

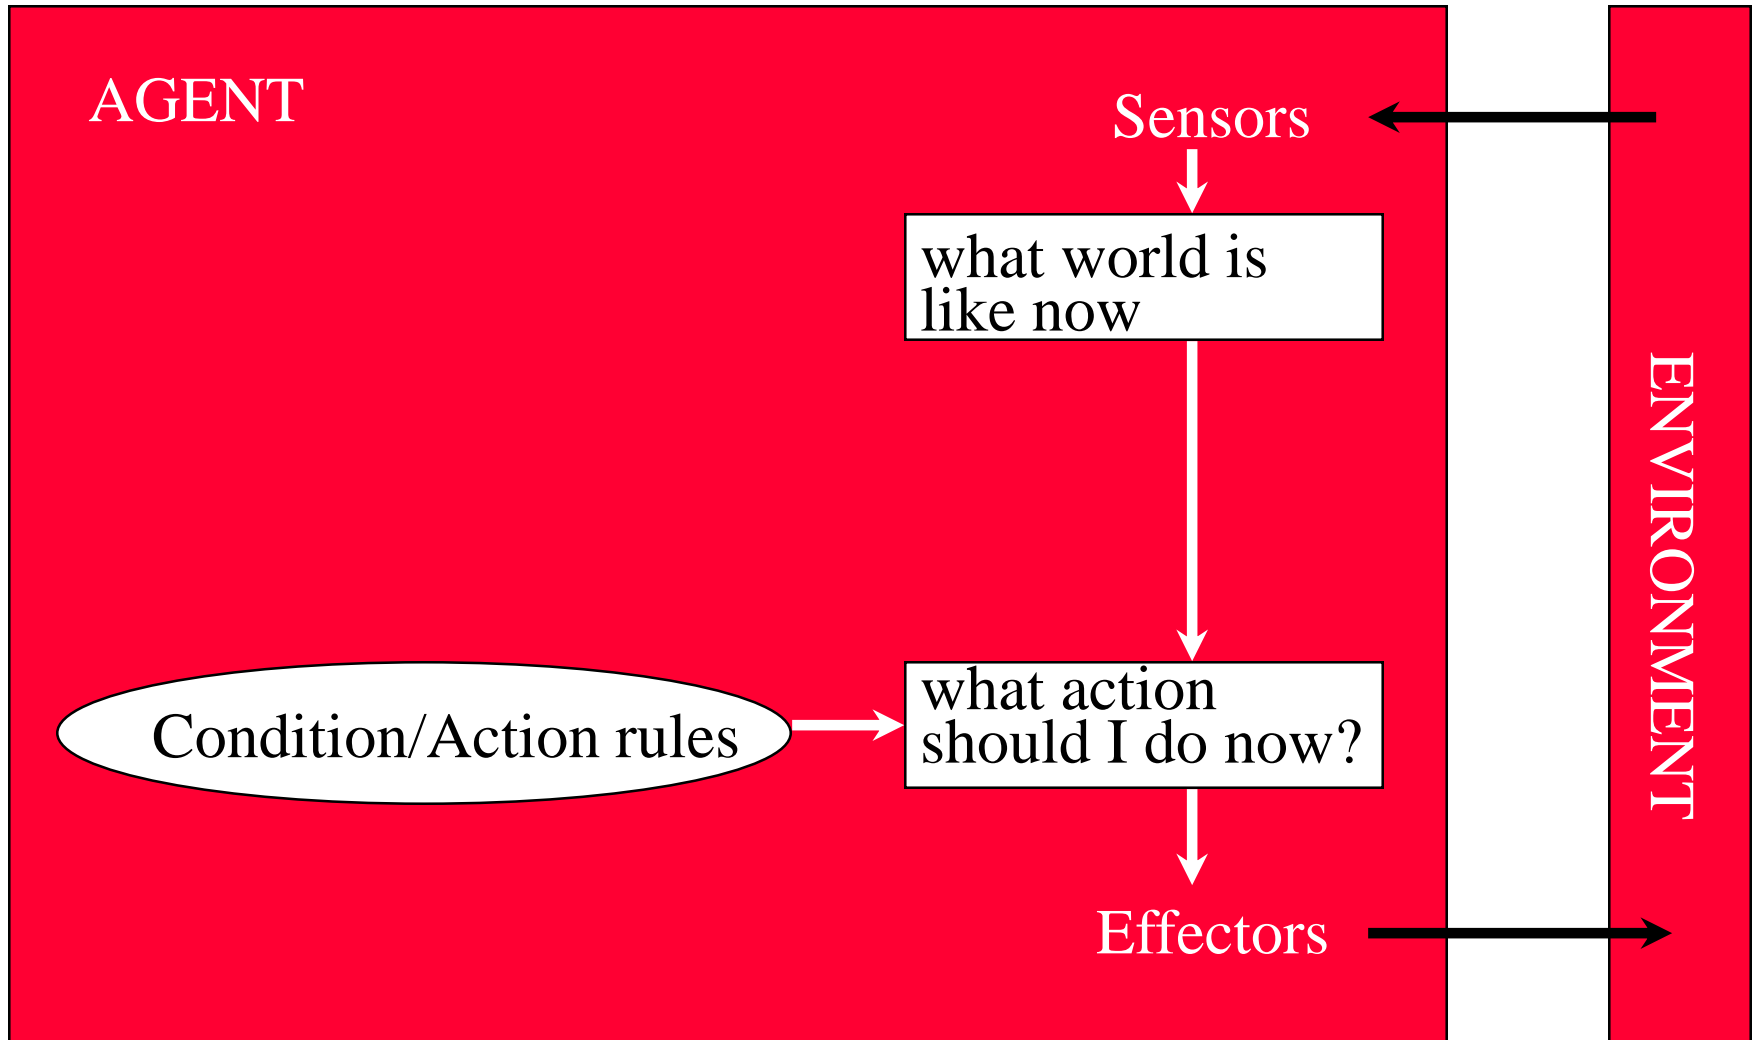
Artificial Intelligence Recap

Mausam

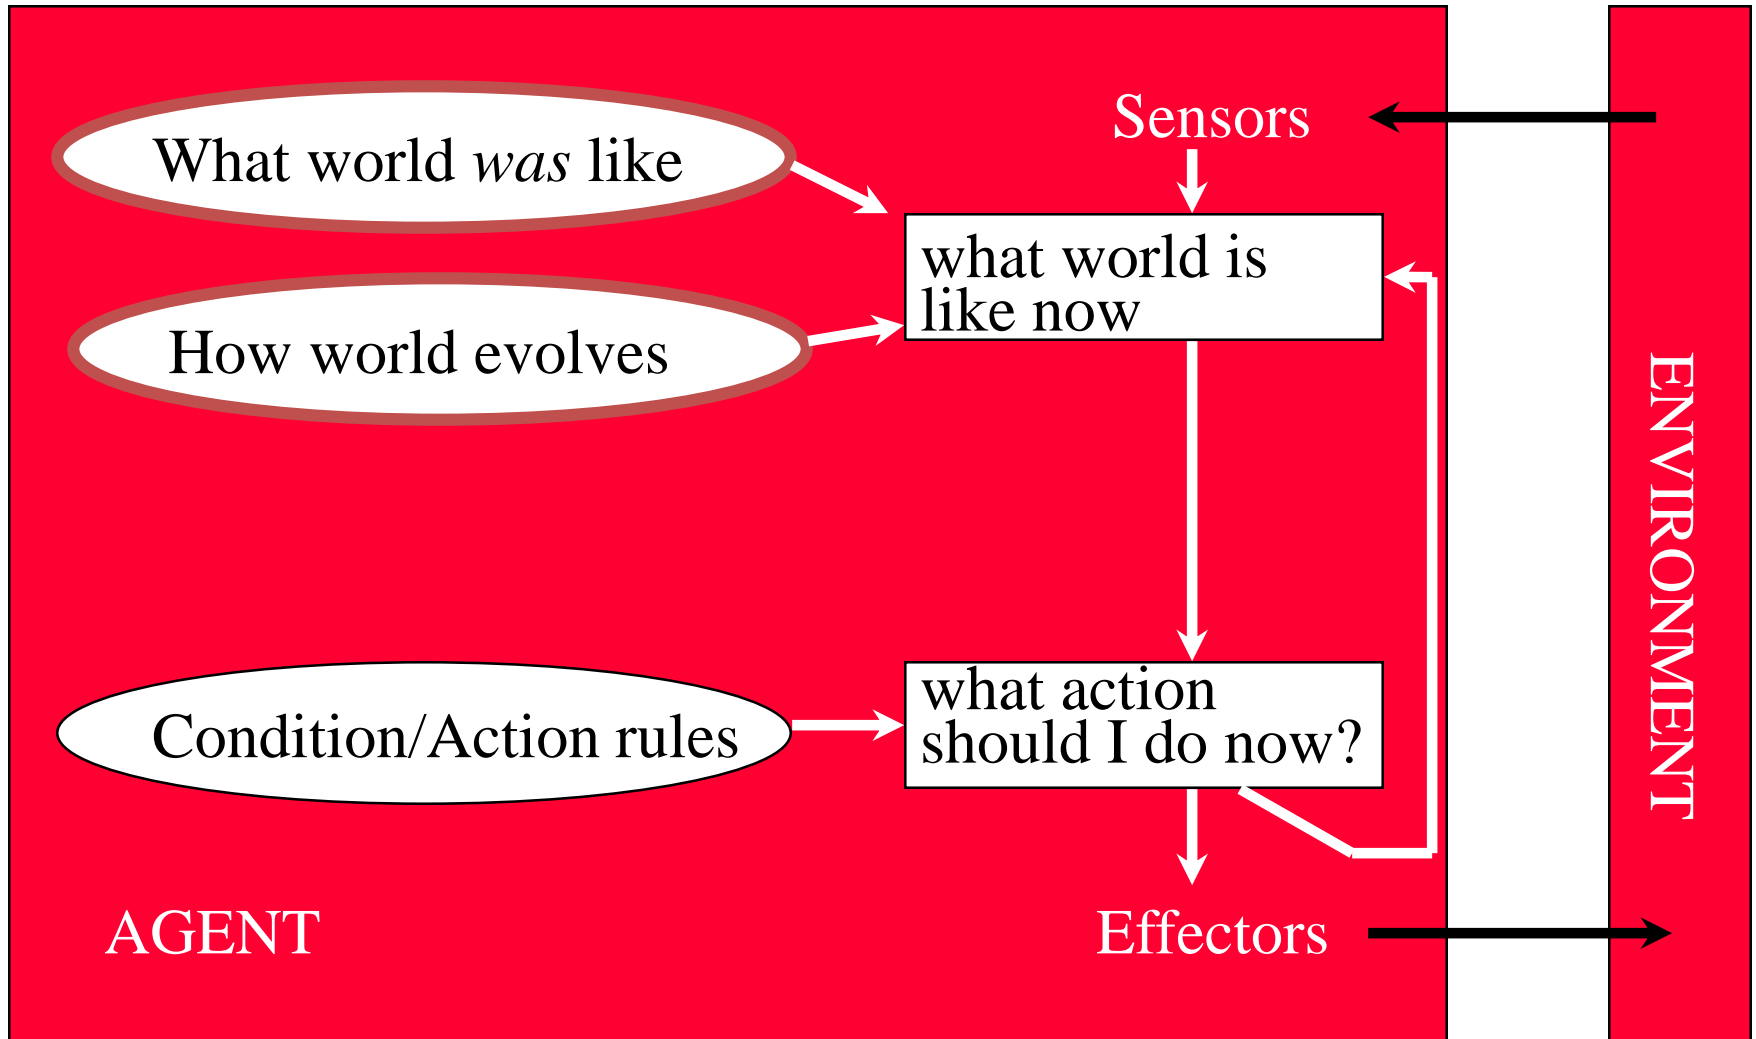
What is intelligence?

- (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

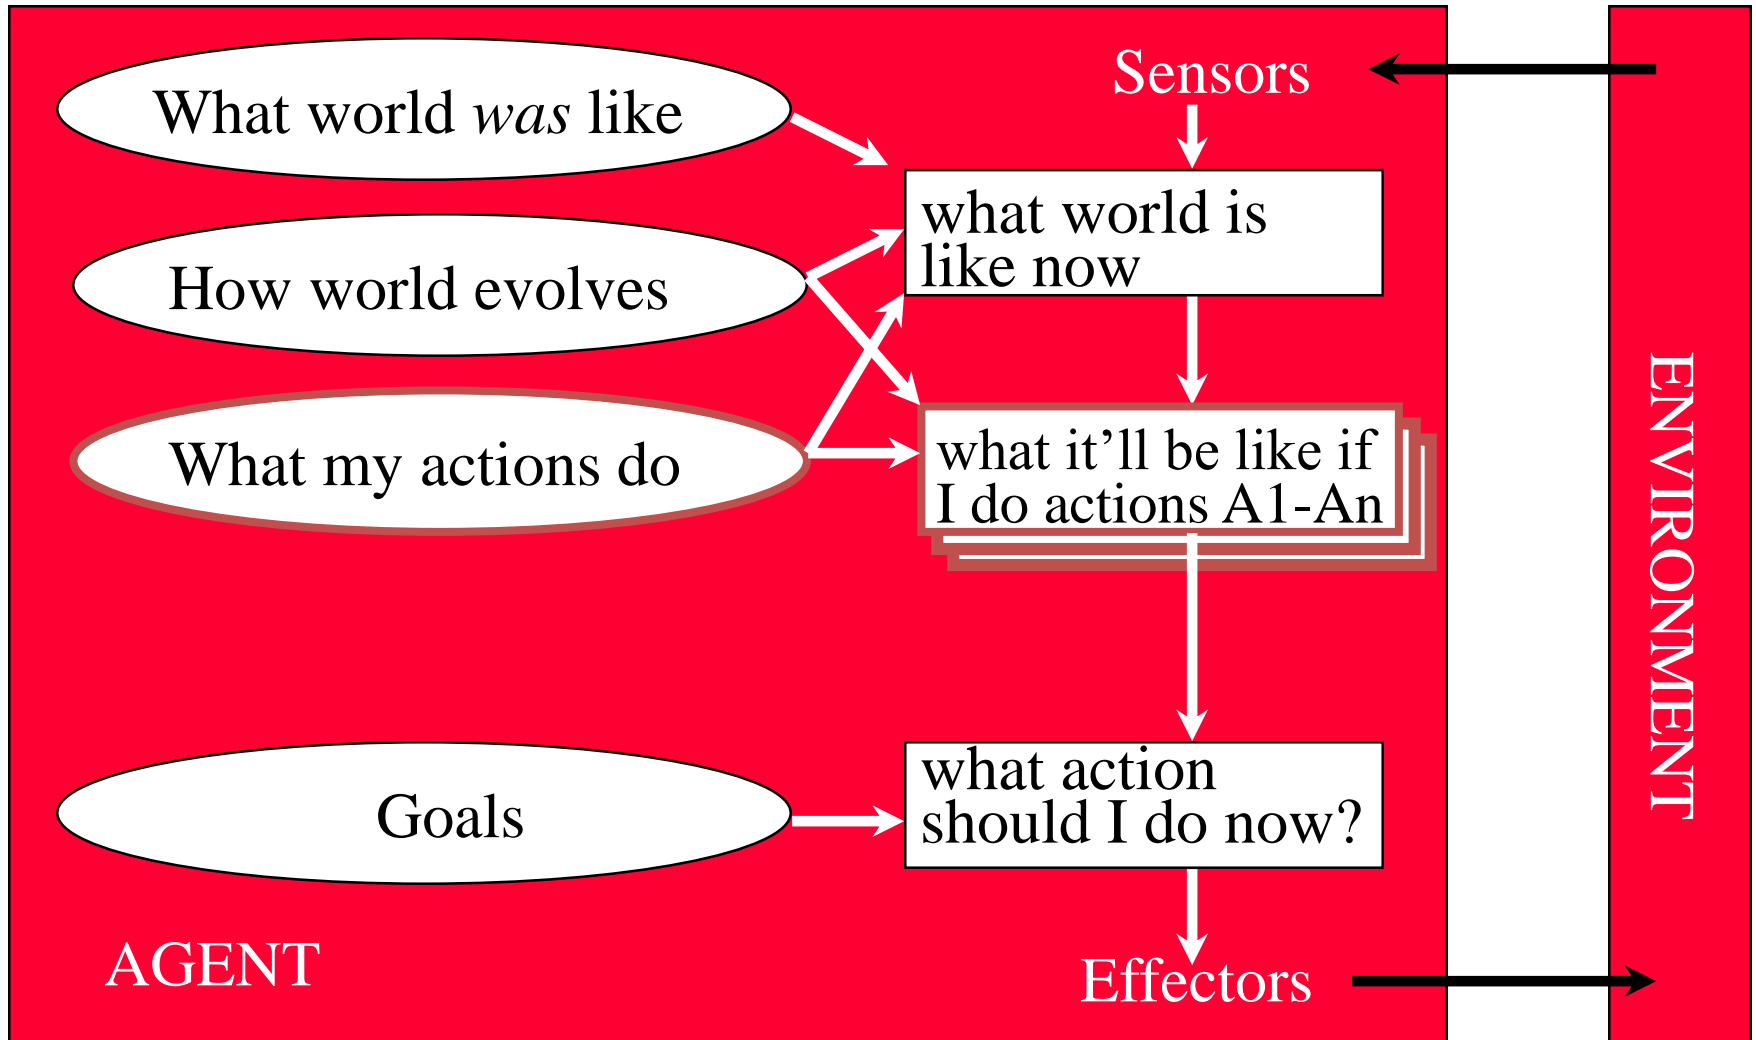
Simple reflex agents



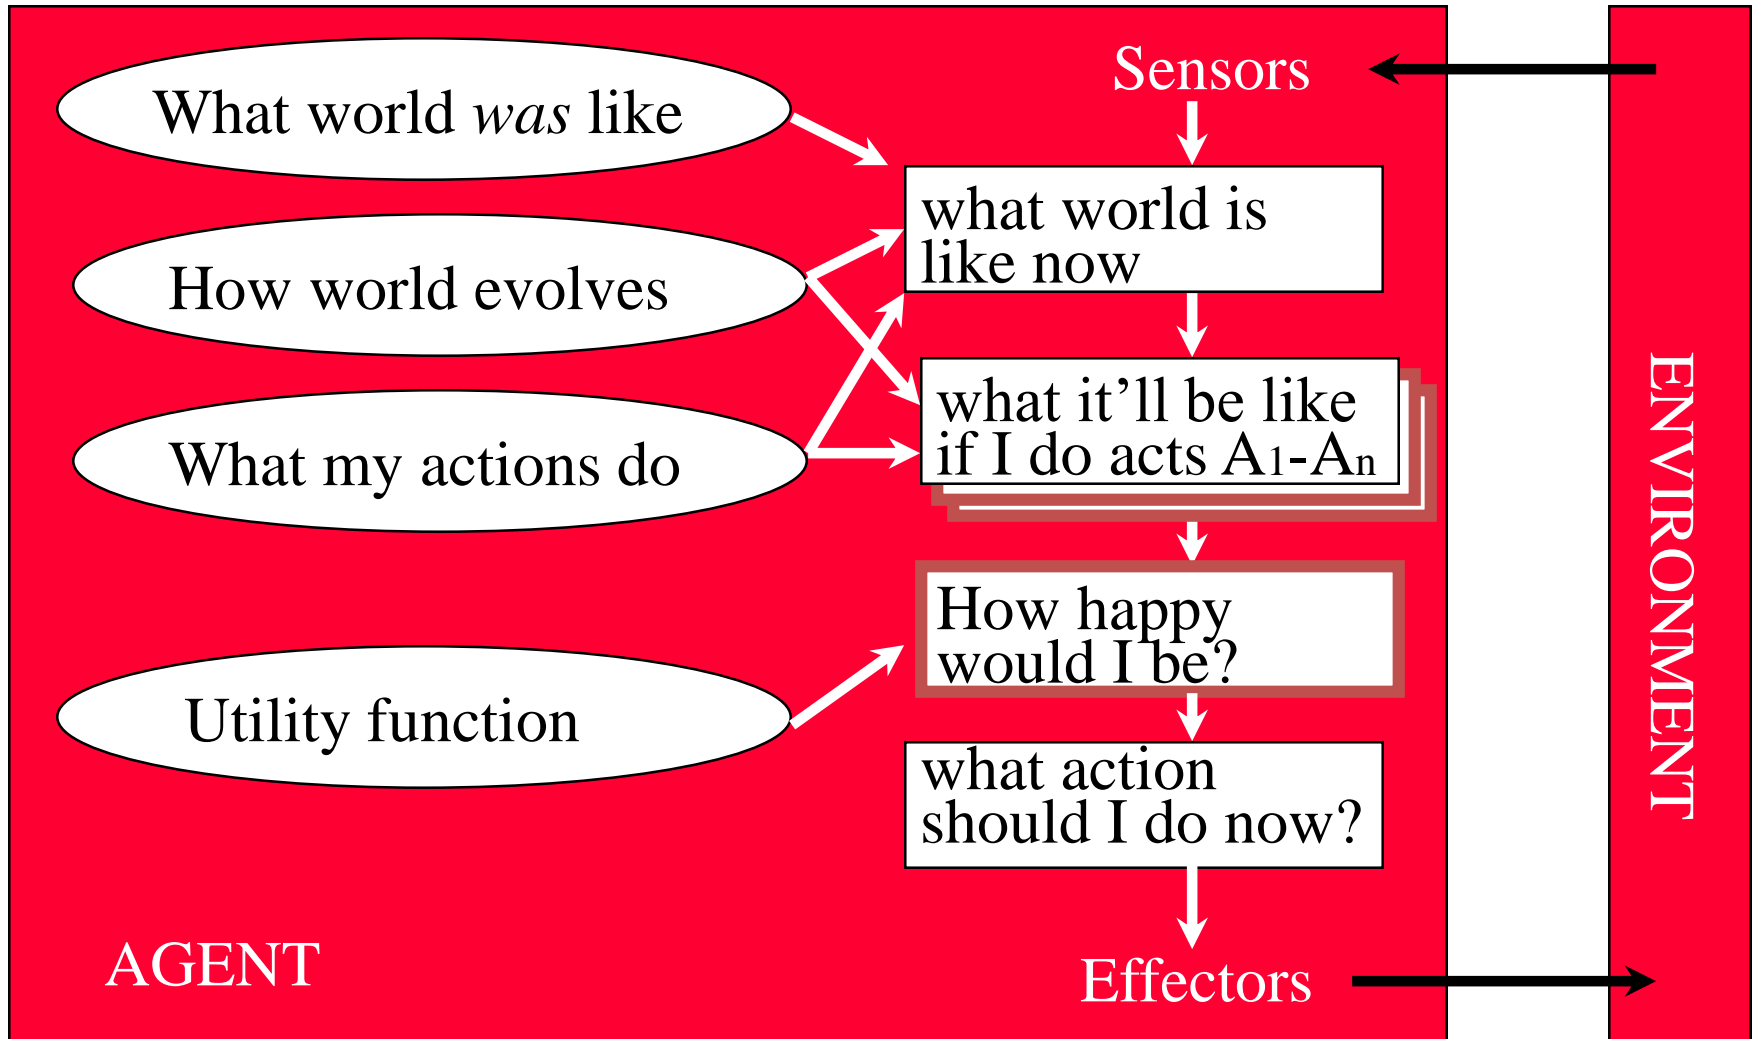
Reflex agent with internal state



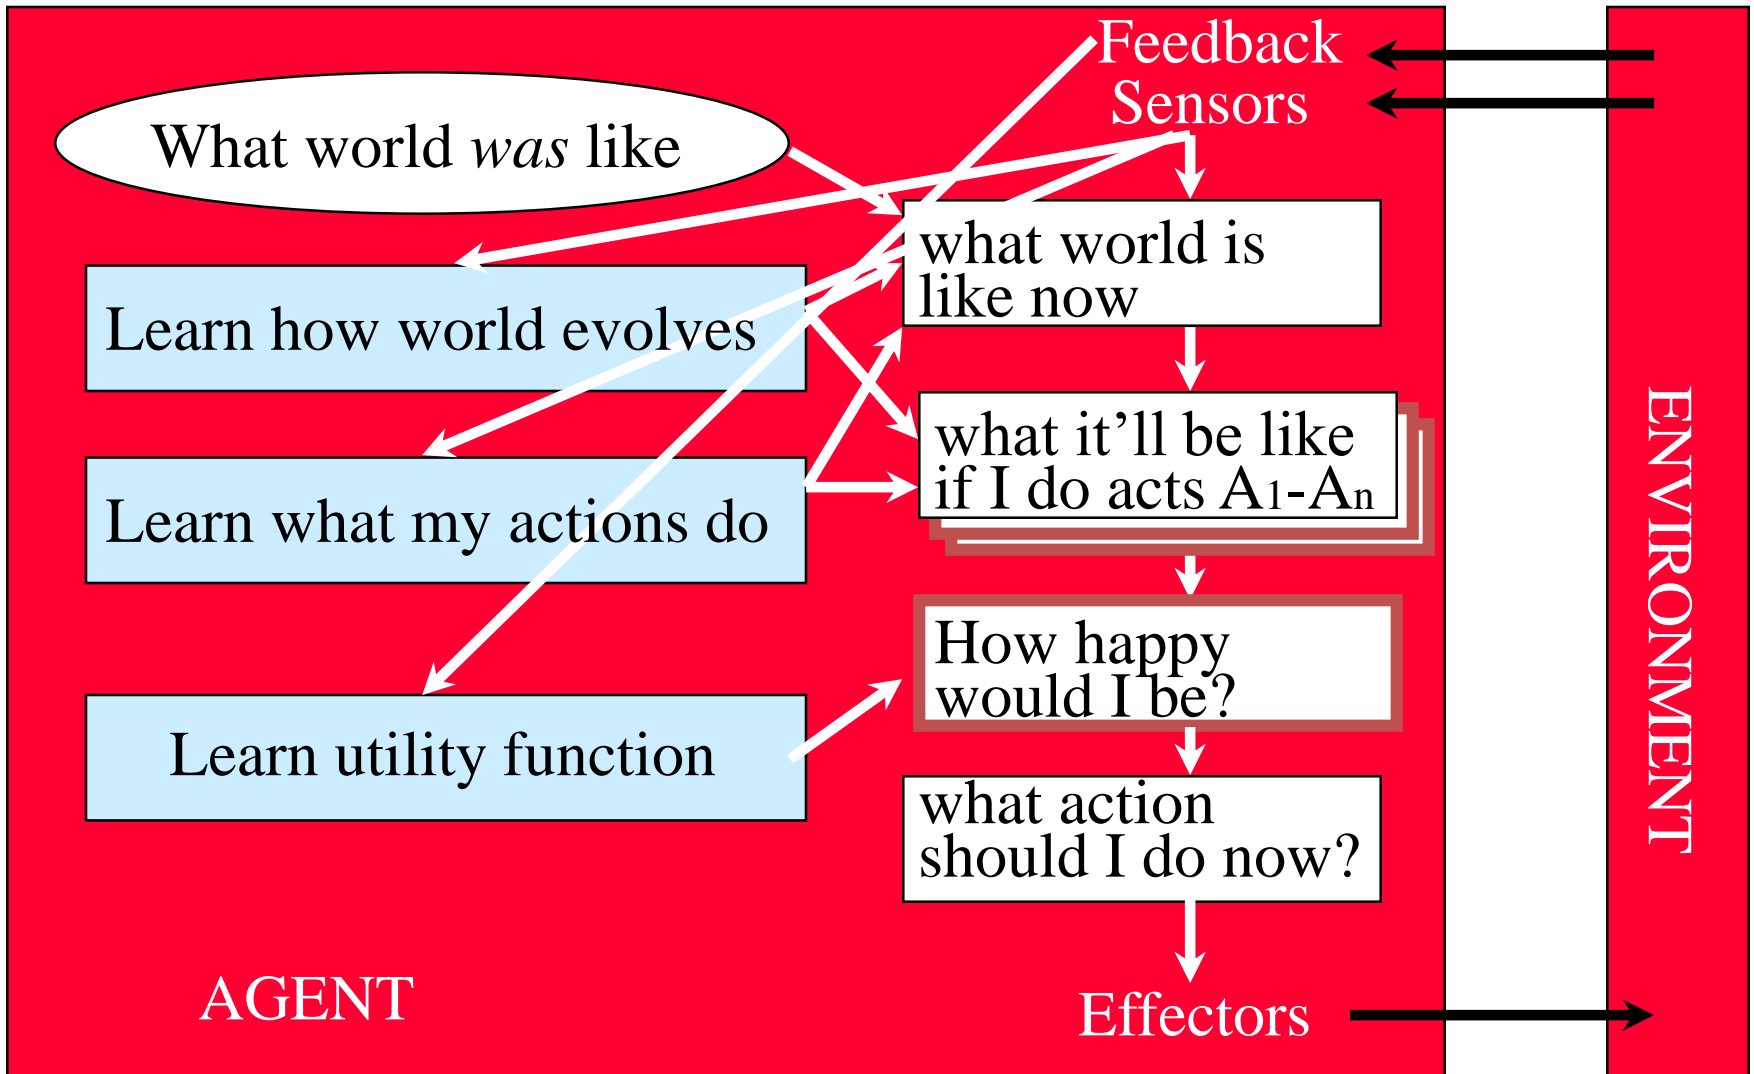
Goal-based agents



Utility-based agents



Learning agents



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

- Logic

- PDDL

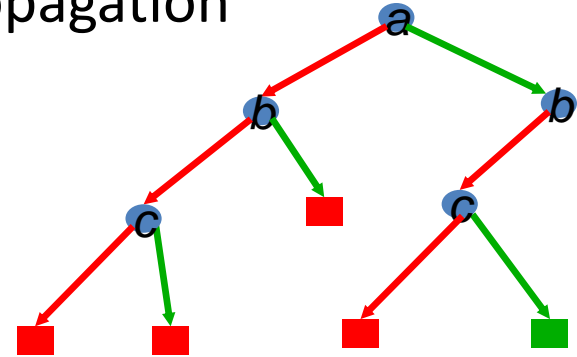
- Bayes Nets

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: 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

Actions	States of Nature	
	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**

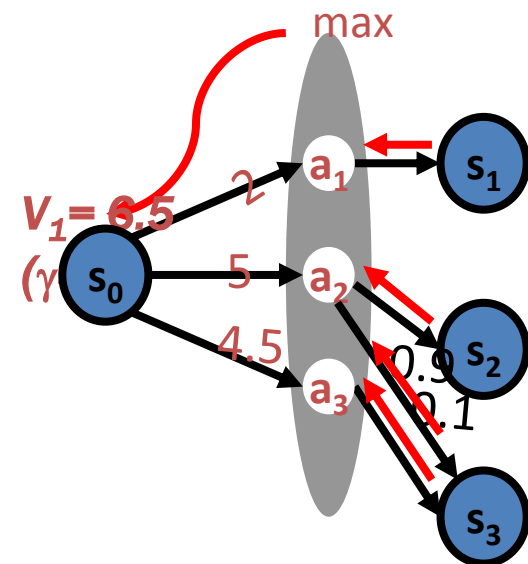
- states, actions, probabilistic outcomes, rewards
- ~AND/OR Graph (sum, max)

- **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/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

Learning: Reinforcement Learning

- Learn model while taking actions
- What to learn
 - T and R: model based
 - Policy: Model free
- Which actions to take
 - Exploration - Exploitation

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
 - $A \wedge B$: 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
- 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 😊

Graphics

Algorithms
Theory

Databases

Operations
Research

AI

Statistics

Linguistics

Robot
Design

Psychology
Neurosc.

...