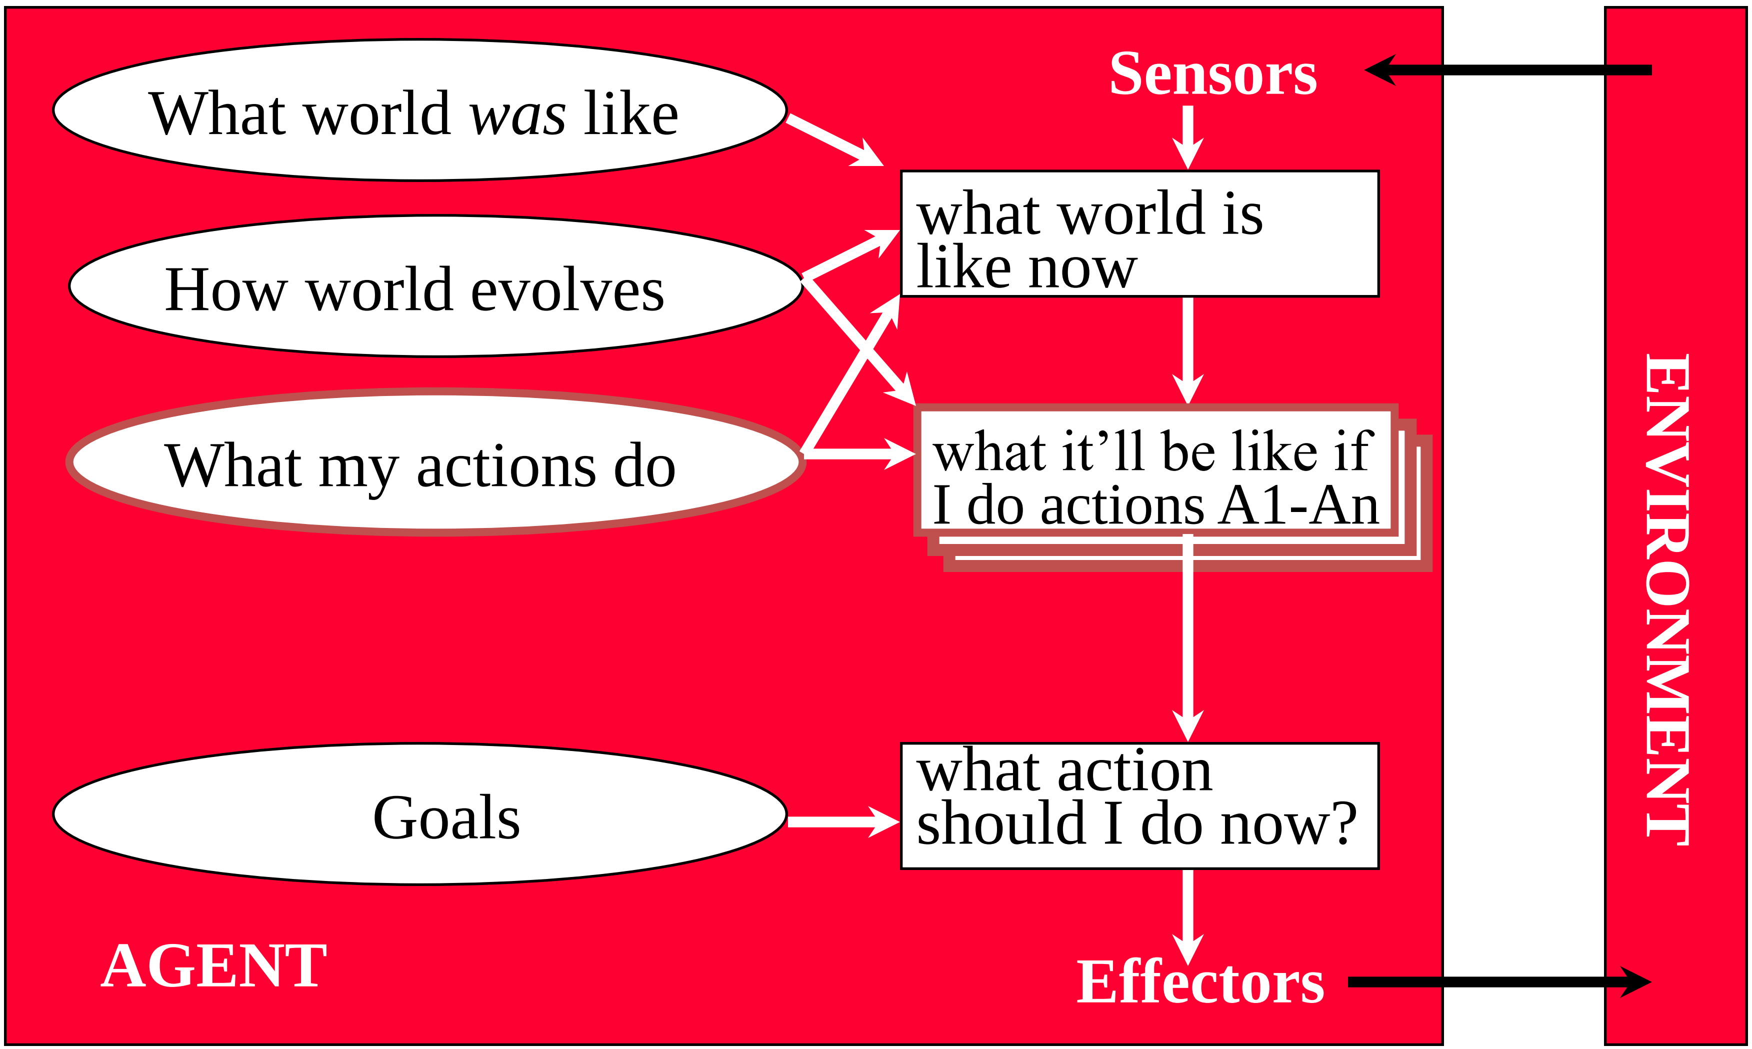

Utility-based agents

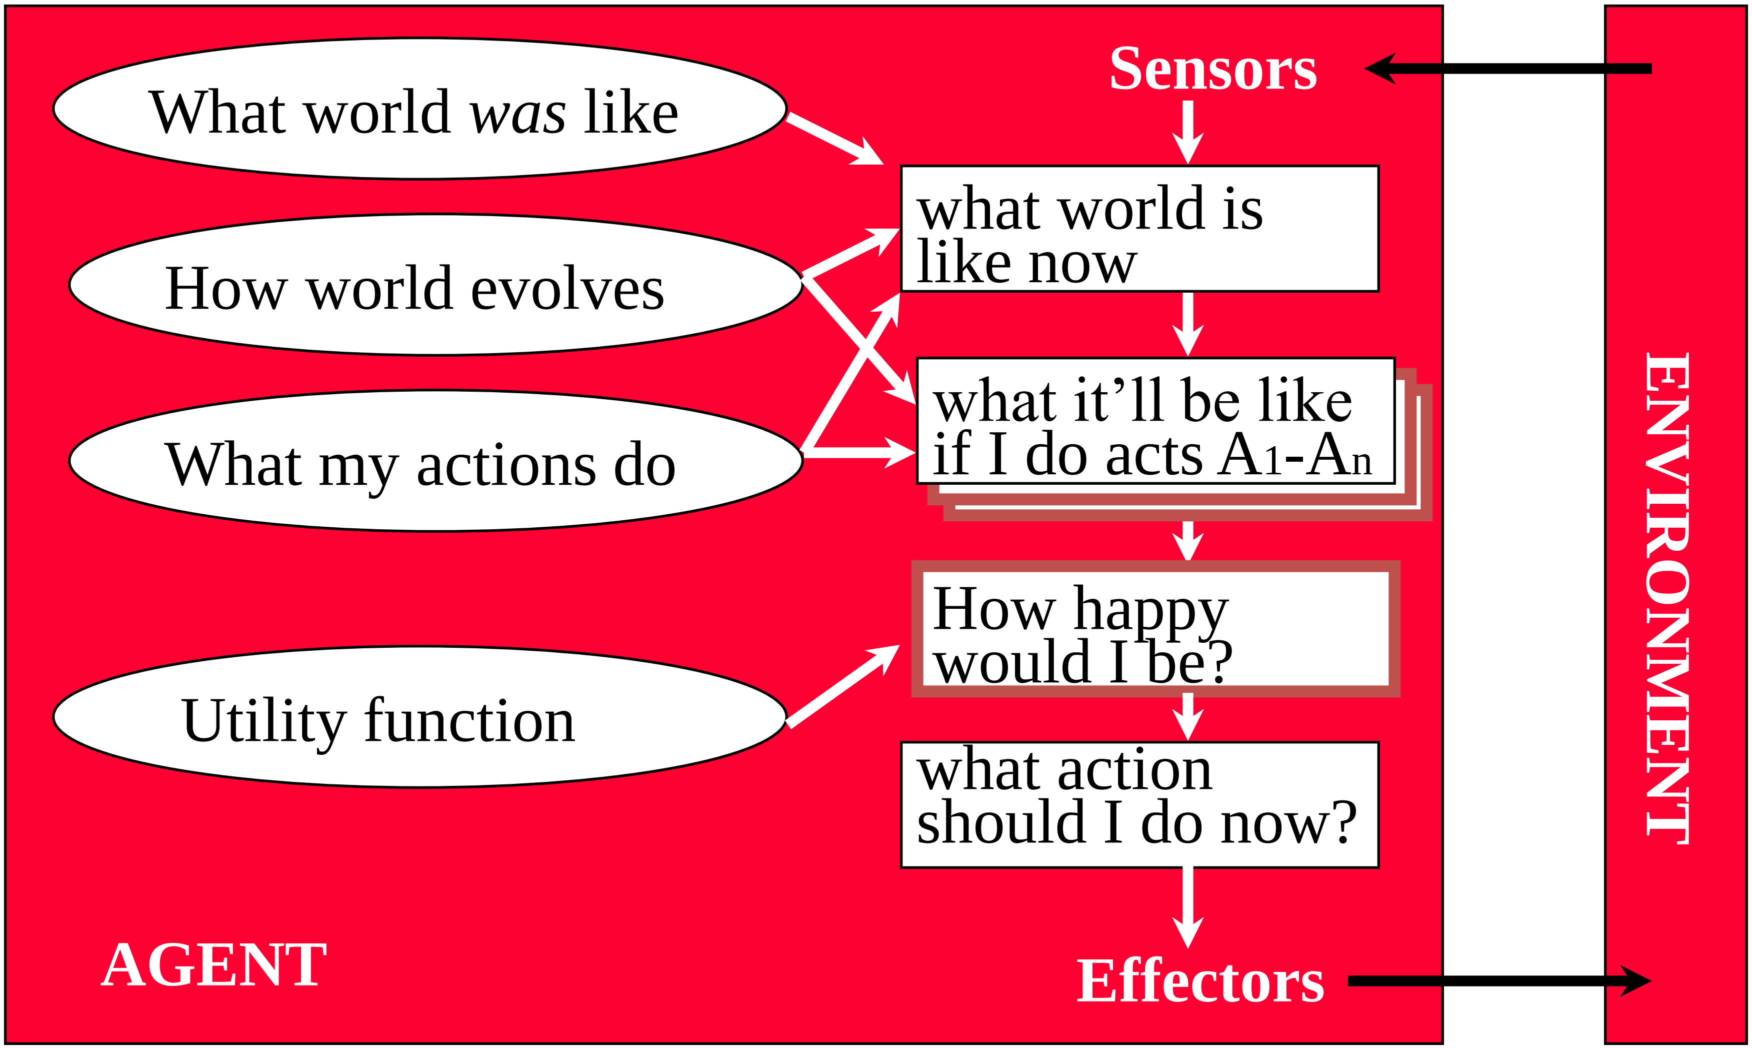

Learning agents

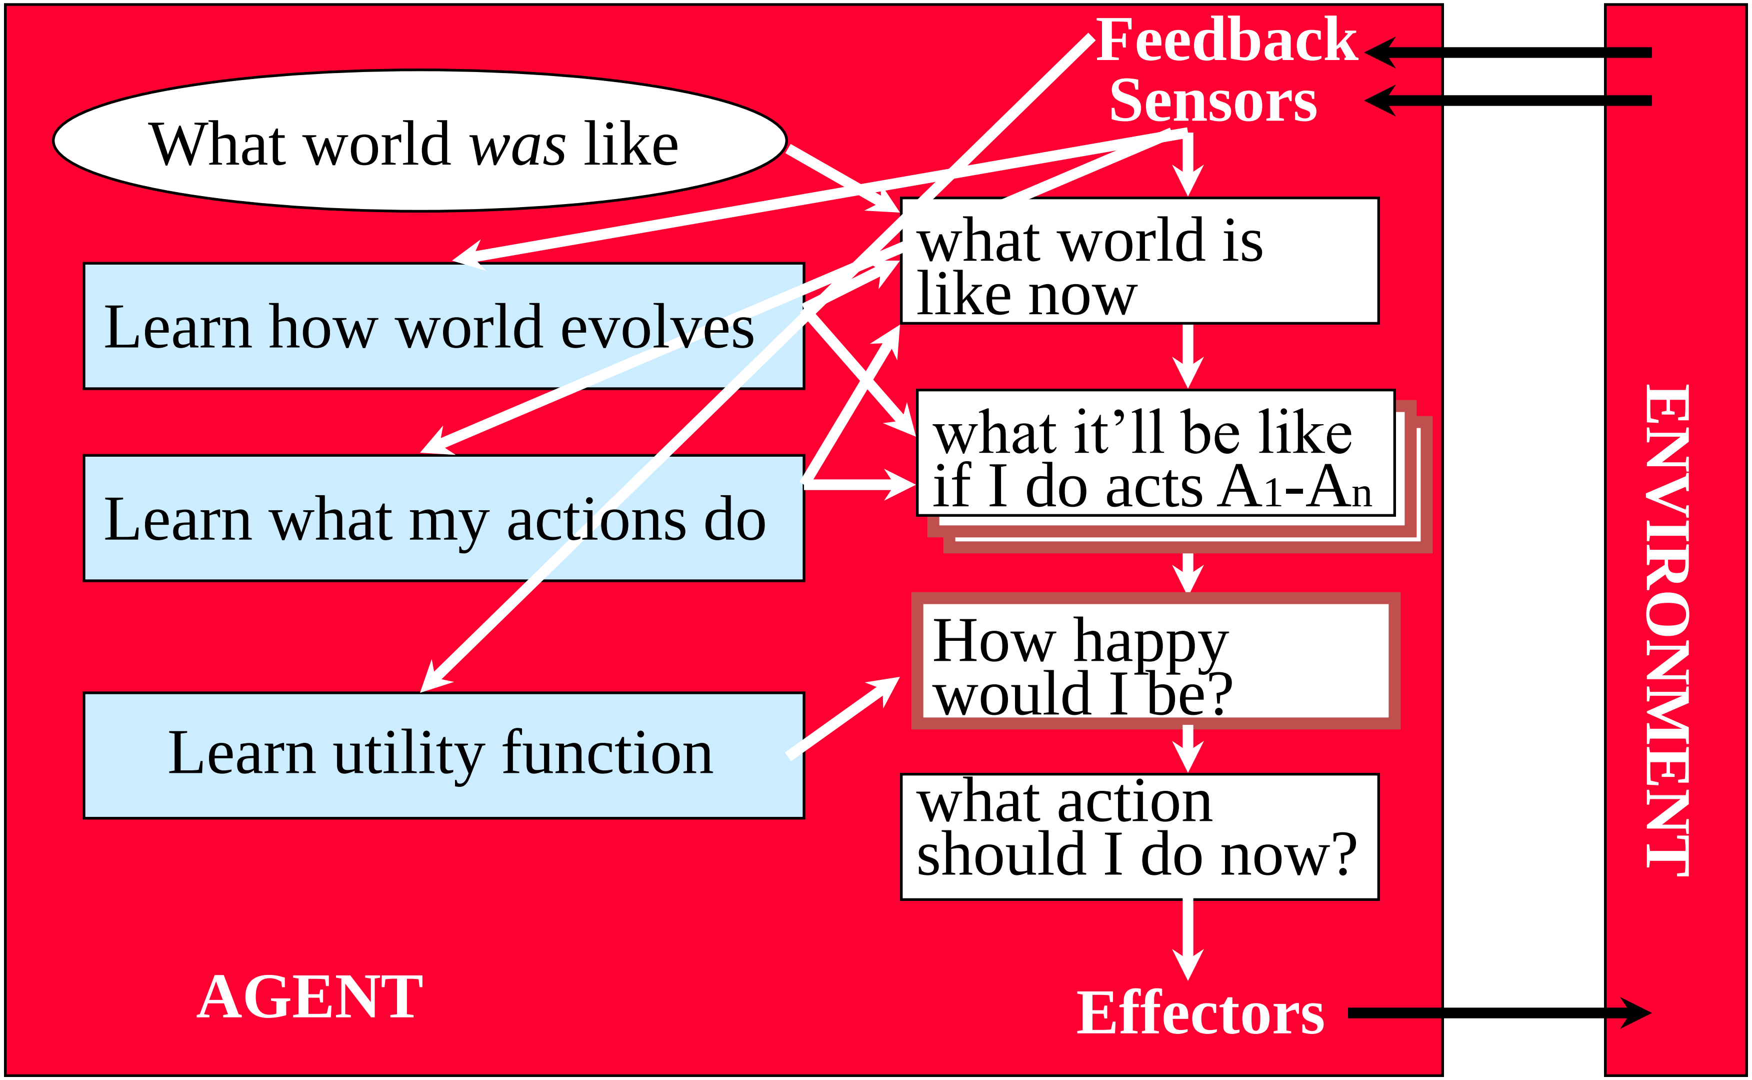

Wrap Up

Mausam

What is intelligence?

- Think like humans
- Act like humans
- Think (bounded) rationally
- Act (bounded) rationally
  - Agent architecture

AI is SEARCH!

- 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

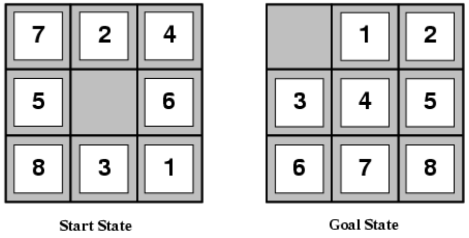

  - Local: fast, fewer guarantees; often local optimal
    - Hill climbing and variations
    - Simulated Annealing: global optimal
    - Genetic algorithms: somewhat non-local due to crossing over
  - (Local) Beam Search

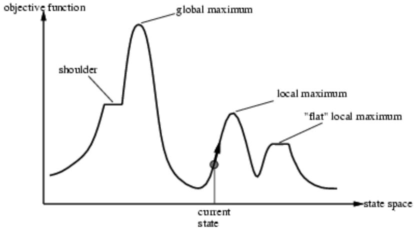

Search Example: Game Playing

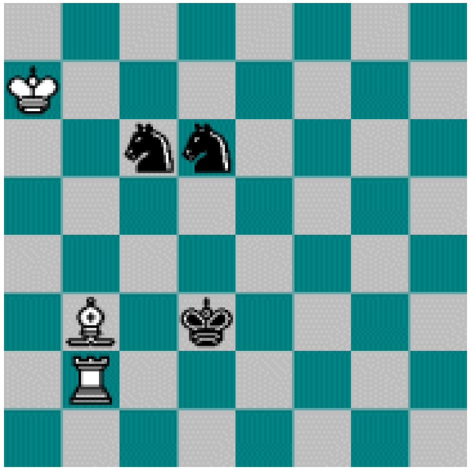

- Game Playing
  - AND/OR search space (max, min)
  - minimax objective function
  - minimax algorithm (~dfs)
    - alpha-beta pruning
  - Utility function for partial search
    - Learning utility functions by playing with itself
  - Openings/Endgame databases
    - Secondary search/Quiescence search

AI is KR&R

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

- CSP
- Logic

- Bayes Nets
- Markov Decision Process

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

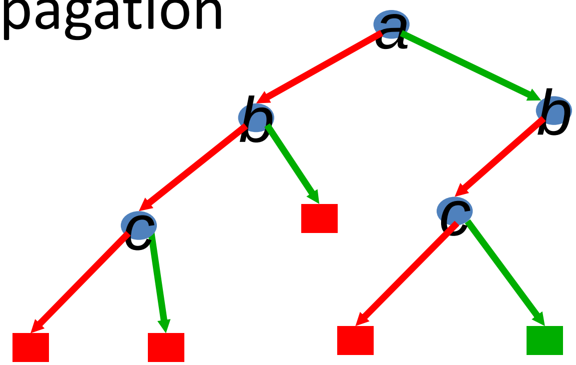

- Advances
  - Clause Learning: learning from mistakes
  - Restarts in systematic search
  - Portfolio of SAT solvers; Parameter tuning

- Phase Transitions in SAT problems

Search+KR&R Example: Planning

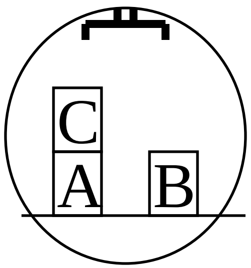

- Representation: STRIPS
- Reasoning: Planning Graph
  - Polynomial data structure
  - reasons about constraints on plans (mutual exclusion)
- Search
  - Forward: state space search
    - planning graph based heuristic
  - Backward: subgoal space search
  - Local: FF (enforced hill climbing)
- Planning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

KR&R: Probability

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

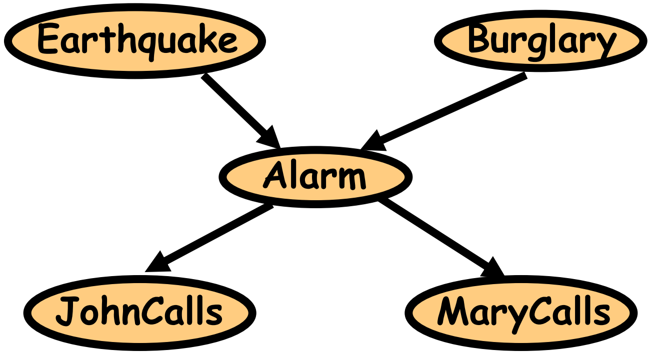

- Reasoning
  - Exact inference: var elimination
  - Approx inference: sampling based methods
    - rejection sampling, likelihood weighting, Gibbs sampling

KR&R: 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*)

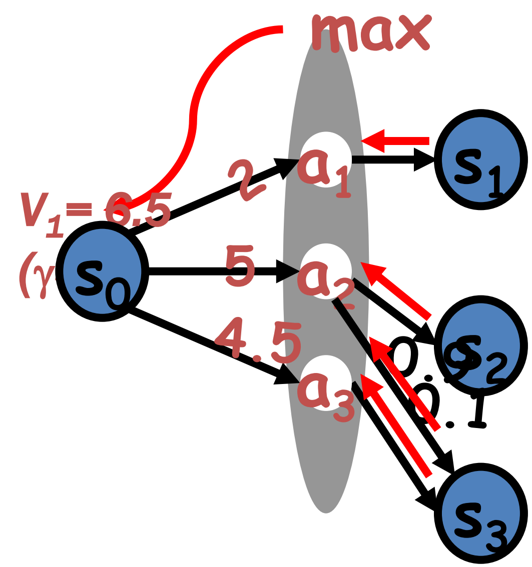

KR&R: Probability

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

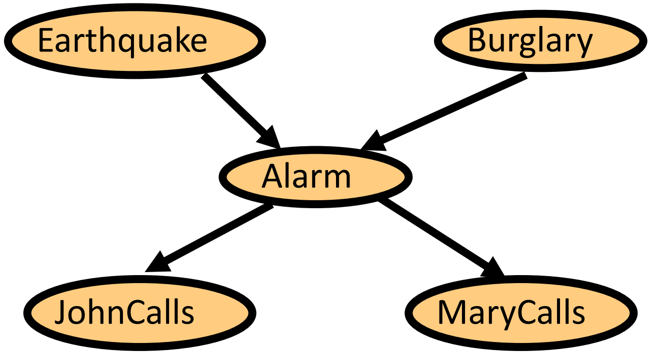

- Reasoning
  - Exact inference: var elimination
  - Approx inference: sampling based methods
    - rejection sampling, likelihood weighting, Gibbs sampling

KR&R: 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
- Approximate Reasoning: monte-carlo planning
  - Policy Rollouts
  - UCT
- State space search
  - LAO* (AND/OR version of A*)

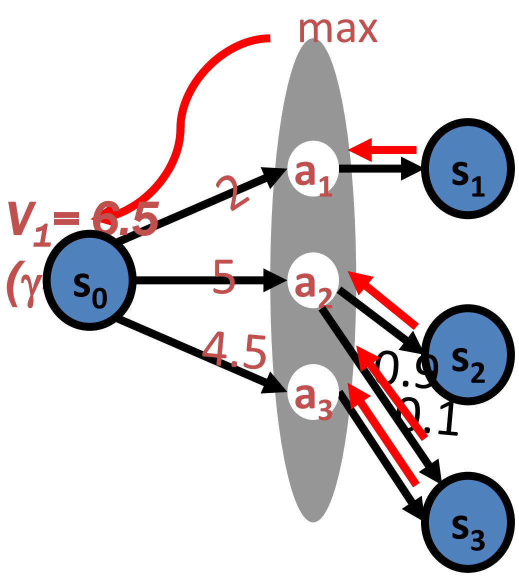

Learning: Bayes Nets

- Learn probabilities from data
- 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

Learning: Reinforcement Learning

- Learning while acting
  - take actions that aid in learning

- Model based vs. Model free learning
- Temporal Difference Learning
  - Q learning
- Exploration/Exploitation tradeoff

Popular Themes

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

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

- Combinatorial auctions
- Subgraph isomorphism
- Blackjack

General Purpose Tools

- CSP solvers
- SAT solvers
- Bayes Net packages

- Algorithm implementations

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
- Sampling based approaches
  - Sampling in systematic search
  - 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

- Temporal models: HMMs, Kalman filters…
- Ontologies
- Robotics
- Vision
- Text processing
- AI for Web
- Mechanism design
- Multi-agent systems
- Sensor Networks
- Computational Neuroscience
- …

AI is about problems.

- It is an application-driven field

- Happy to beg, borrow, steal ideas from anywhere

- Traditionally discrete … more and more cont.
- Traditionally logic… almost all probability
  - Recent close connections with EE/Stat due to ML

- HUGE field

Applications of AI

- Mars rover: planning
- Jeopardy: NLP, info retrieval, machine learning
- Puzzles: search, CSP, logic
- Chess: search
- Web search: IR
- 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 ☺