Intelligent Agents

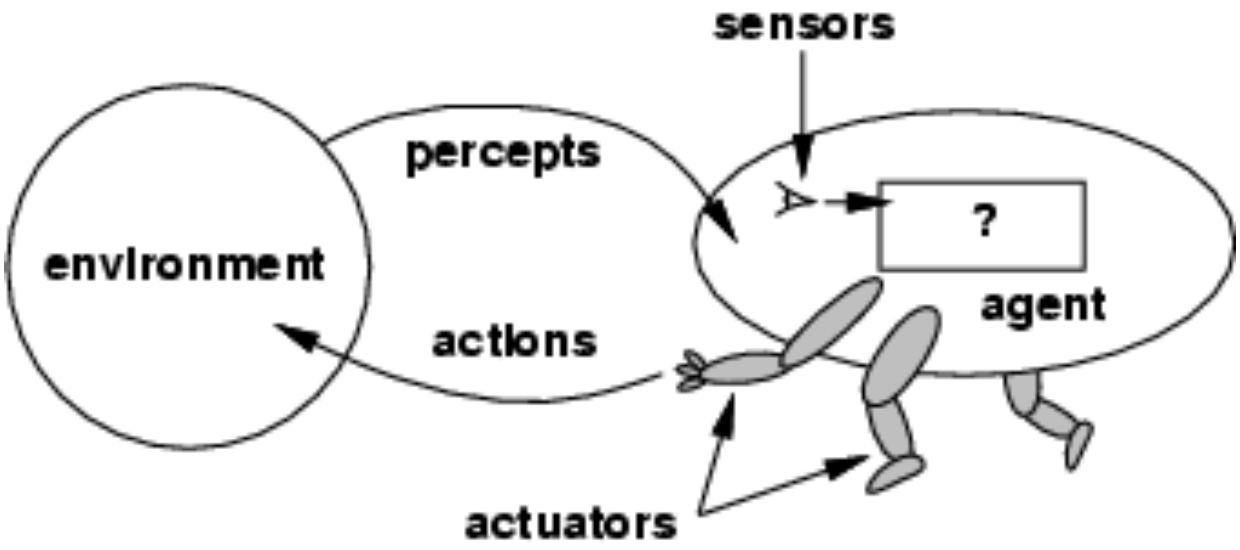

What is an agent?

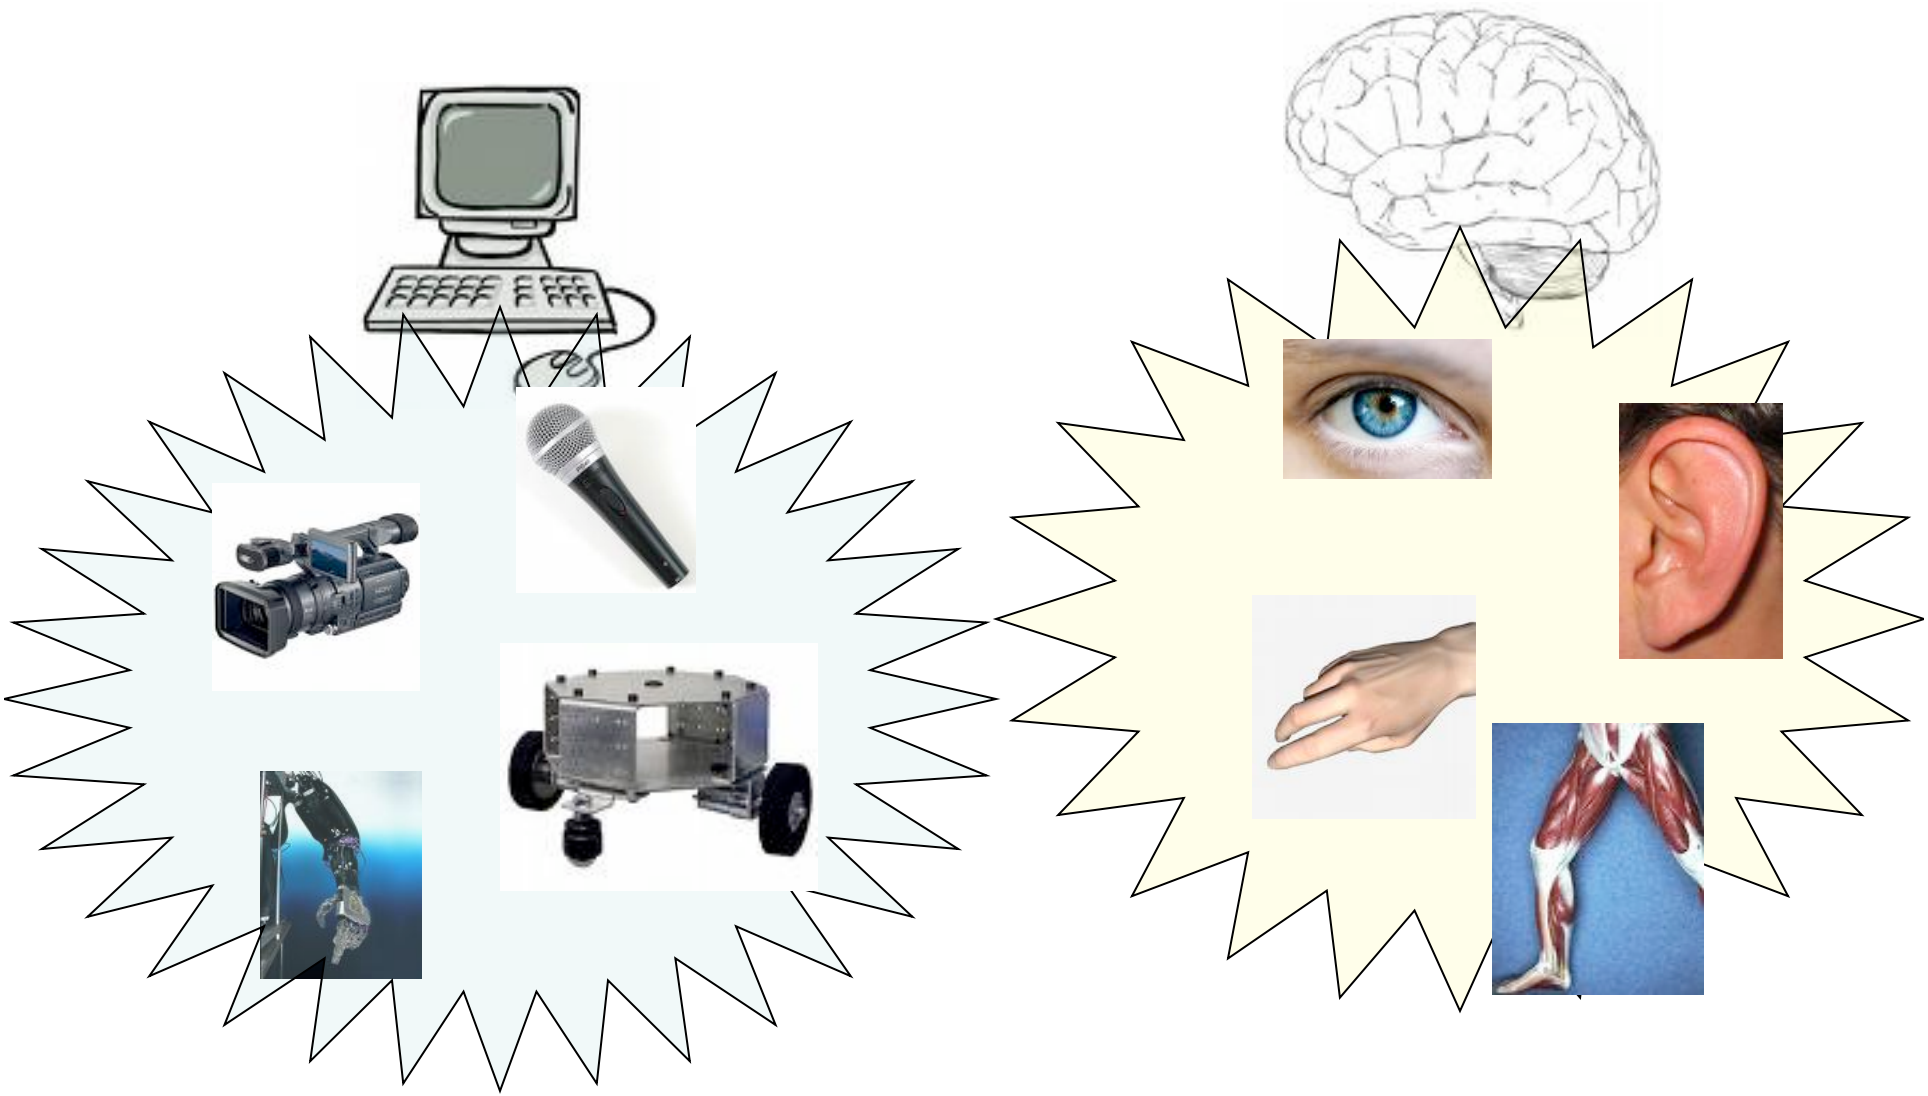

Vacuum-cleaner world

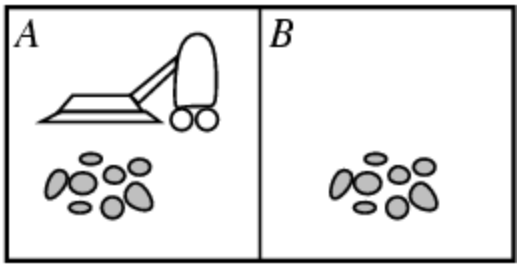

- percepts: location and contents
  e.g., [A,Dirty]
- actions: *Left*, *Right*, *Suck*, *NoOp*

The Lowly Dung Beetle

After digging its nest and laying its eggs, the dung beetle fetches a ball of dung from a nearby en route heap to plug the entrance; if the ball of dung is removed from its grasp, the beetle continues on and pantomimes plugging the nest with the nonexistent dung ball.

from Hanski & Cambefort, 1991

The Sphex Wasp

The female will dig a burrow, go out and sting a caterpillar and drag it to the burrow, enter the burrow again to check all is well, drag the caterpillar inside, and lay its eggs… but if an entomologist moves the caterpillar a few inches away while the sphex is doing the check, it will revert back to the "drag" step of its plan, and will continue the plan without modification, even after dozens of caterpillar-moving interventions.

Autonomy

- rational agent relies on its own percepts and experience, not just prior “hardwired” knowledge of its designer
- what is required:
  - perception
  - learning
- why is autonomy important?

Environment types

	Chess	Taxi driving
Fully vs partially observable	fully	partially
Deterministic vs stochastic	strategic*	stochastic
Episodic vs sequential	sequential	sequential
Static vs. dynamic	static/semi	dynamic
Discrete vs continuous	discrete	continuous
Single vs multi-agent	multi	multi

*strategic = deterministic except for actions of other agents

Agent functions and programs

- an agent is completely specified by the *agent function* mapping percept sequences to actions
- we are interested in an efficient implementation of a *rational* agent function

Table-lookup agent

- could construct a look-up table with appropriate action to take for any percept sequence

Simple reflex agents

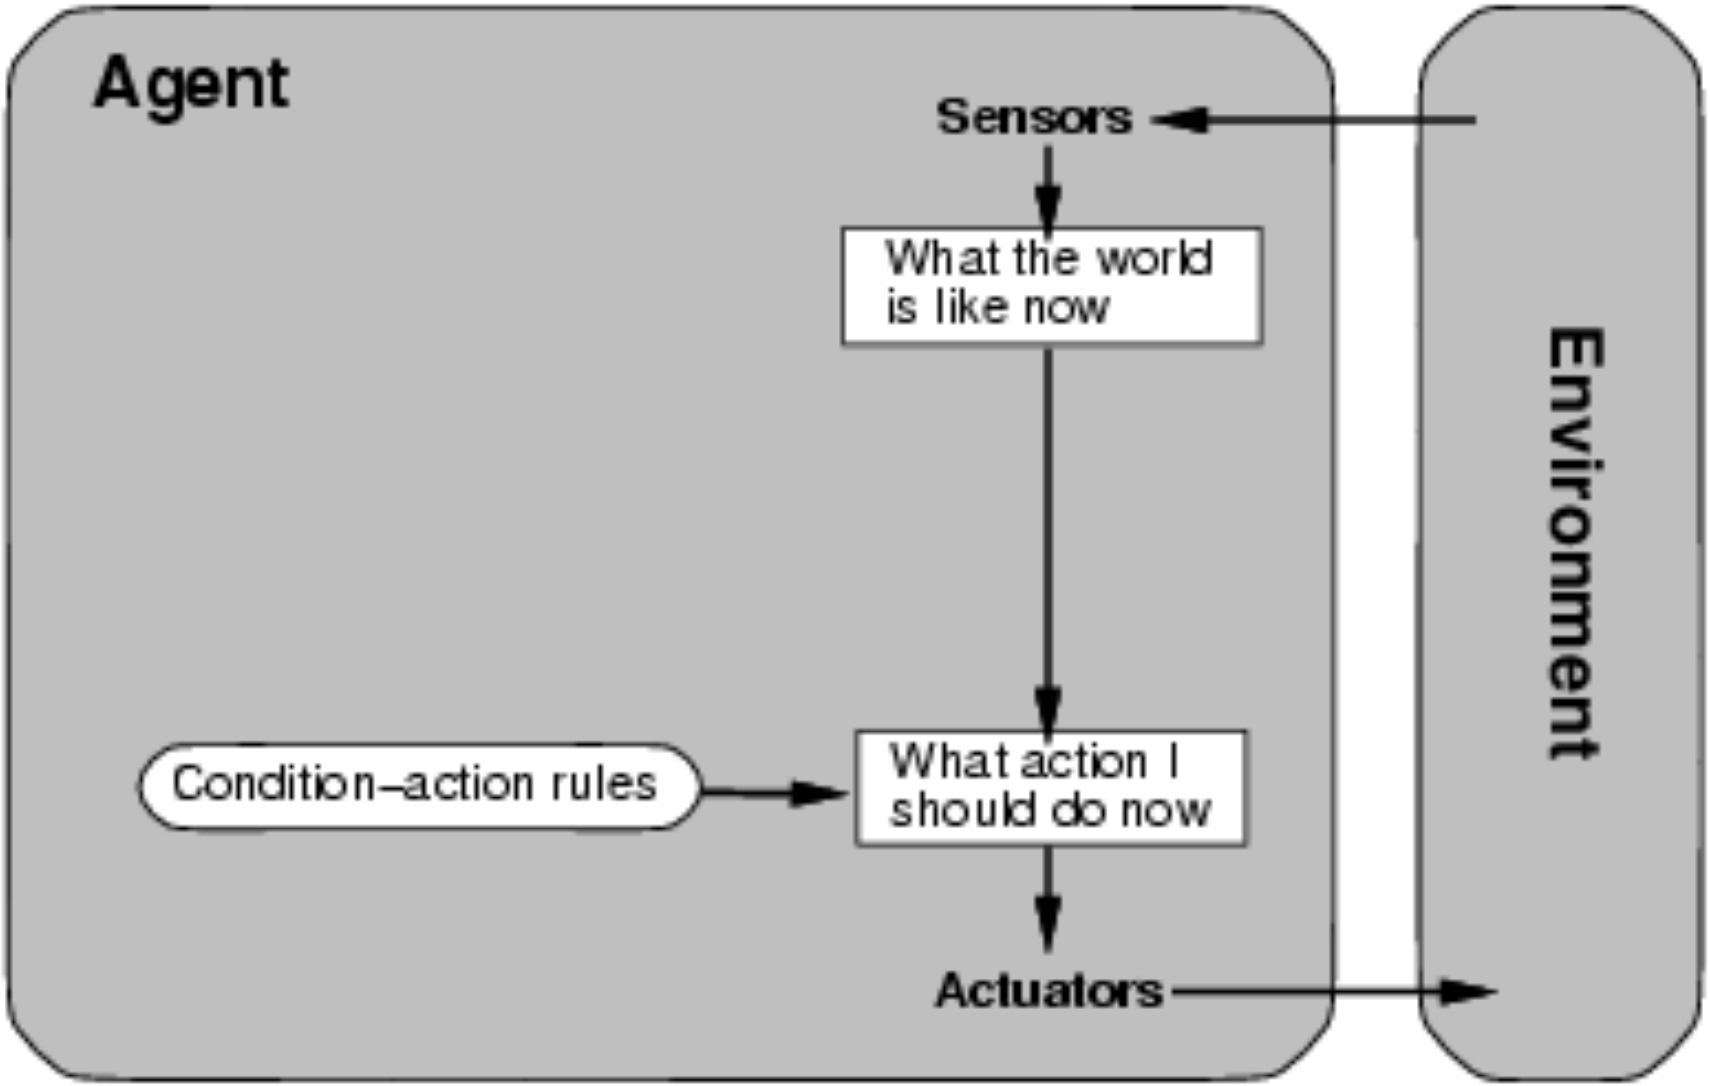

Model-based reflex agents

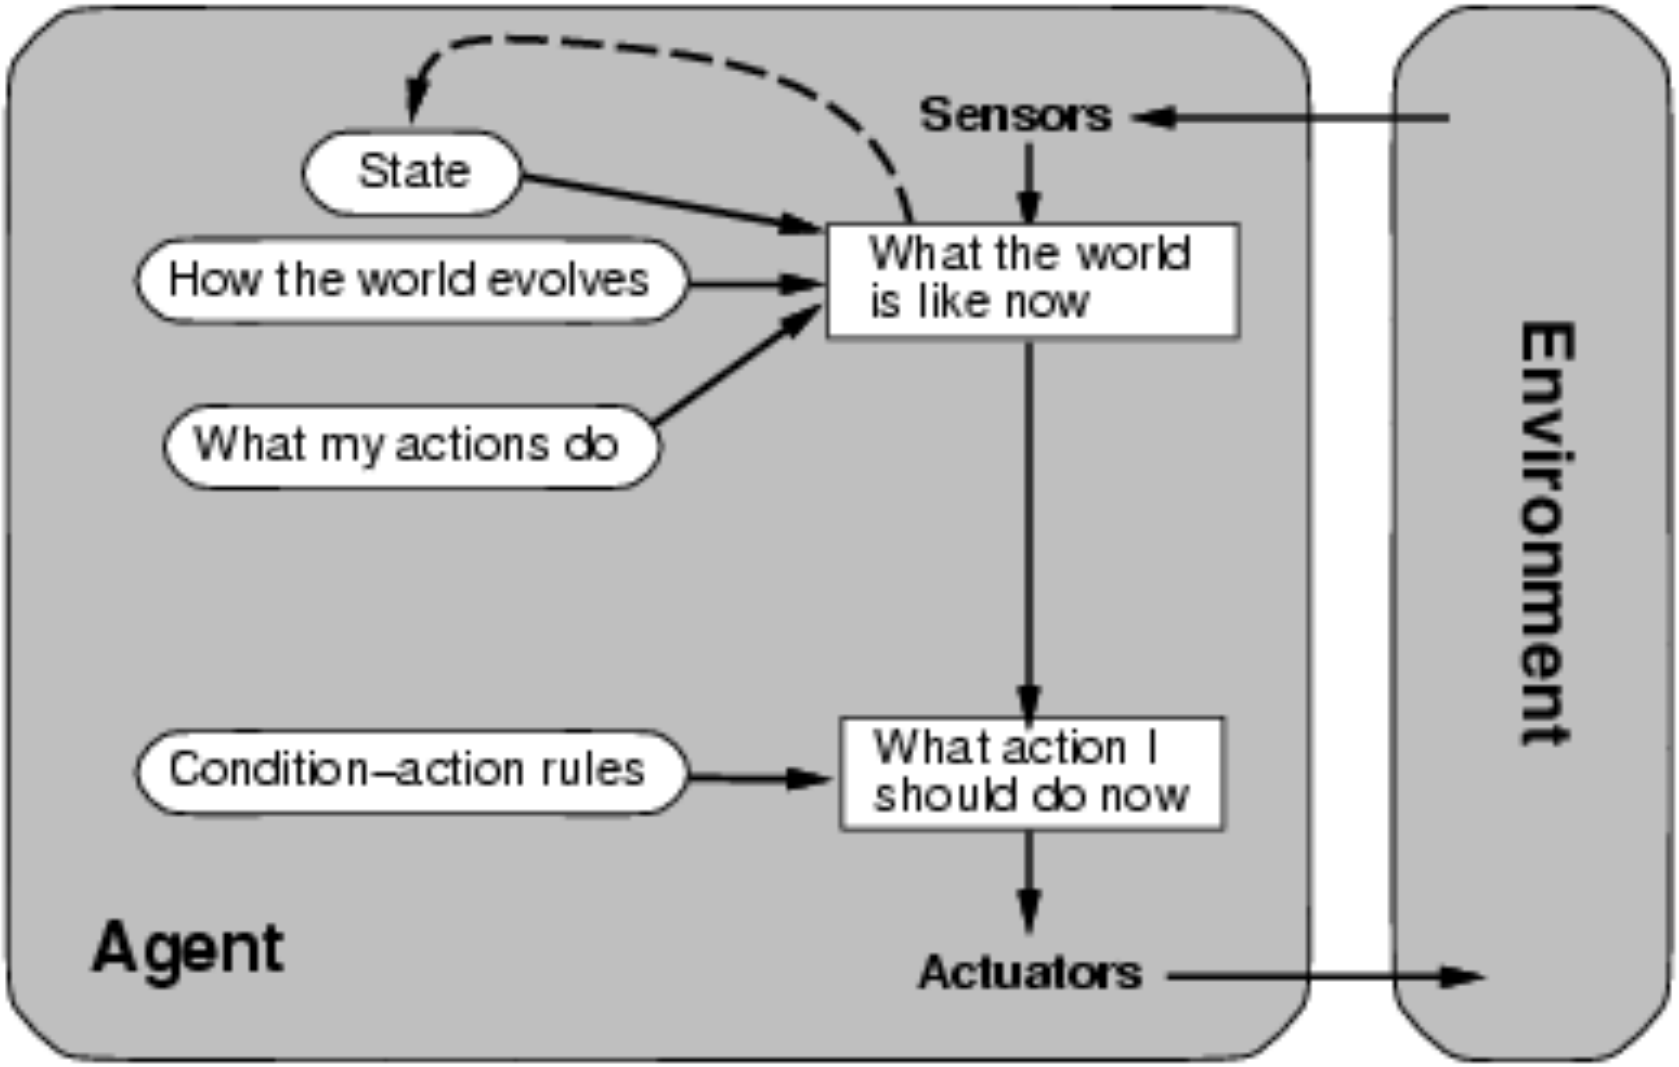

Goal-based agents

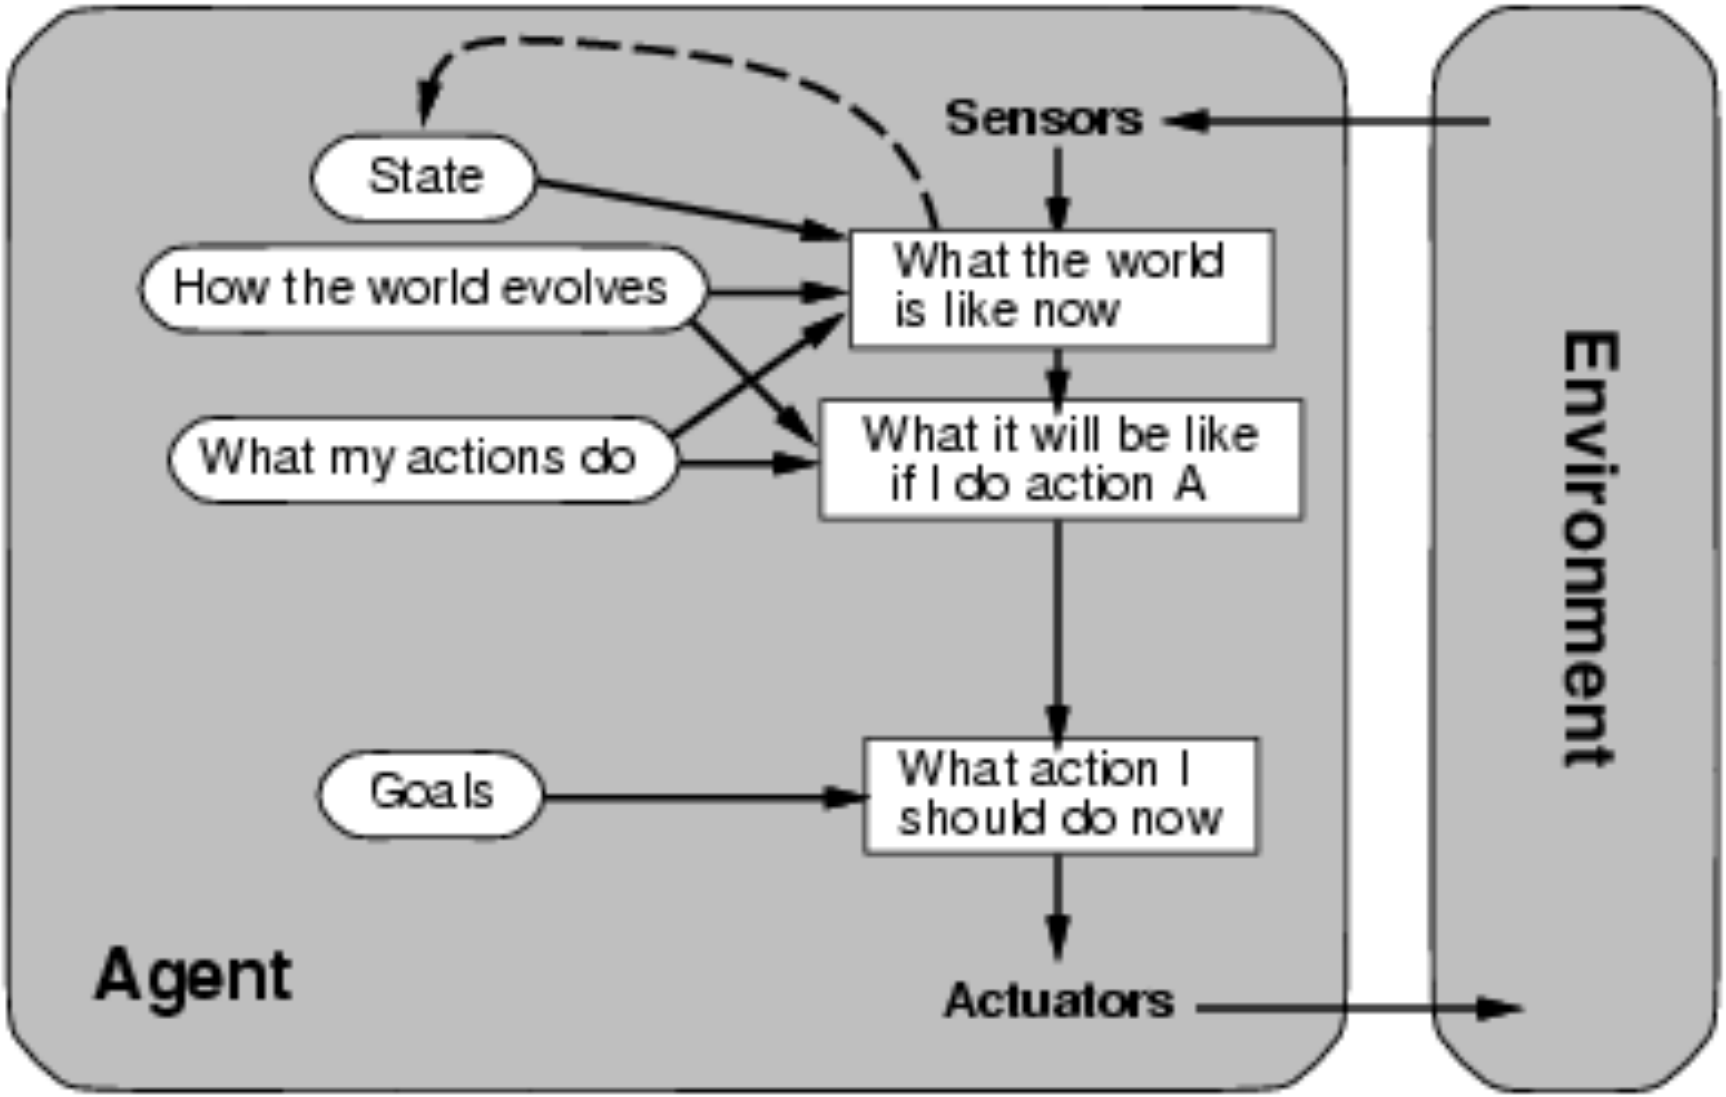

Utility-based agents

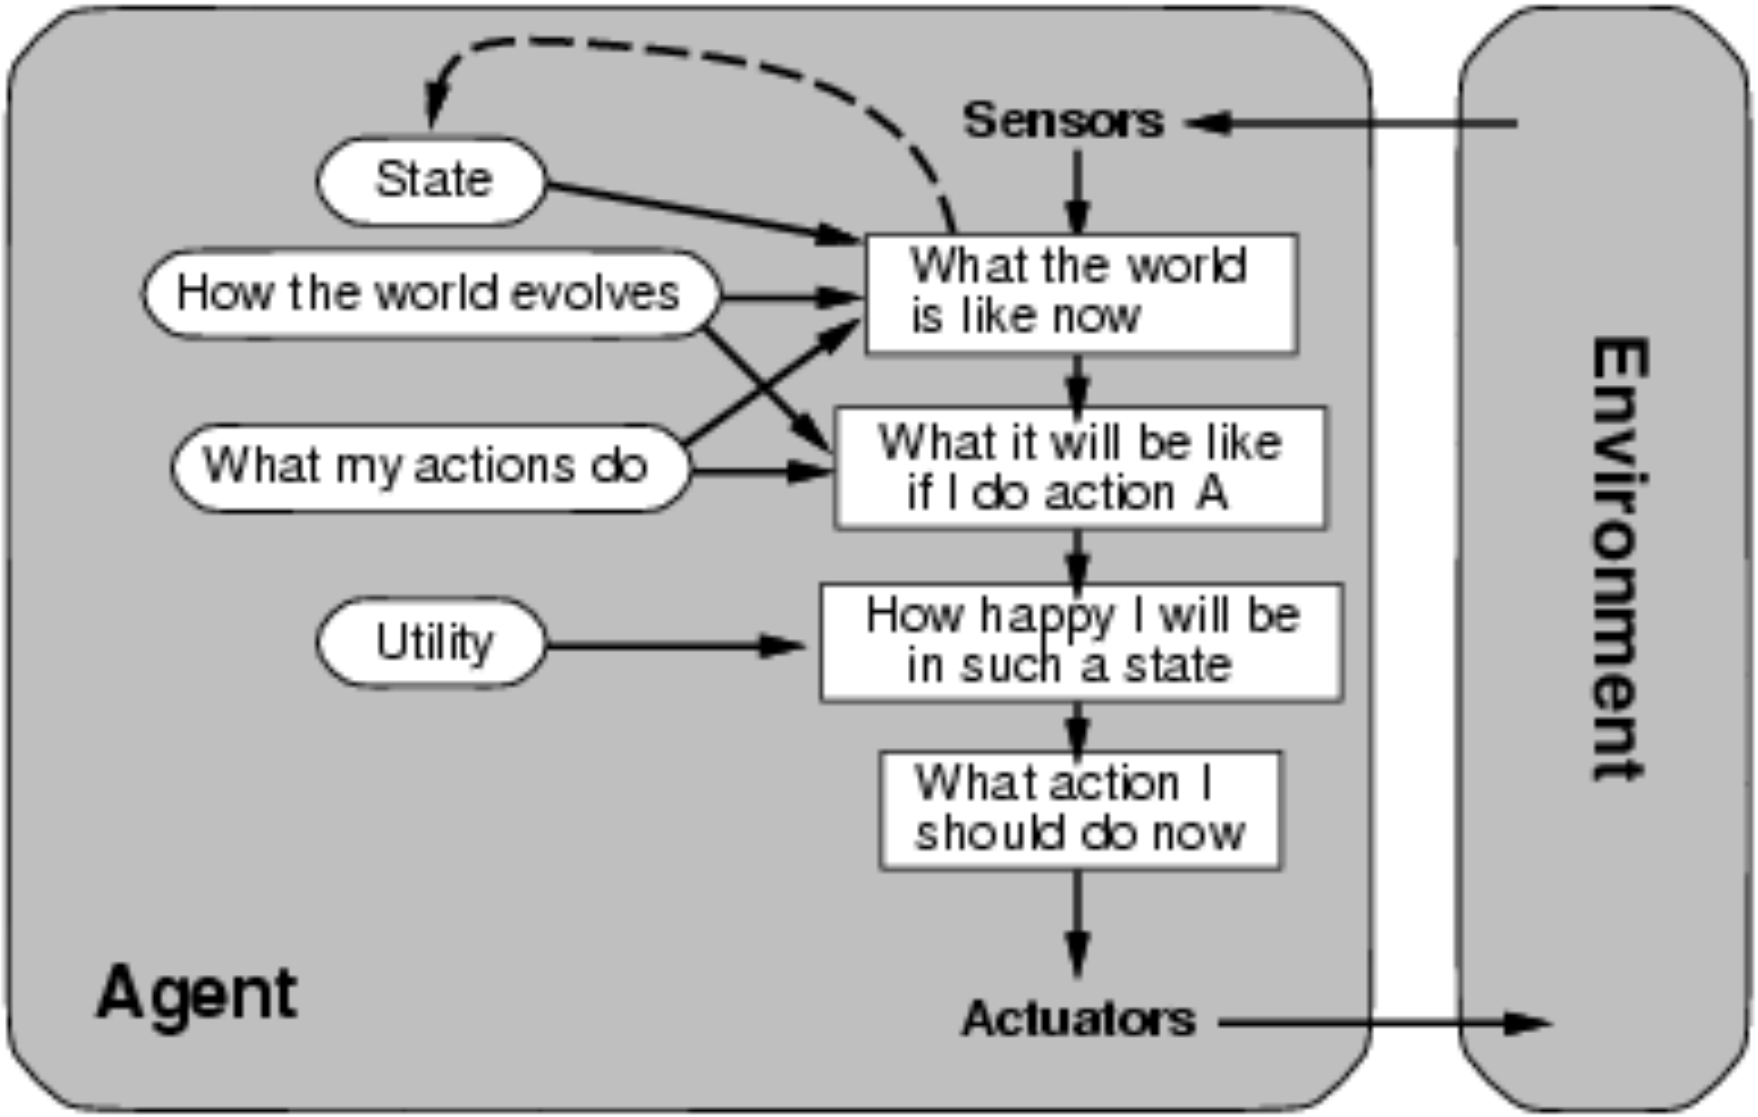

Learning agents

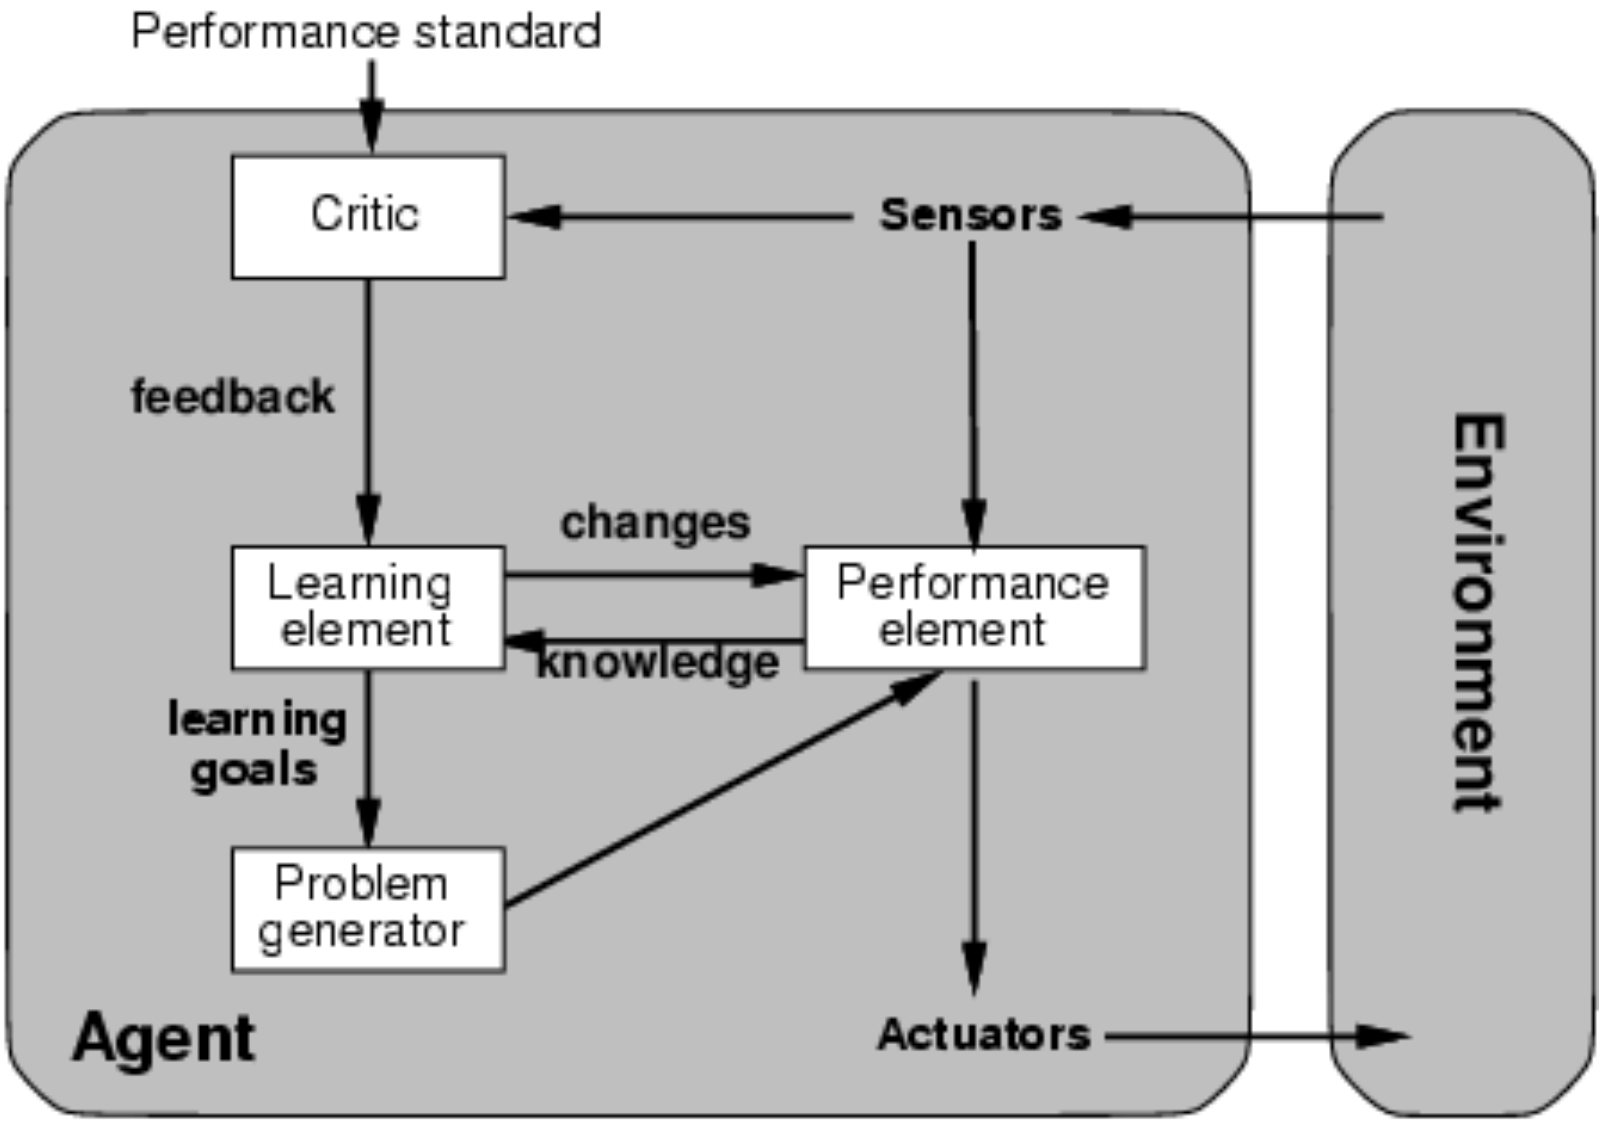

Learning

- off-line learning: collect data, then learn
- on-line learning: learn while doing
- why do we need to learn?
  - complex skill acquisition
  - dynamic environment
  - interaction strategies

Deliberative Architectures

- SMPA: sense, model, plan, act
- use internal, symbolic representation (model) of world for reasoning/planning

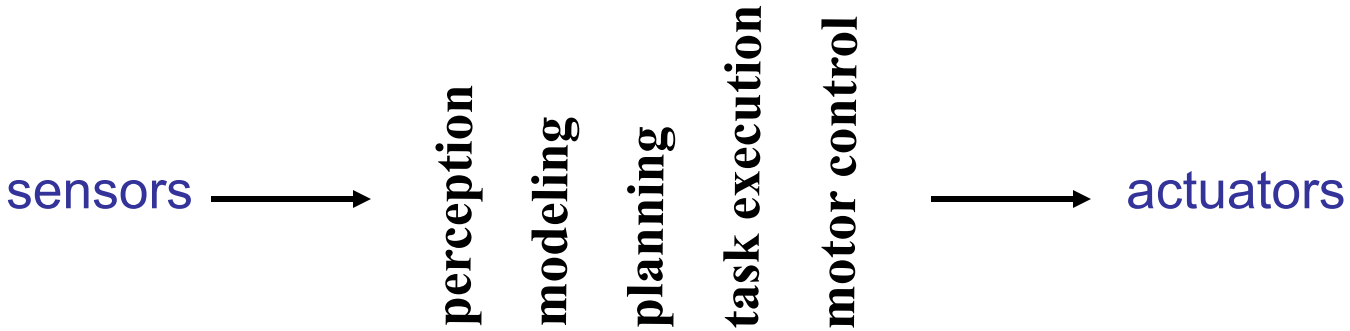

Shakey (1966-1972)

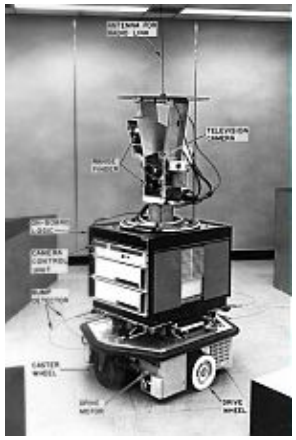

- PLANEX: accepts goals, maintains world state, calls planner, executes best plan
- STRIPS: propositional logic planning system
- ILA: intermediate-level actions
- LLA: low-level actions

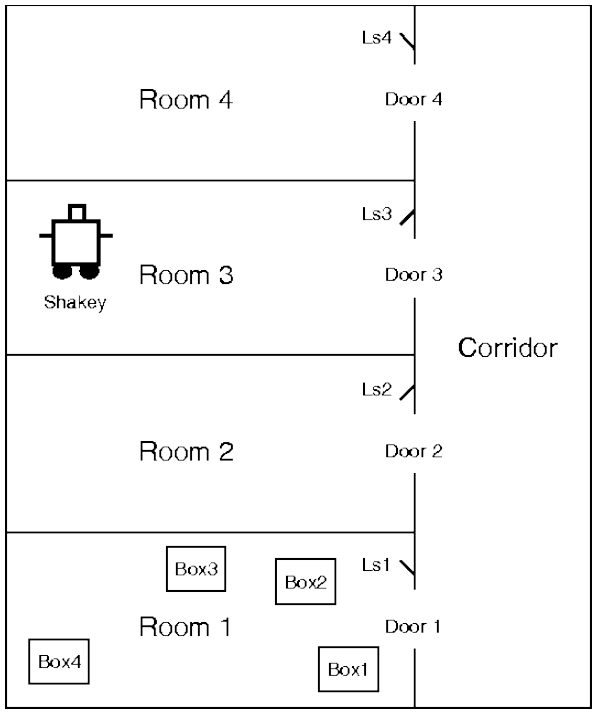

Strengths of Deliberation

- allows agent to look-ahead at possible outcomes of action (search) without actually performing action
- powerful learning (adaptation, compilation) schemes can be incorporated

Weaknesses

- cost of generating model, inaccuracies
- cost of reasoning about low-level actions
- too expensive for real-time behaviour
- the “microworlds” problem
  - AI programs operated in small domains
  - relied on very simple world models
  - couldn't scale to real world scenarios

Which robot is more intelligent?

Brooks' criticism

- failure of simulated systems in past -> only way to do AI is with mobile robots
- *"the world is its own best model"*
- intelligence does not require:
  - explicit representation
  - abstract, symbolic reasoning
- real intelligence:
  - arises from systems situated in the real world
  - is an emergent property of systems, as agents interact with a complex environment

Reactive Architectures

- behaviour-based robotics
- act using stimulus/response behaviour
- no centralized control
- layered task decomposition, subsumption

Subsumption Architecture

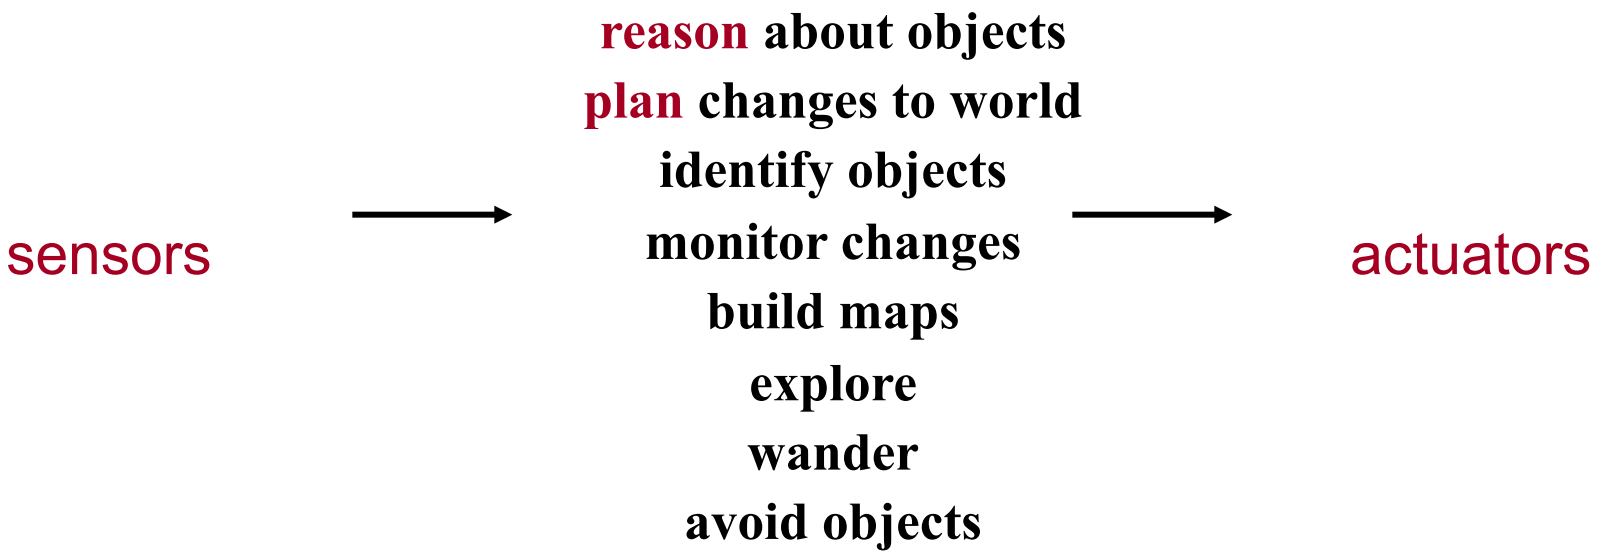

Strengths

- competence can be achieved without explicit reasoning or rule following [Dreyfus, 1972]
- internal models are expensive and inaccurate so best to avoid them
- complex behaviour emerges through interaction of simple behaviours

How far can it go?