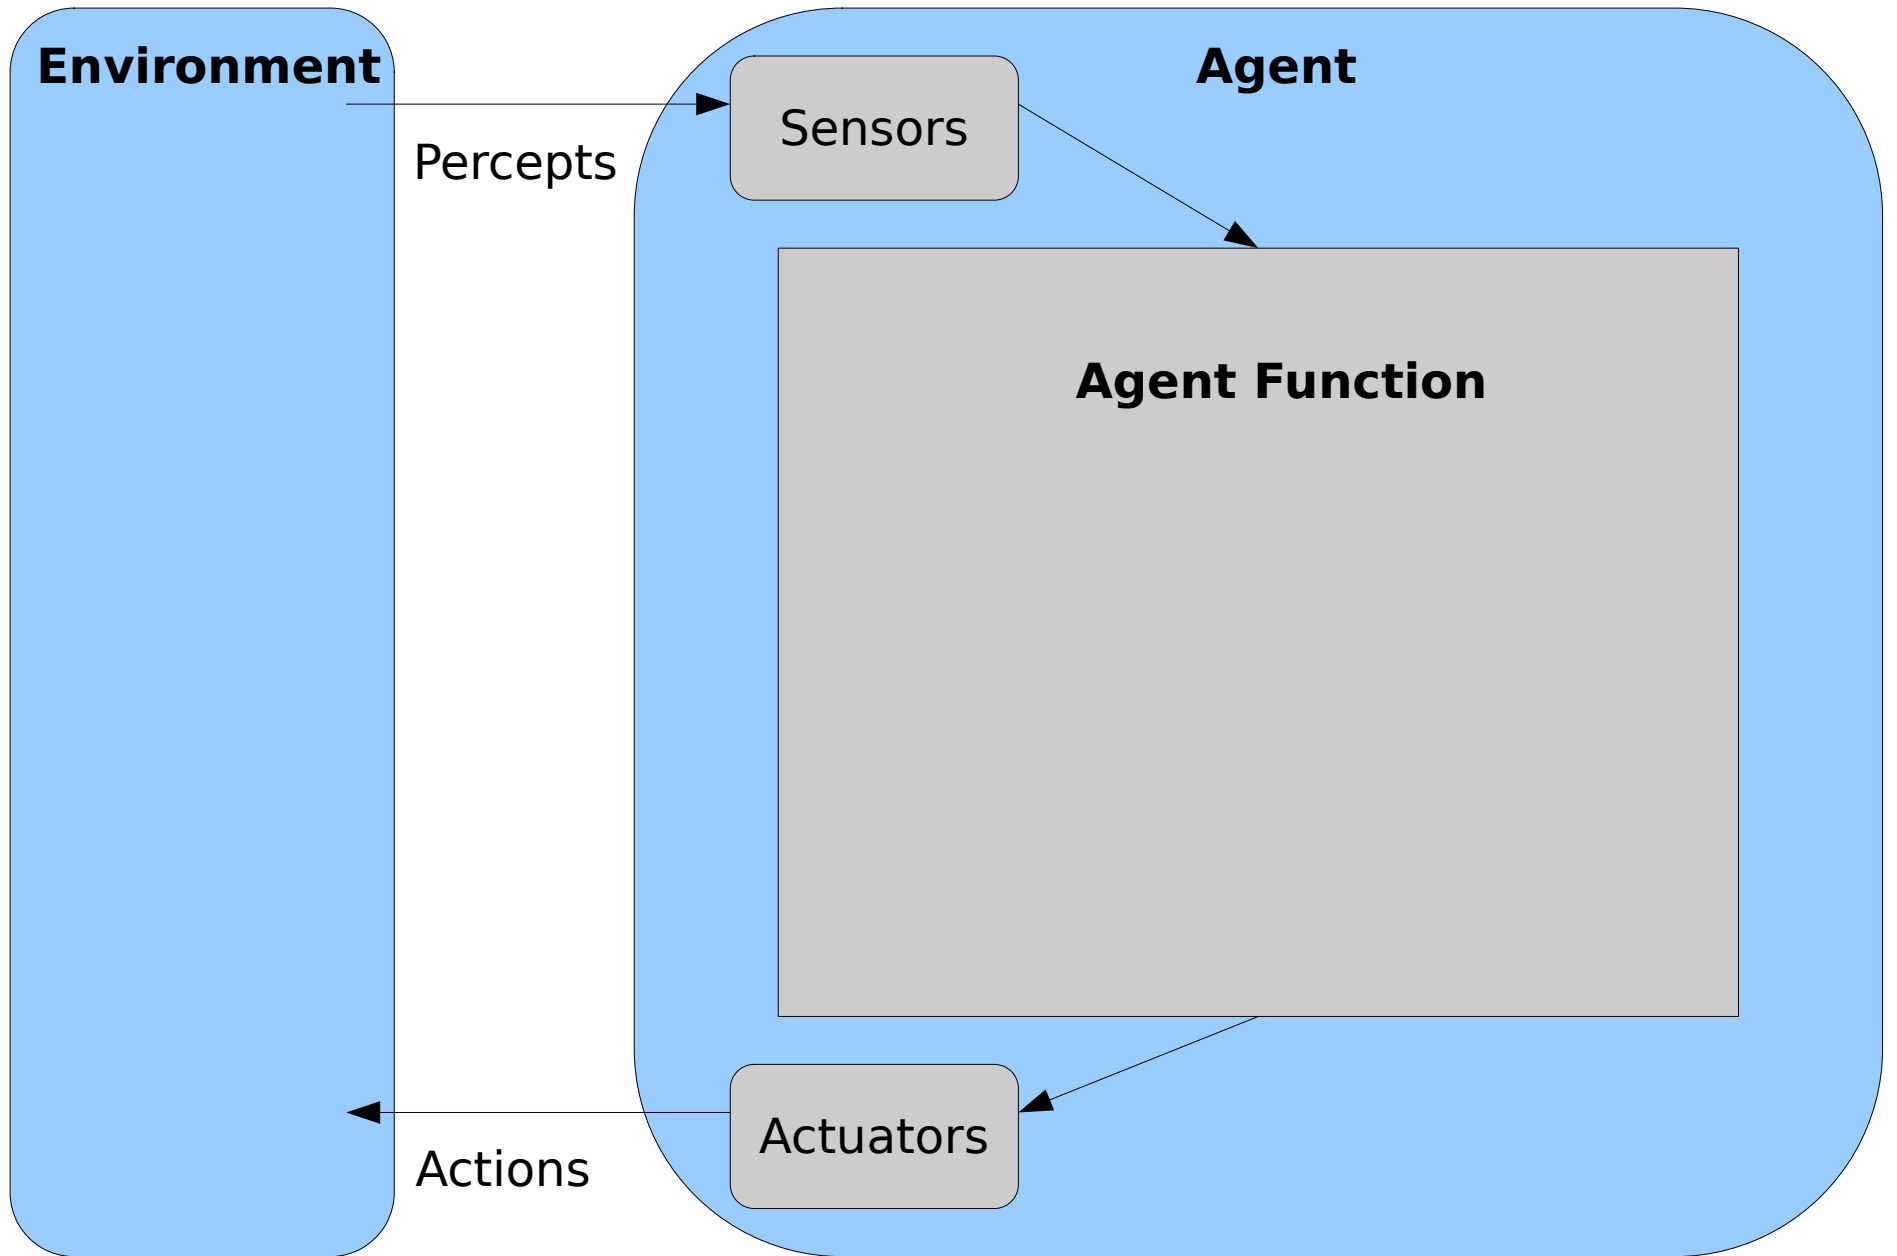


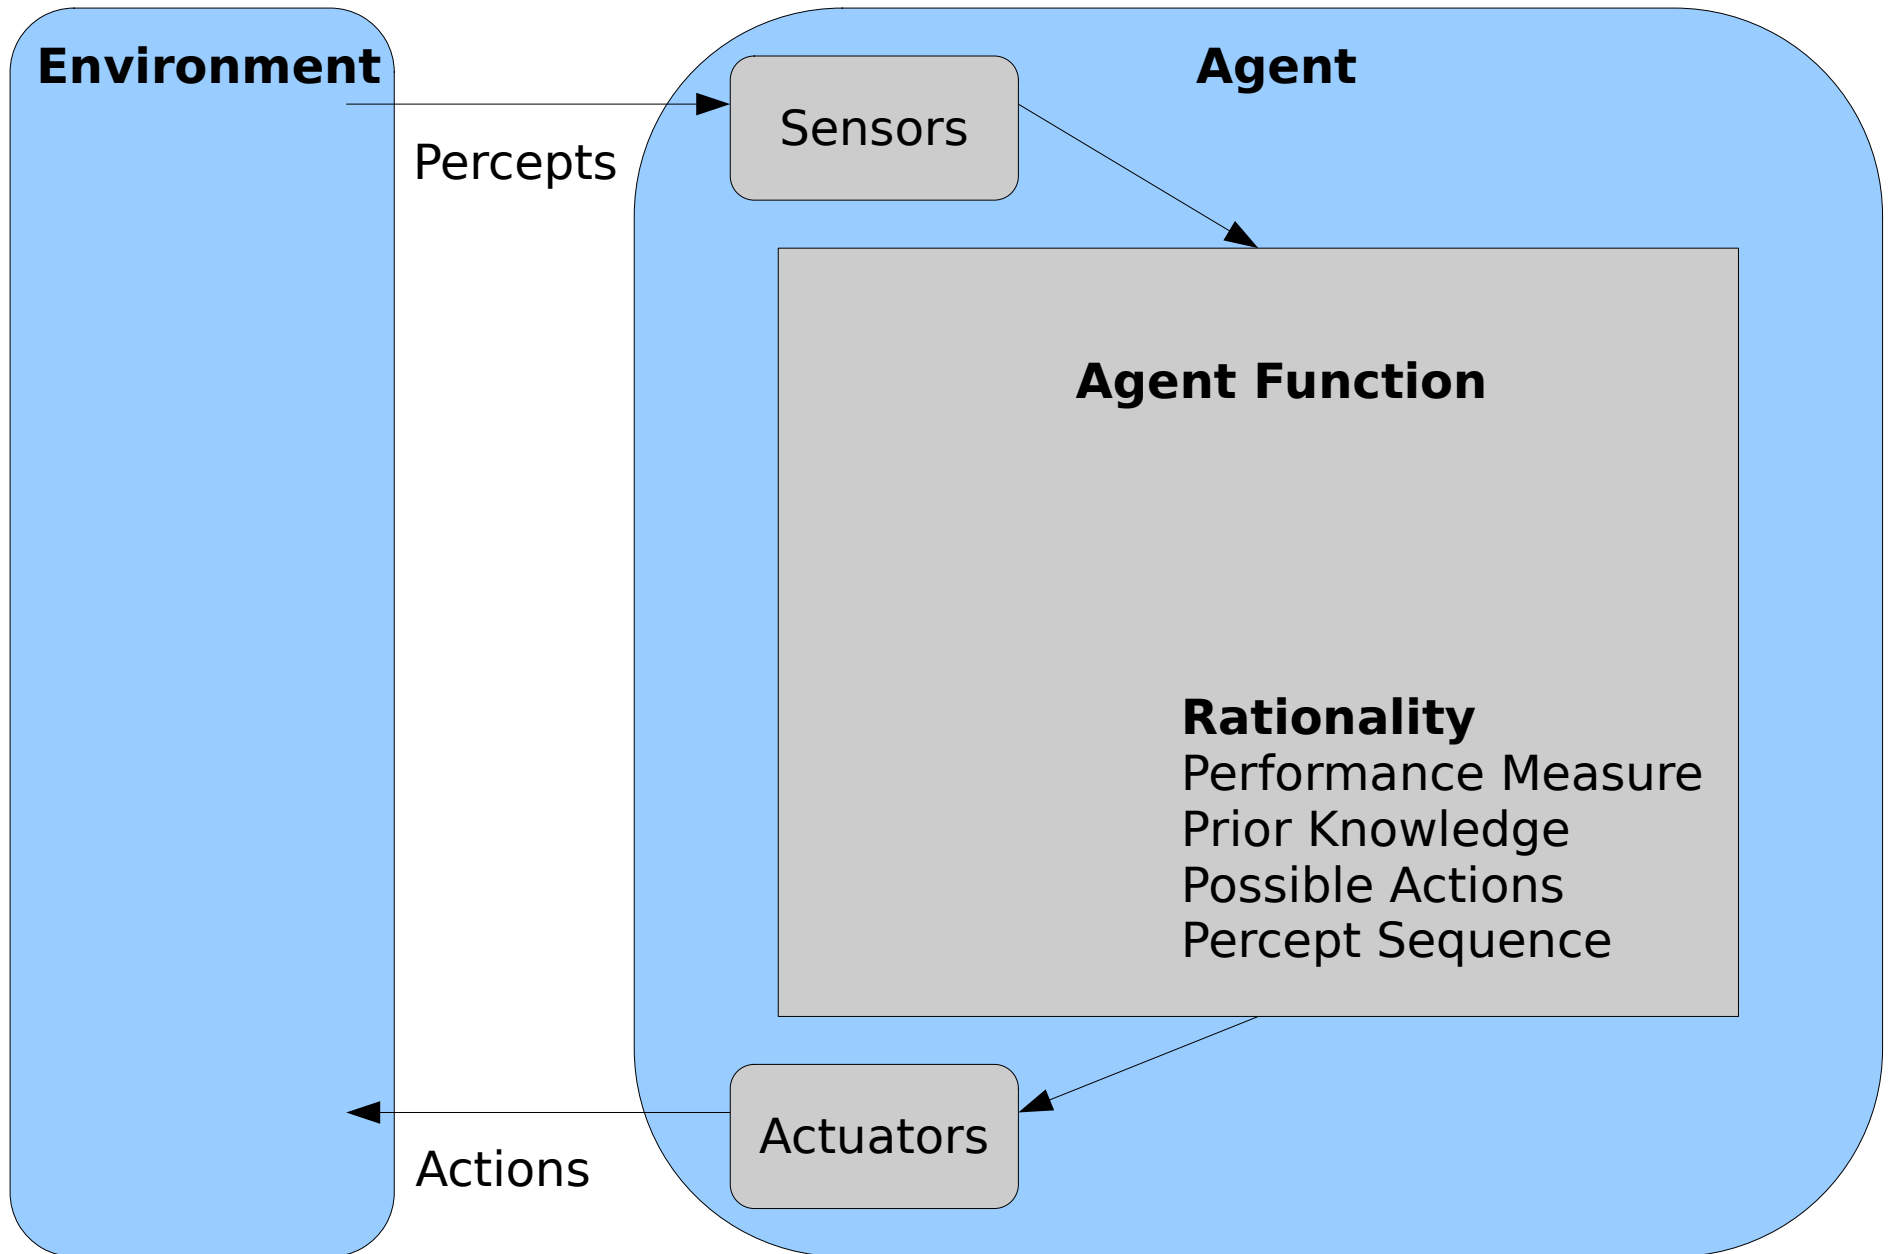
Intelligent Agents

- Agents
- Environments
- Rationality



Example: Vacuum World

- Vacuum agent to clean a pair of rooms
- Can move left, move right or suck dirt
- Can sense location and cleanliness

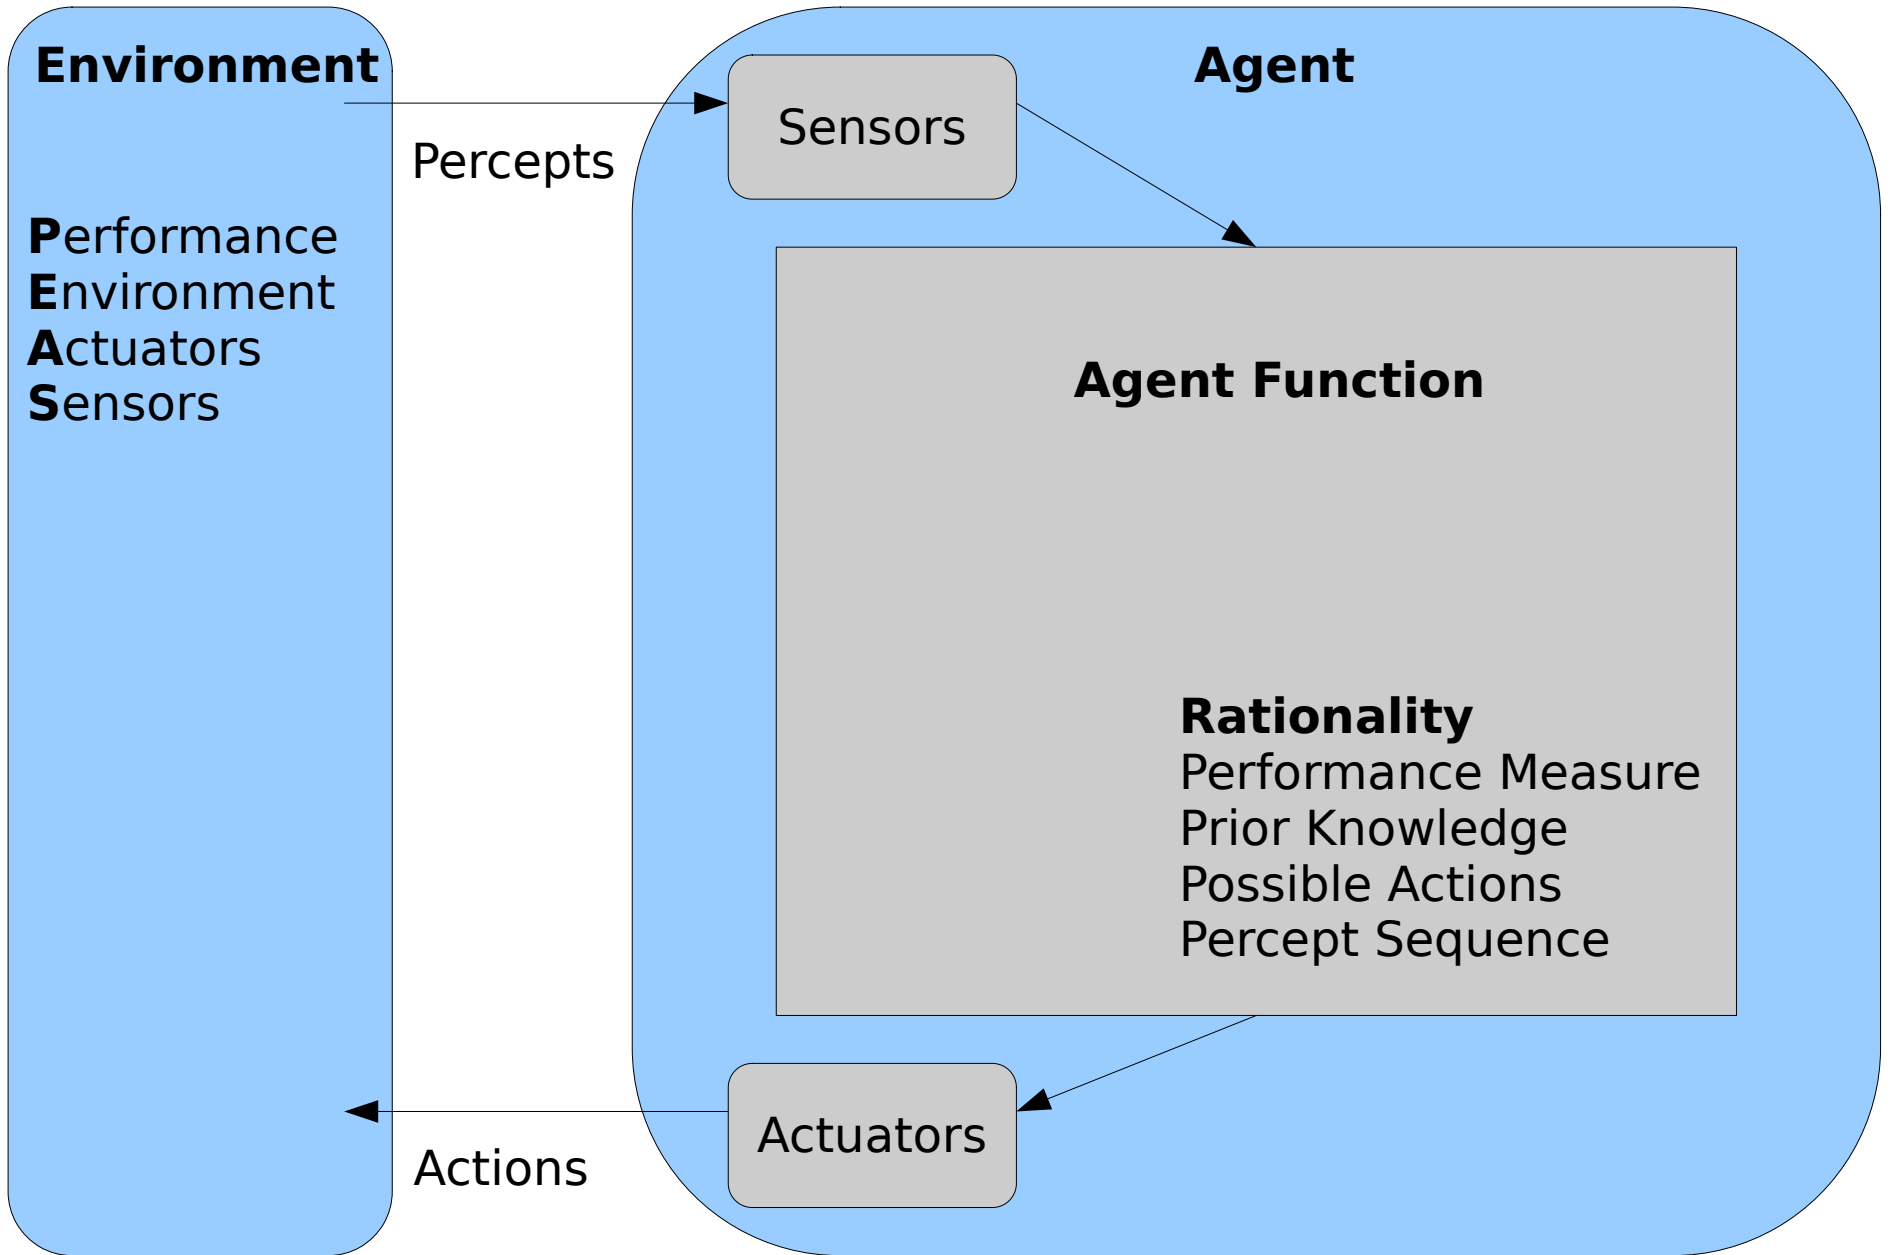


Example: VW Performance Measure

- Average cleanliness
- End cleanliness
- Noise levels
- Power consumption

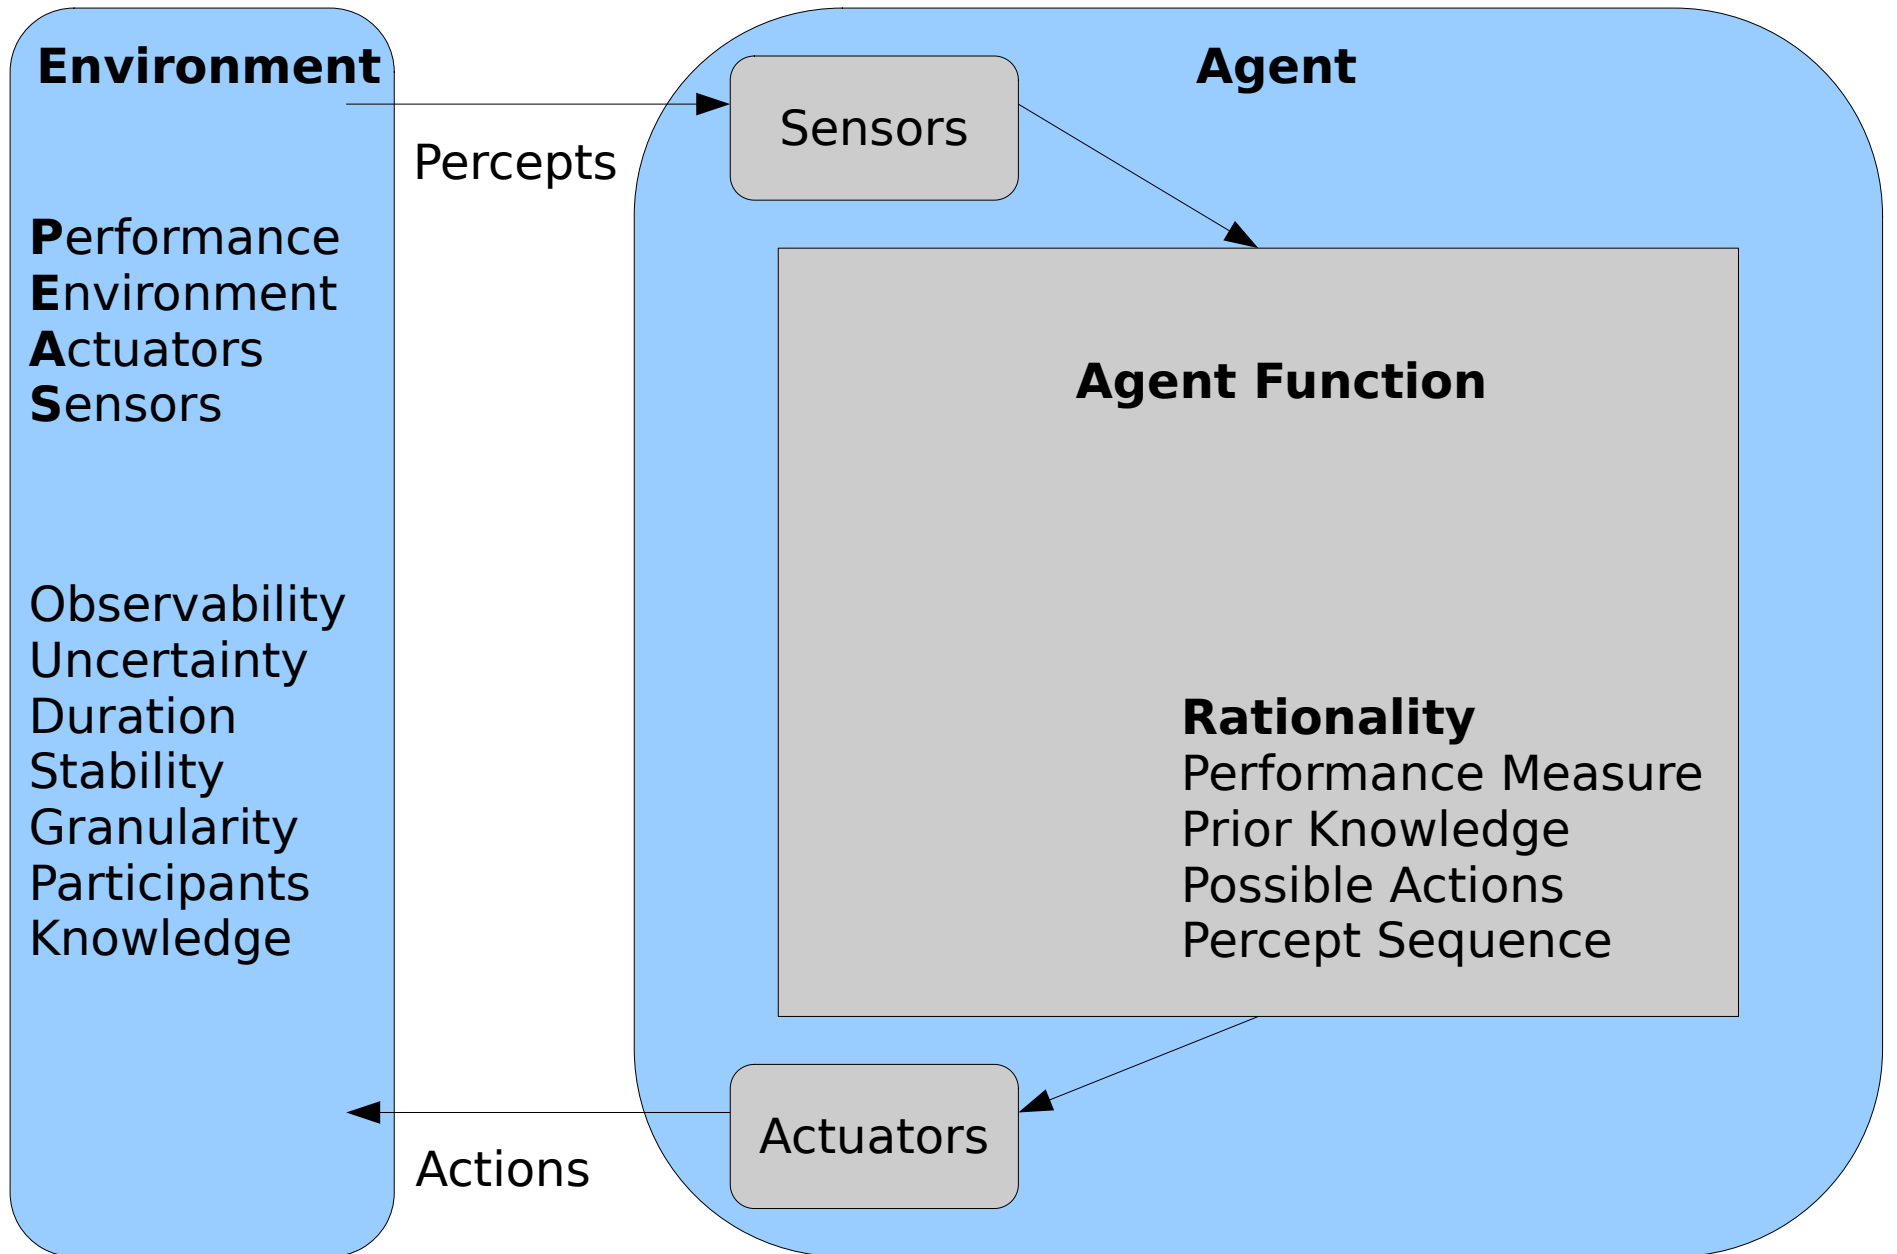
Rationality Gotchas

- Agents not expected to be omniscient (expected, not actual performance).
- Information gathering and exploration may improve performance measure.
- Learn from percepts.
- Autonomy and designer bias.



Environment Properties

- Fully vs partially observable
- Deterministic vs non-deterministic (stochastic)
- Episodic vs sequential
- Static vs dynamic
- Discrete vs continuous
- Single vs multiple agent
- Known vs unknown



Rationality

- Rationality should be measured as the average performance measure of an agent over a large number of environments from an environment class.

Agent Program Implementations

- Table driven
- Simple reflex
- Model based reflex
- Goal based
- Utility based
- Learning

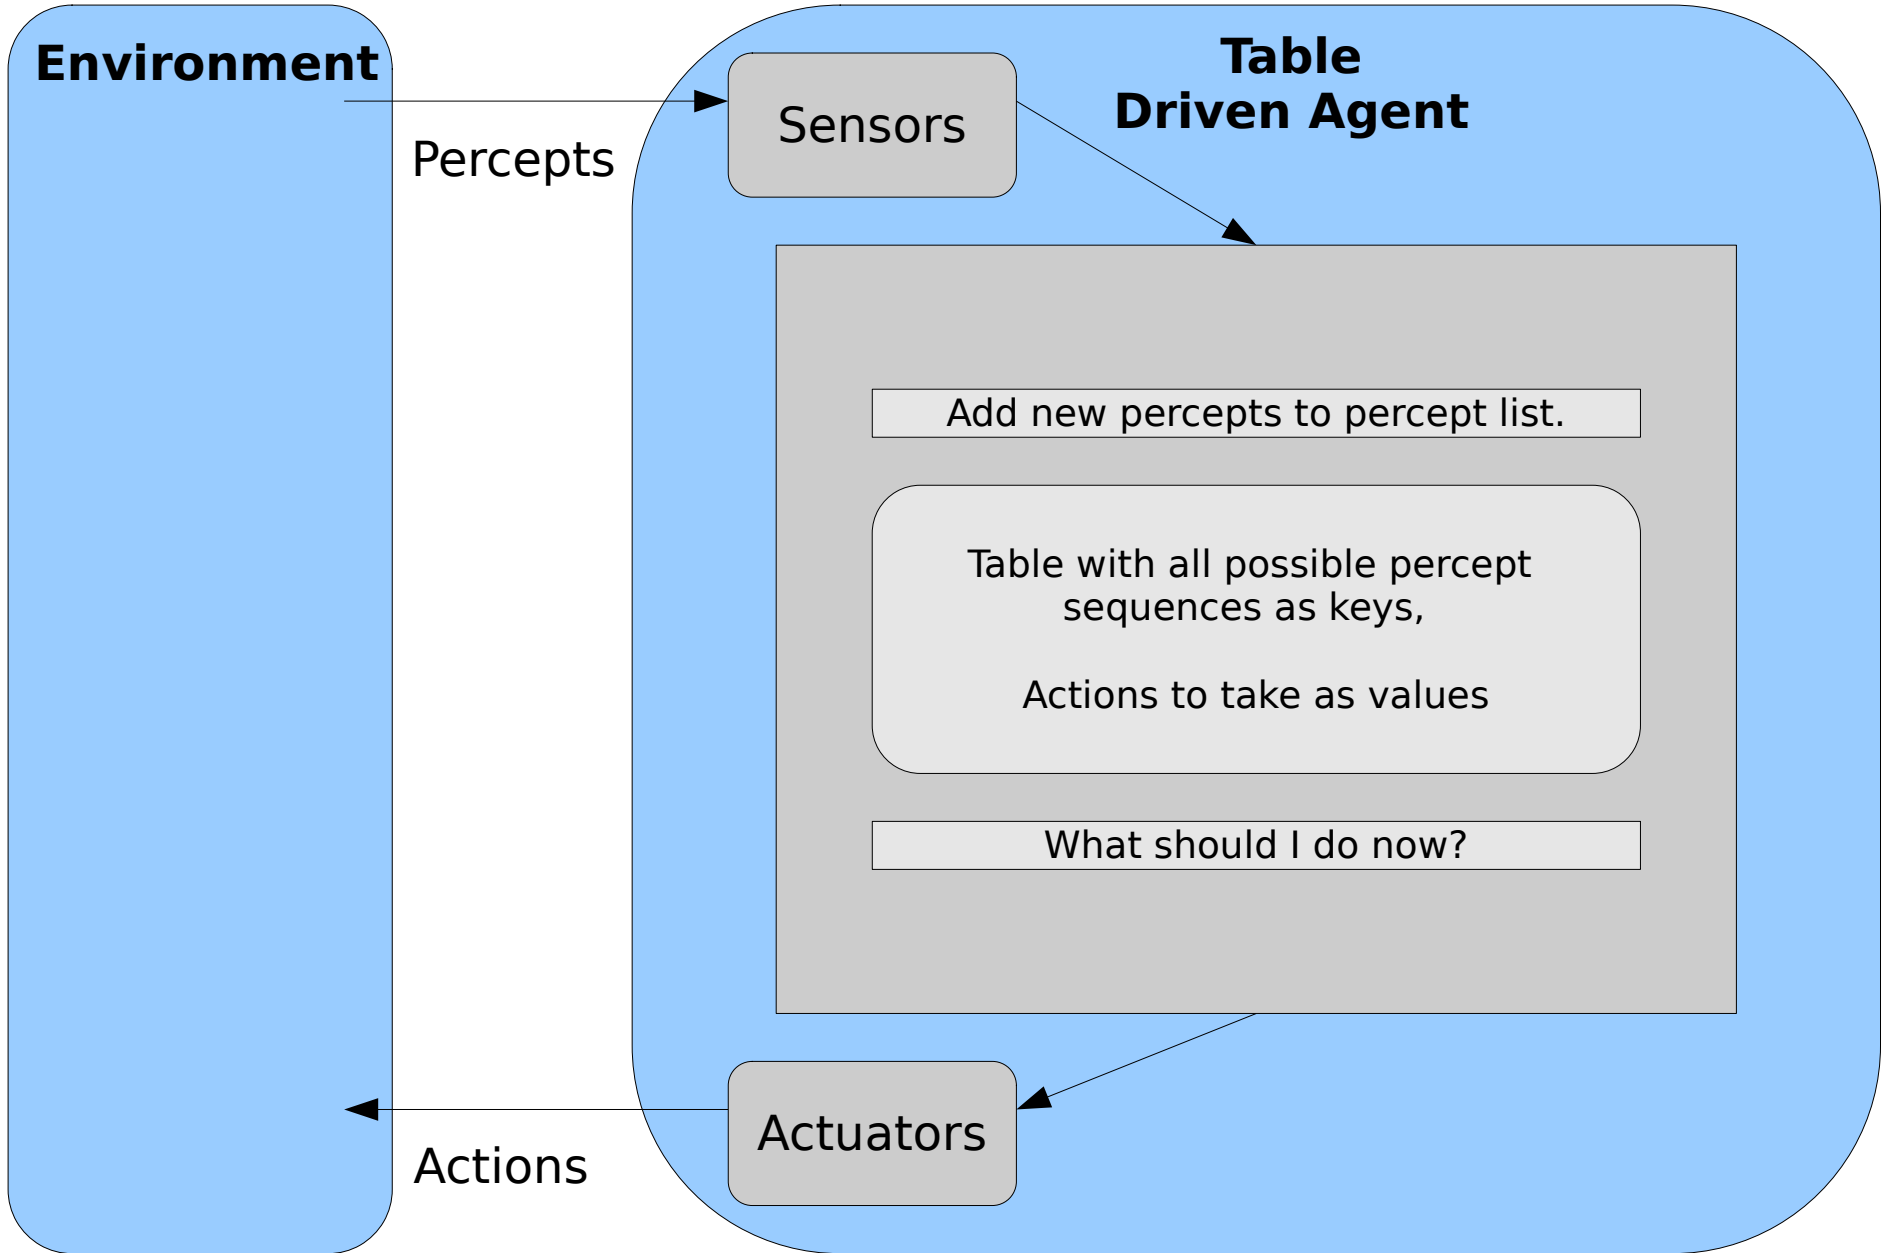
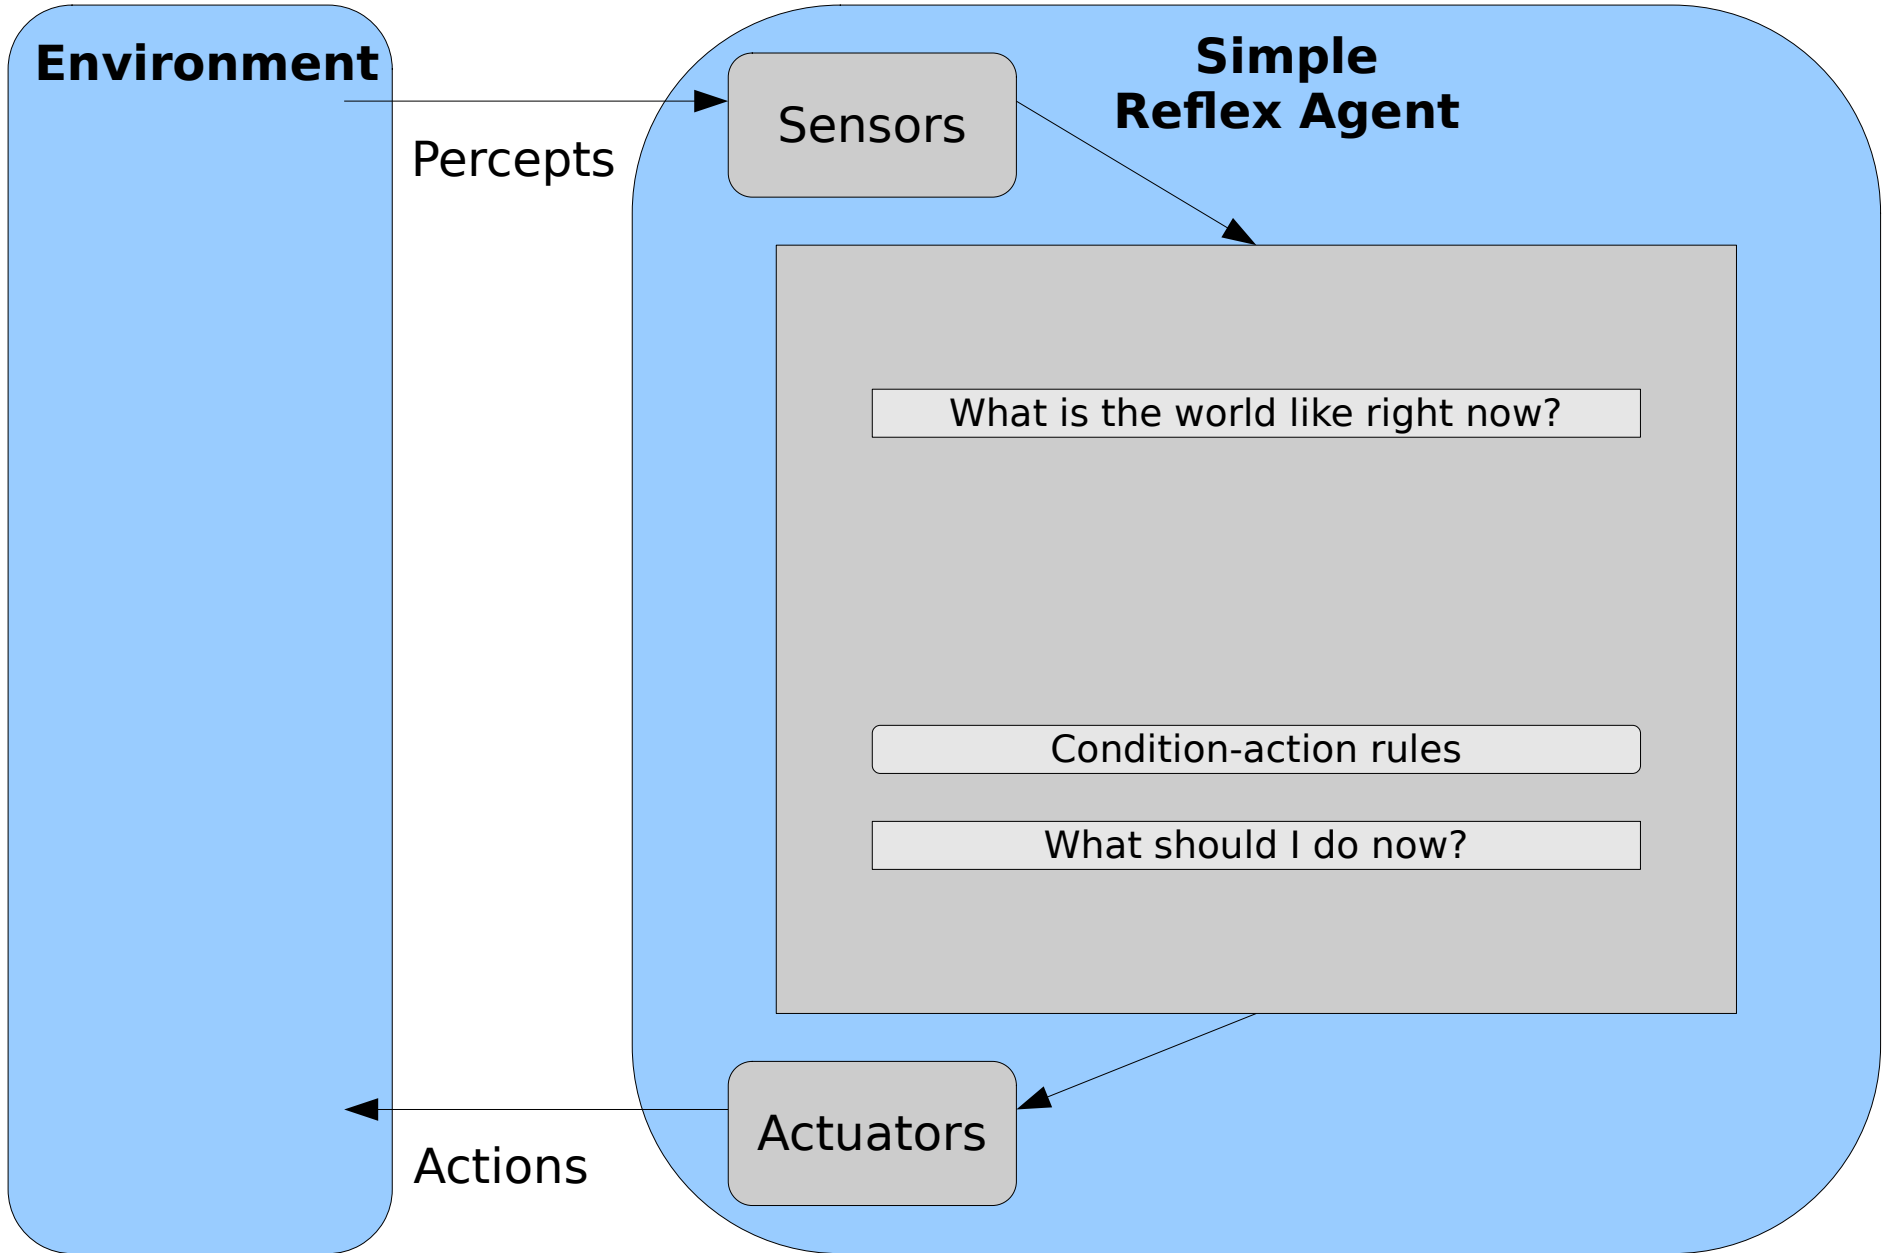


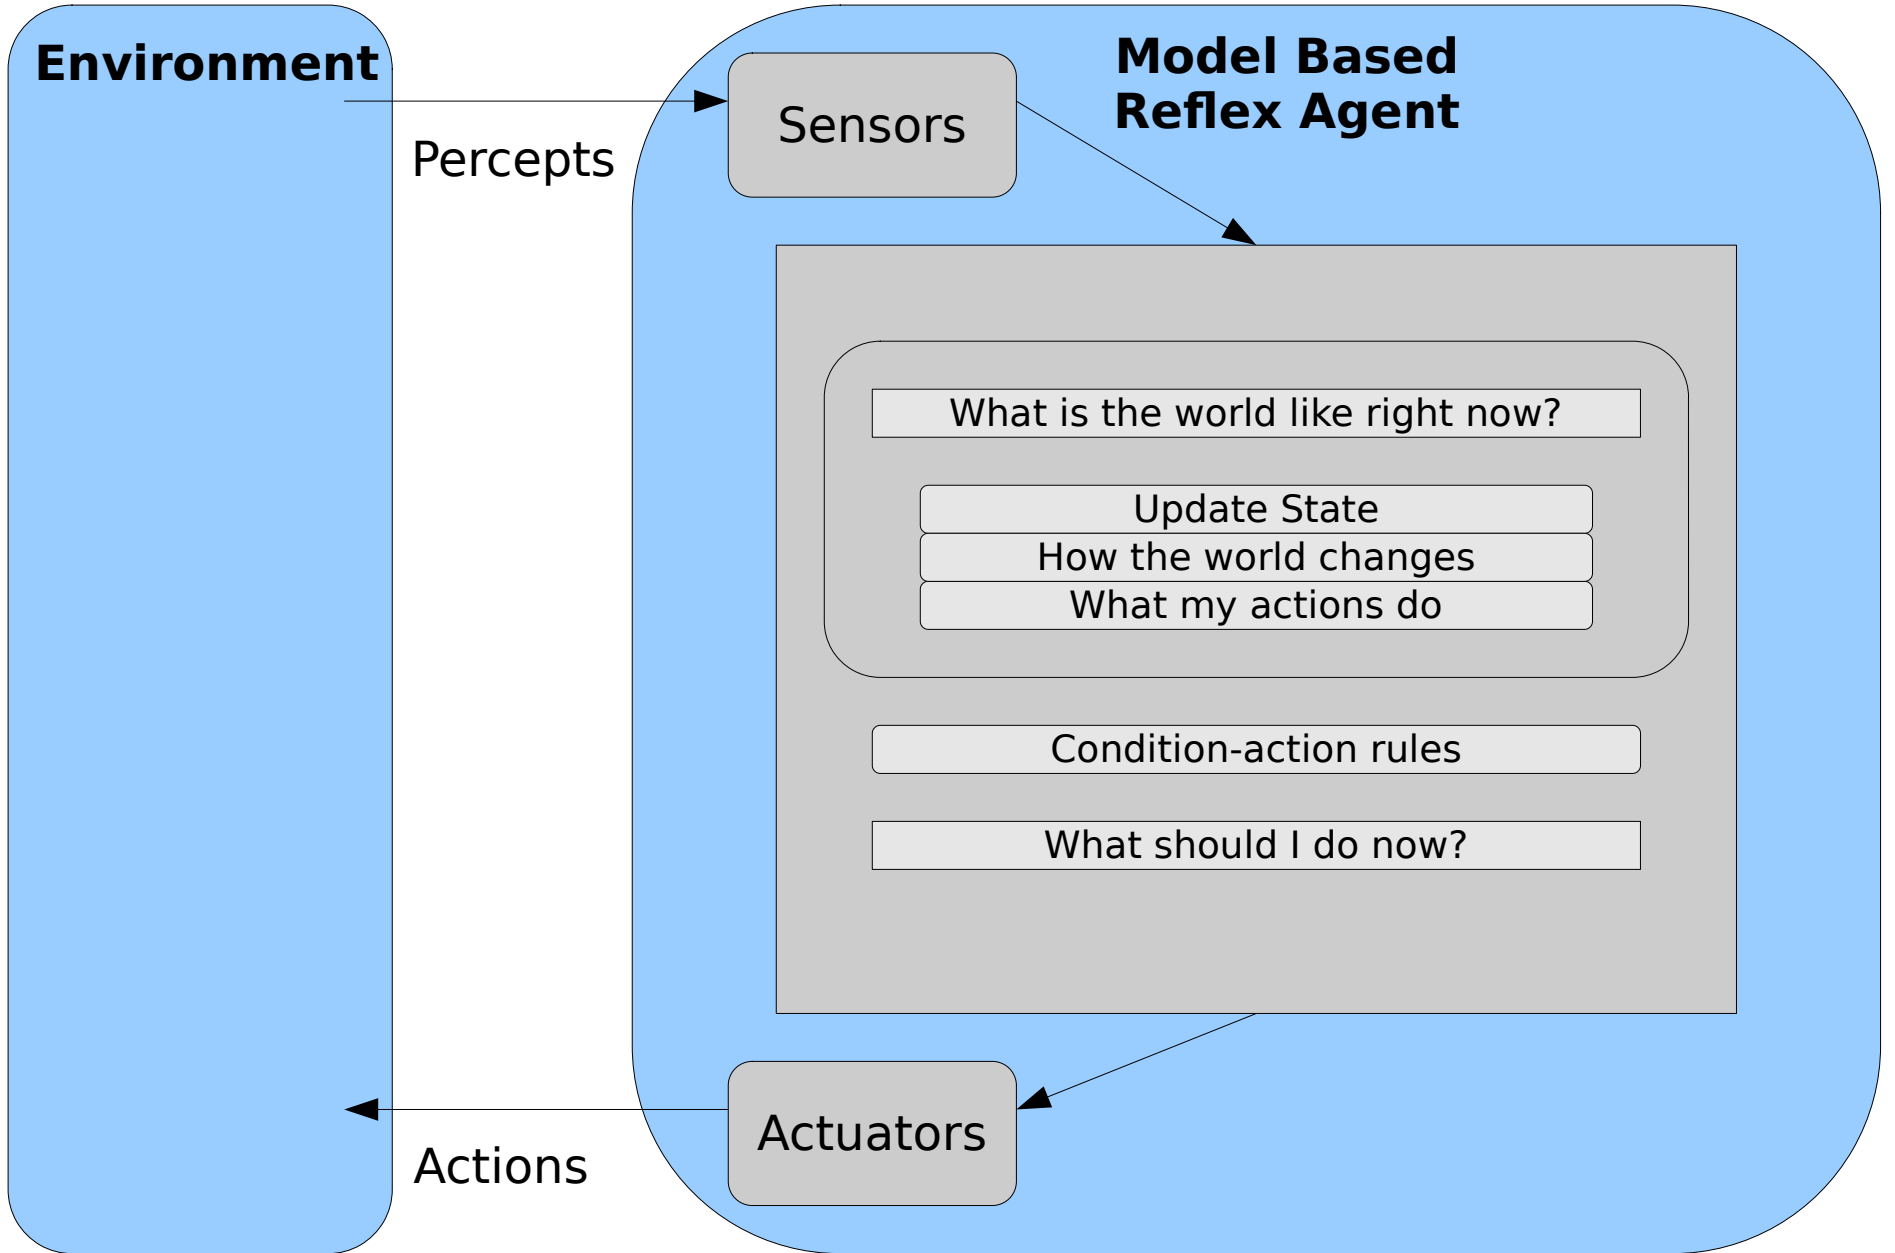
Table Driven Agents

- Entry for every possible percept sequence
- Translates percept sequence into action
- How does this rely on designer bias?
- For what kinds of environments is this a good choice? Bad choice?
- Very large tables, usually not feasible



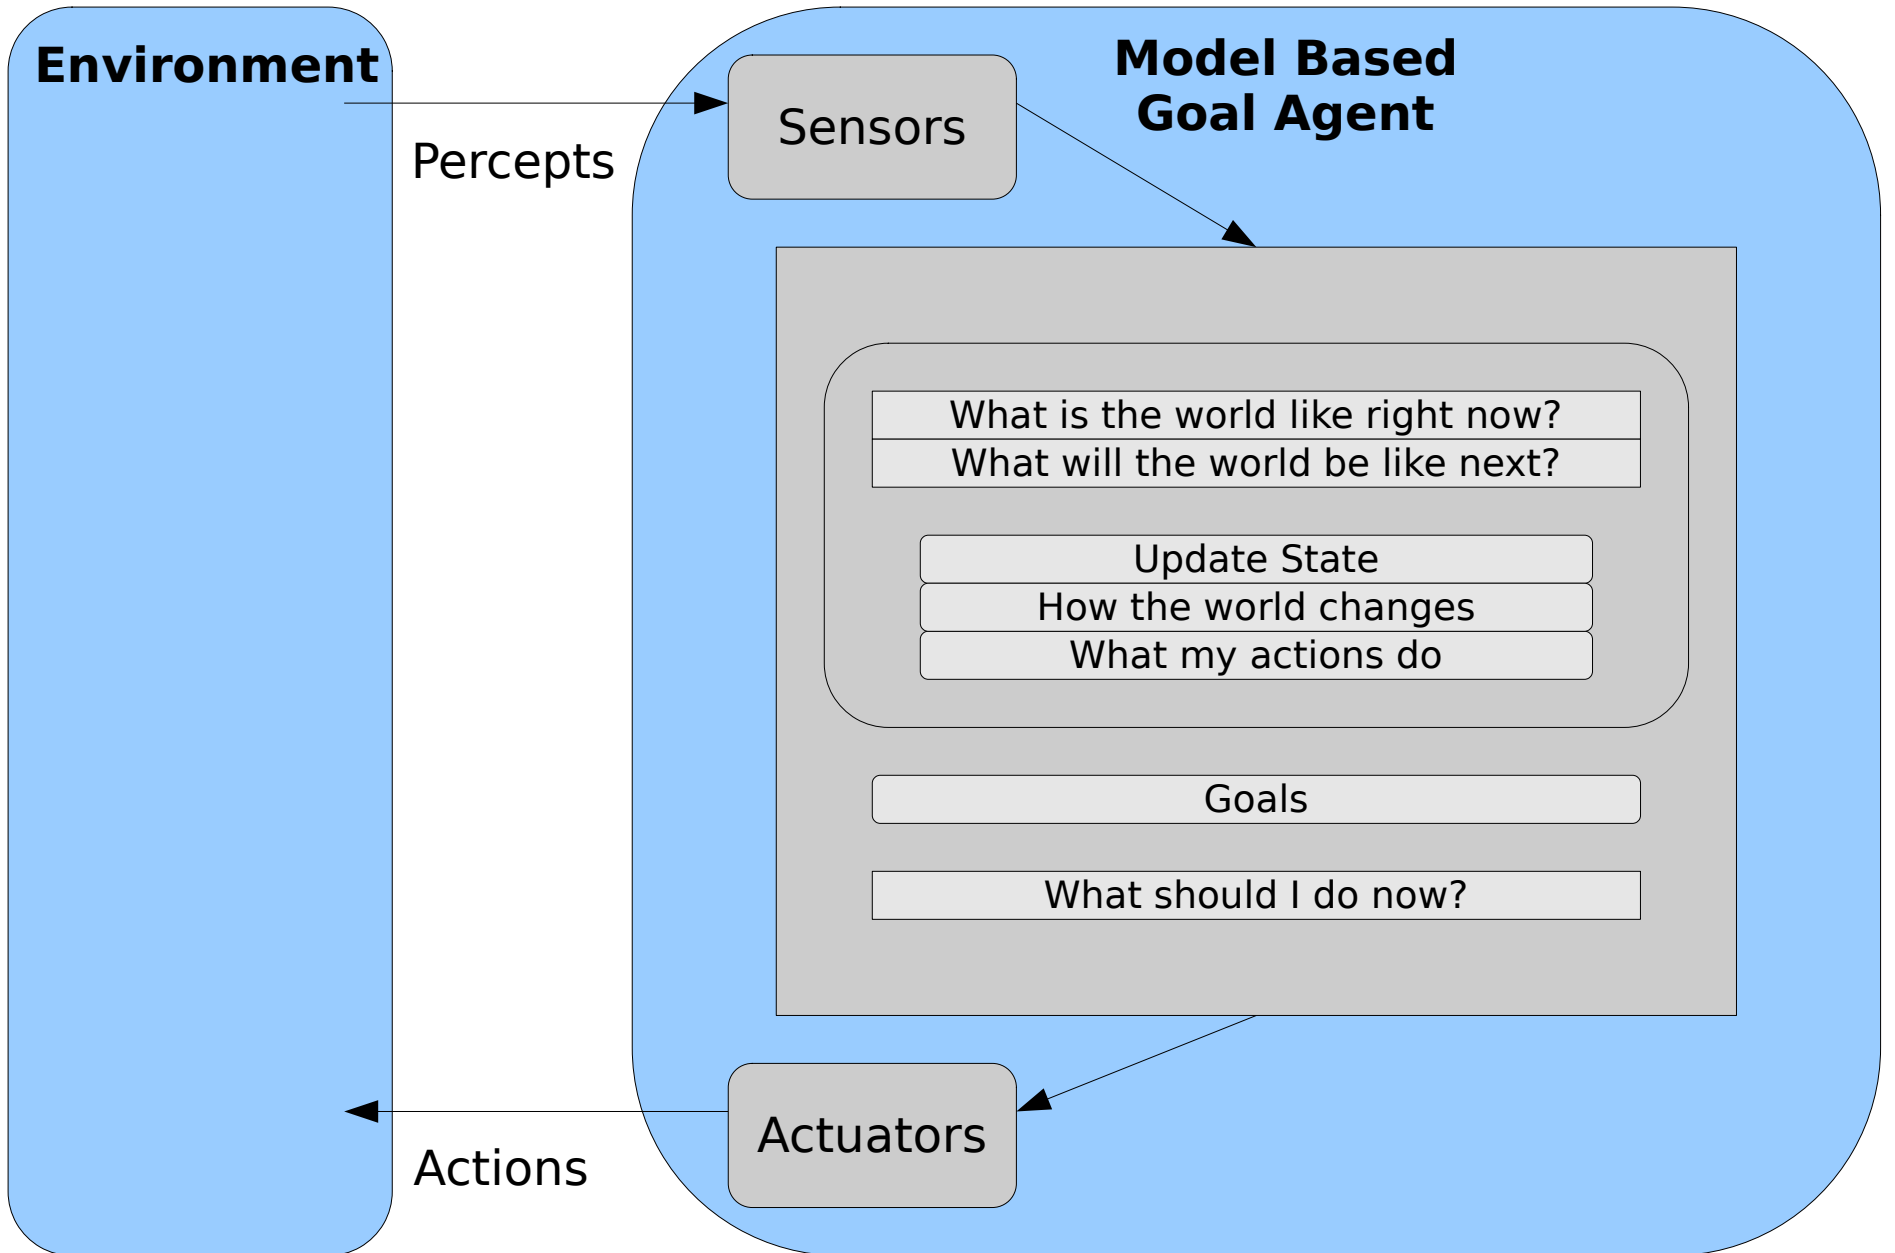
Simple Reflex Agent

- Only looks at current percepts, not entire percept sequence.
- Table only needs to store entry for each possible percept set.
- How much designer bias?
- What if designer misses entries?
- For what kinds of environments is this a good choice? Bad choice?



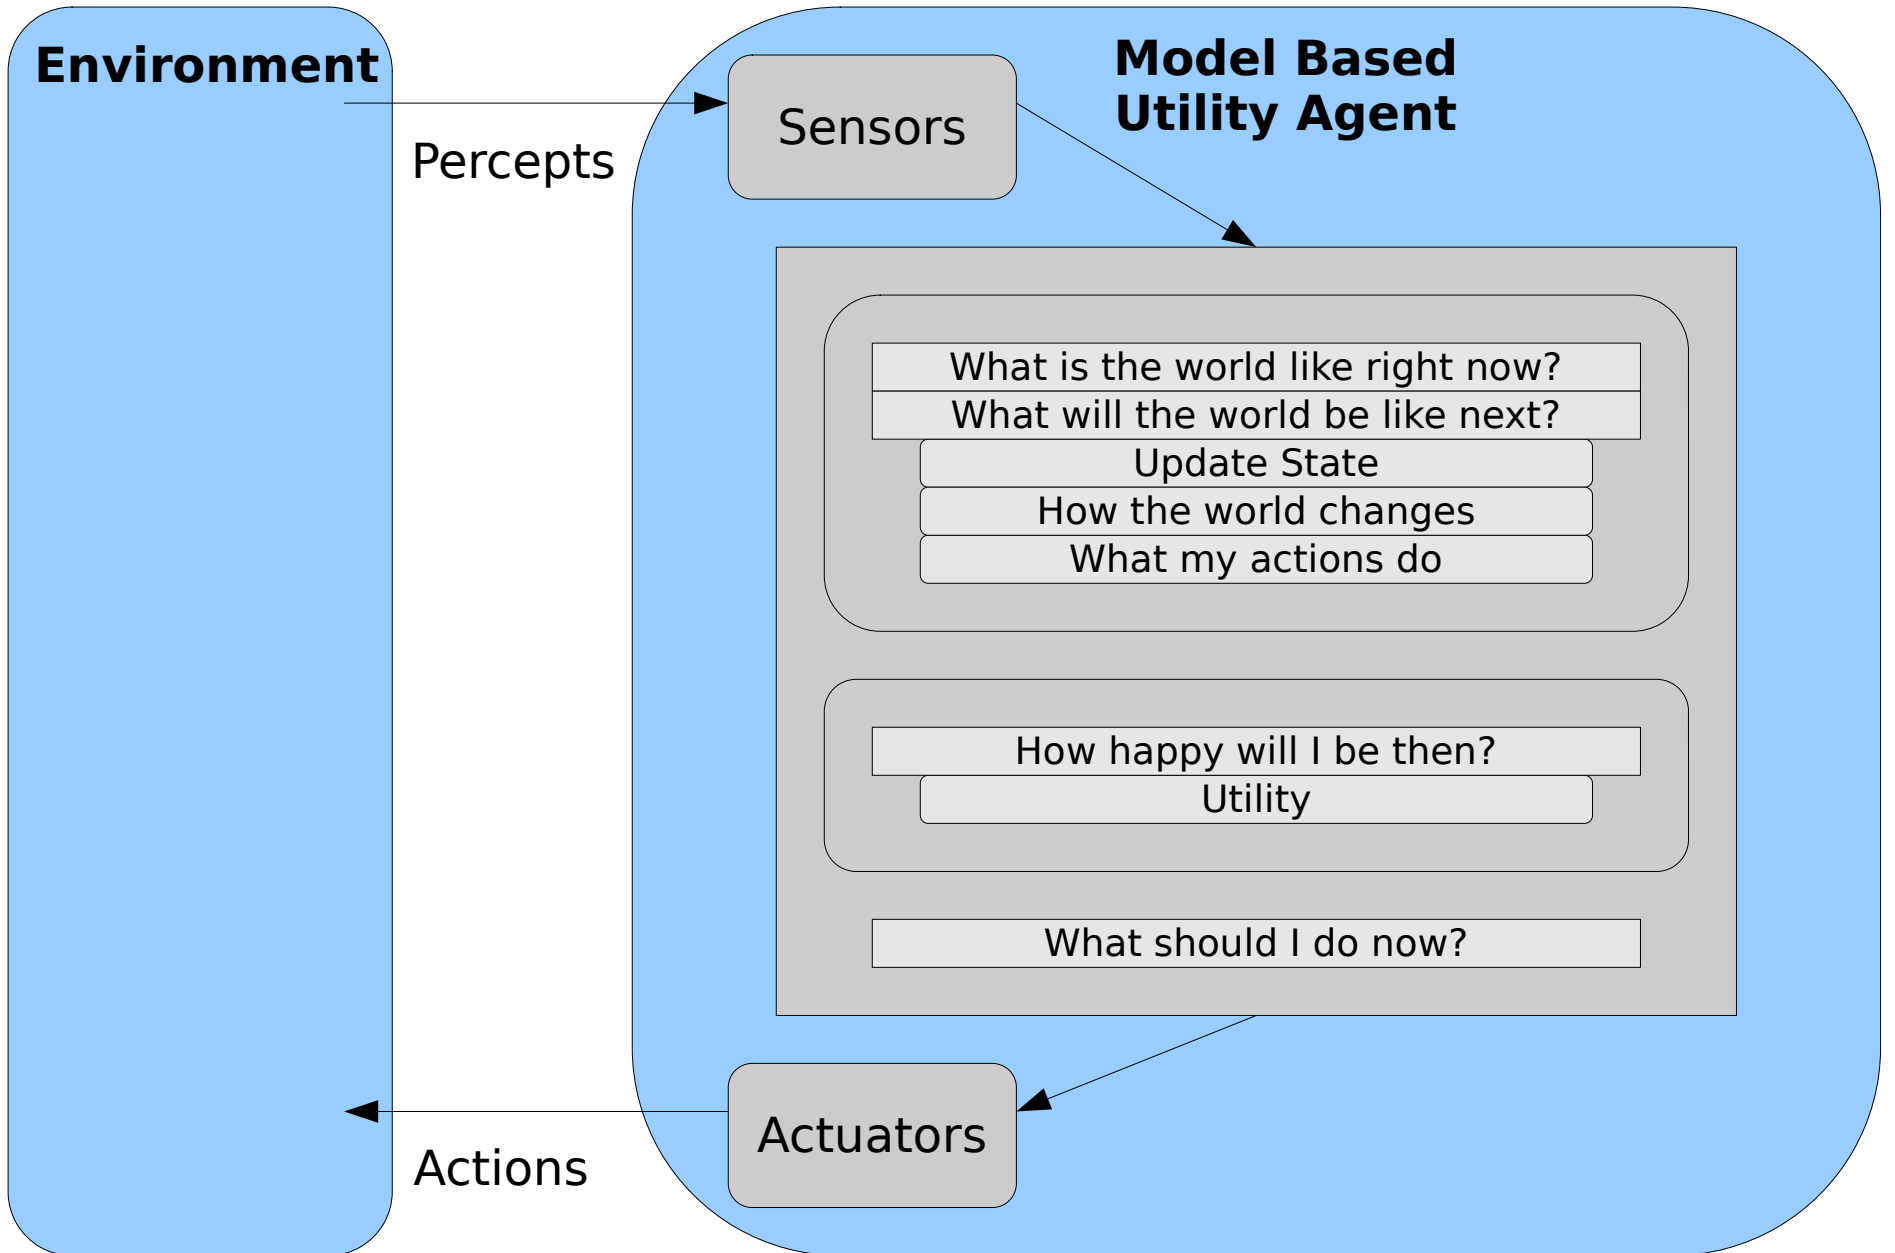
Model Based Reflex Agent

- Updates current state based on percepts
- Table needs to store entry for each possible state.
- How much designer bias?
- What if designer misses entries?
- For what kinds of environments is this a good choice? Bad choice?



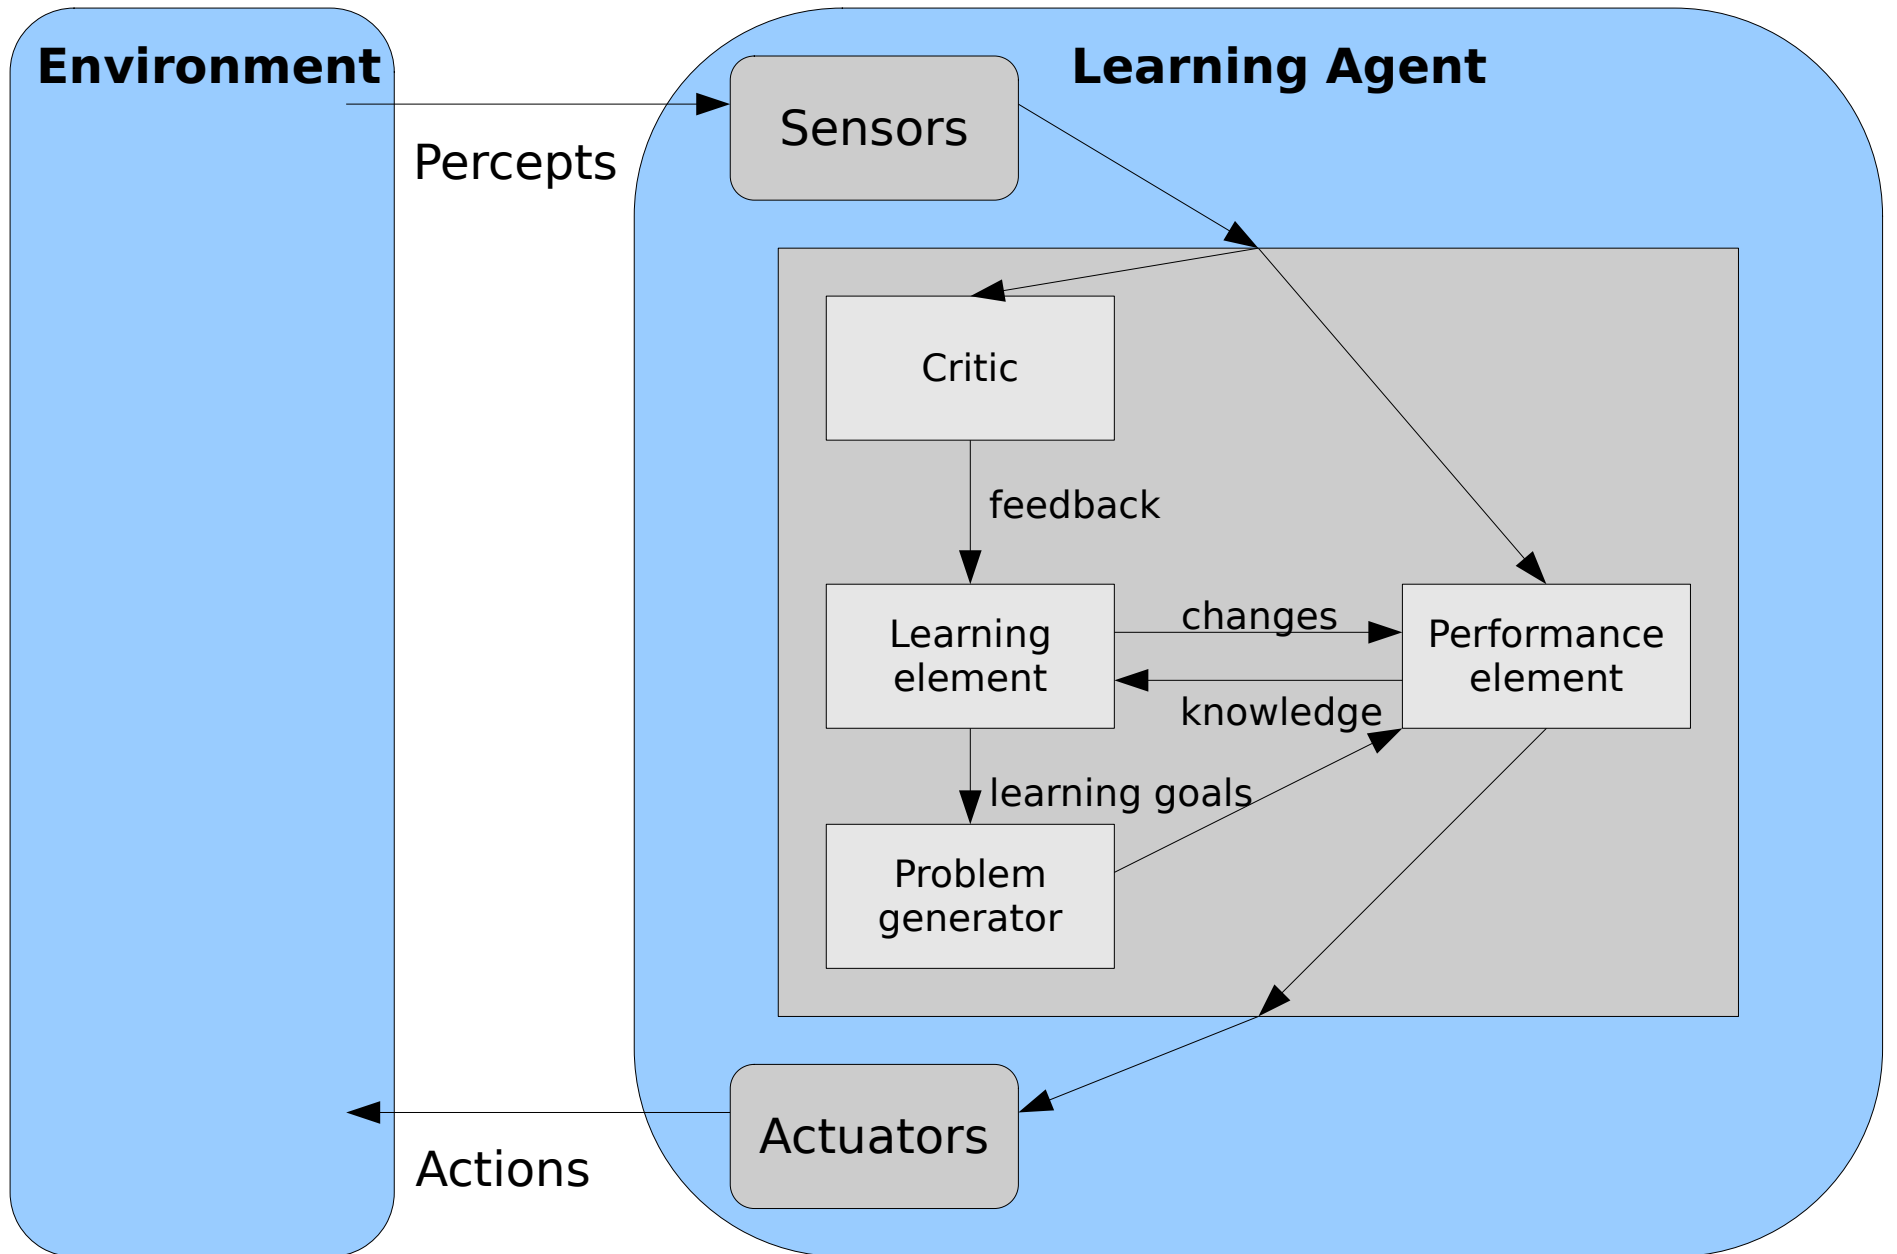
Model Based Goal Agent

- Updates current state based on percepts
- Searches for multi-step path to a goal state.
- How does designer contribute?
- How much designer bias?
- For what kinds of environments is this a good choice? Bad choice?



Model Based Utility Agent

- Updates current state based on percepts
- Searches for multi-step path to goal states. Chooses “best” path.
- How does designer contribute?
- How much designer bias?
- For what kinds of environments is this a good choice? Bad choice?



Learning Agent

- Learning element updates performance element.
- Uses critic to decide what needs to change, and what needs to stay the same.
- Problem generator causes exploration actions to occur, just for learning.
- For what kinds of environments is this a good choice? Bad choice?

Agent Representations

- Atomic
- Factored
- Structured
- Benefits and costs of expressiveness