March 14, 2024

Intelligent Agents, Multi-Agent Systems, & Cybernetics

Dubberly Design Office

PART ONE

The structure of intelligent agents according to Stuart Russell and Peter Norvig

In ‘Artificial Intelligence: A Modern Approach’, Russell & Norvig define and classify ‘Intelligent Agents’

In the field of artificial intelligence, an ‘intelligent agent’ can be defined as an entity that:

1 perceives its environment

2 takes actions autonomously
in order to achieve goals

an intelligent agent might also:

3 improve its performance
with learning or acquiring knowledge

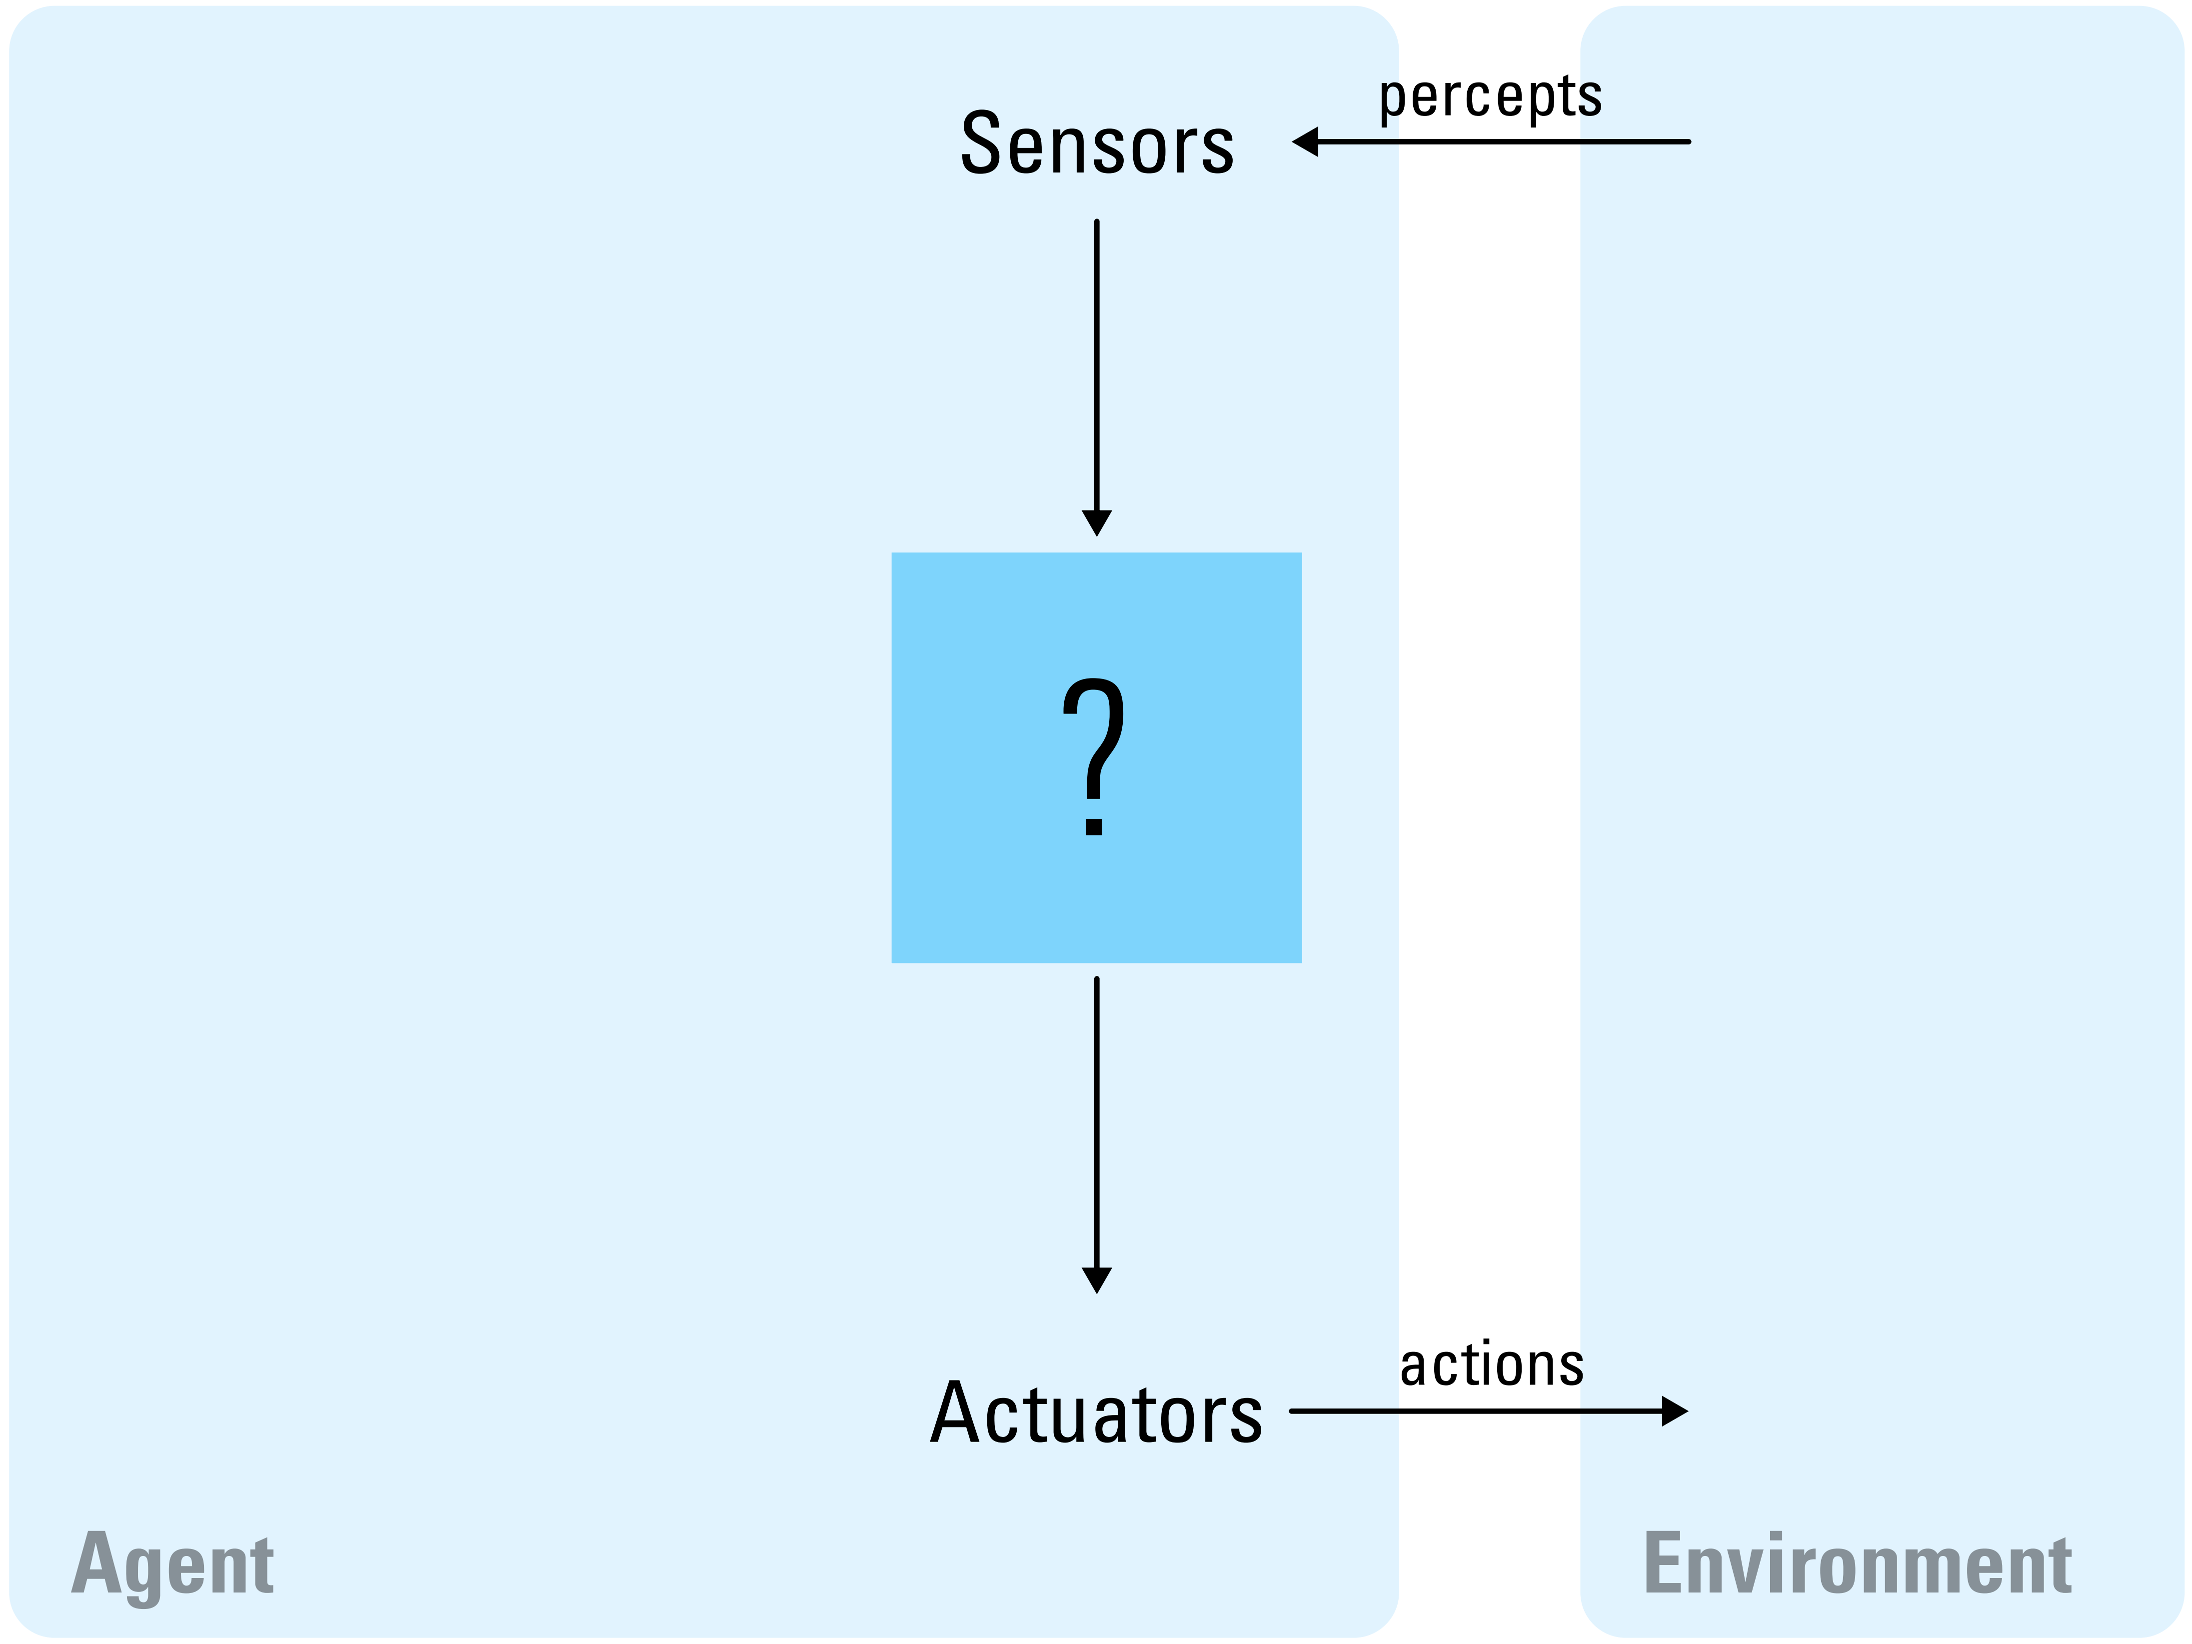

Russell & Norvig, ‘Artificial Intelligence: A Modern Approach’ (2003)
http://aima.cs.berkeley.edu/figures.pdf

Stuart Russell and Peter Norvig group agents into five classes based on their degree of perceived intelligence and capability:

1. Simple reflex agents
2. Model-based reflex agents
3. Goal-based agents
4. Utility-based agents
5. Learning agents

Russell & Norvig, ‘Artificial Intelligence: A Modern Approach’ (2003)
http://aima.cs.berkeley.edu/figures.pdf

Simple reflex agent

'Schematic diagram of a simple reflex agent. We use rectangles to denote the current internal state of the agent's decision process, and ovals to represent the background information used in the process."

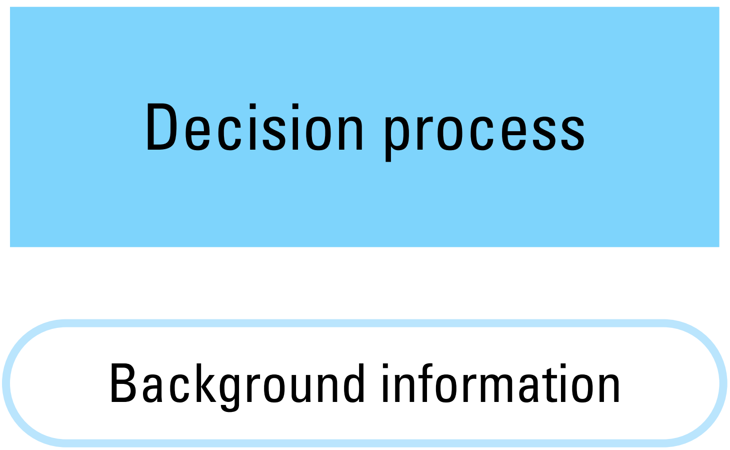

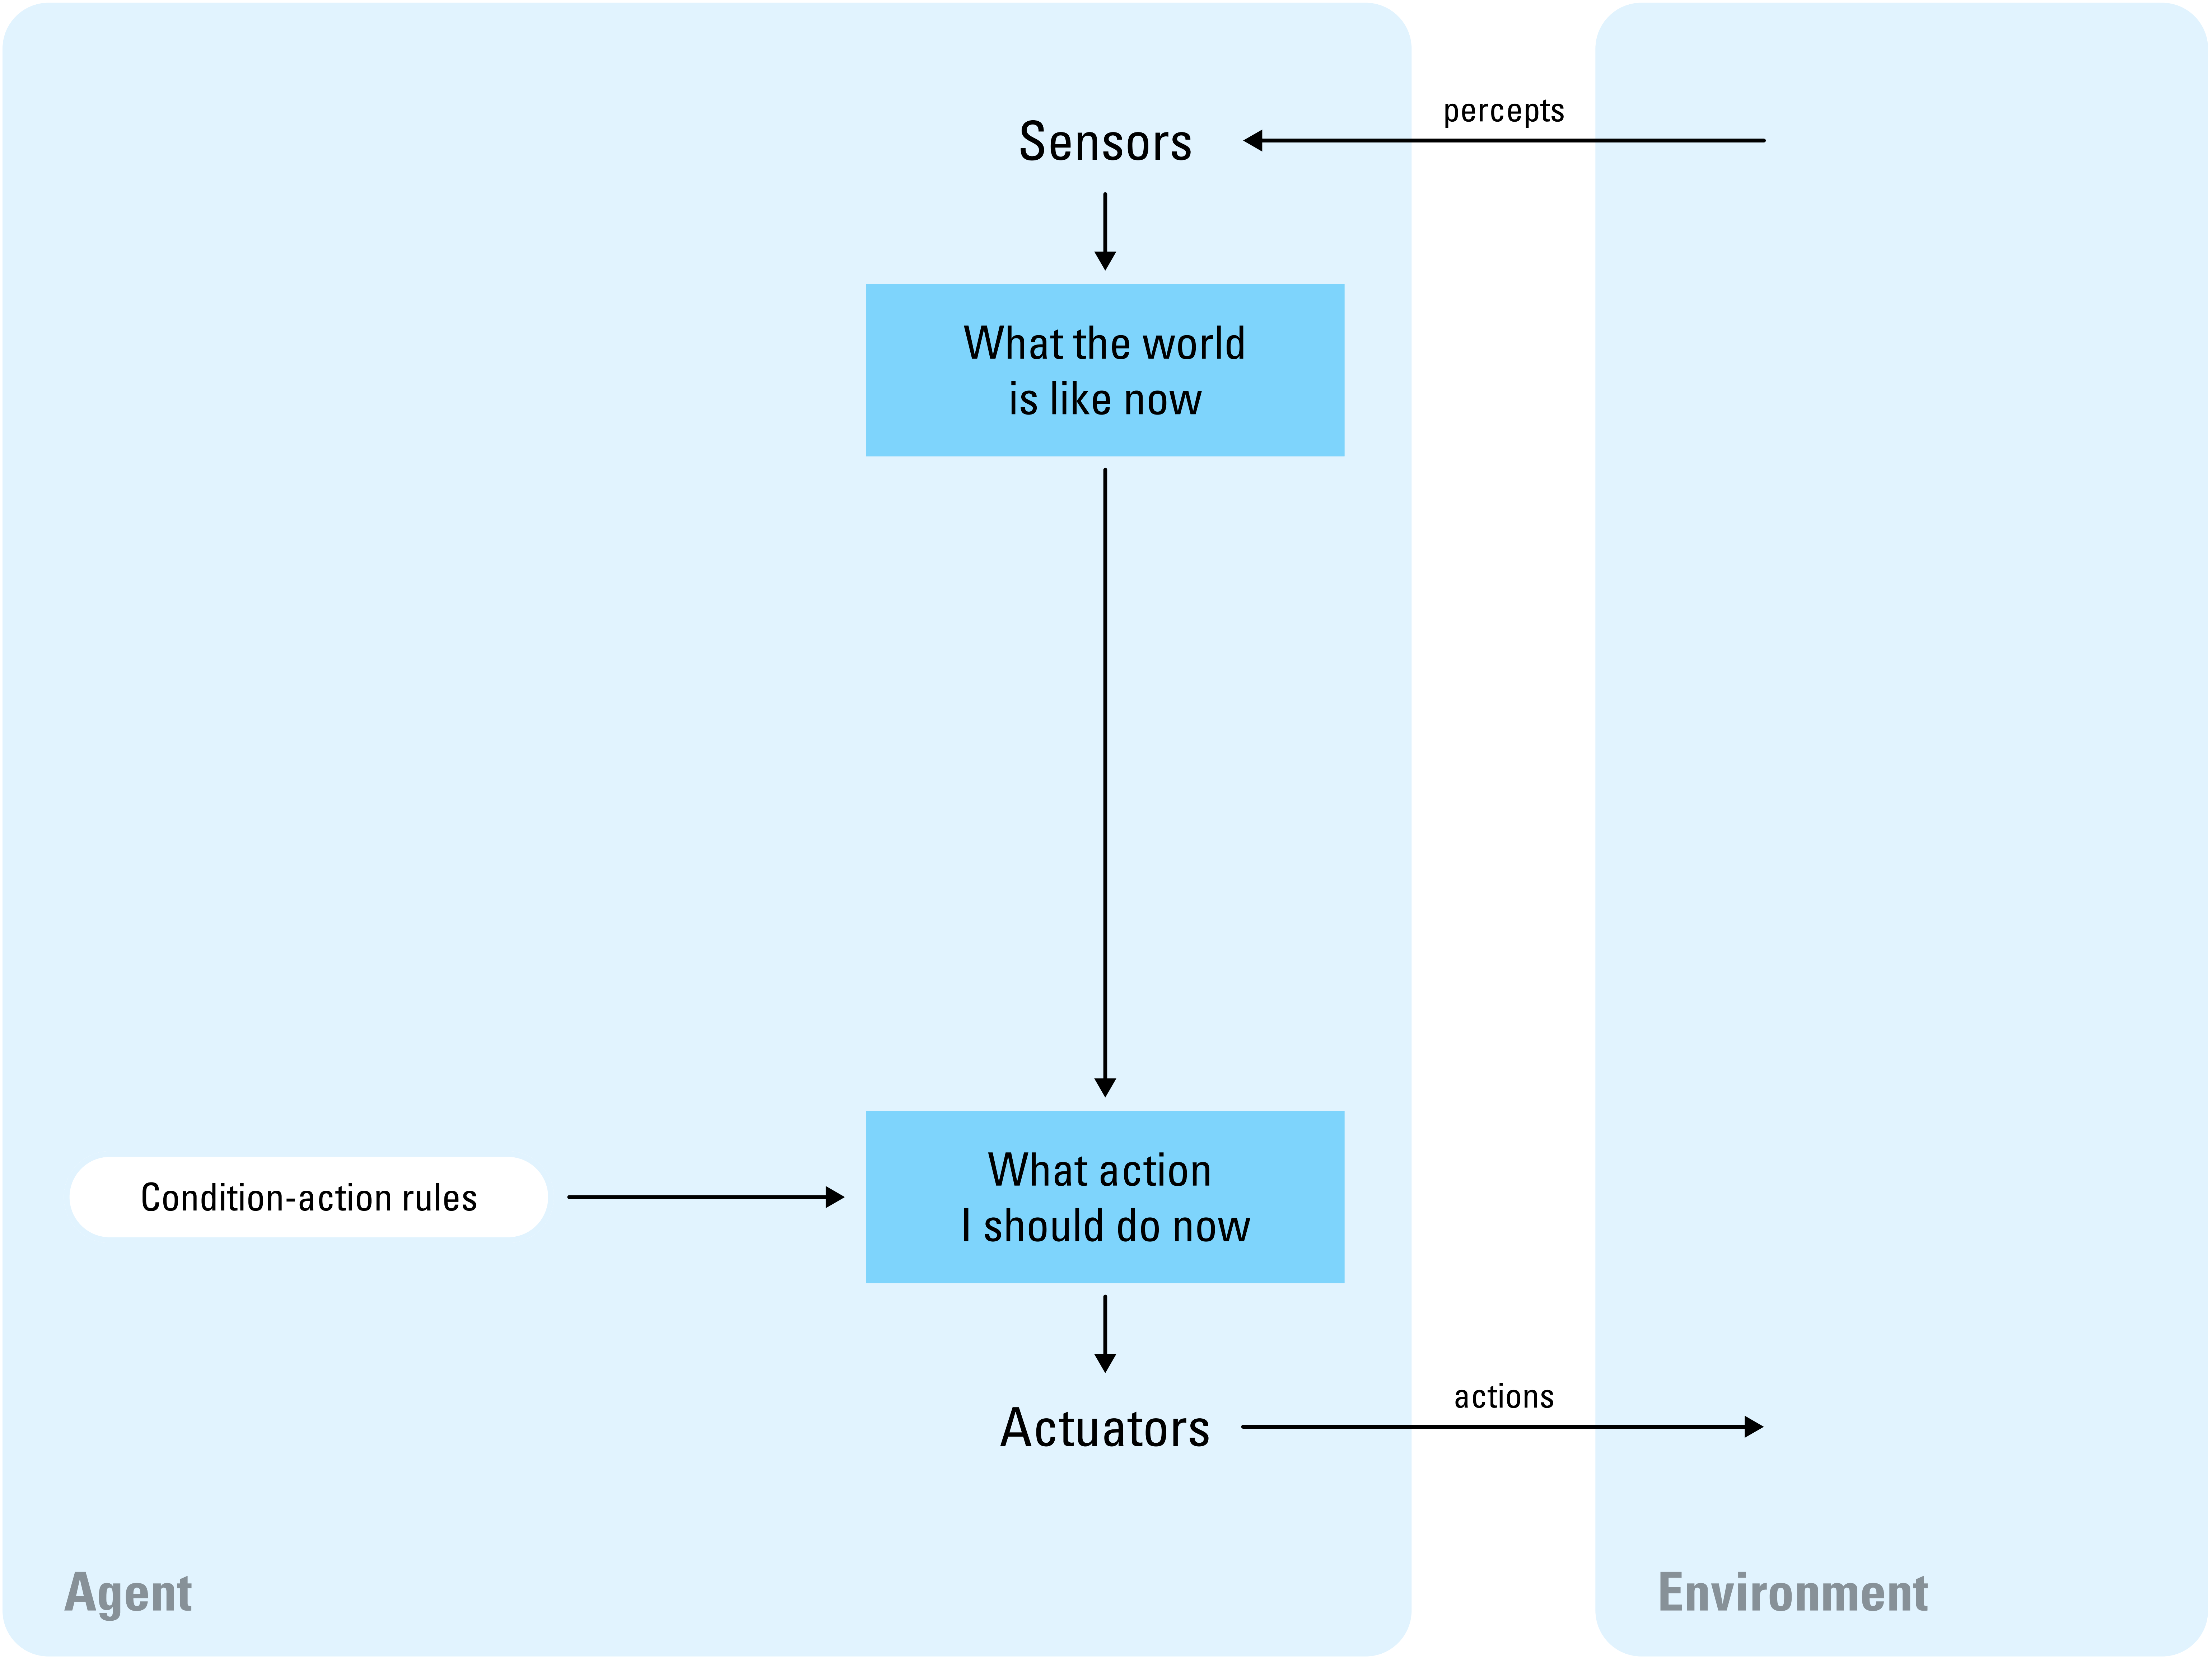

Russell & Norvig, 'Artificial Intelligence: A Modern Approach' (2003)
http://aima.cs.berkeley.edu/figures.pdf

Model-based reflex agent

'A model-based reflex agent. It keeps track of the current state of the world[1], using an internal model. It then chooses an action in the same way as the reflex agent."

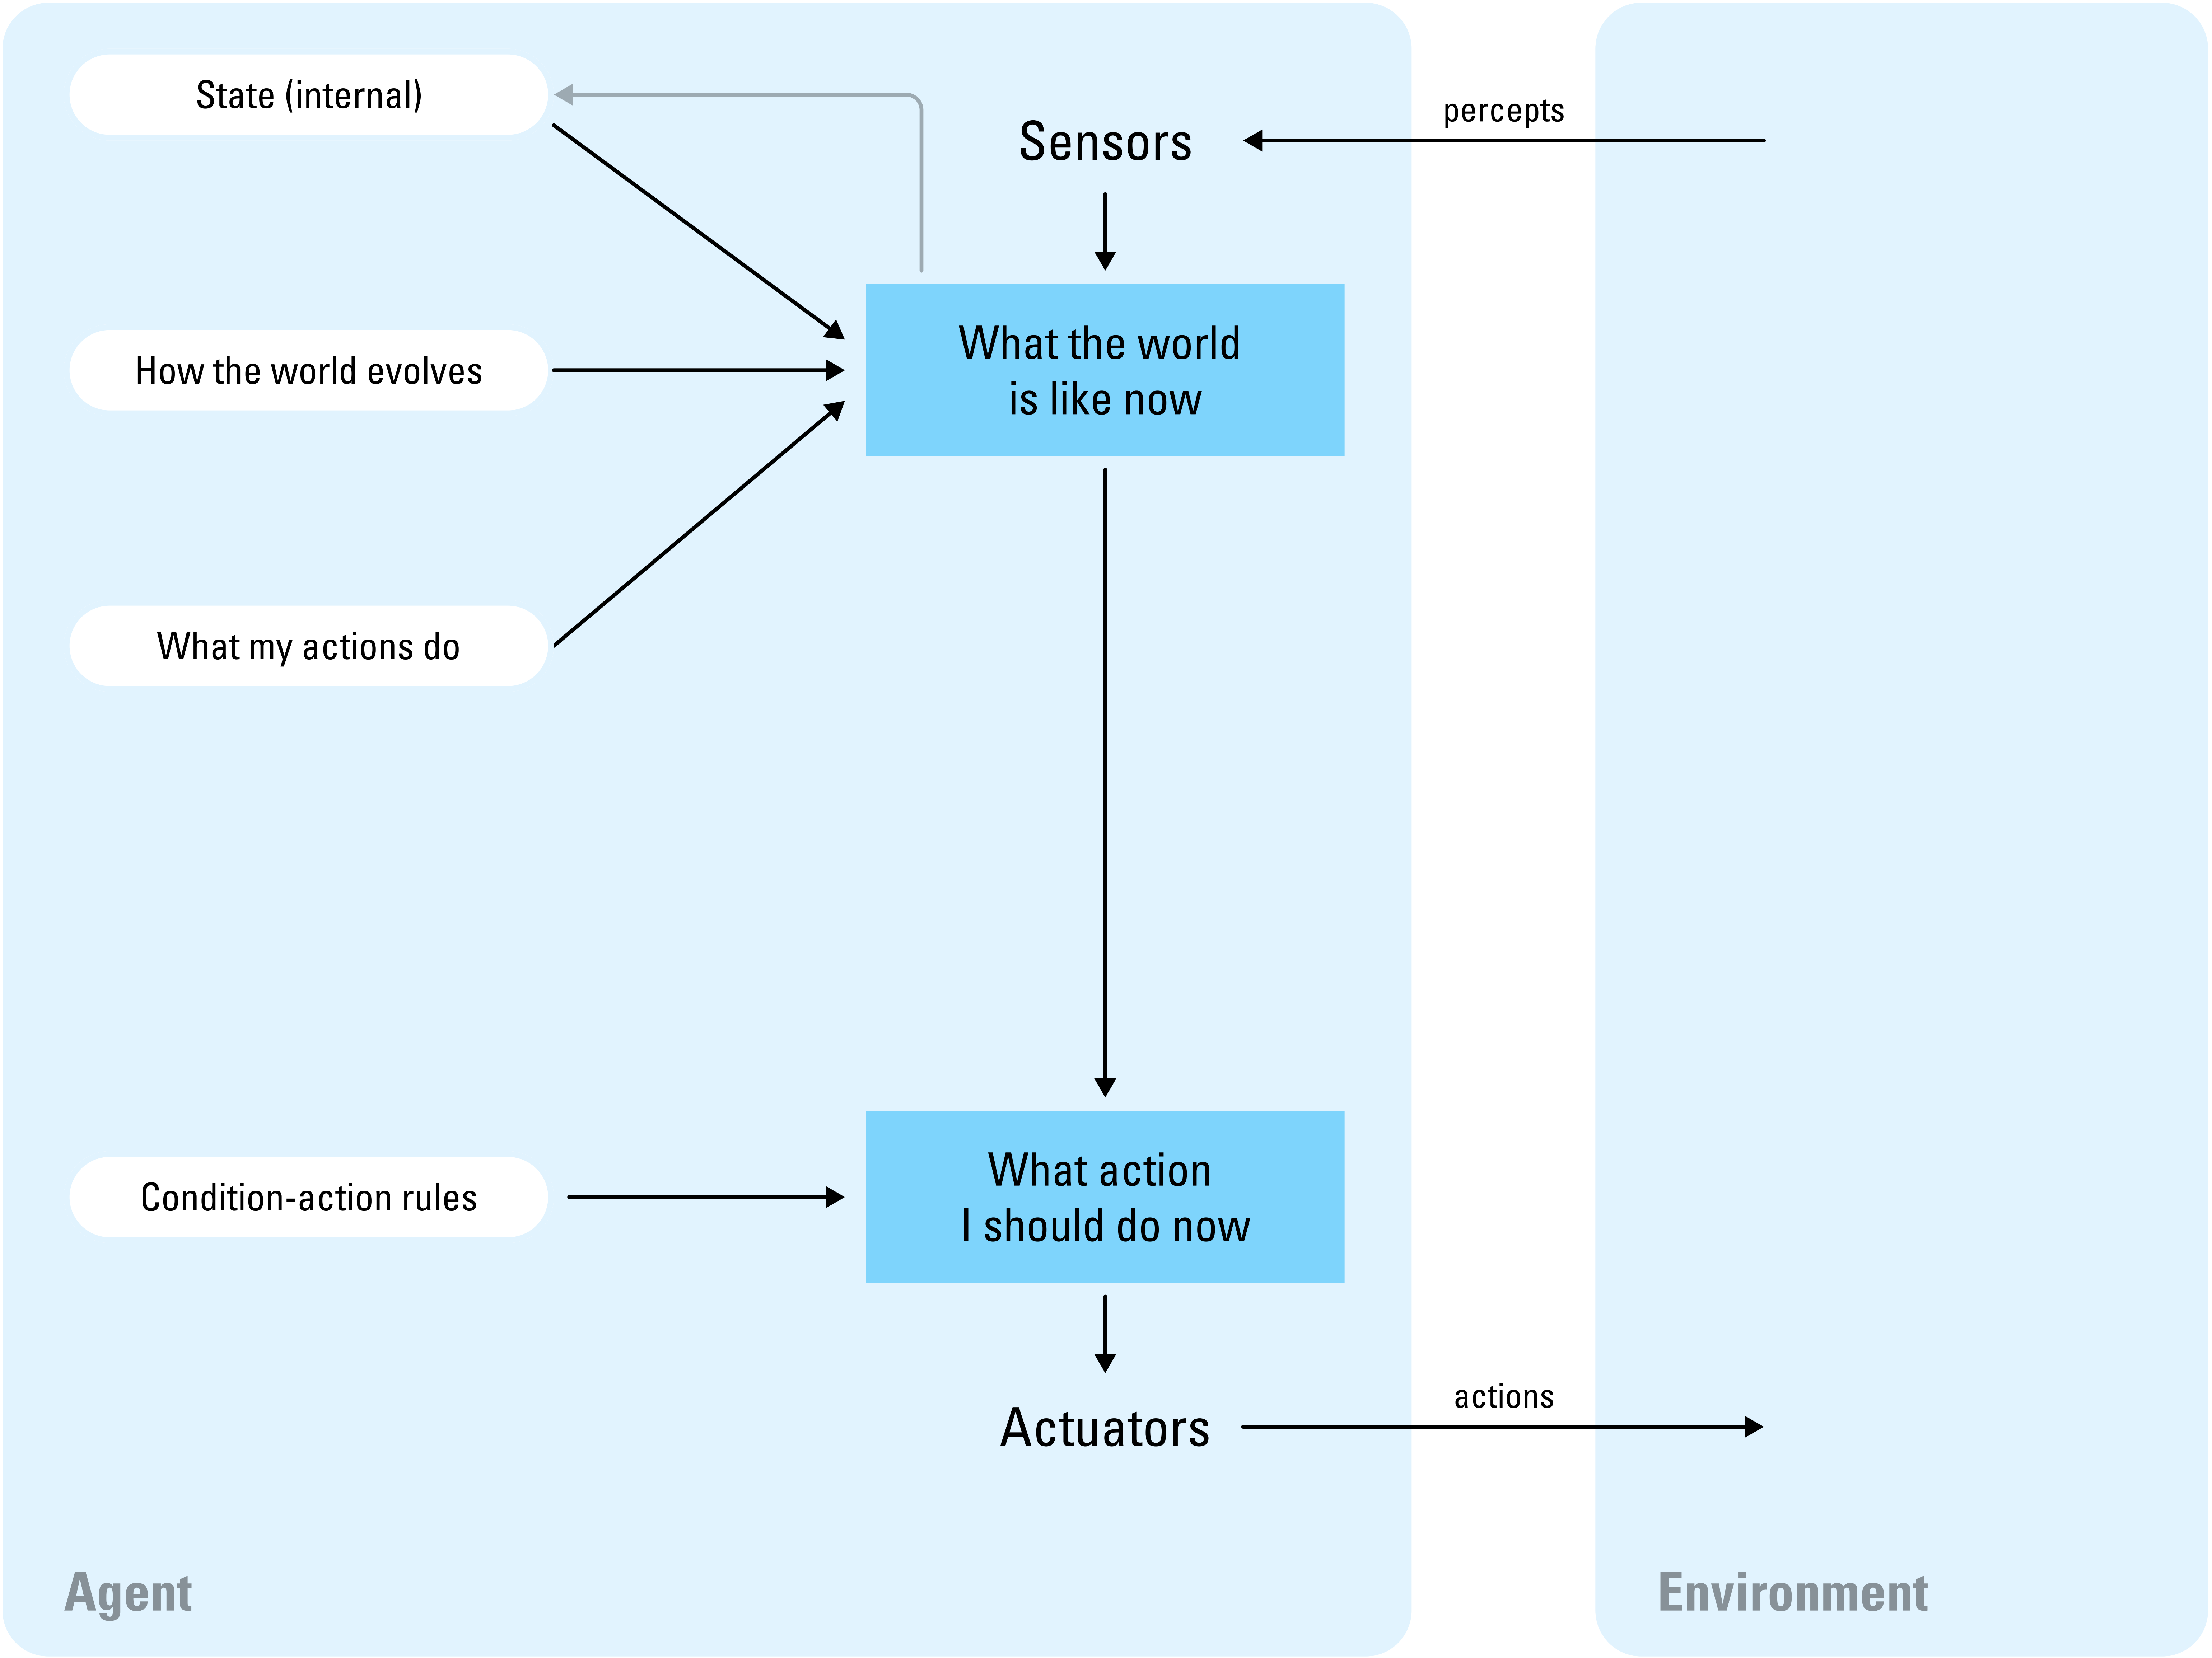

In these contexts, the term 'world' is used synonymously with the agent's 'task environment'.

Russell & Norvig, 'Artificial Intelligence: A Modern Approach' (2003)
http://aima.cs.berkeley.edu/figures.pdf

Model-based, goal-based ag

'A model-based, goal-based agent. It keeps track of the world state as well as a set of goals it is trying to achieve, and chooses an action that will (eventually) lead to the achievement of its goals."

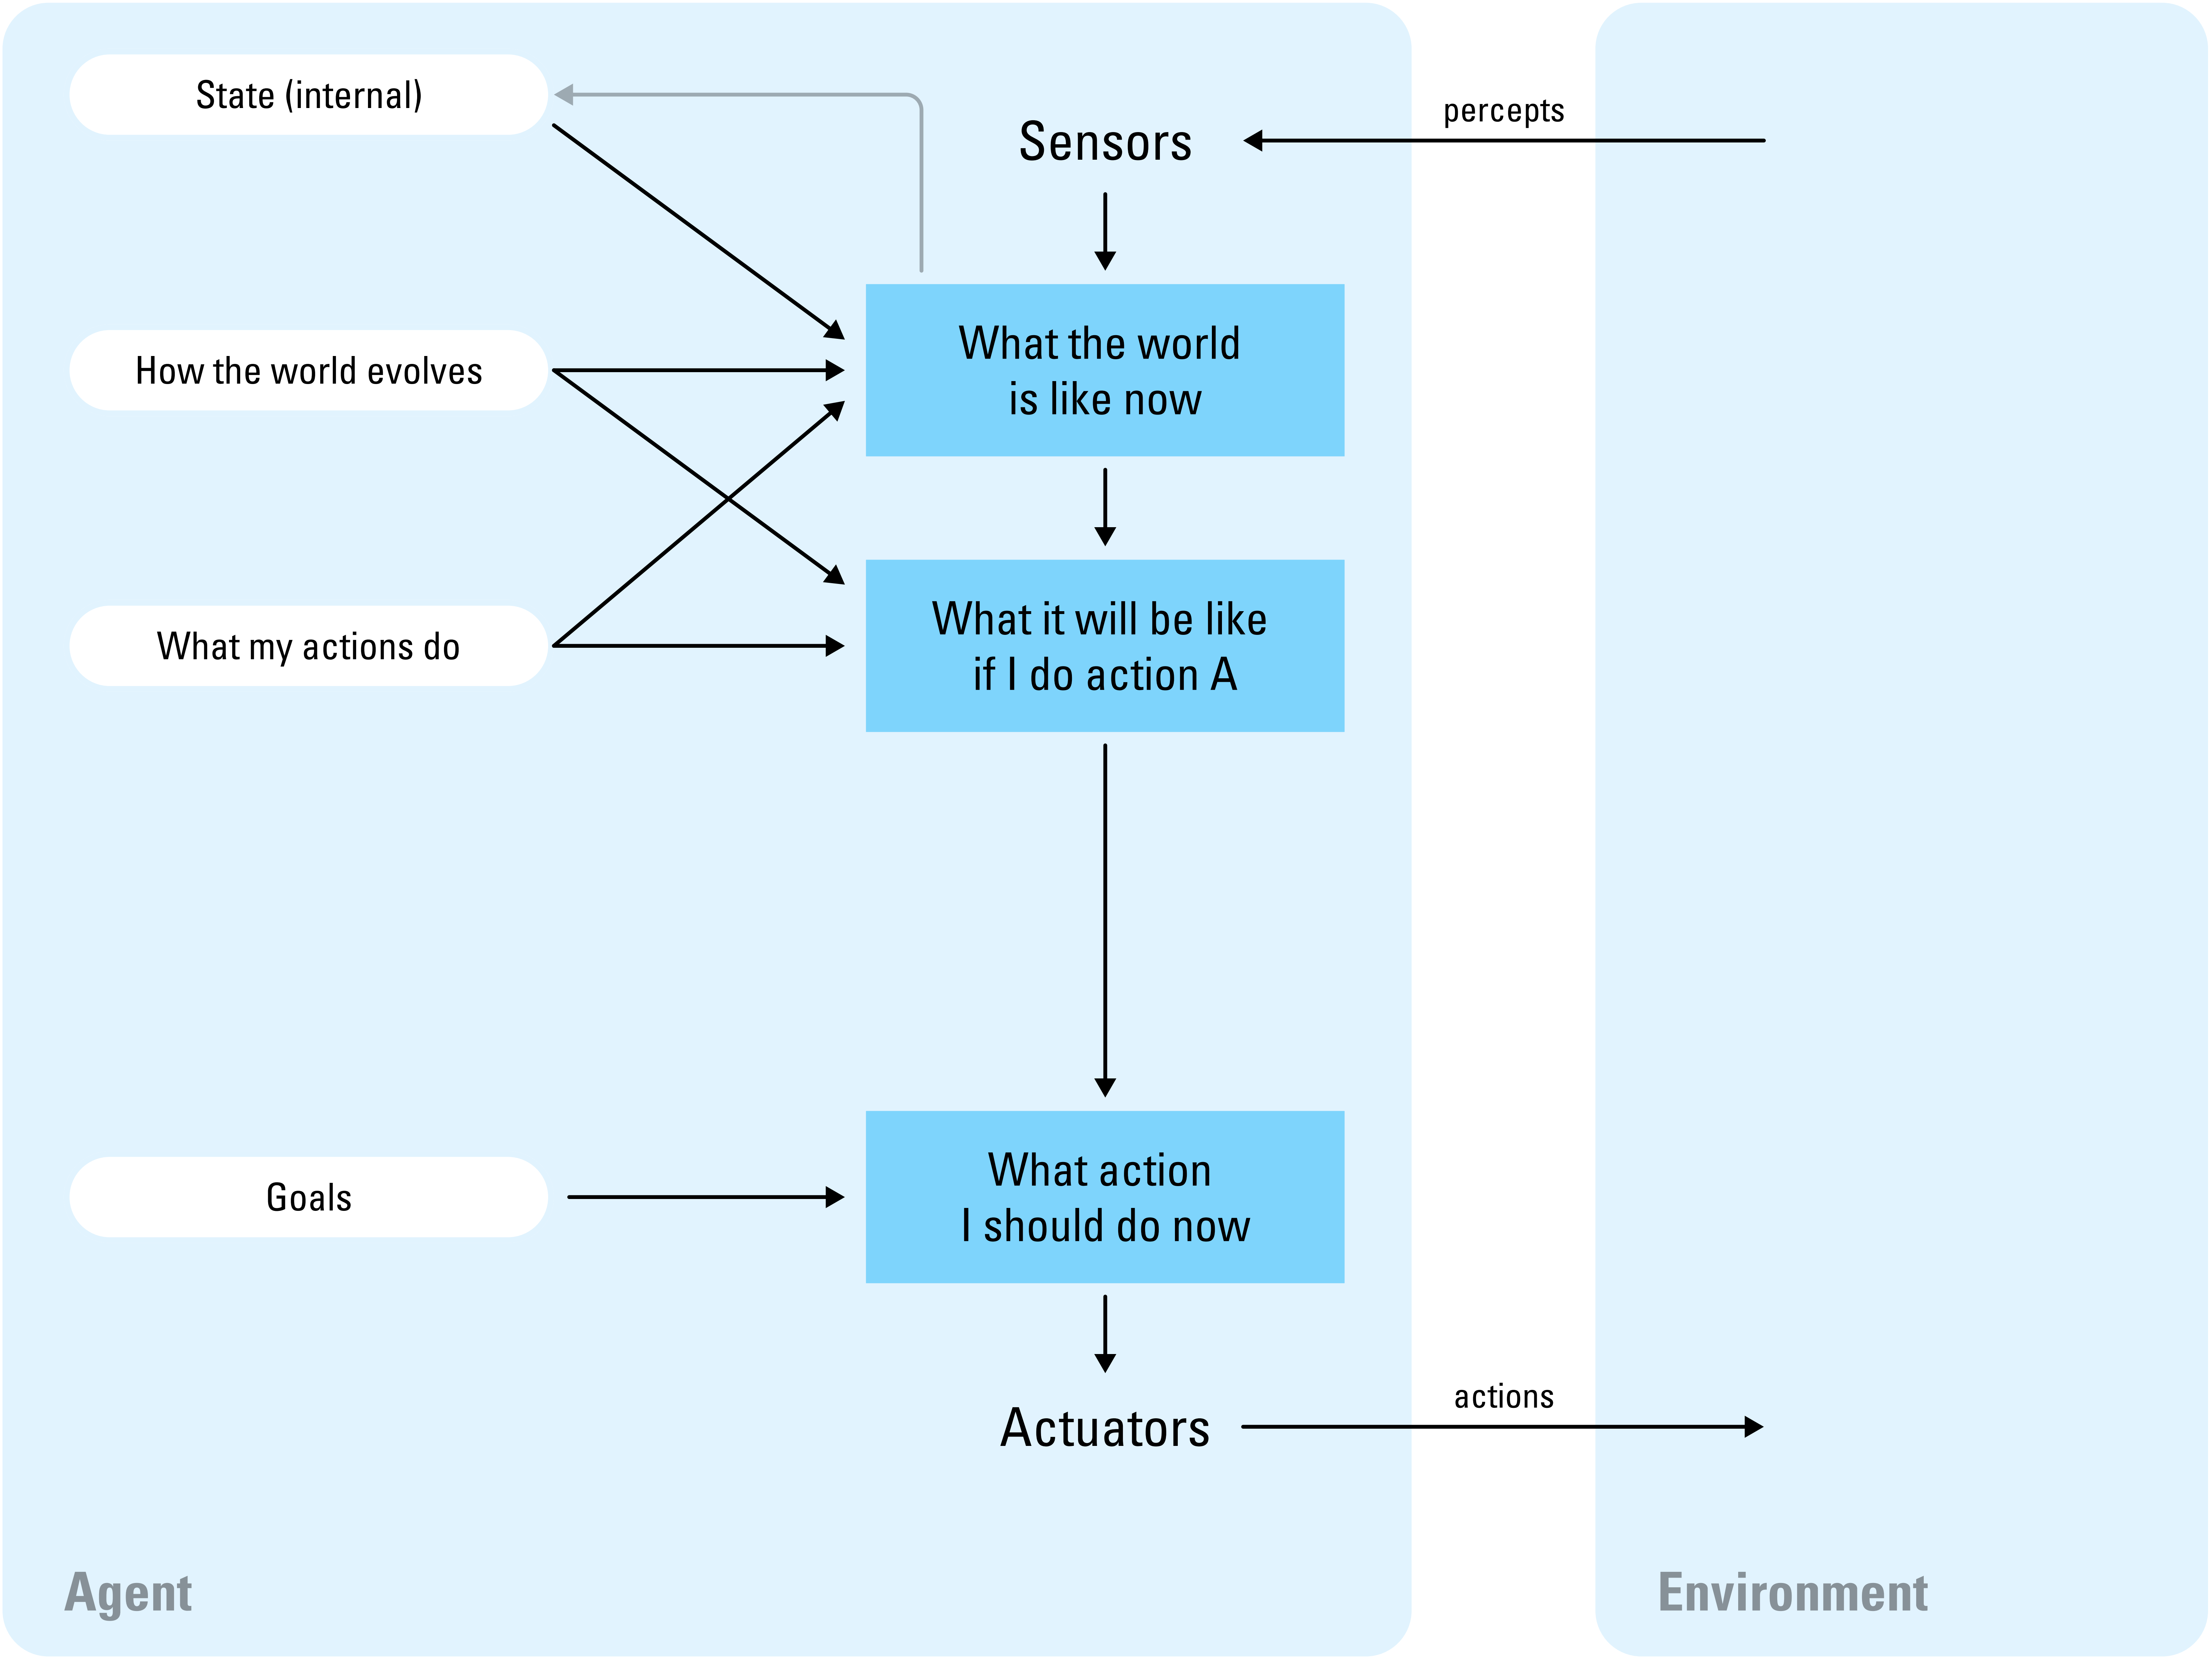

Russell & Norvig, 'Artificial Intelligence: A Modern Approach' (2003)
http://aima.cs.berkeley.edu/figures.pdf

Model-based, utility-based a

'A model-based, utility-based agent. It uses a model of the world, along with a utility function that measures its preferences among states of the world. Then it chooses the action that leads to the best expected utility, where expected utility is computed by averaging over all possible outcome states, weighted by the probability of the outcome."

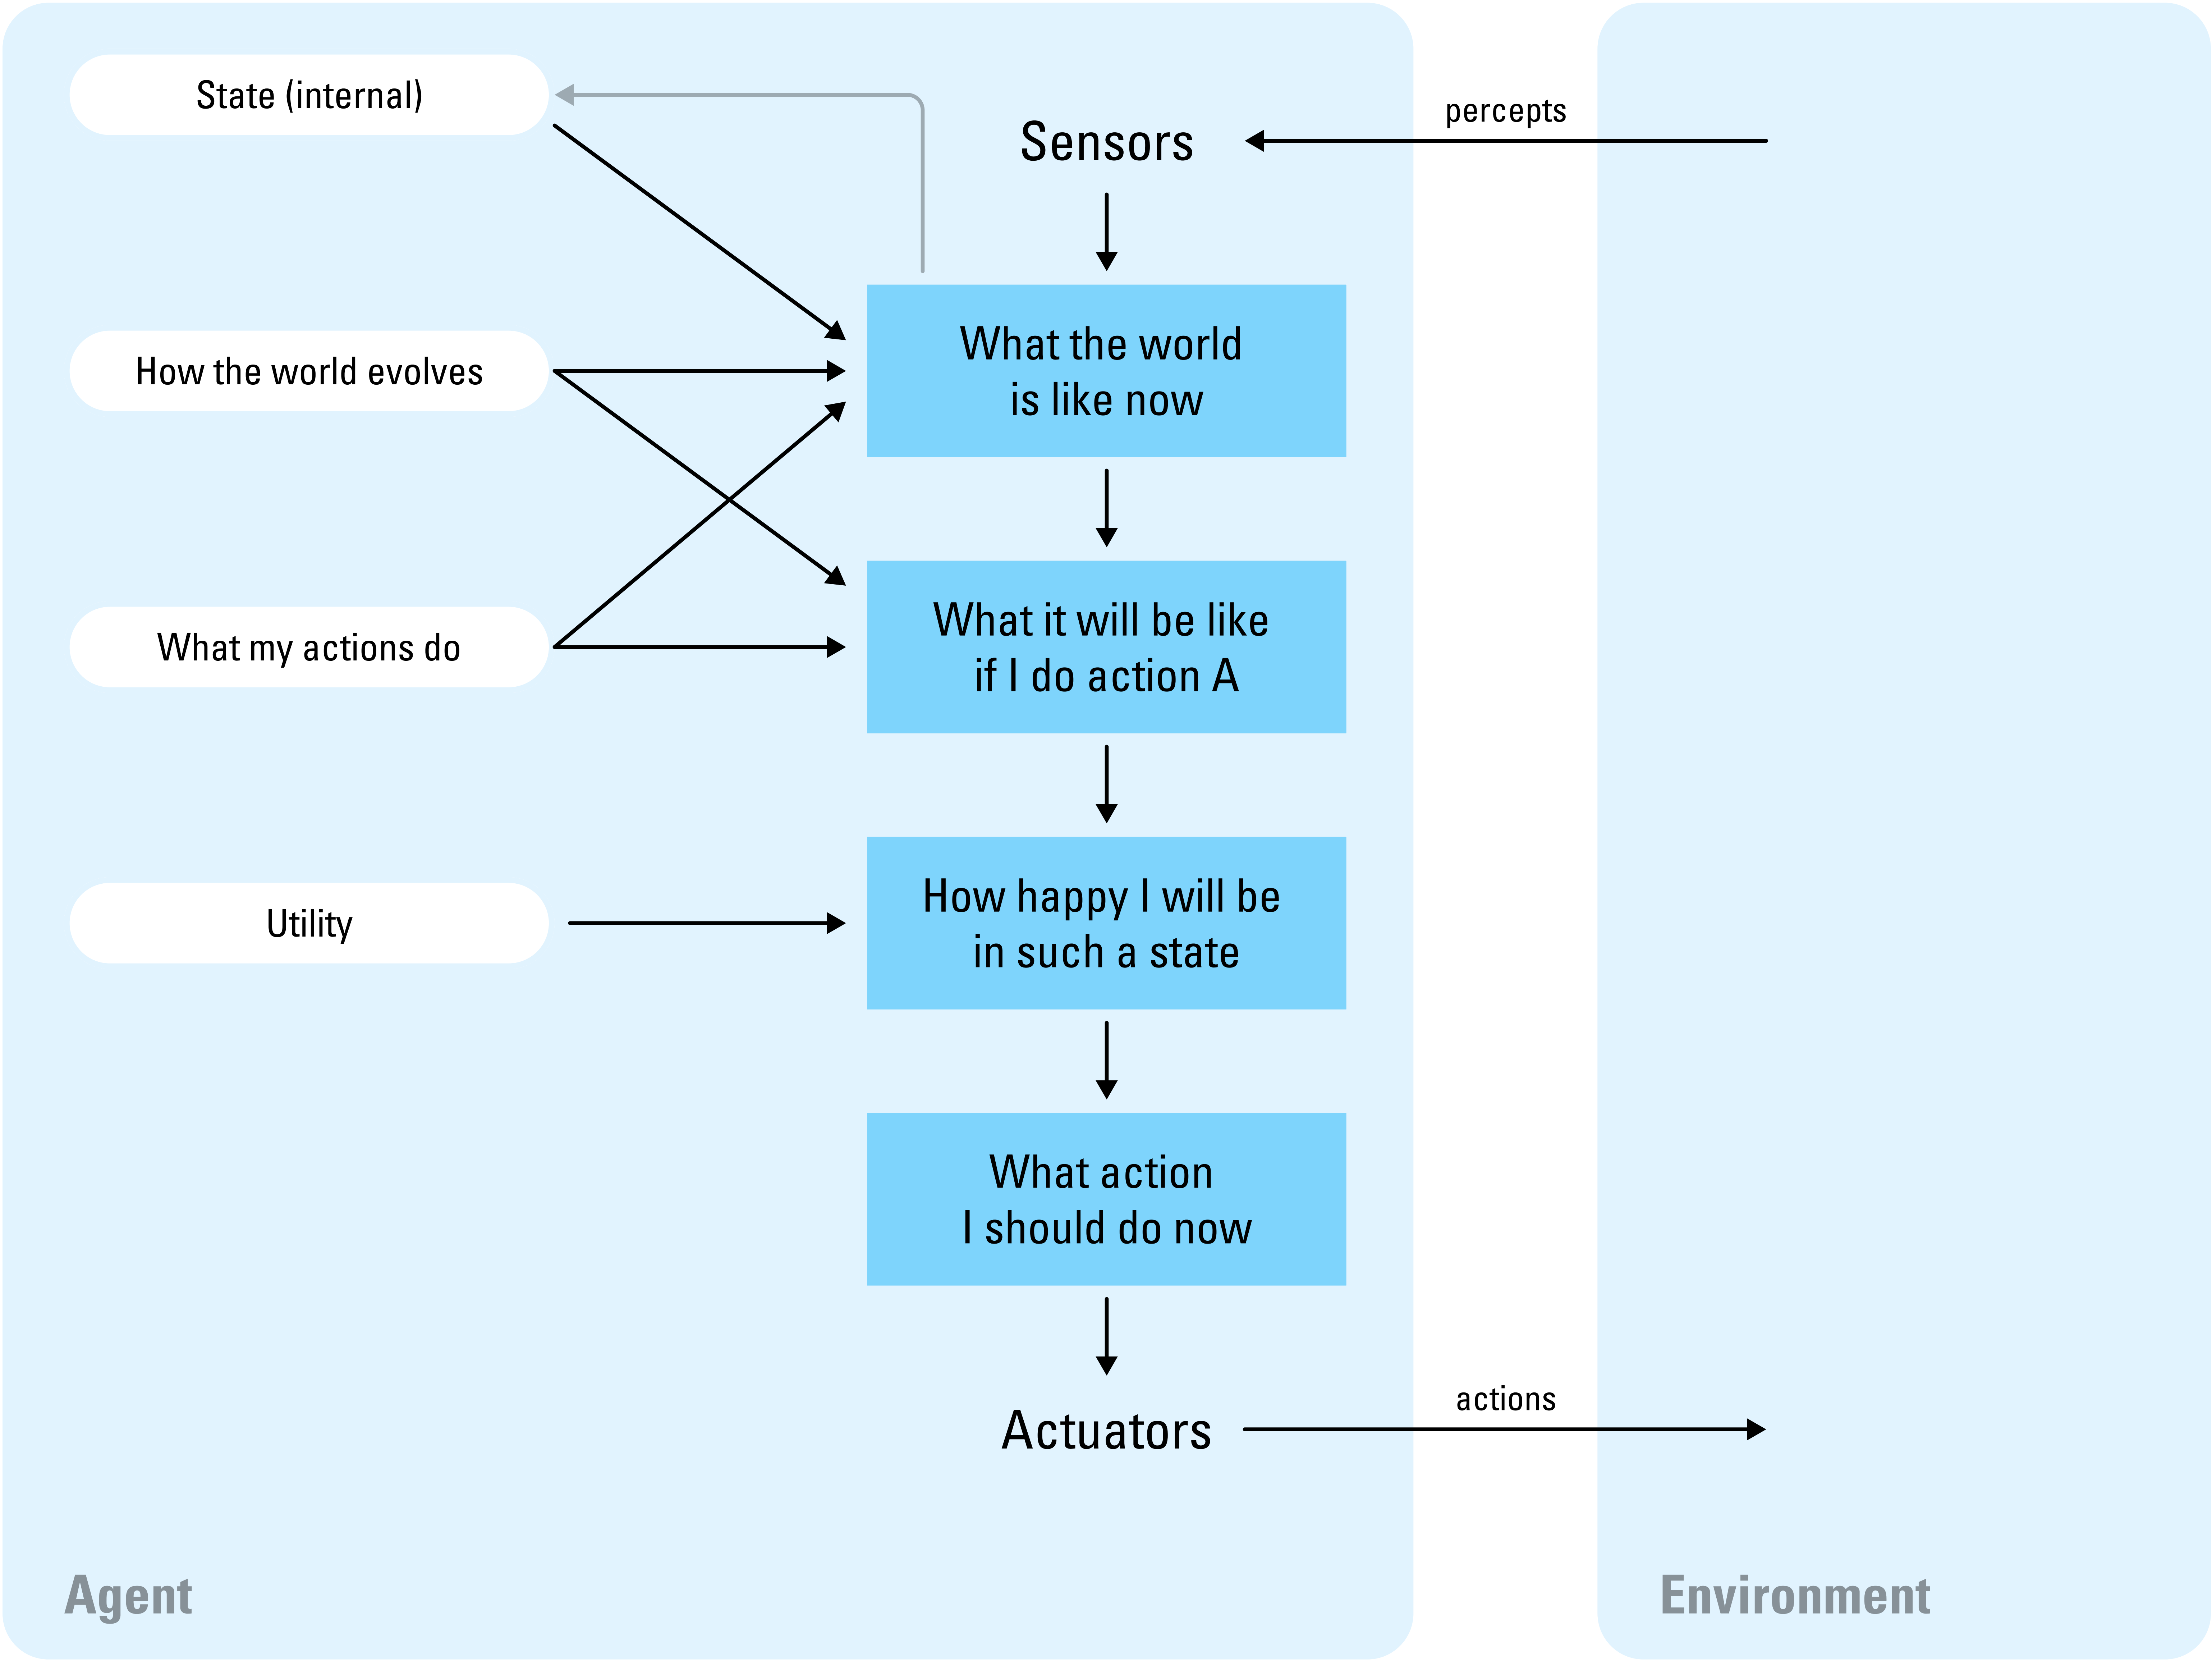

Russell & Norvig, 'Artificial Intelligence: A Modern Approach' (2003)
http://aima.cs.berkeley.edu/figures.pdf

General learning agent

'A general learning agent. The 'performance element' box represents what we have previously considered to be the whole agent program. Now, the 'learning element' box gets to modify that program to improve its performance."

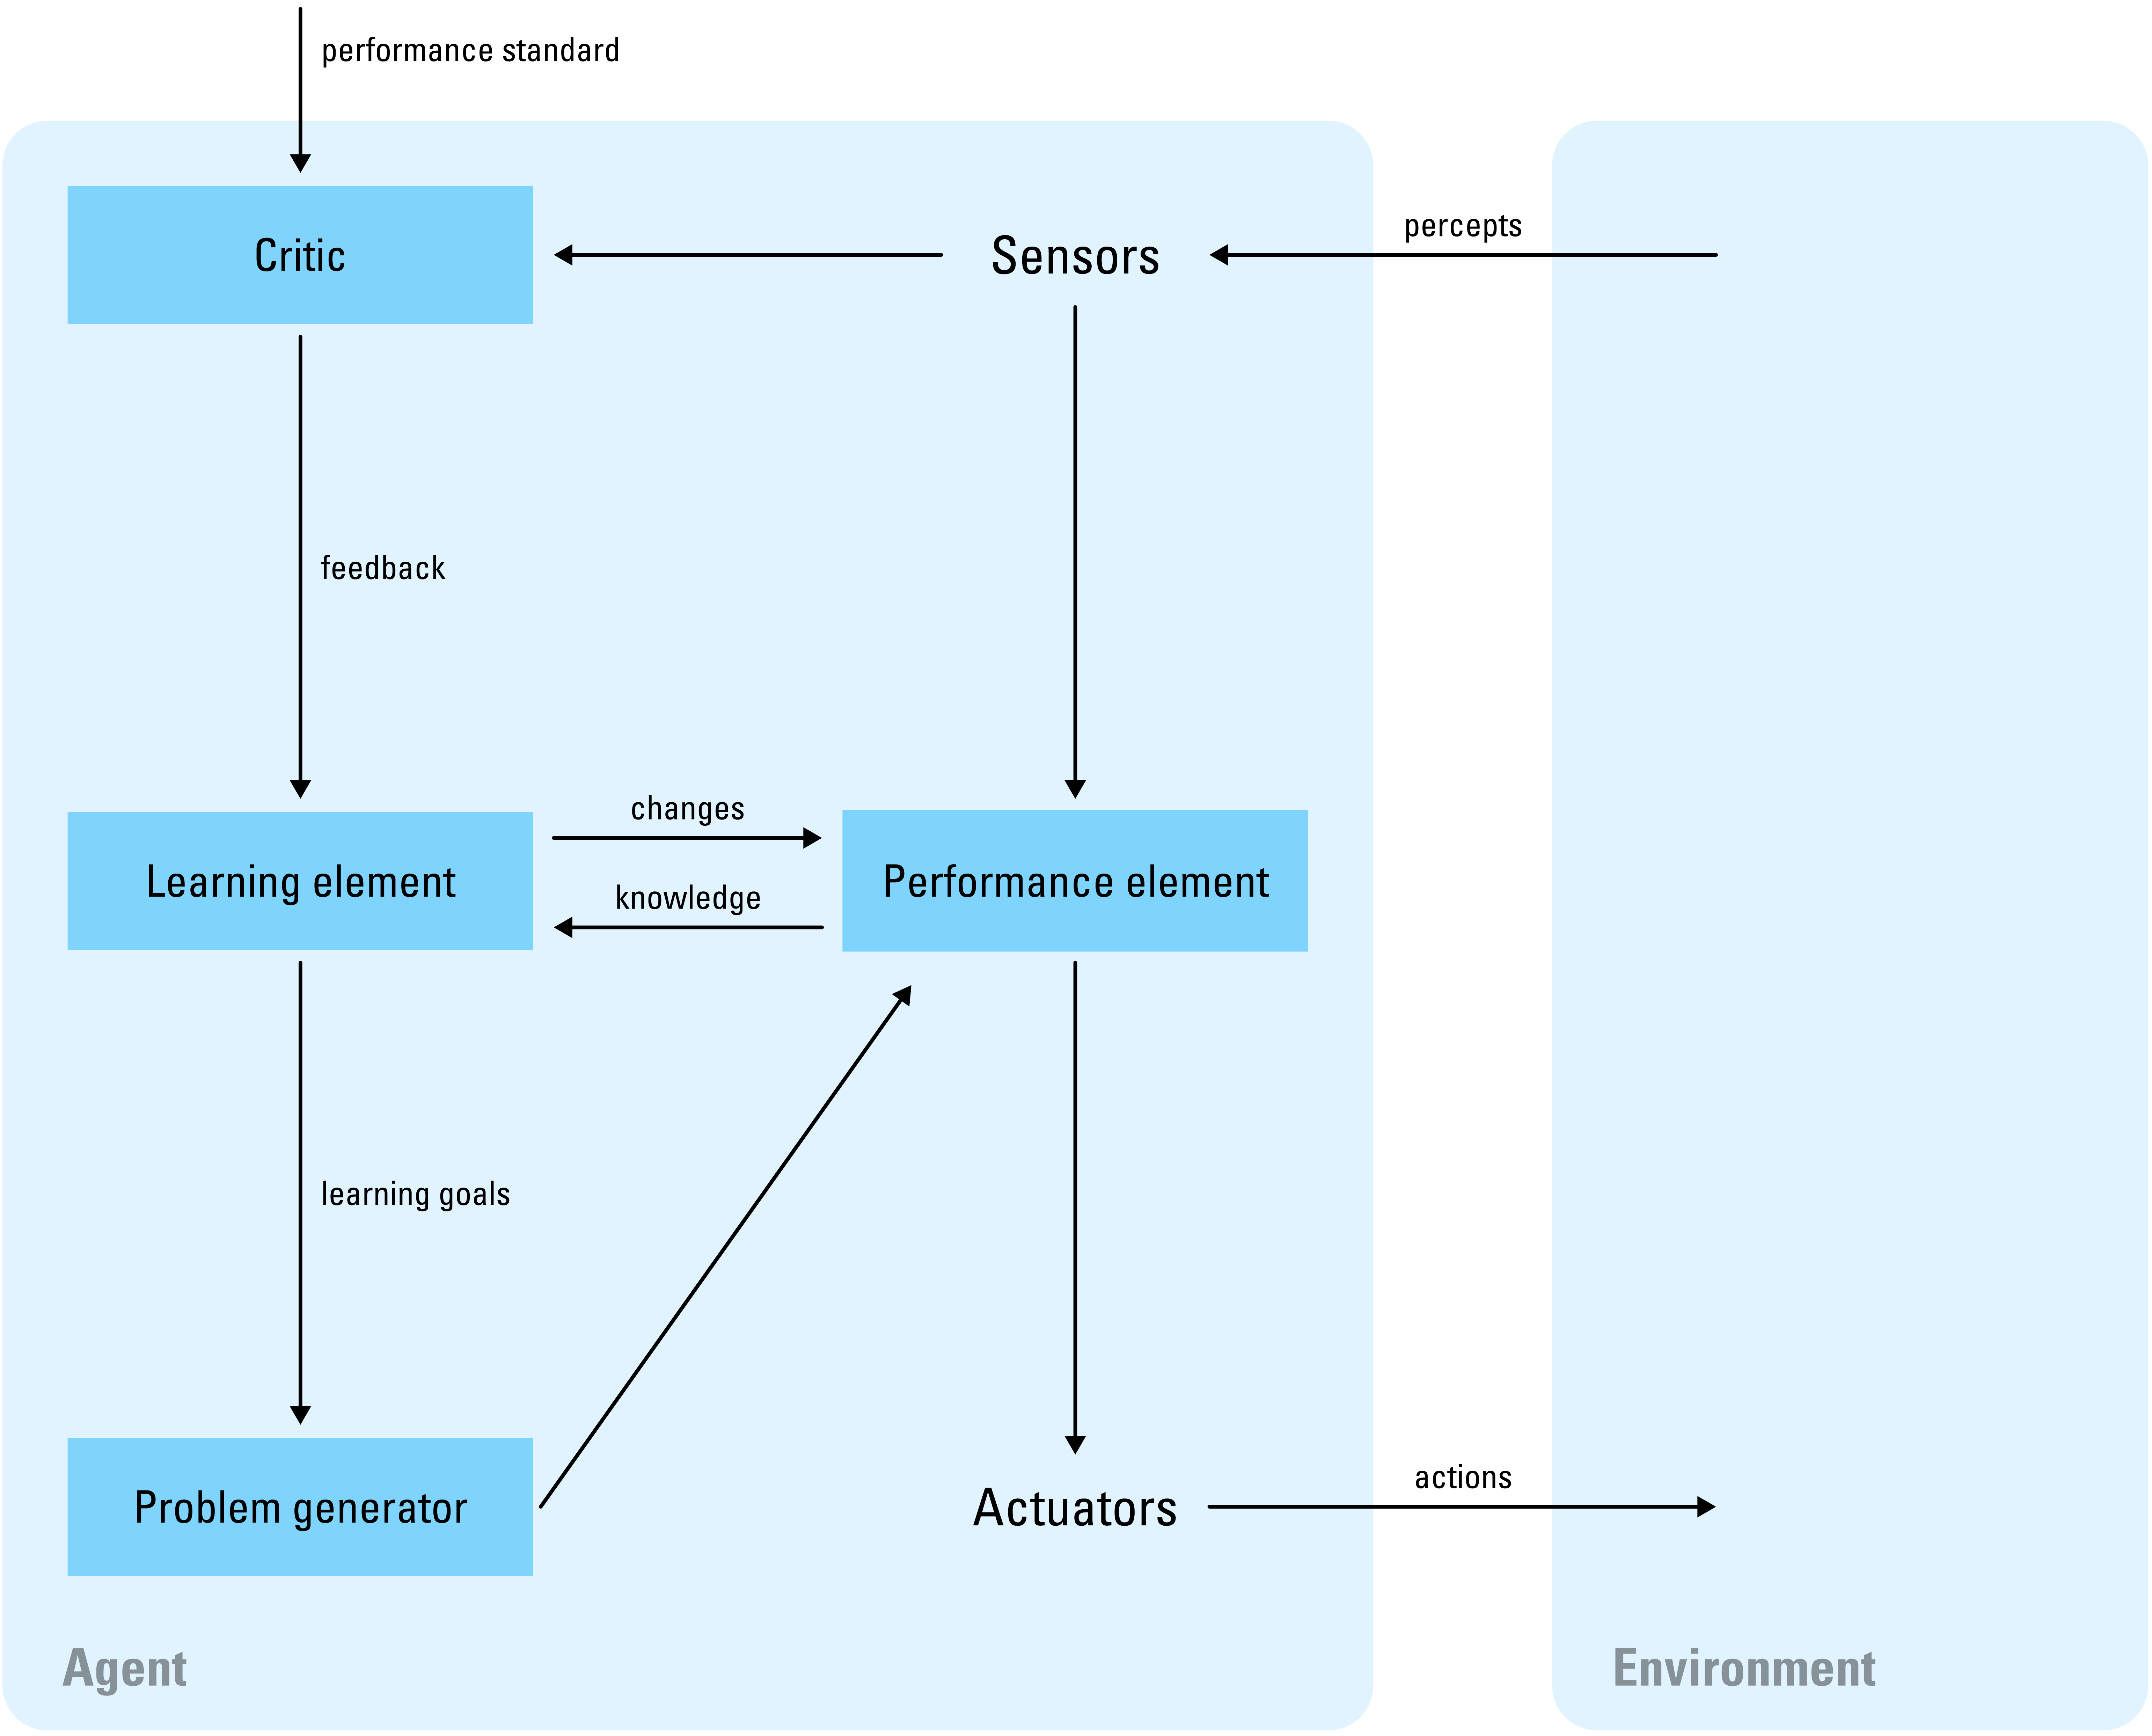

Russell & Norvig, 'Artificial Intelligence: A Modern Approach' (2003)
http://aima.cs.berkeley.edu/figures.pdf

General learning agent

'A general learning agent. ***The 'performance element' box represents what we have previously considered to be the whole agent program.*** *Now, the 'learning element' box gets to modify that program to improve its performance."*

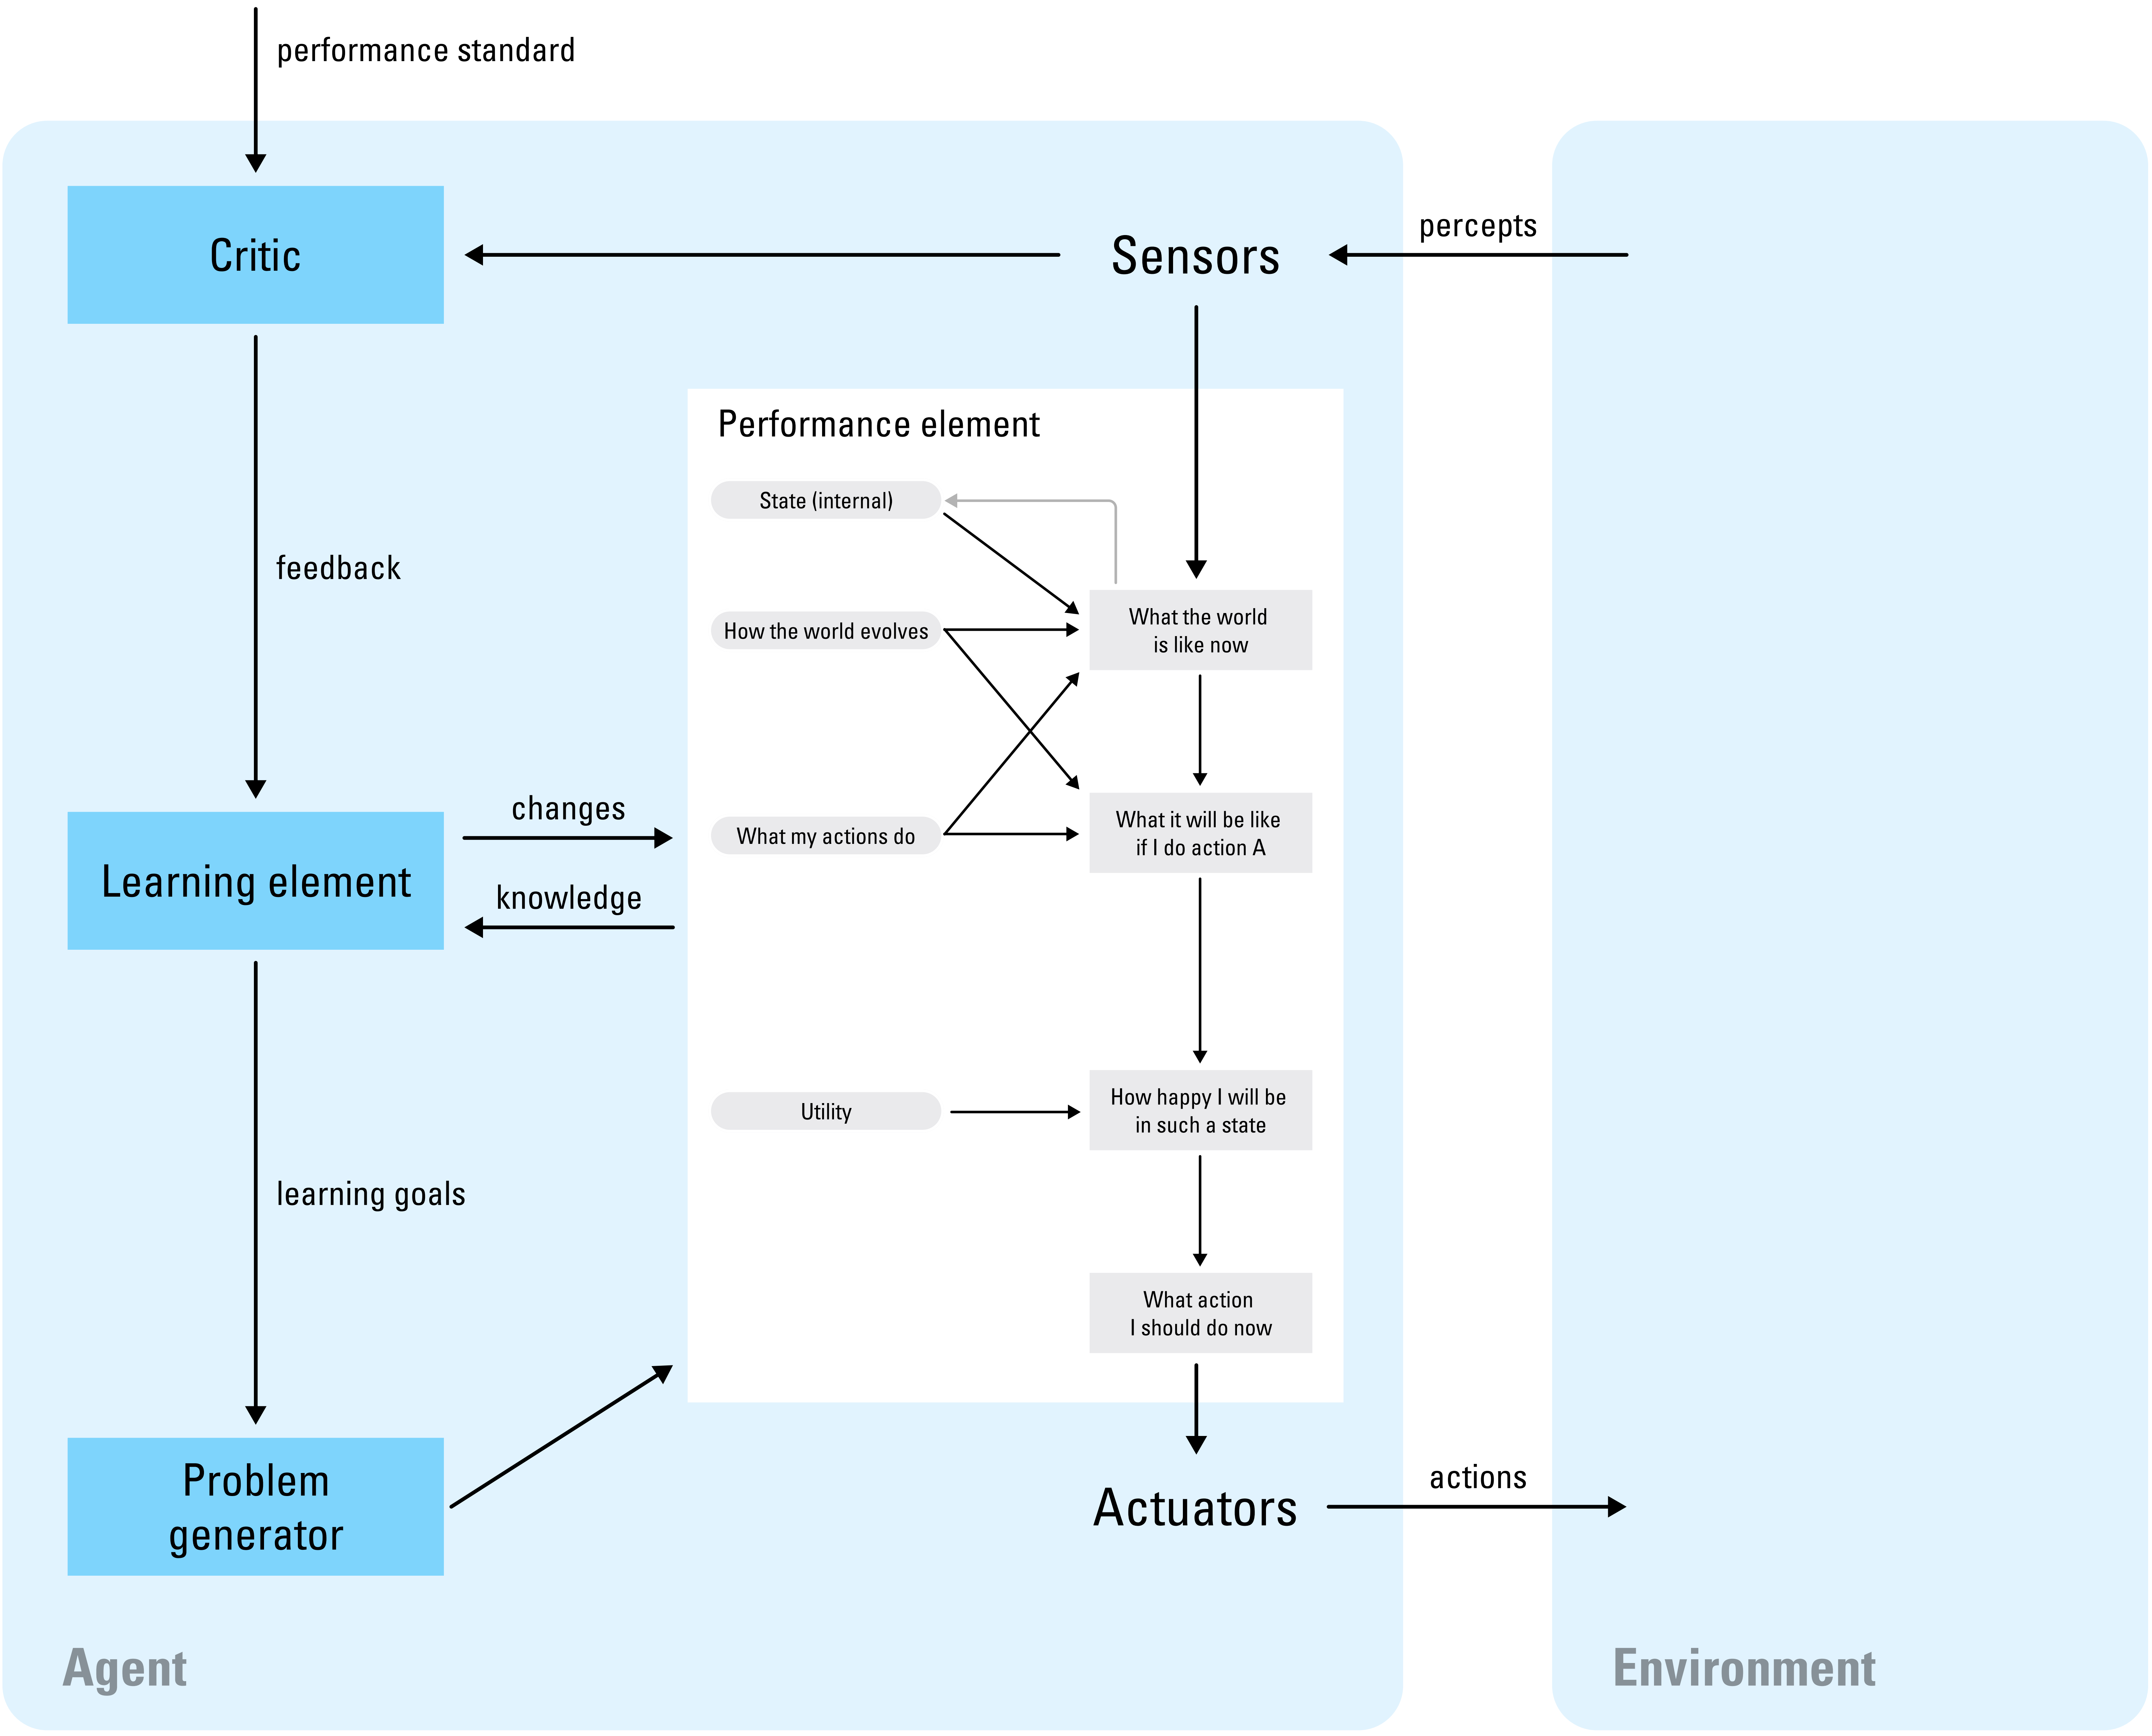

Russell & Norvig, 'Artificial Intelligence: A Modern Approach' (2003)
http://aima.cs.berkeley.edu/figures.pdf

Key terms Russell & Norvig use in describing the structure of intelligent agents:

Term	Definition
Rules	Condition-action rules describe a direct connection between an external state and an action that should be taken in response.
Model	Knowledge about “how the world works”, including: • how the world evolves independently of the agent’s actions • how the agent’s actions affect the world
Goals	Information that describes desirable situations. Goals provide a binary between “happy” and “unhappy” states.
Utility	Utility allows a comparison of different world states according to exactly how happy they would make the agent.
Performance element	What we have previously considered to be the entire agent. It takes in percepts and decides on actions.
Critic	Provides feedback to the learning element on how the agent is doing by observing the world and passing information along to the learning element.
Learning element	Responsible for making improvements. Uses feedback from the critic and determines how the performance element should be modified to do better in the future.
Problem generator	Suggests actions that will lead to new and informative experiences — exploratory actions that might be suboptimal in the short-term, but could lead to discovering better actions in the future.

Russell & Norvig, ‘Artificial Intelligence: A Modern Approach’ (2003)
http://aima.cs.berkeley.edu/figures.pdf

An intelligent agent uses sensors and actuators to interact with their environment, and makes decisions about which actions to take in order to achieve certain outcomes.

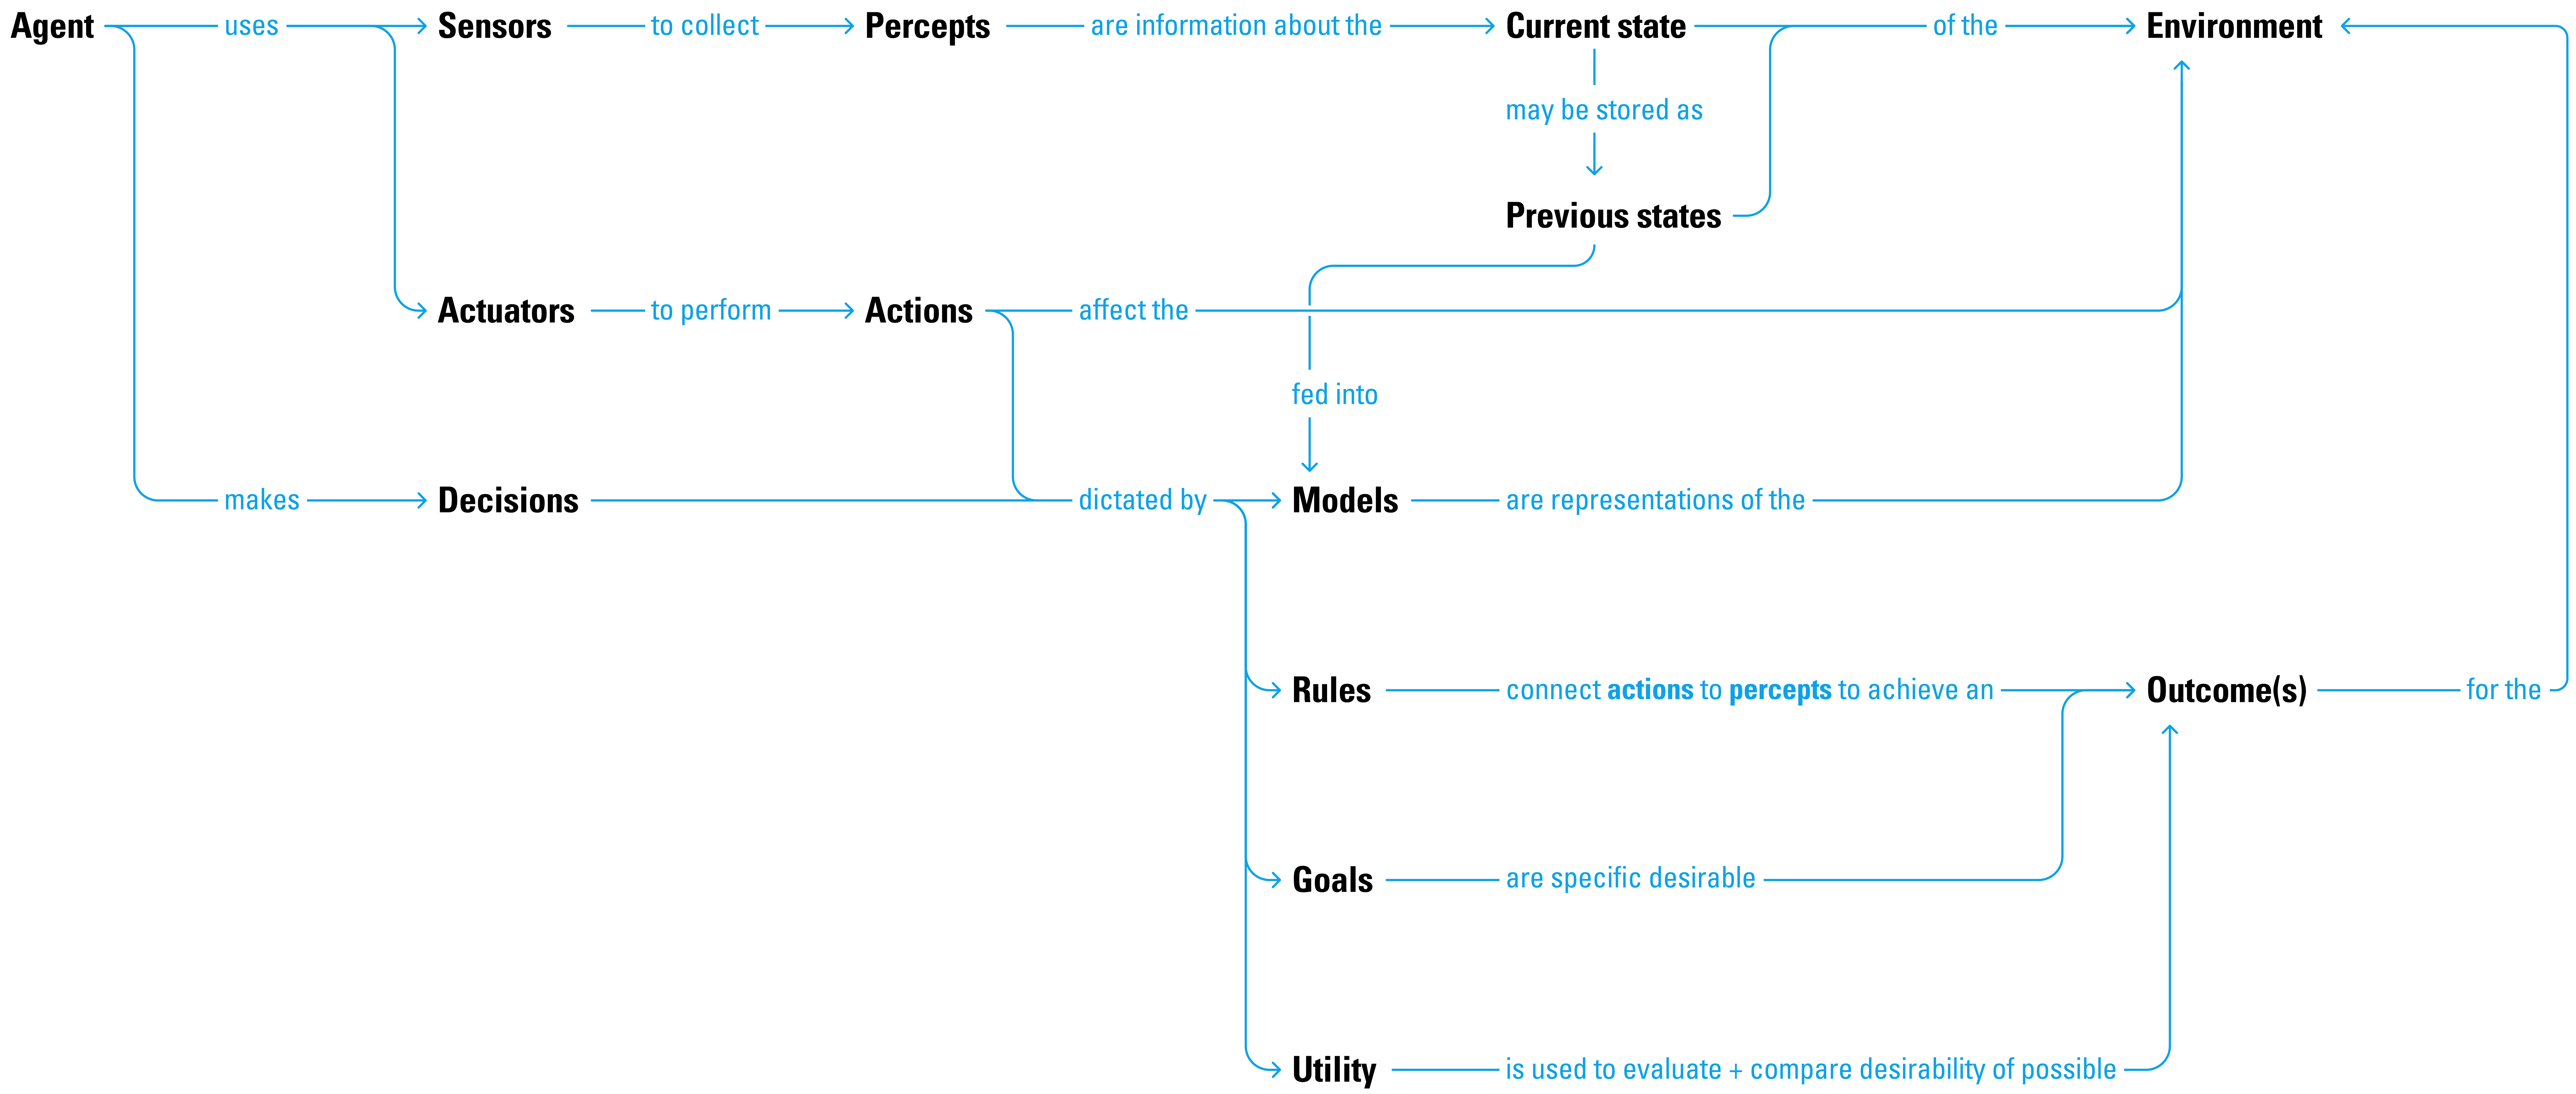

Russell & Norvig illustrate these terms with an example of a self-driving taxi

Imagine yourself as the driving agent of a self-driving taxi.

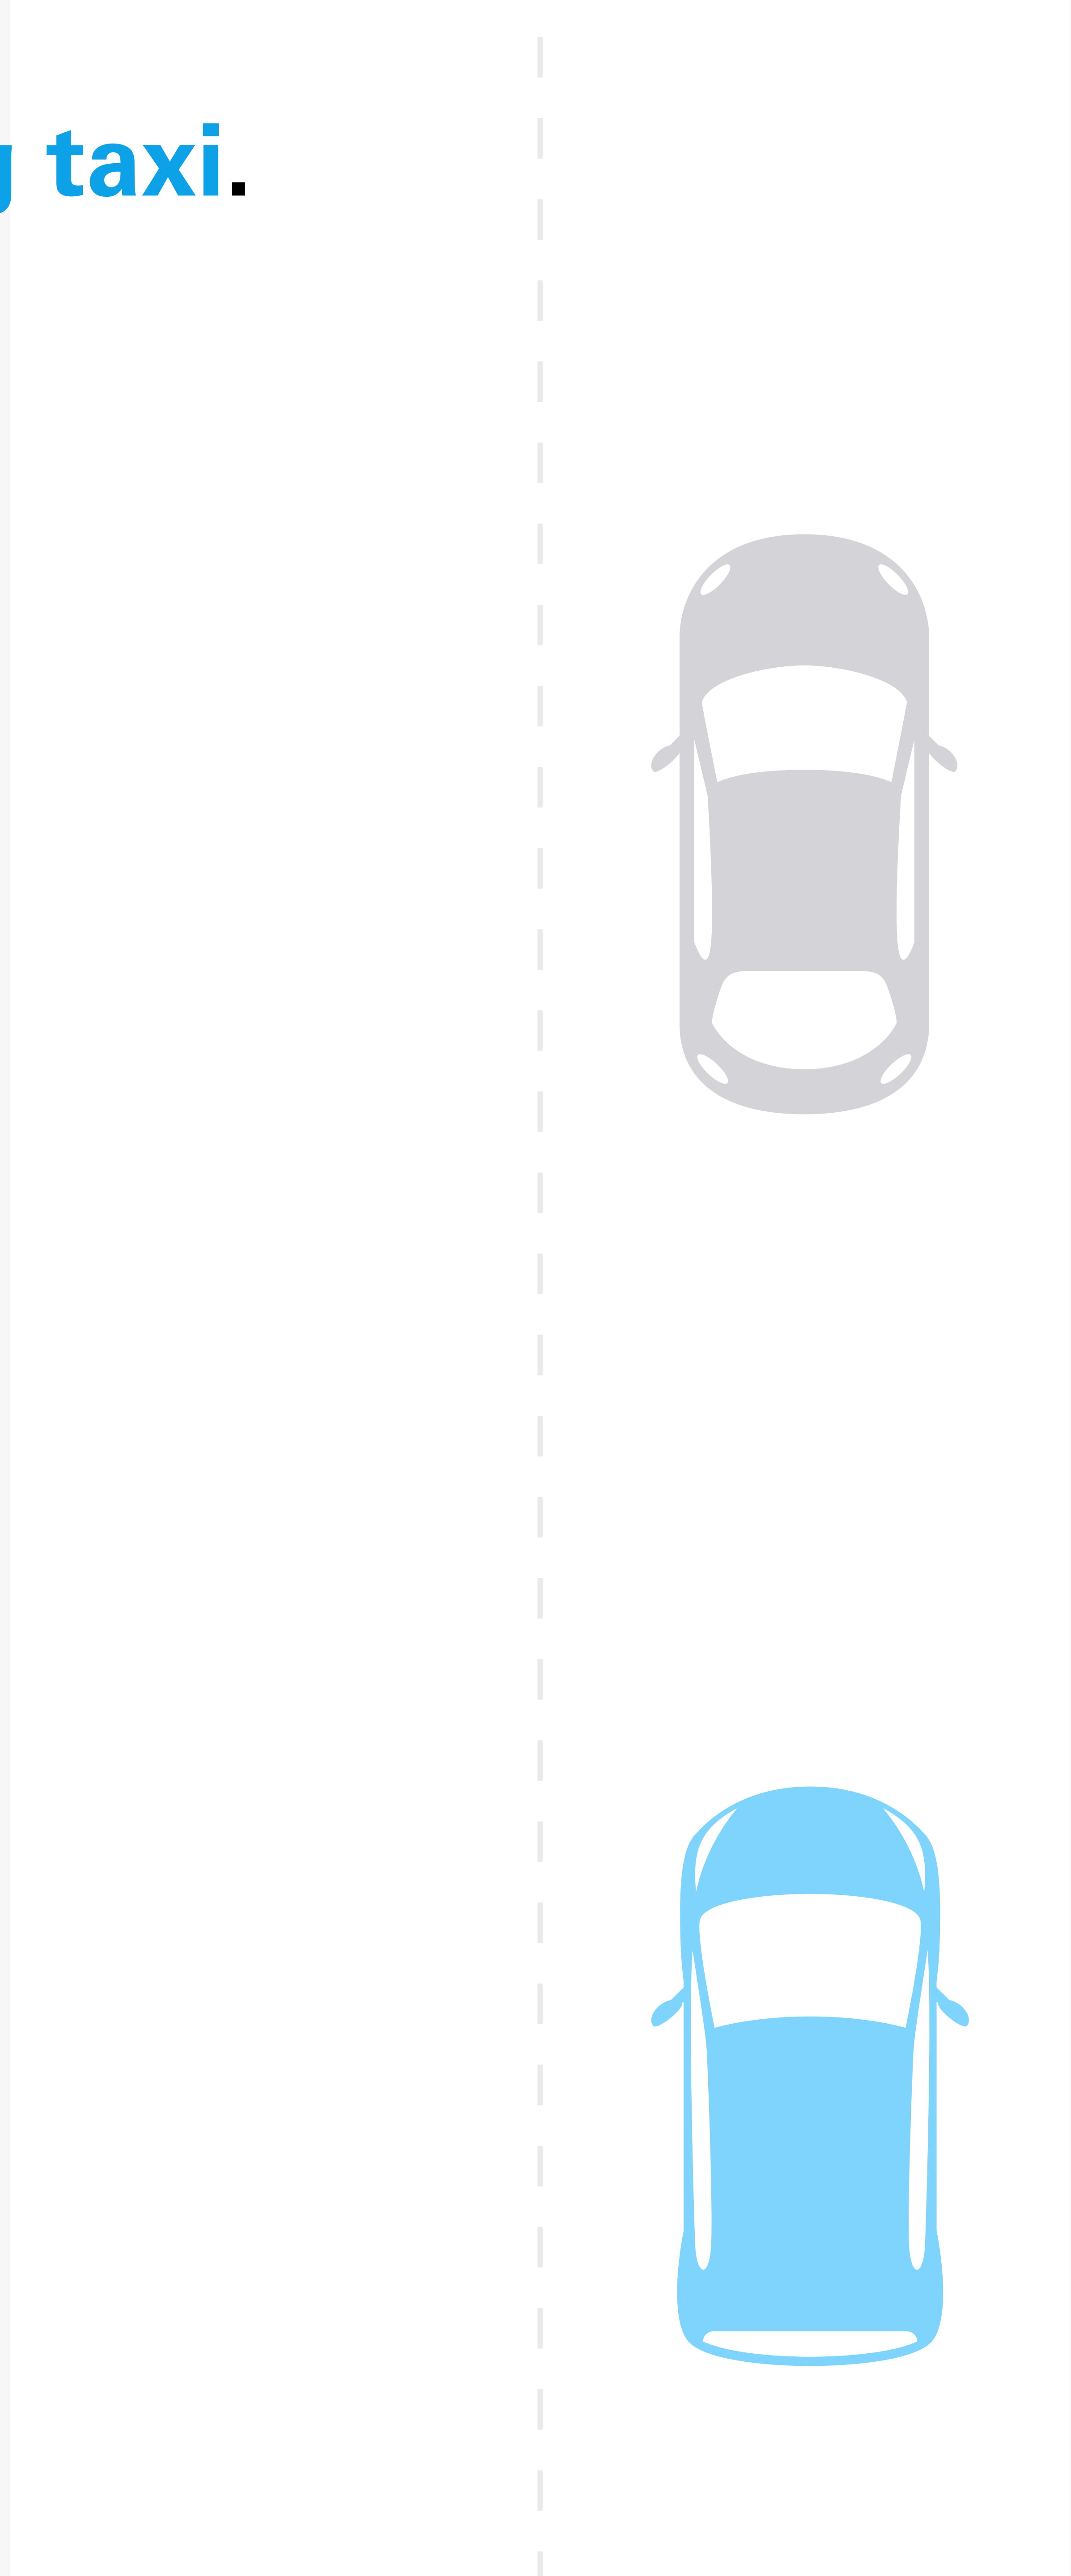

If the car in front brakes and its brake lights come on, then you should notice this and initiate braking.

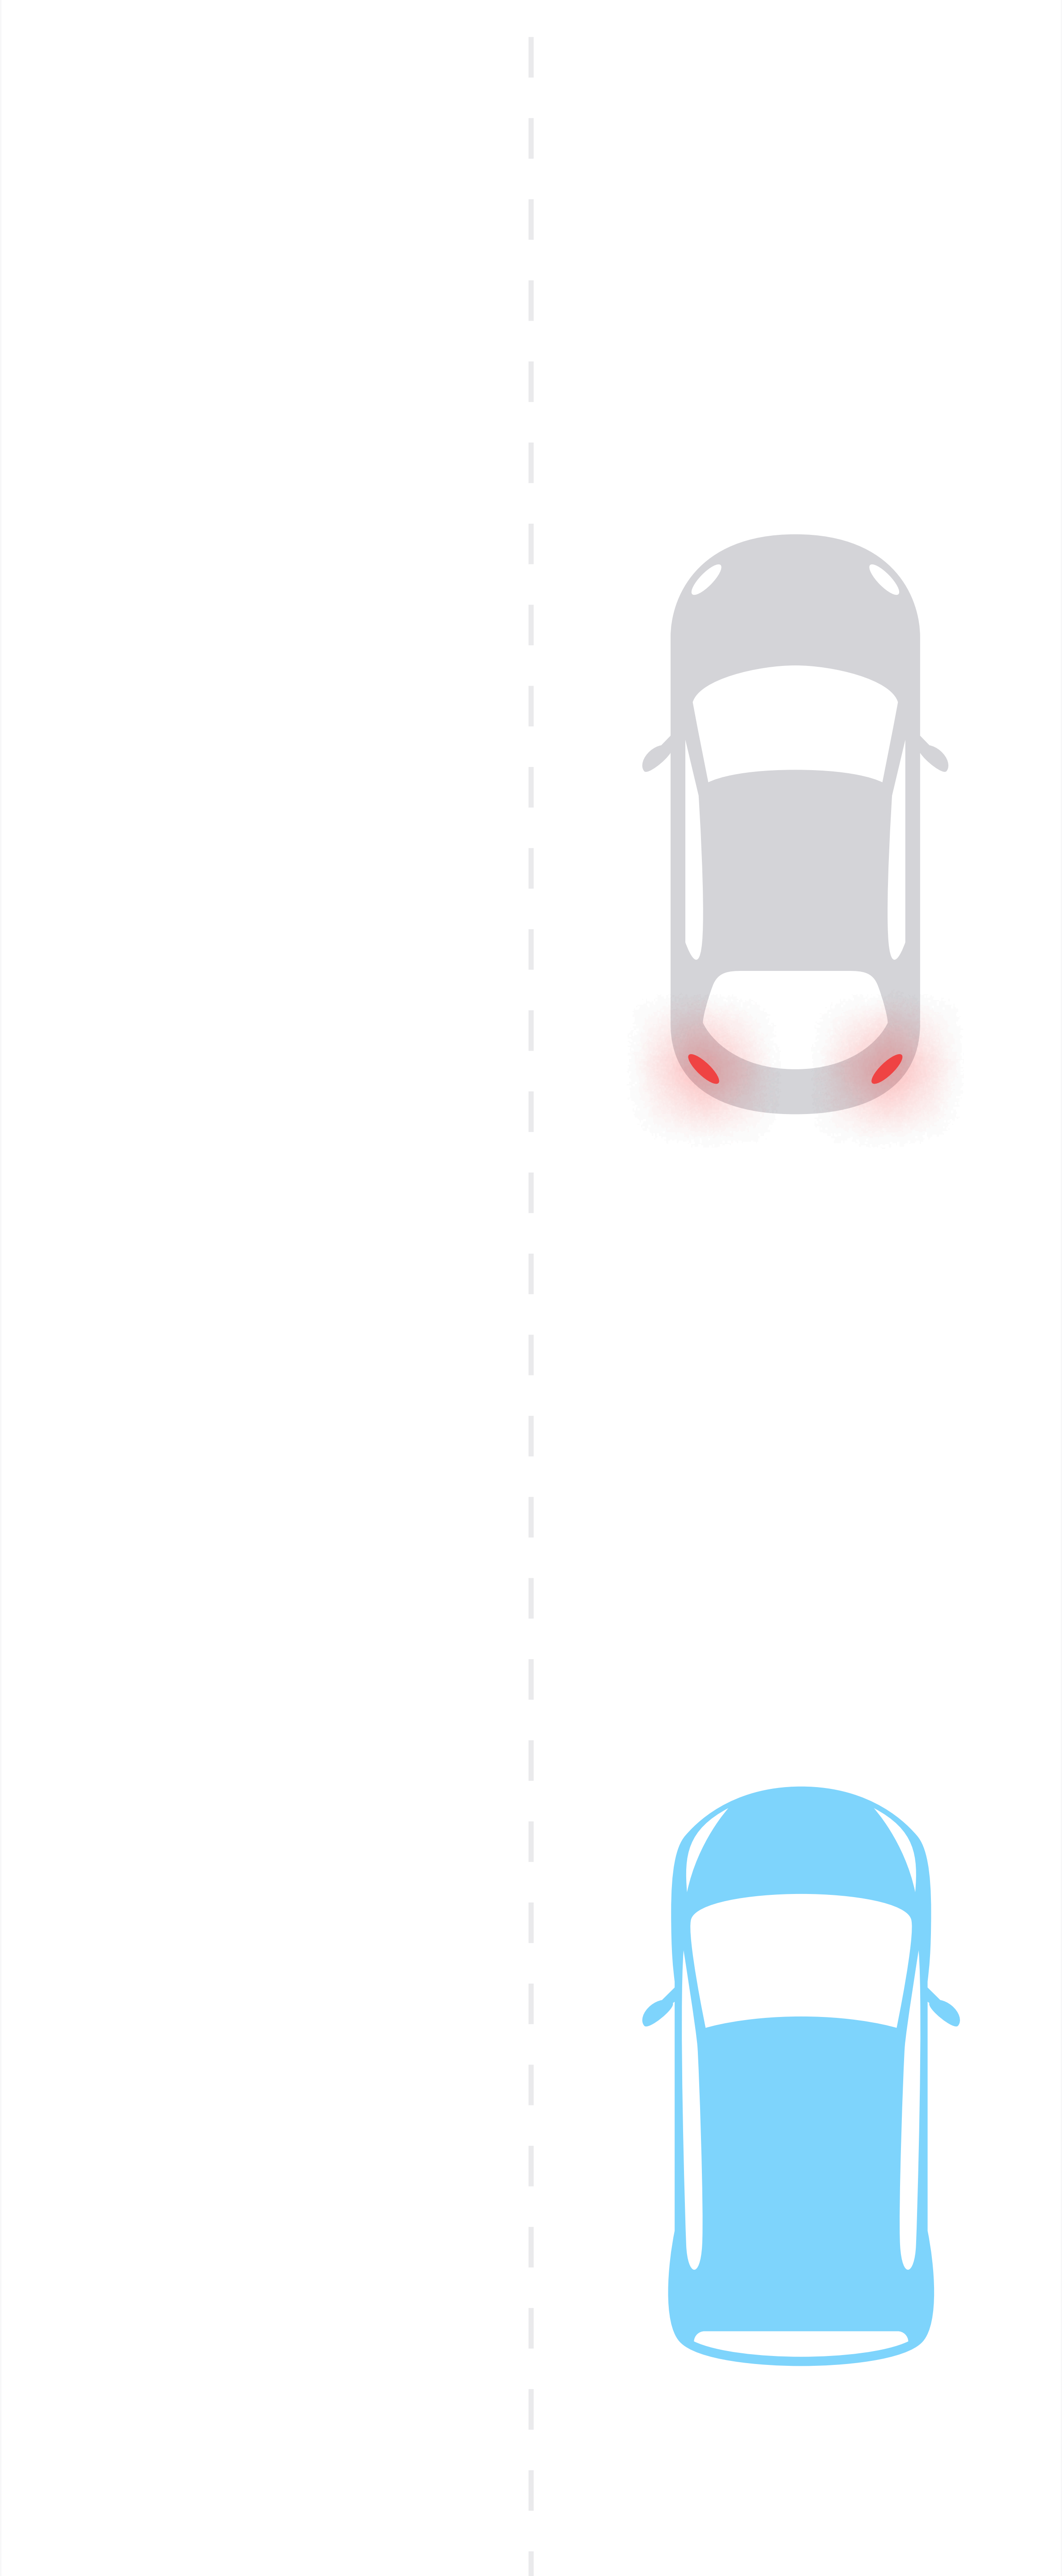

Some processing is done on the visual input to establish the condition 'The car in front is braking.' This triggers some established connection in the agent program to the action 'initiate braking.'

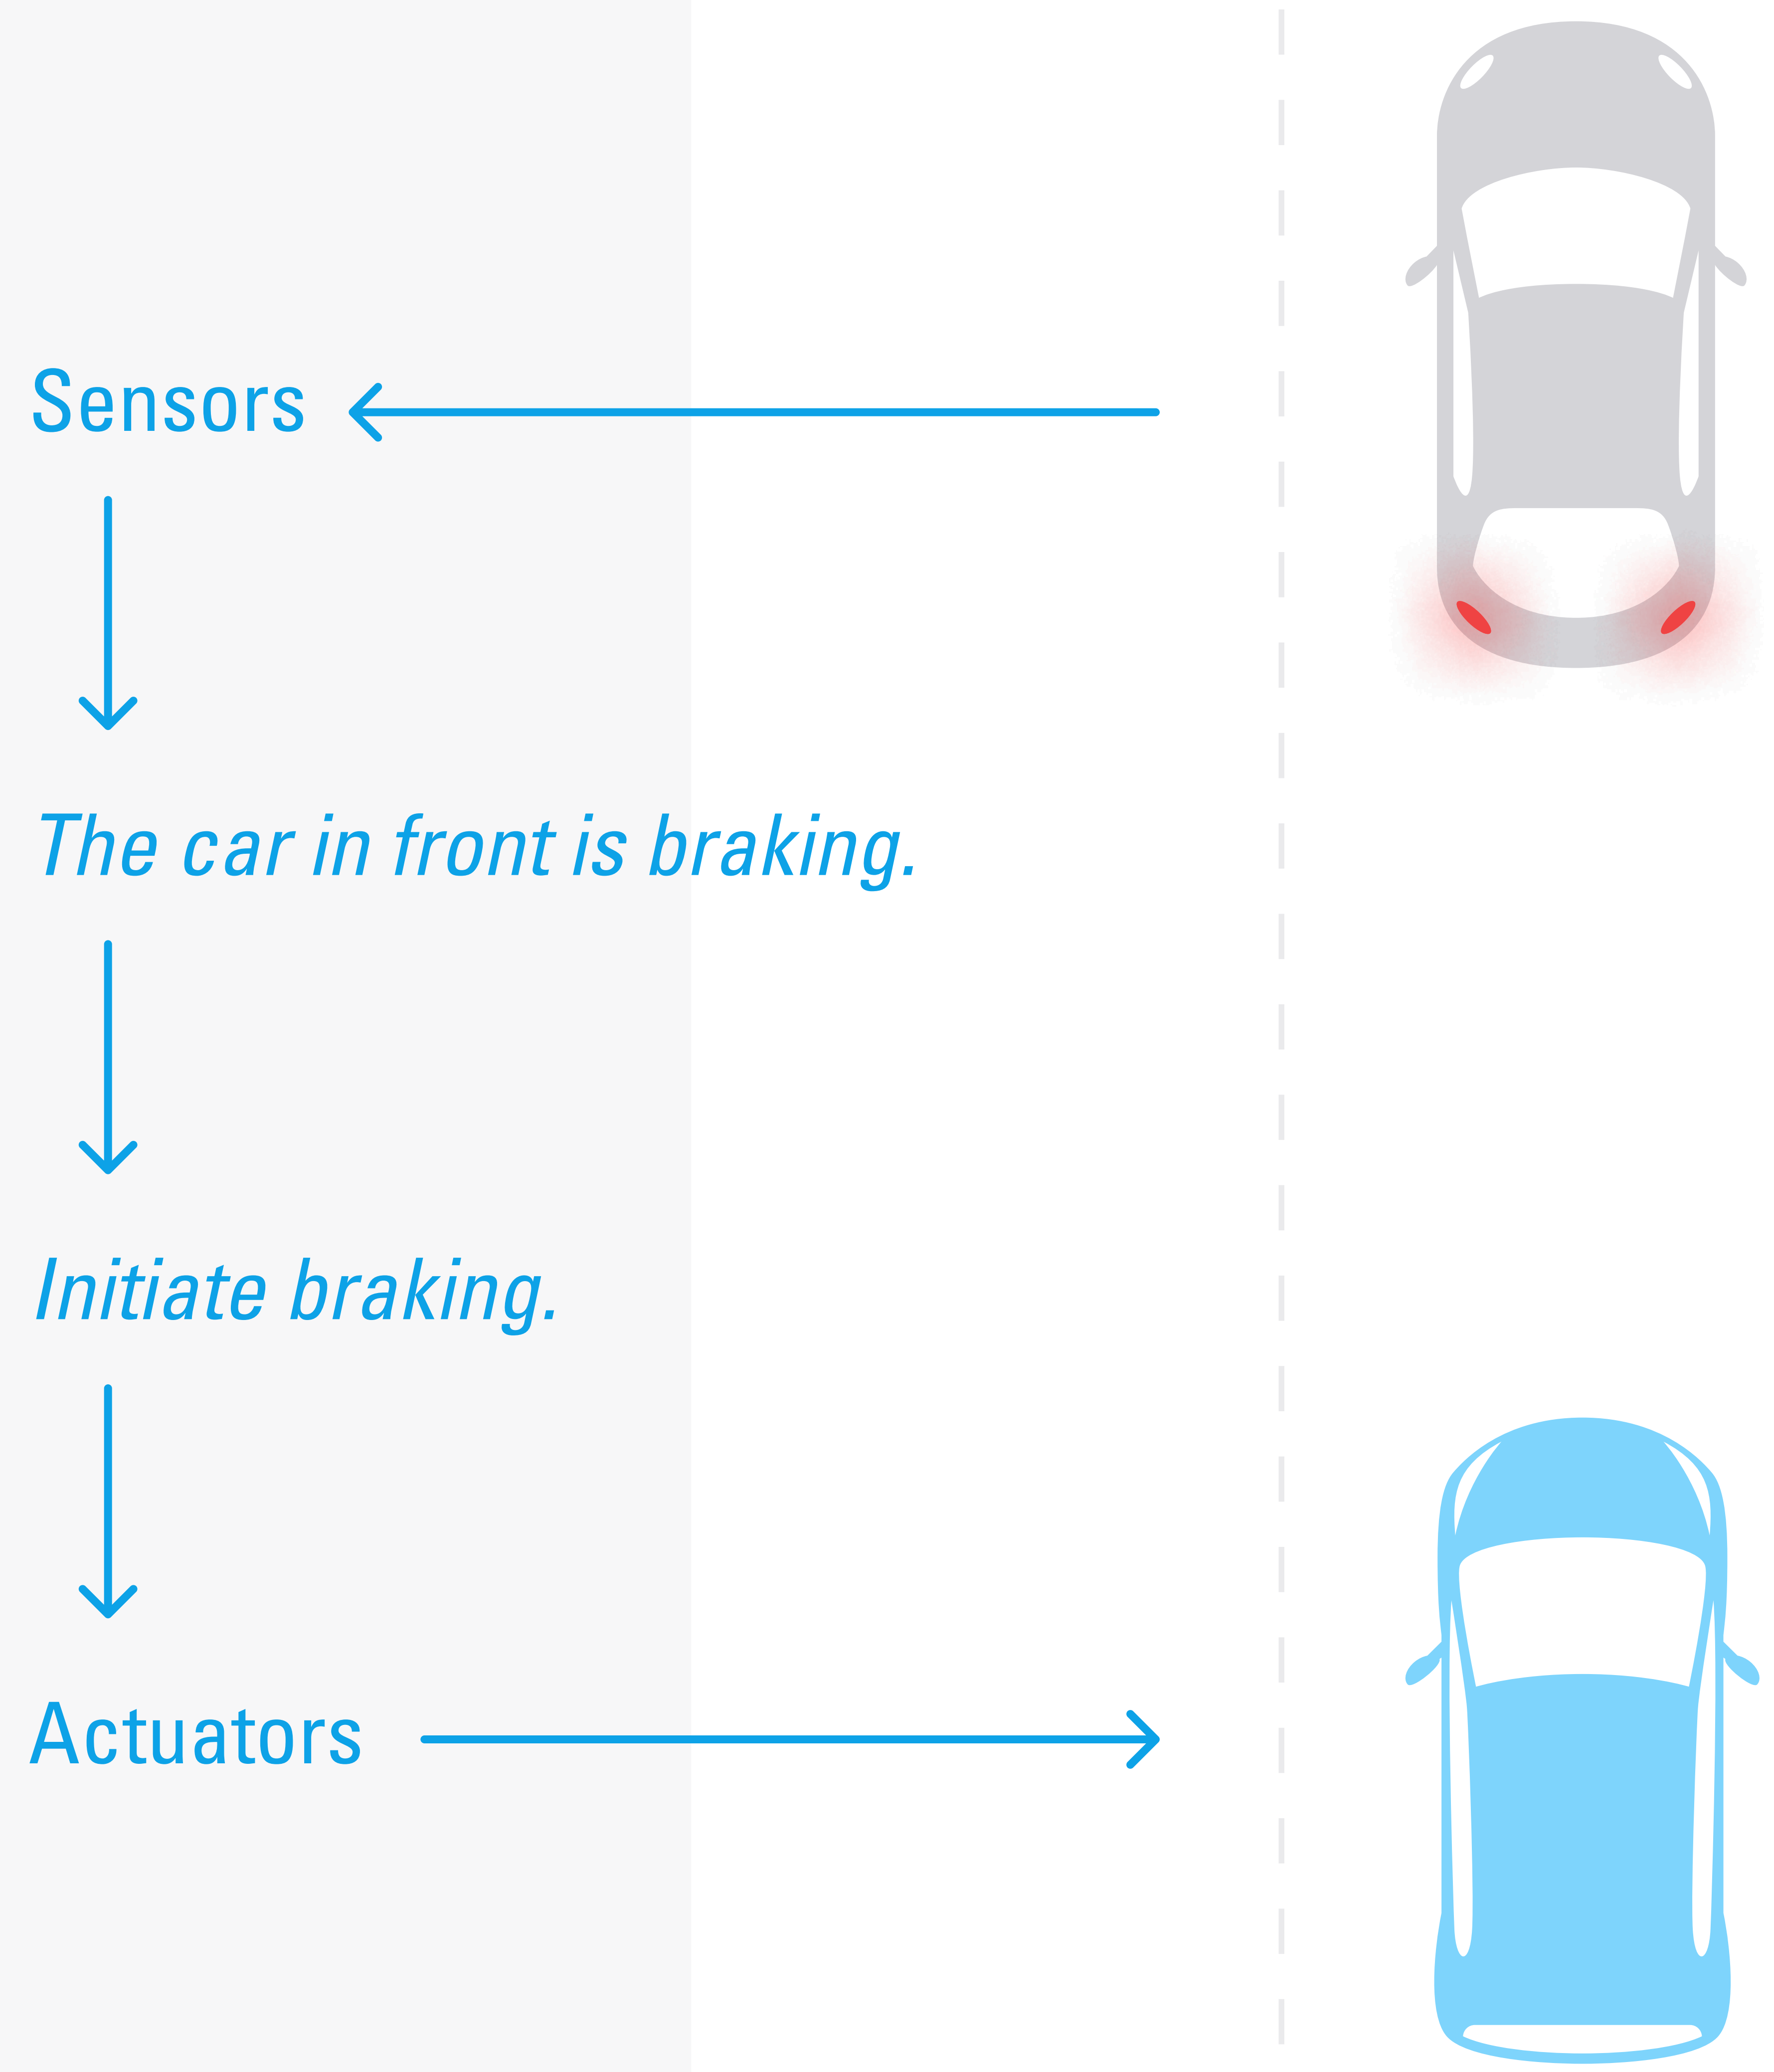

This is a condition-action rule.
Condition-action rules describe a direct connection between an external state and an action that should be taken in response.

Condition-action rule:

if:

The car in front is braking.

then:

Initiate braking.

In a model-based agent, a model is some knowledge of 'how the world works' which includes:

1 How the world evolves independently of the agent's actions

and

2 How the agent's actions affect the world:

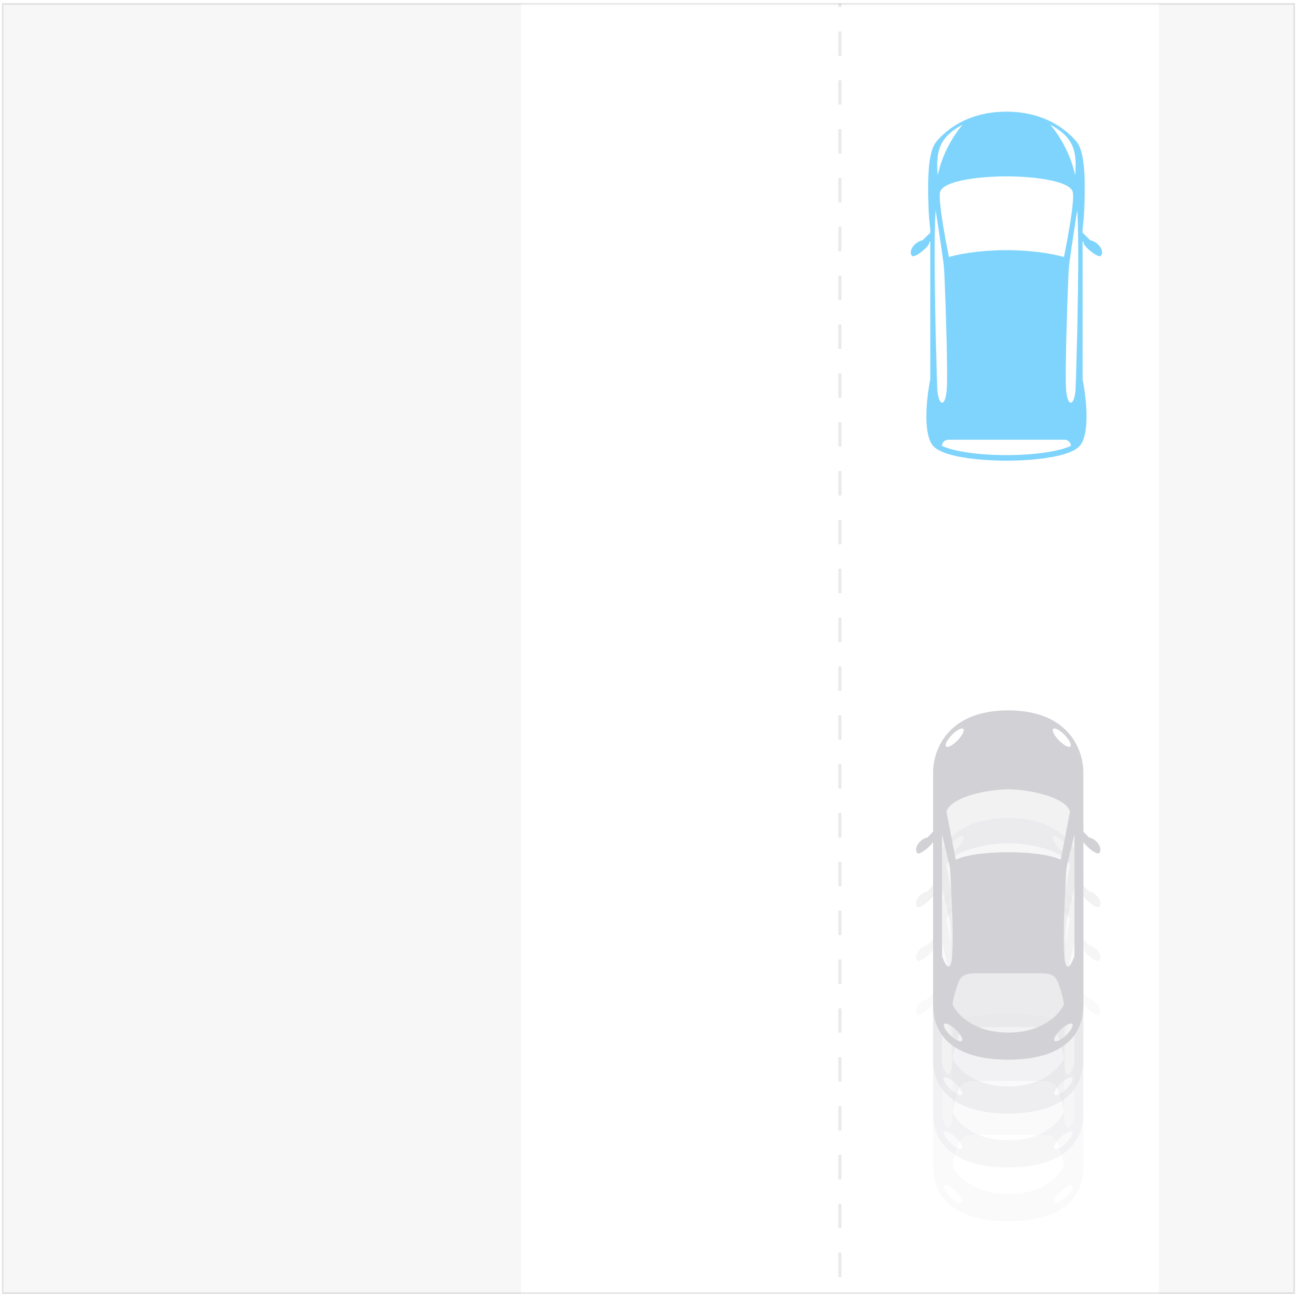

1A

An overtaking car generally will be closer behind than it was a moment ago.

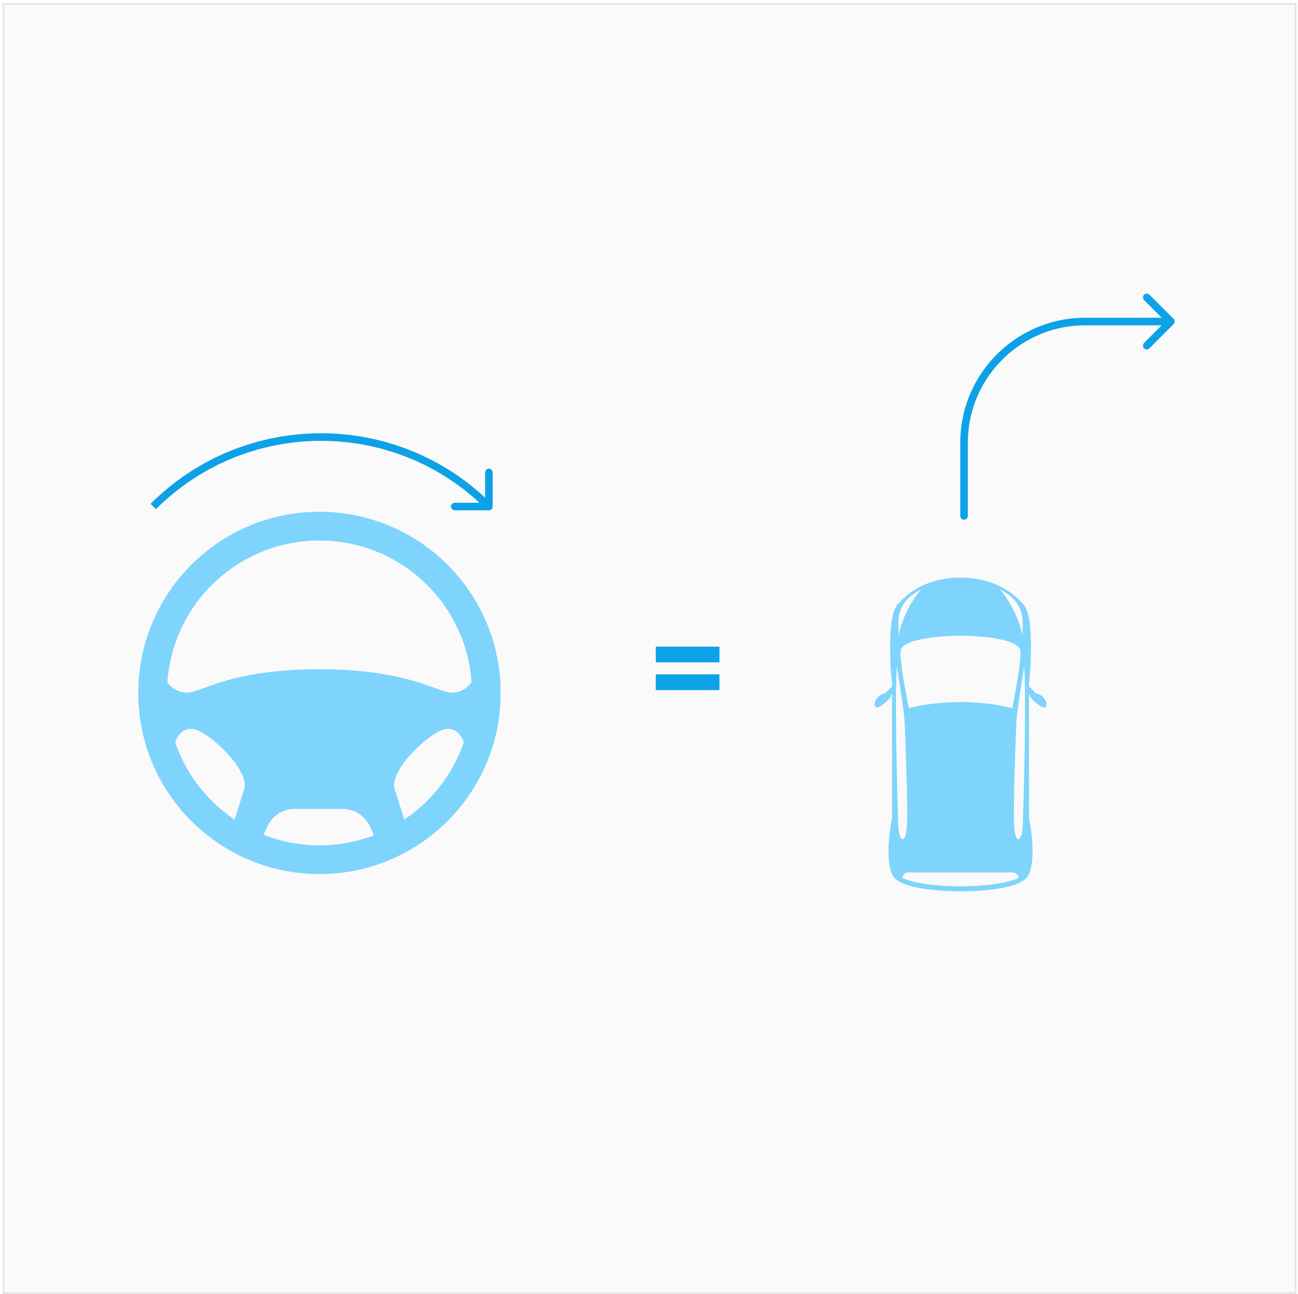

2A

When the agent turns the steering wheel clockwise, the car turns to the right.

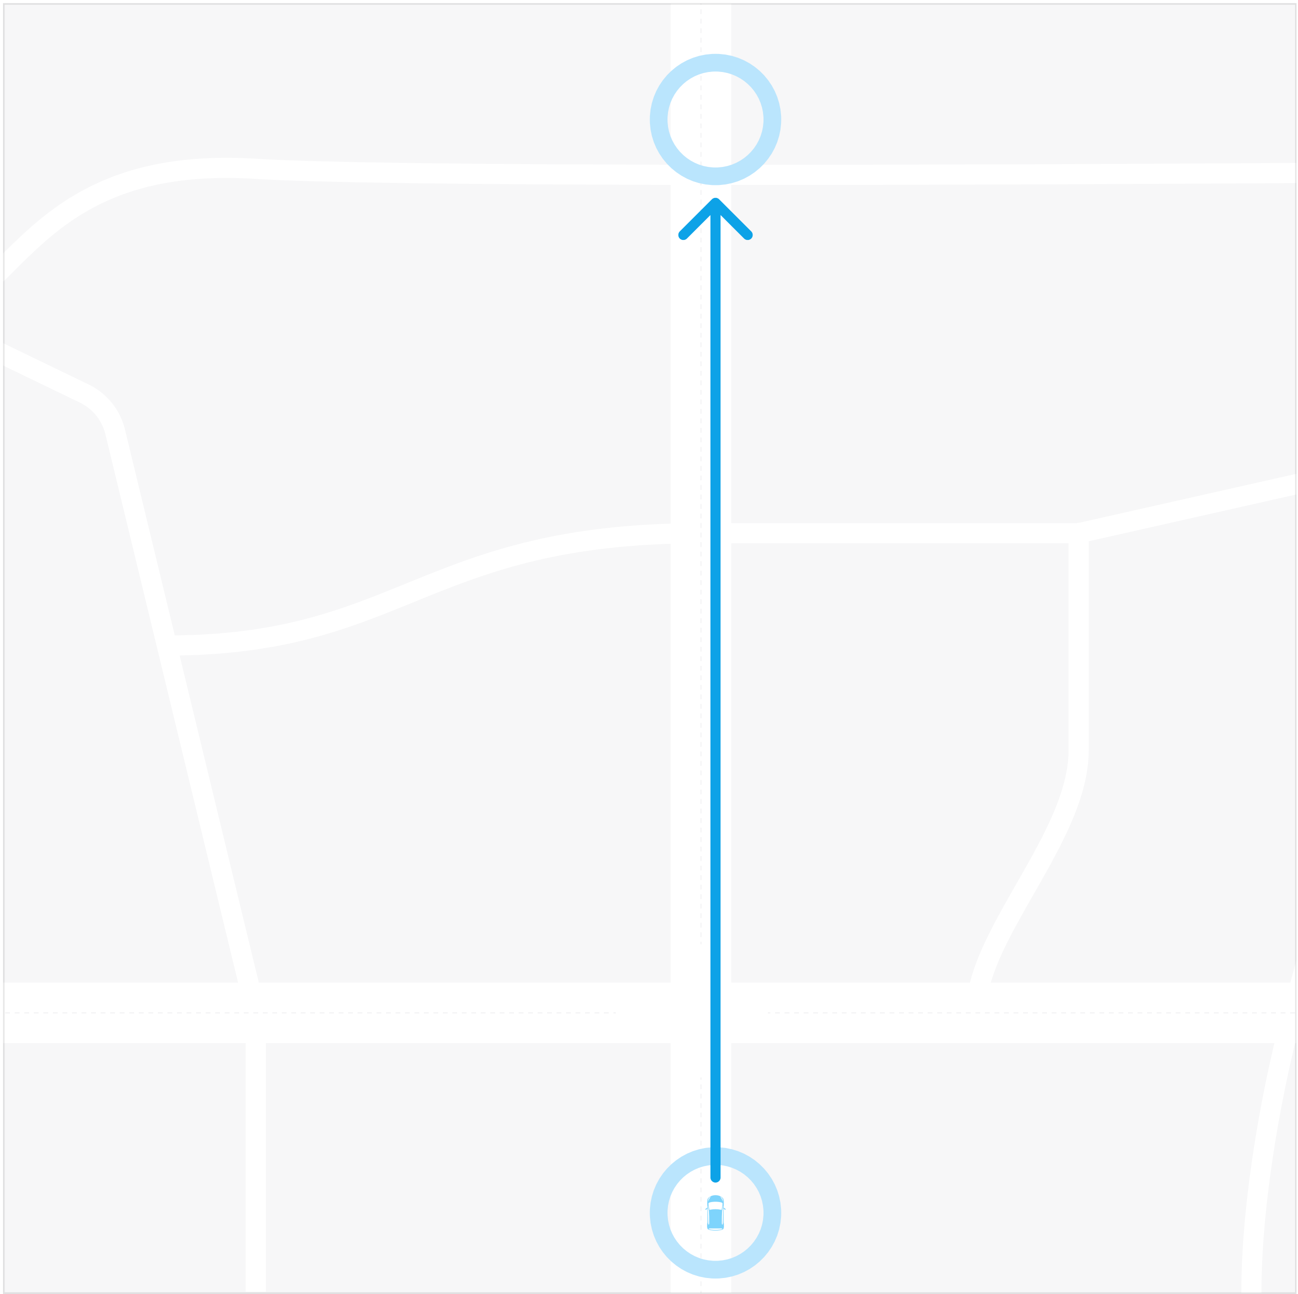

2B

After driving for five minutes northbound on the freeway, one is usually about five miles north of where one was five minutes ago.

Condition-action rules based on the current state of the environment are not always enough to decide what to do.

For example, at a road junction, the taxi can turn left, turn right, or go straight on. The correct decision depends on where the taxi is trying to get to.

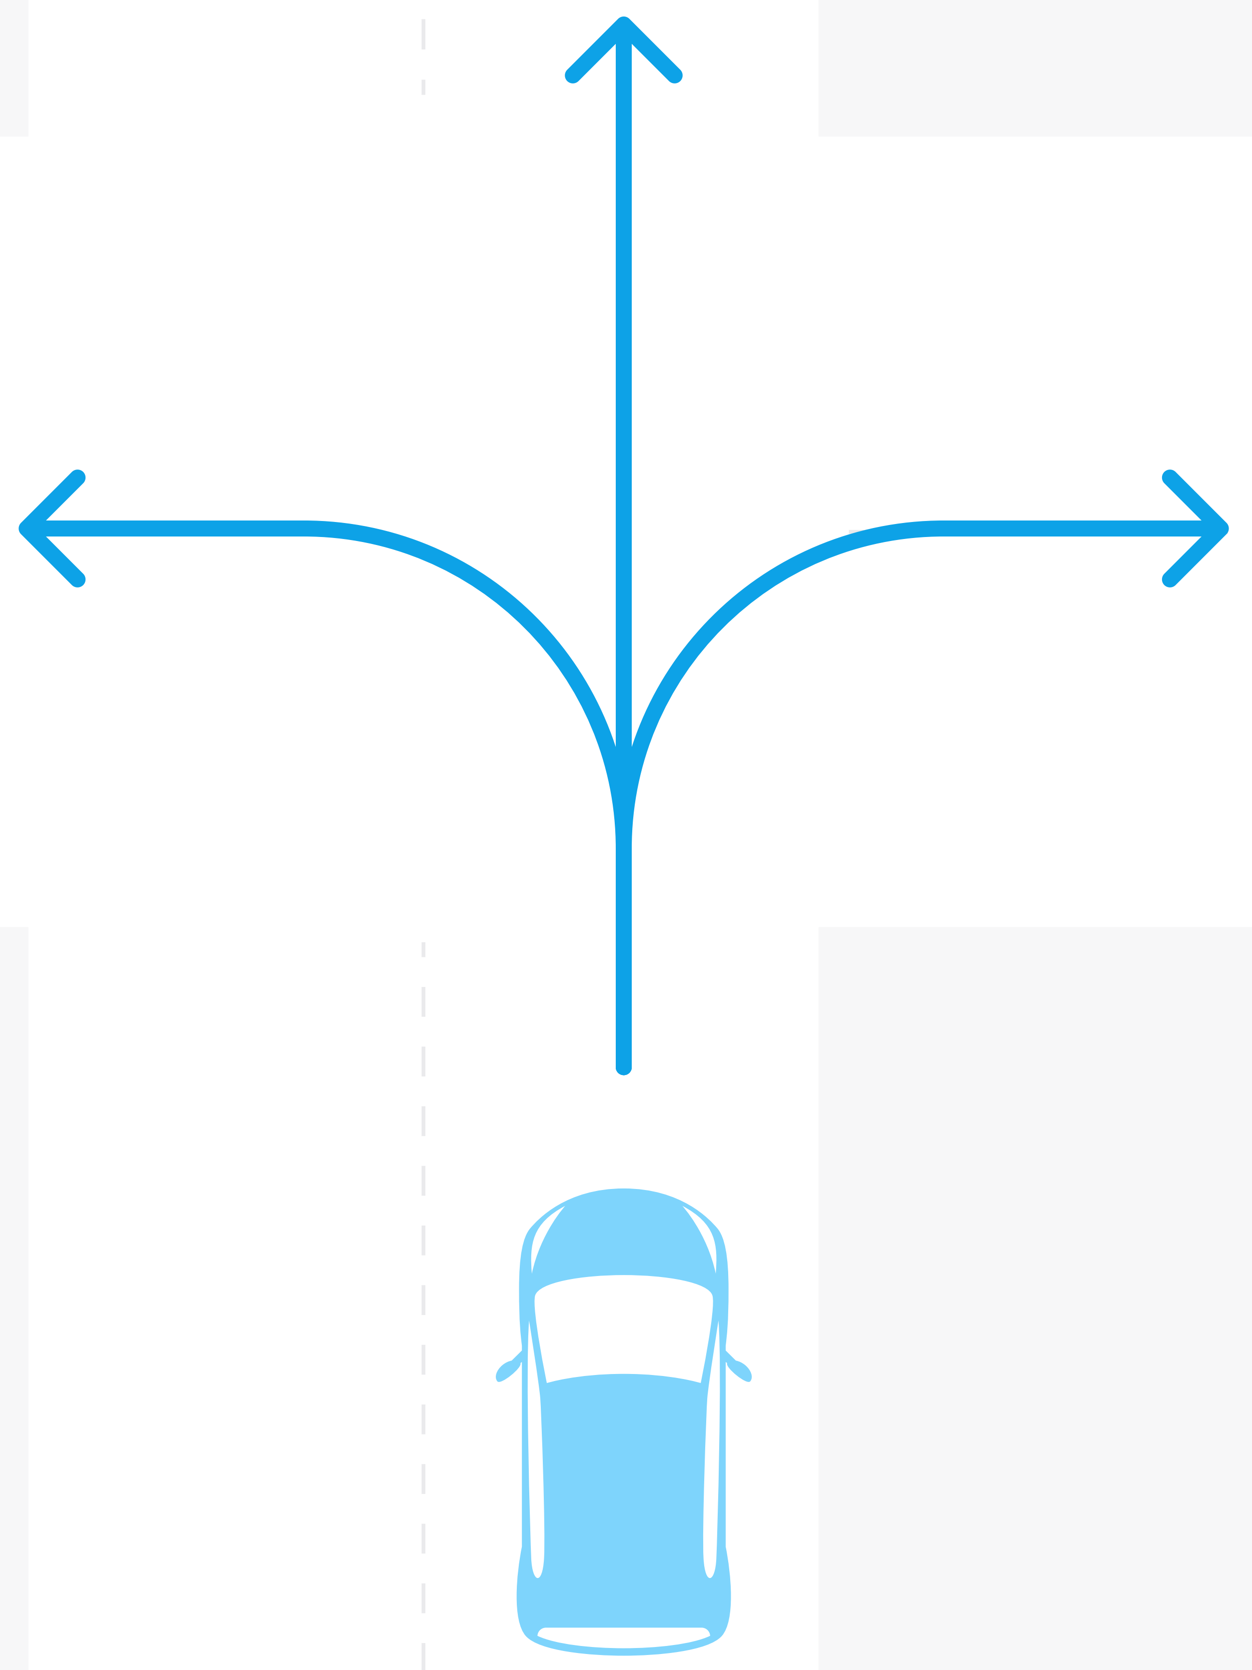

The agent needs some sort of goal — information that describes desirable outcomes. In this case, the goal would be the passenger's destination.

Goals only provide a crude distinction between “happy” and “unhappy” states.

Many action sequences will get the taxi to its destination (thereby achieving the goal) — but some are quicker, safer, more reliable, or cheaper than others.

Utility allows a comparison of different world states according to exactly how "happy" they would make the agent.

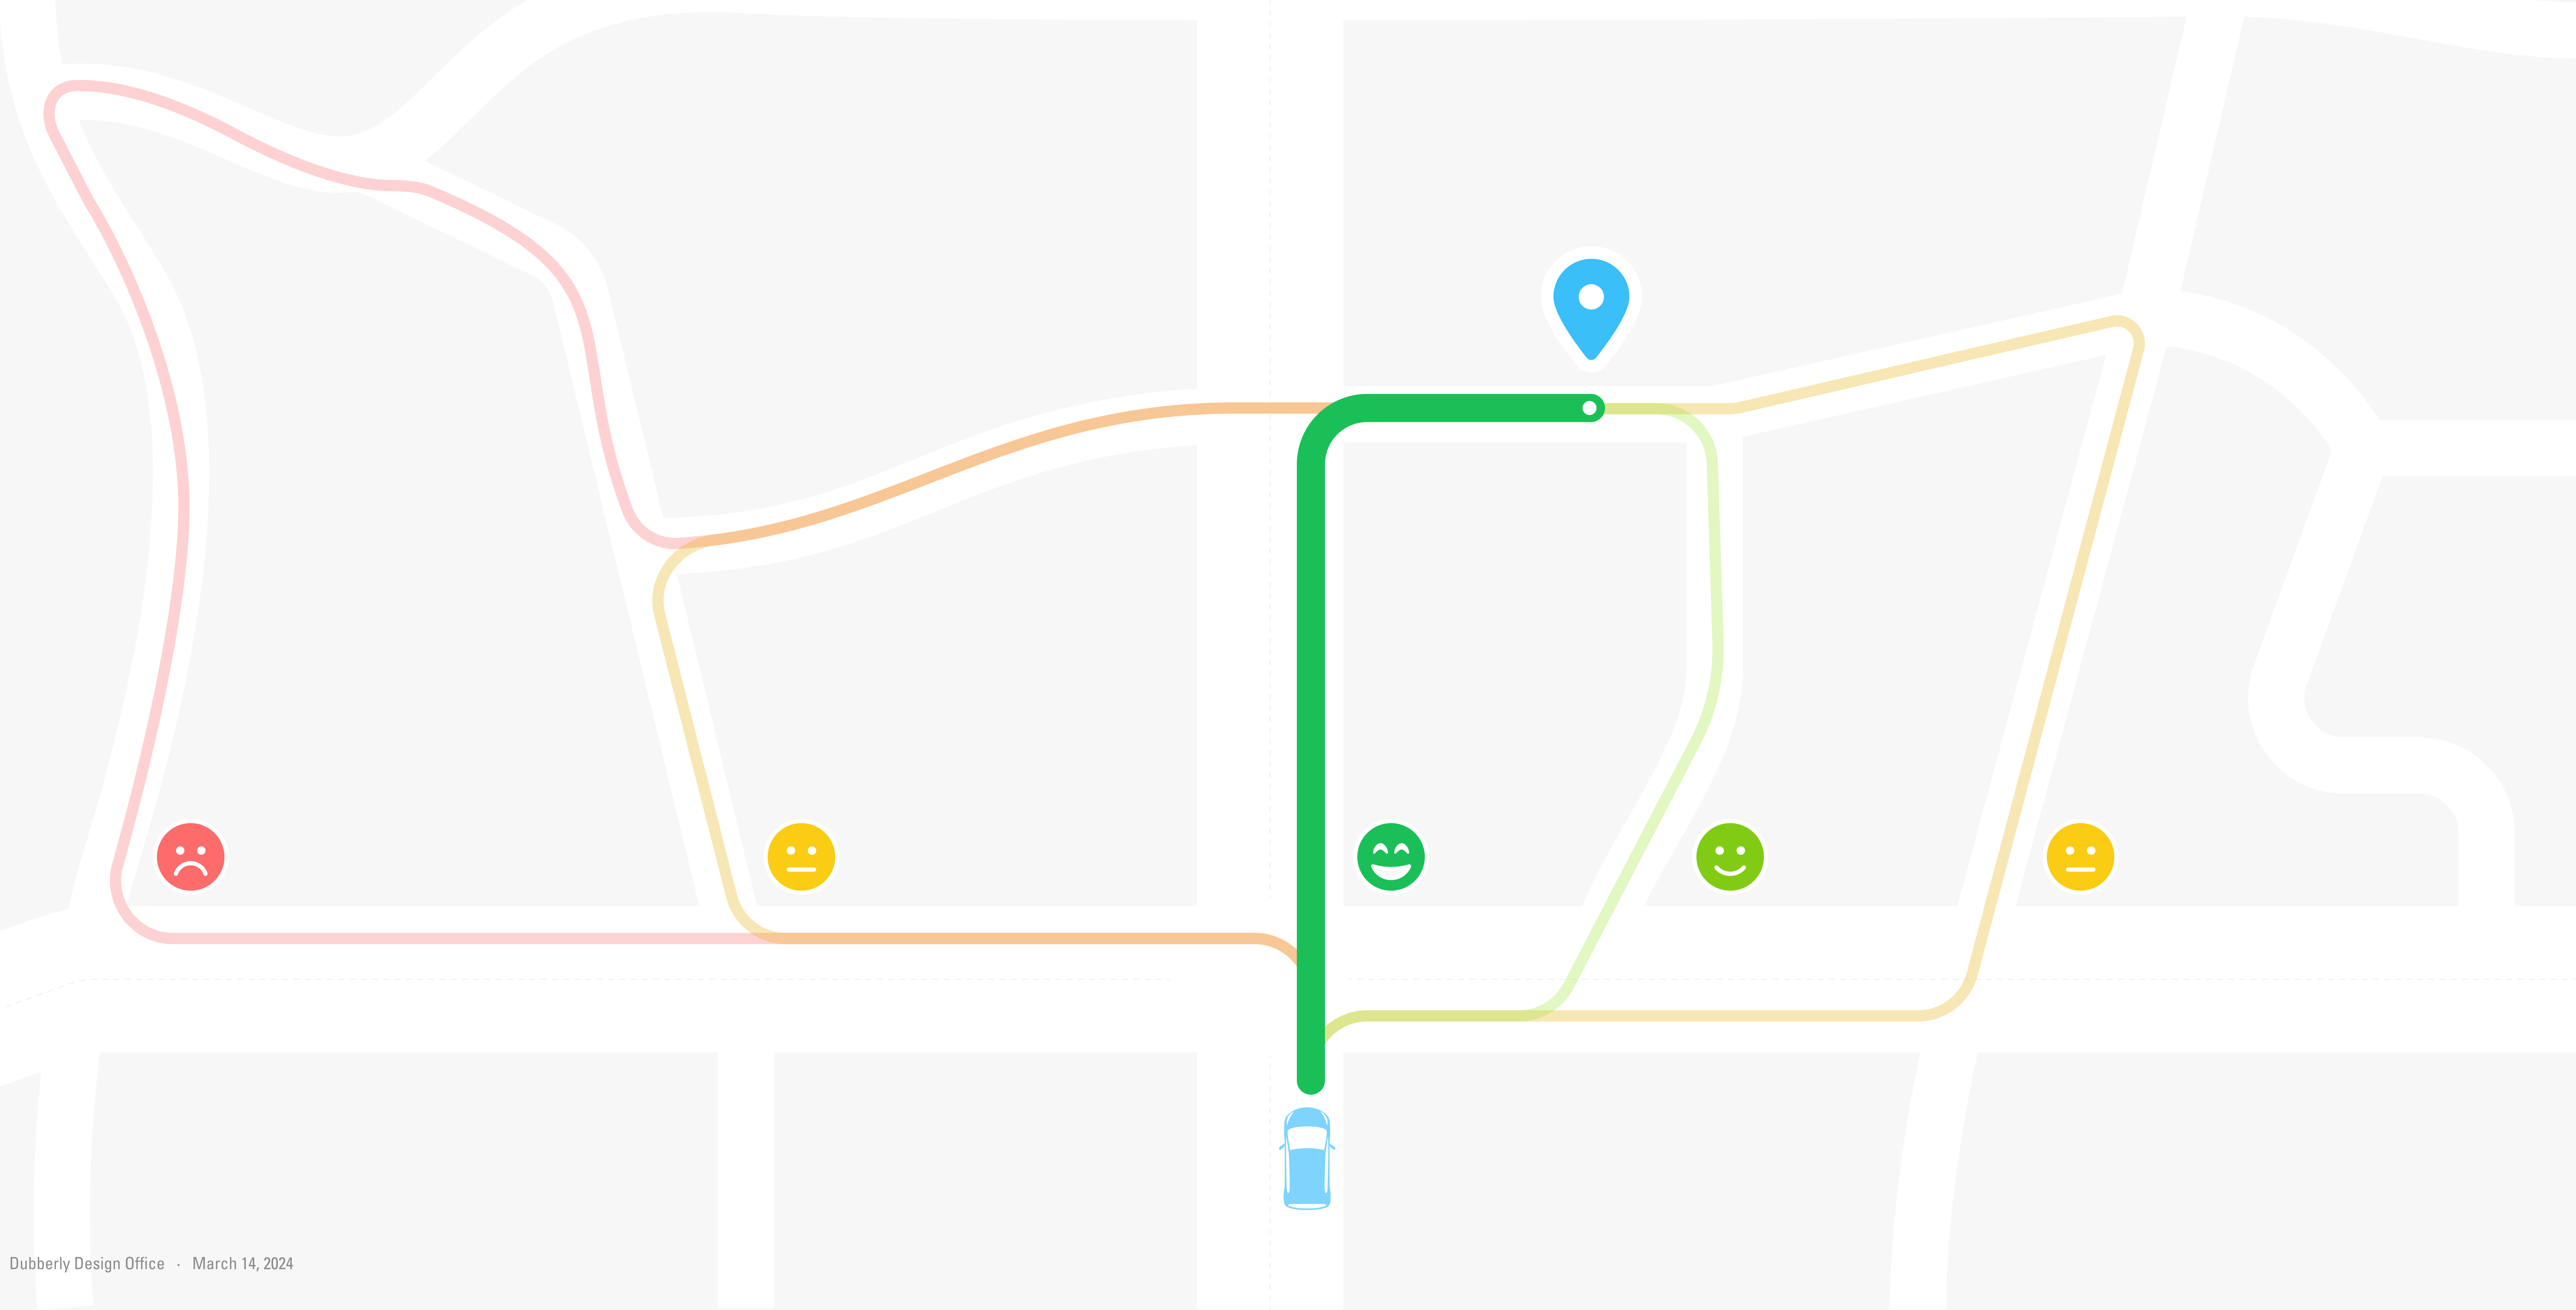

PART TWO

Multi-Agent Systems

Single- vs Multi-Agent systems

If each entity sends its perceptions to and receives its actions from a single central process, then there is only a single agent: the central process. The central agent models all of the entities as a single "self".

Multiagent Systems: A Survey from a Machine Learning Perspective
https://www.cs.cmu.edu/afs/cs/usr/pstone/public/papers/97MAS-survey/revised-survey.html

Single- vs Multi-Agent systems

Multiagent systems differ from single-agent systems in that several agents exist which model each other's goals and actions.

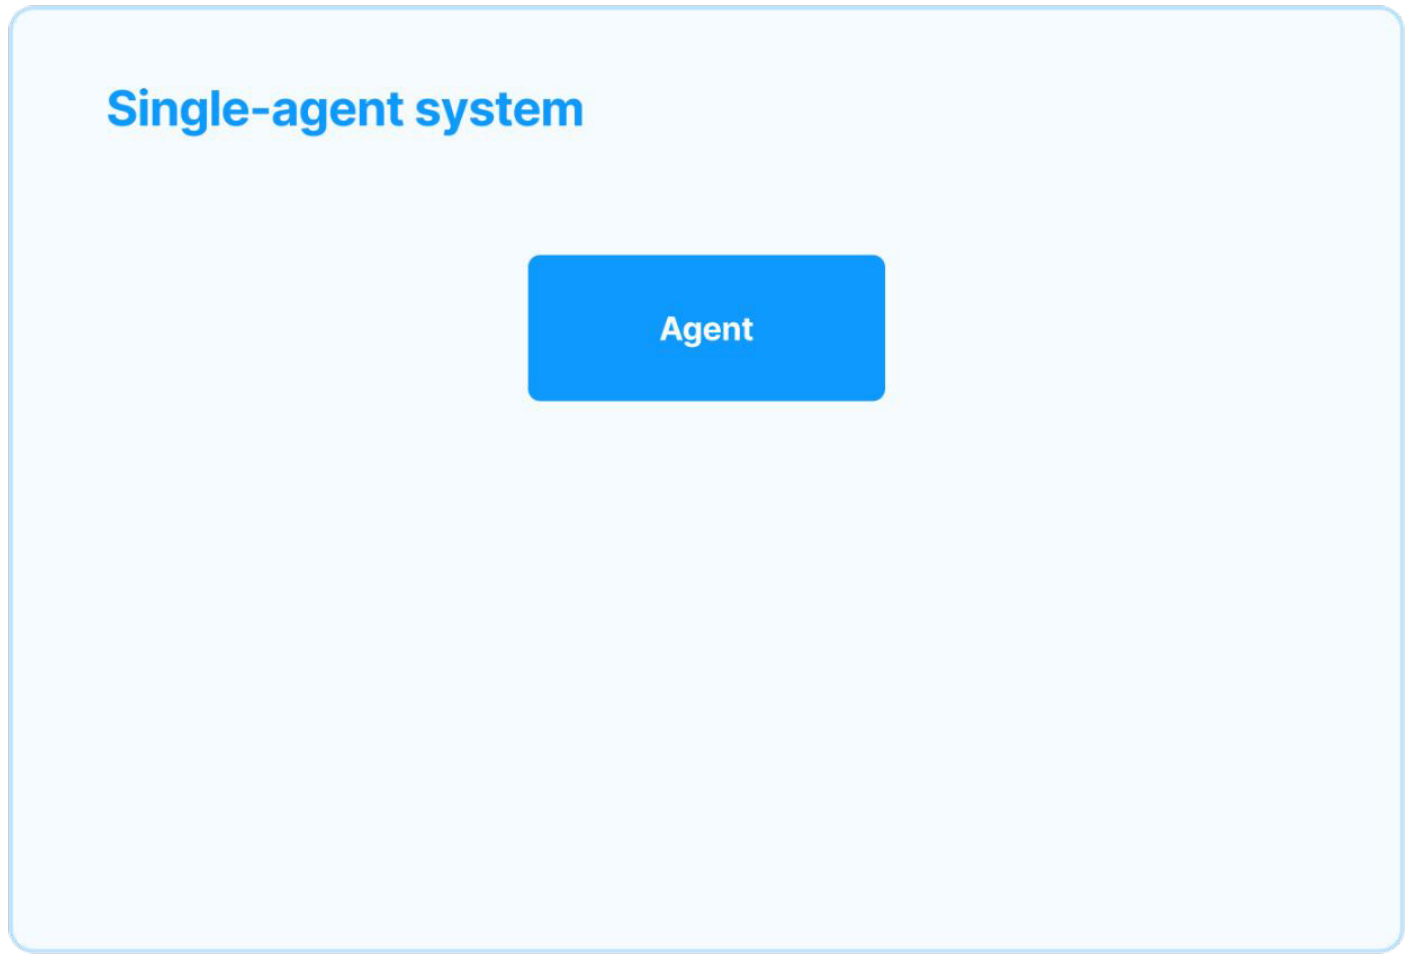

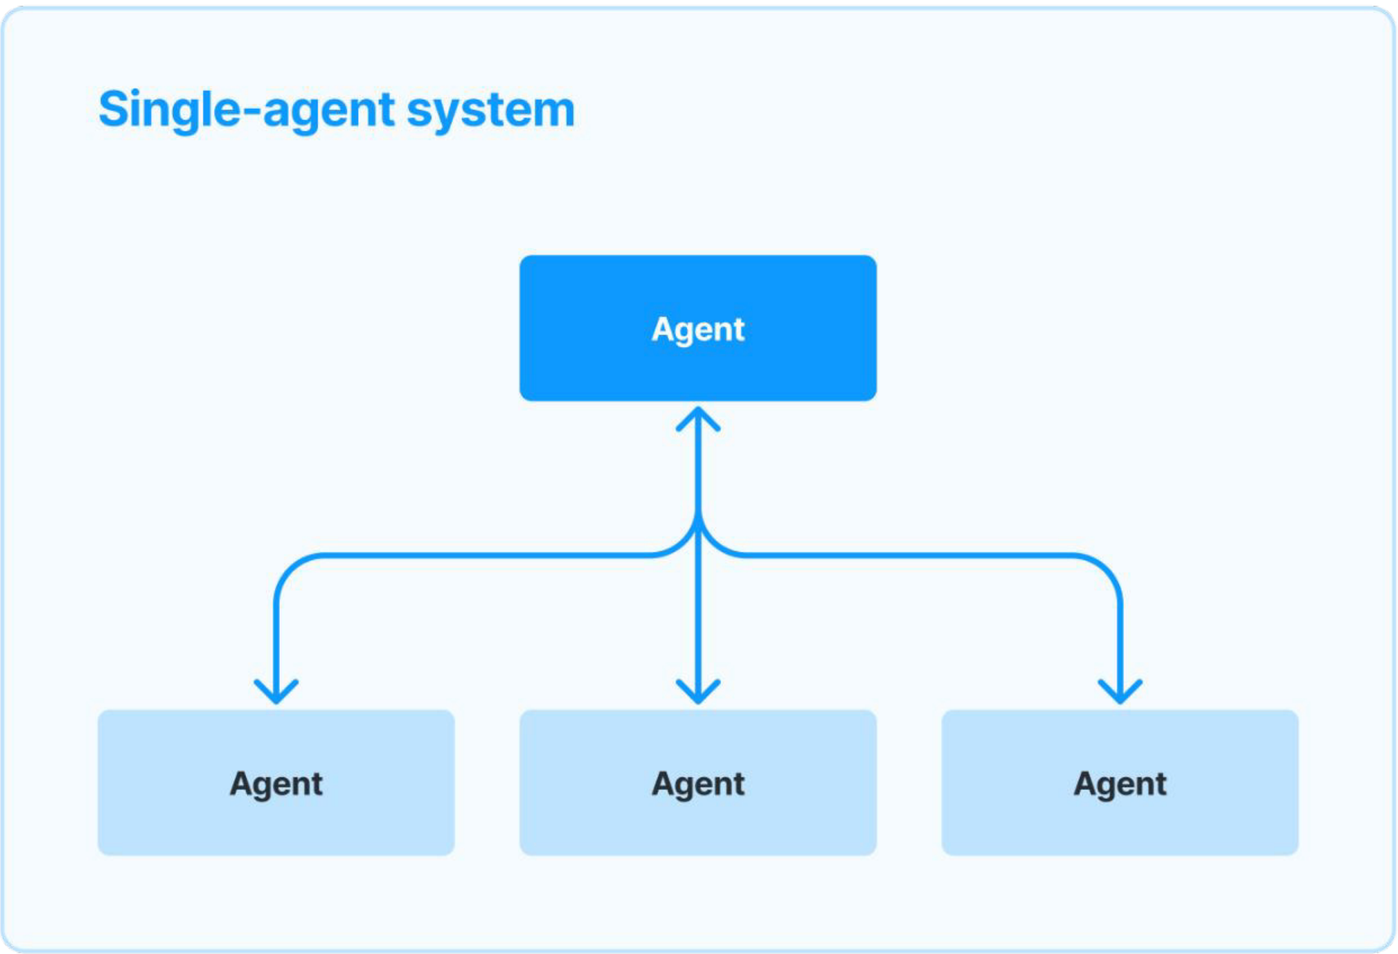

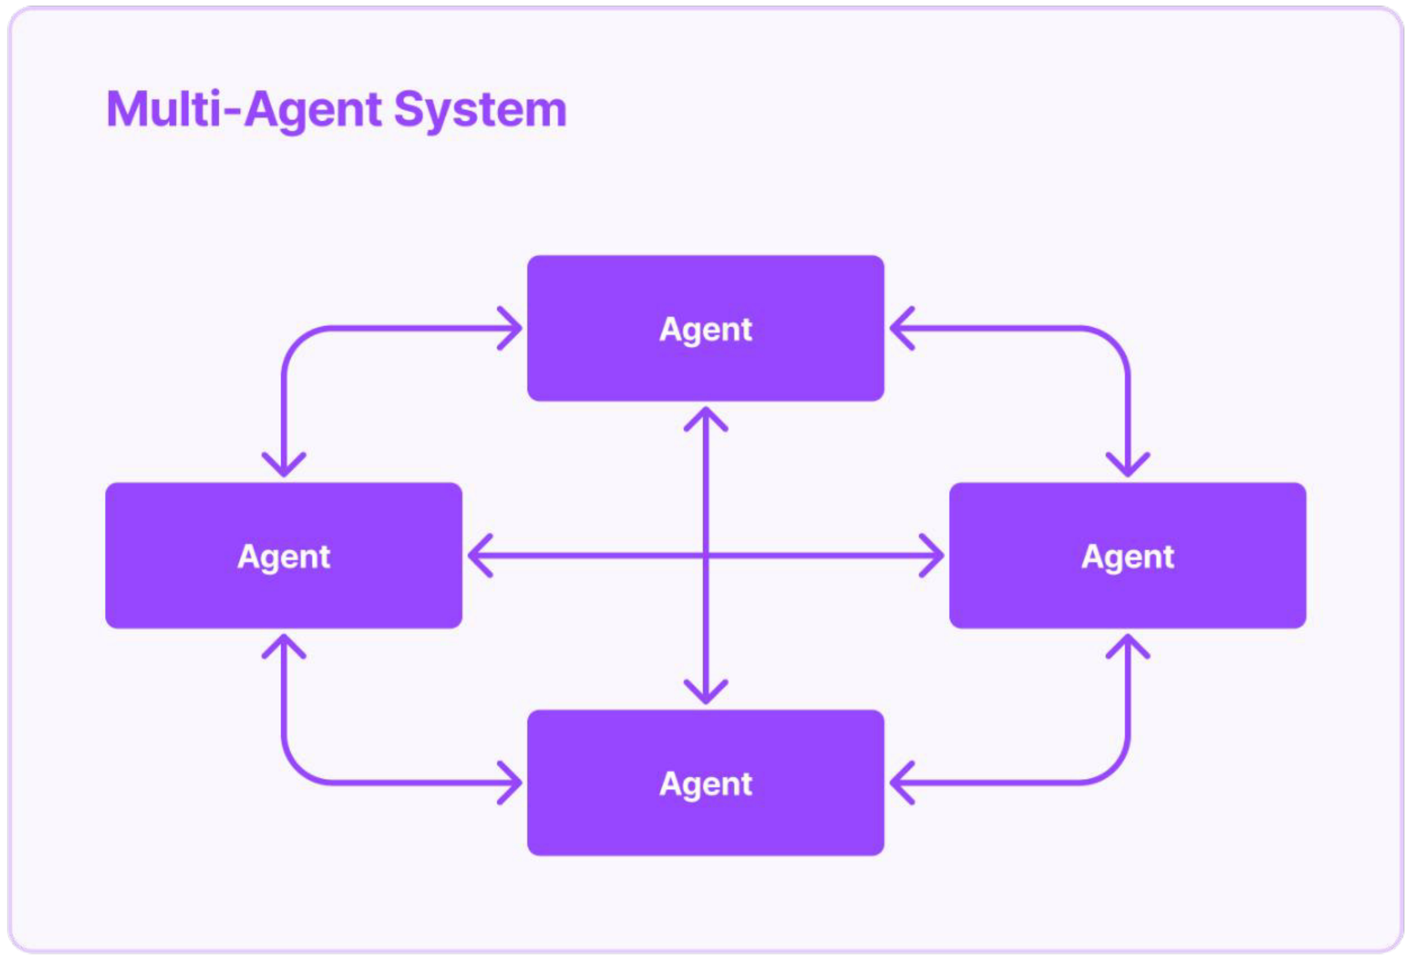

Multiagent Systems: A Survey from a Machine Learning Perspective
https://www.cs.cmu.edu/afs/cs/usr/pstone/public/papers/97MAS-survey/revised-survey.html

Multi-Agent Systems

Intelligent agent systems may act as a nested hierarchy. Multiple individual agents organized within a single environment acting as a single entity.

Multi-Agent Systems

**Each individual agent within the system
has the intelligent agent structure which includes:
sensors (percepts),
goals,
a decision process,
and actuators
which act on the environment.**

A general structure of a Multi-Agent system.

A Multi-Agent System (MAS)
is constituted by a network of agents,
which typically interact with each other.

A MAS represents a society of
autonomous interacting components
dedicated to provide solutions
for larger-scale problems.

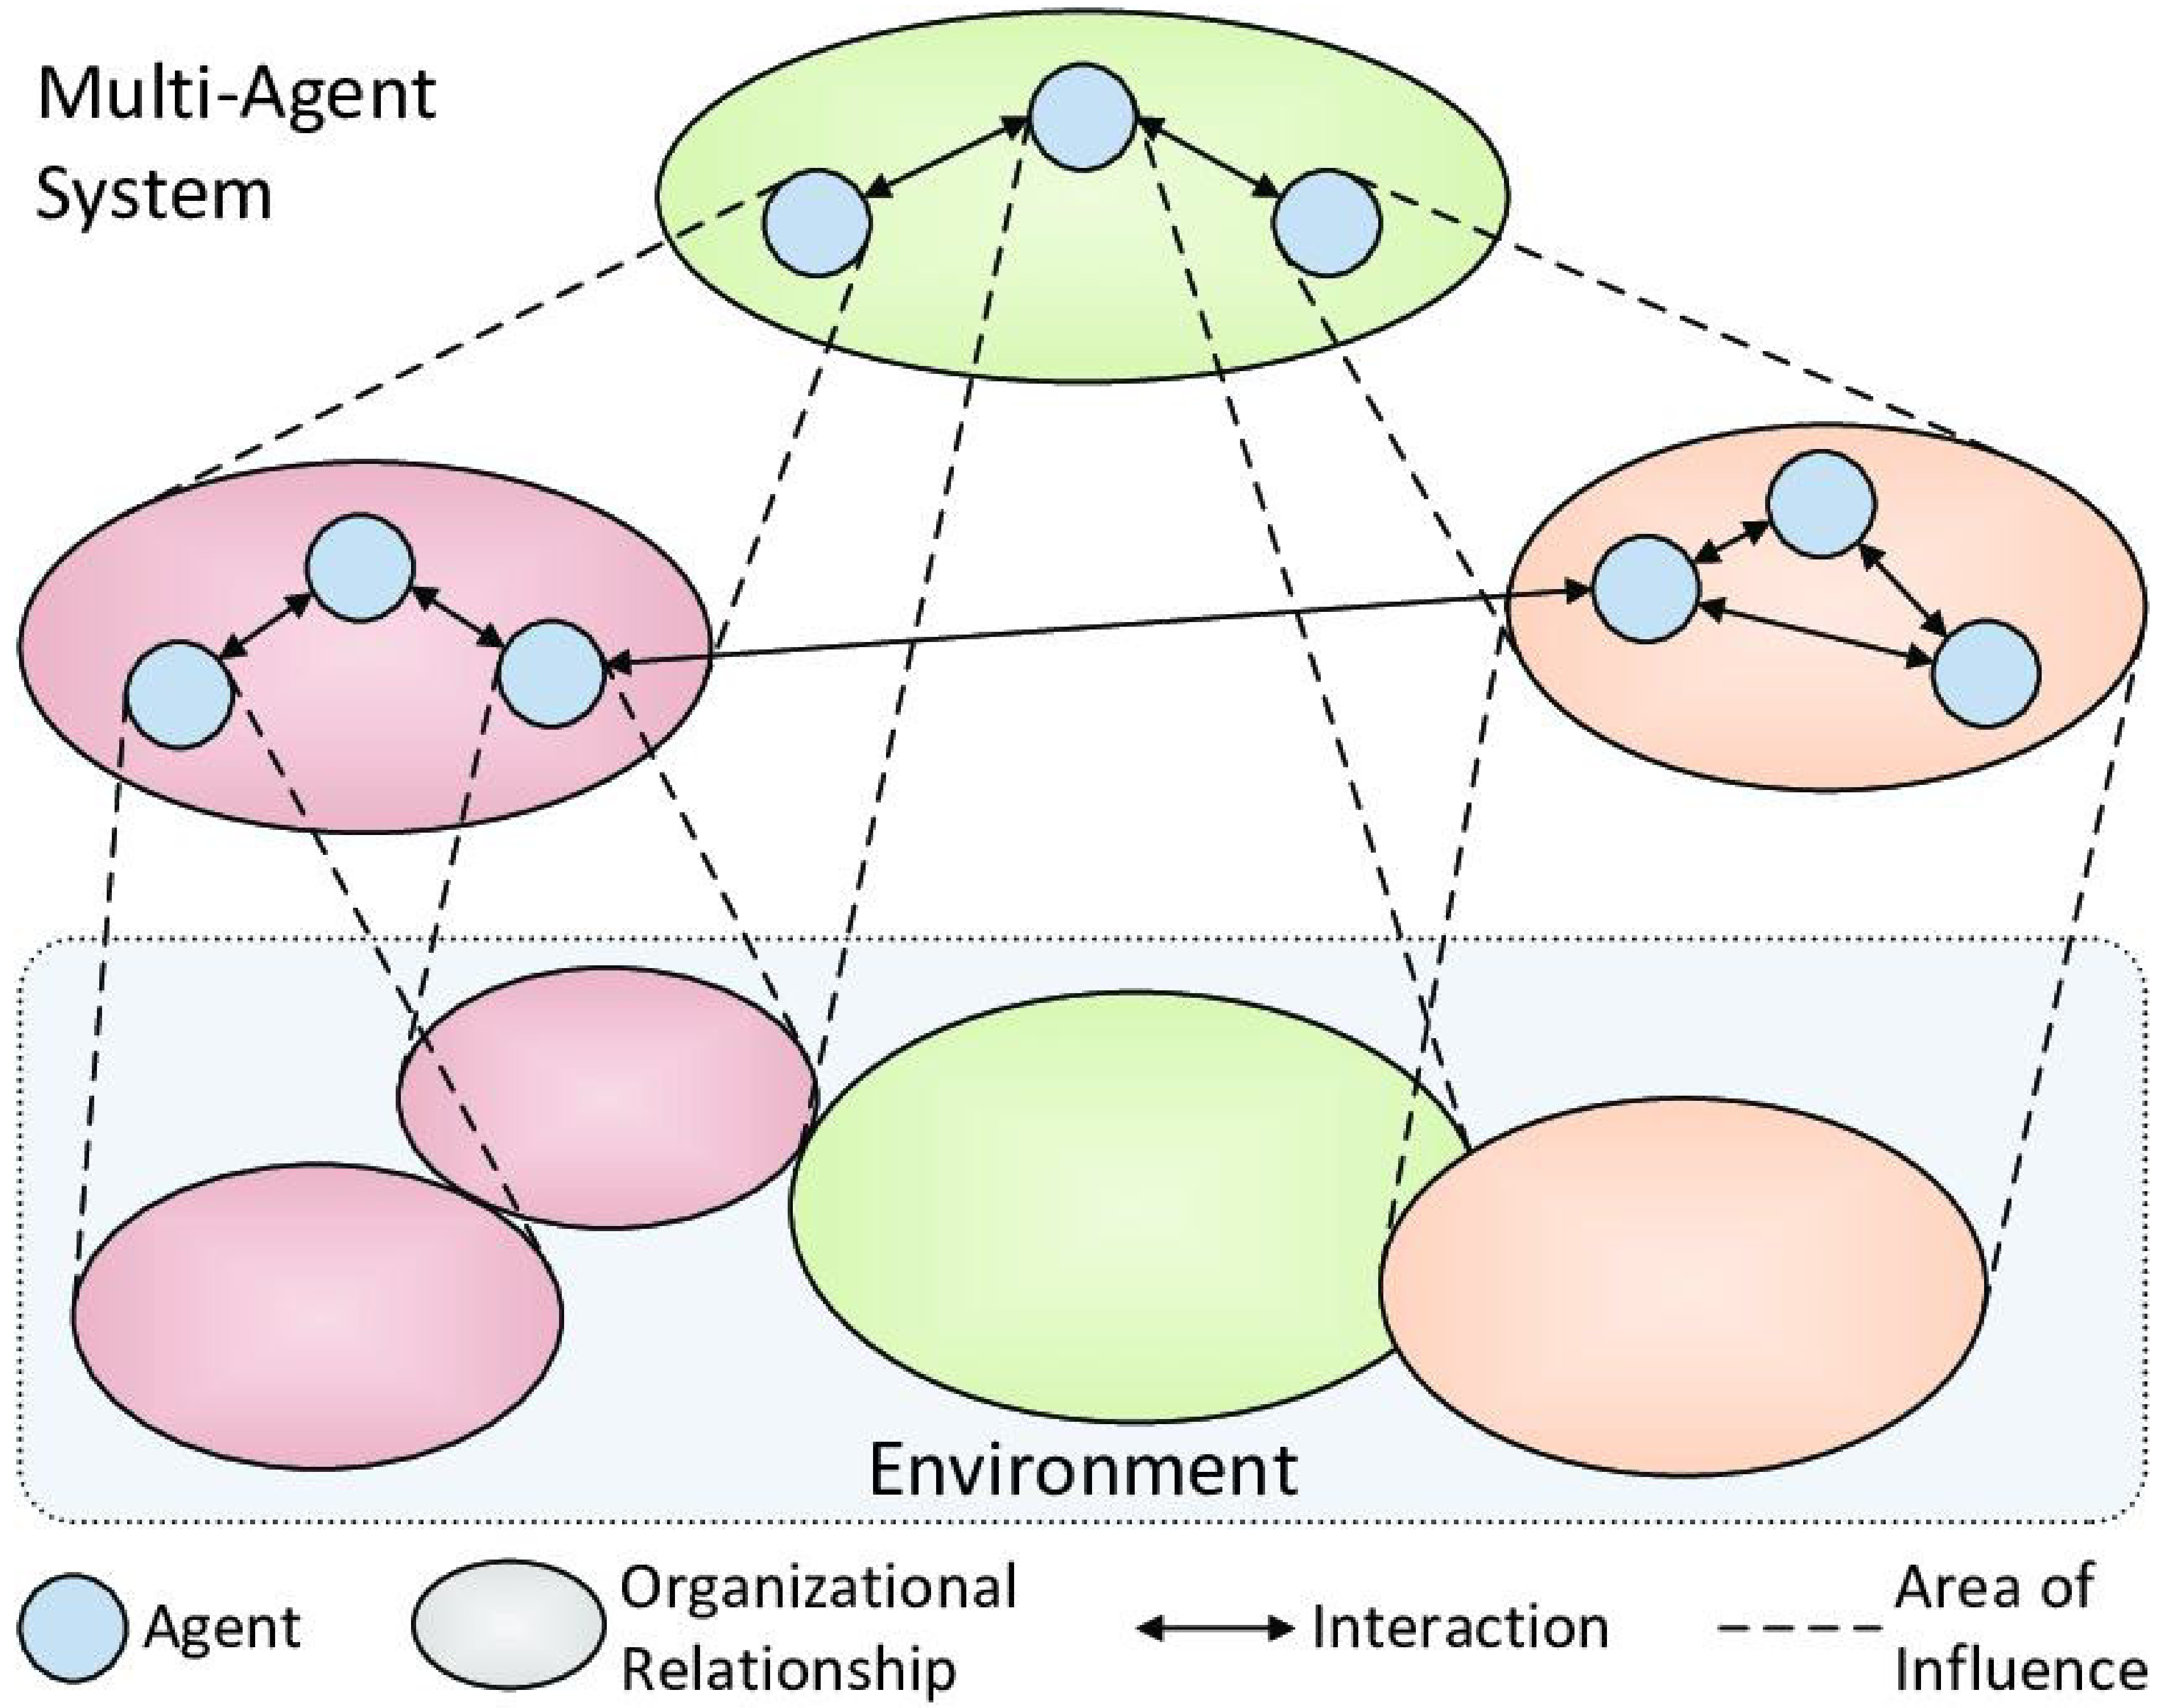

Lepuschitz, Wilfried. (2018). Self-Reconfigurable Manufacturing Control based on Ontology-Driven Automation Agents.
https://www.researchgate.net/figure/General-structure-of-a-multi-agent-system-42_fig4_324808821

The software development SOPs between MetaGPT and real-world human teams.

In software engineering,
SOPs promote collaboration among various roles.
MetaGPT showcases its ability
to decompose complex tasks
into specific actionable procedures
assigned to various roles
(e.g., Product Manager, Architect, Engineer, etc.).

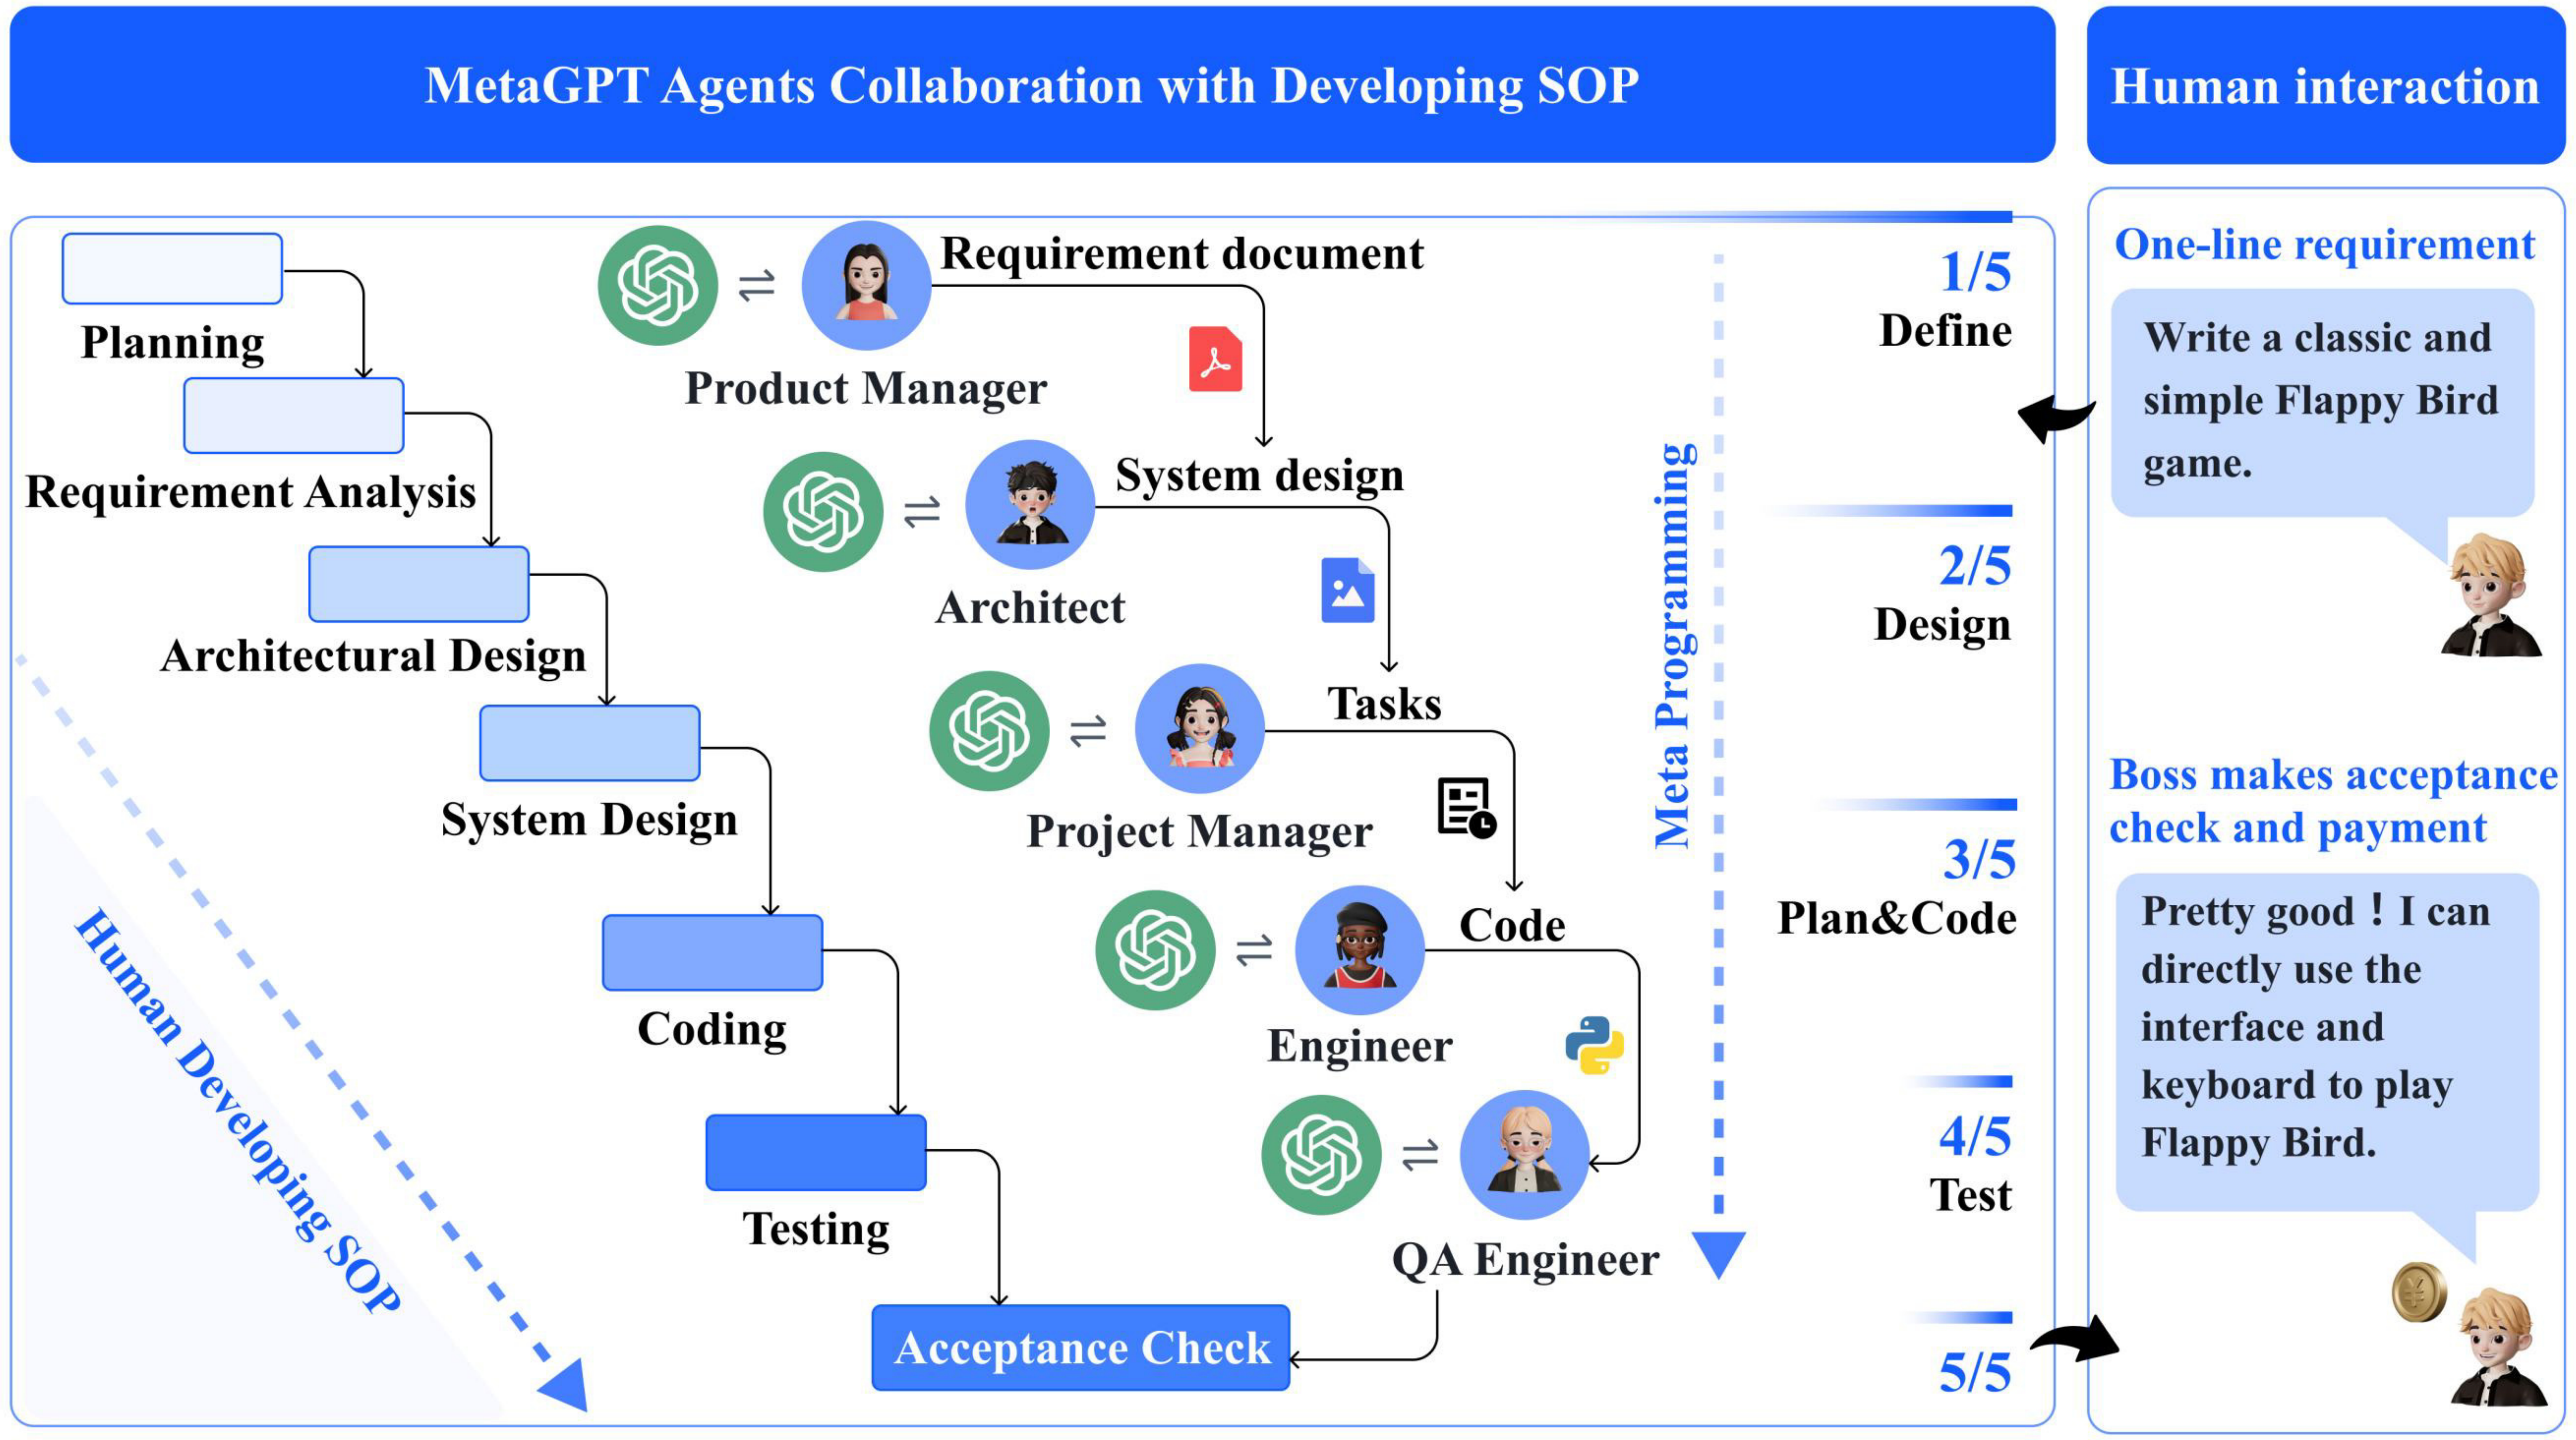

MetaGPT: Meta Programming for a Multi-Agent Collaborative Framework
https://arxiv.org/pdf/2308.00352

PART THREE

Intelligent agents as cybernetic control systems

How does Russell & Norvig's framework map to cybernetic systems?

The ‘decision process’ in cybernetic systems

A first-order cybernetic control system perceives its environment, then takes actions based on a goal.

This model maps to Russell & Norvig's model of intelligent agents.

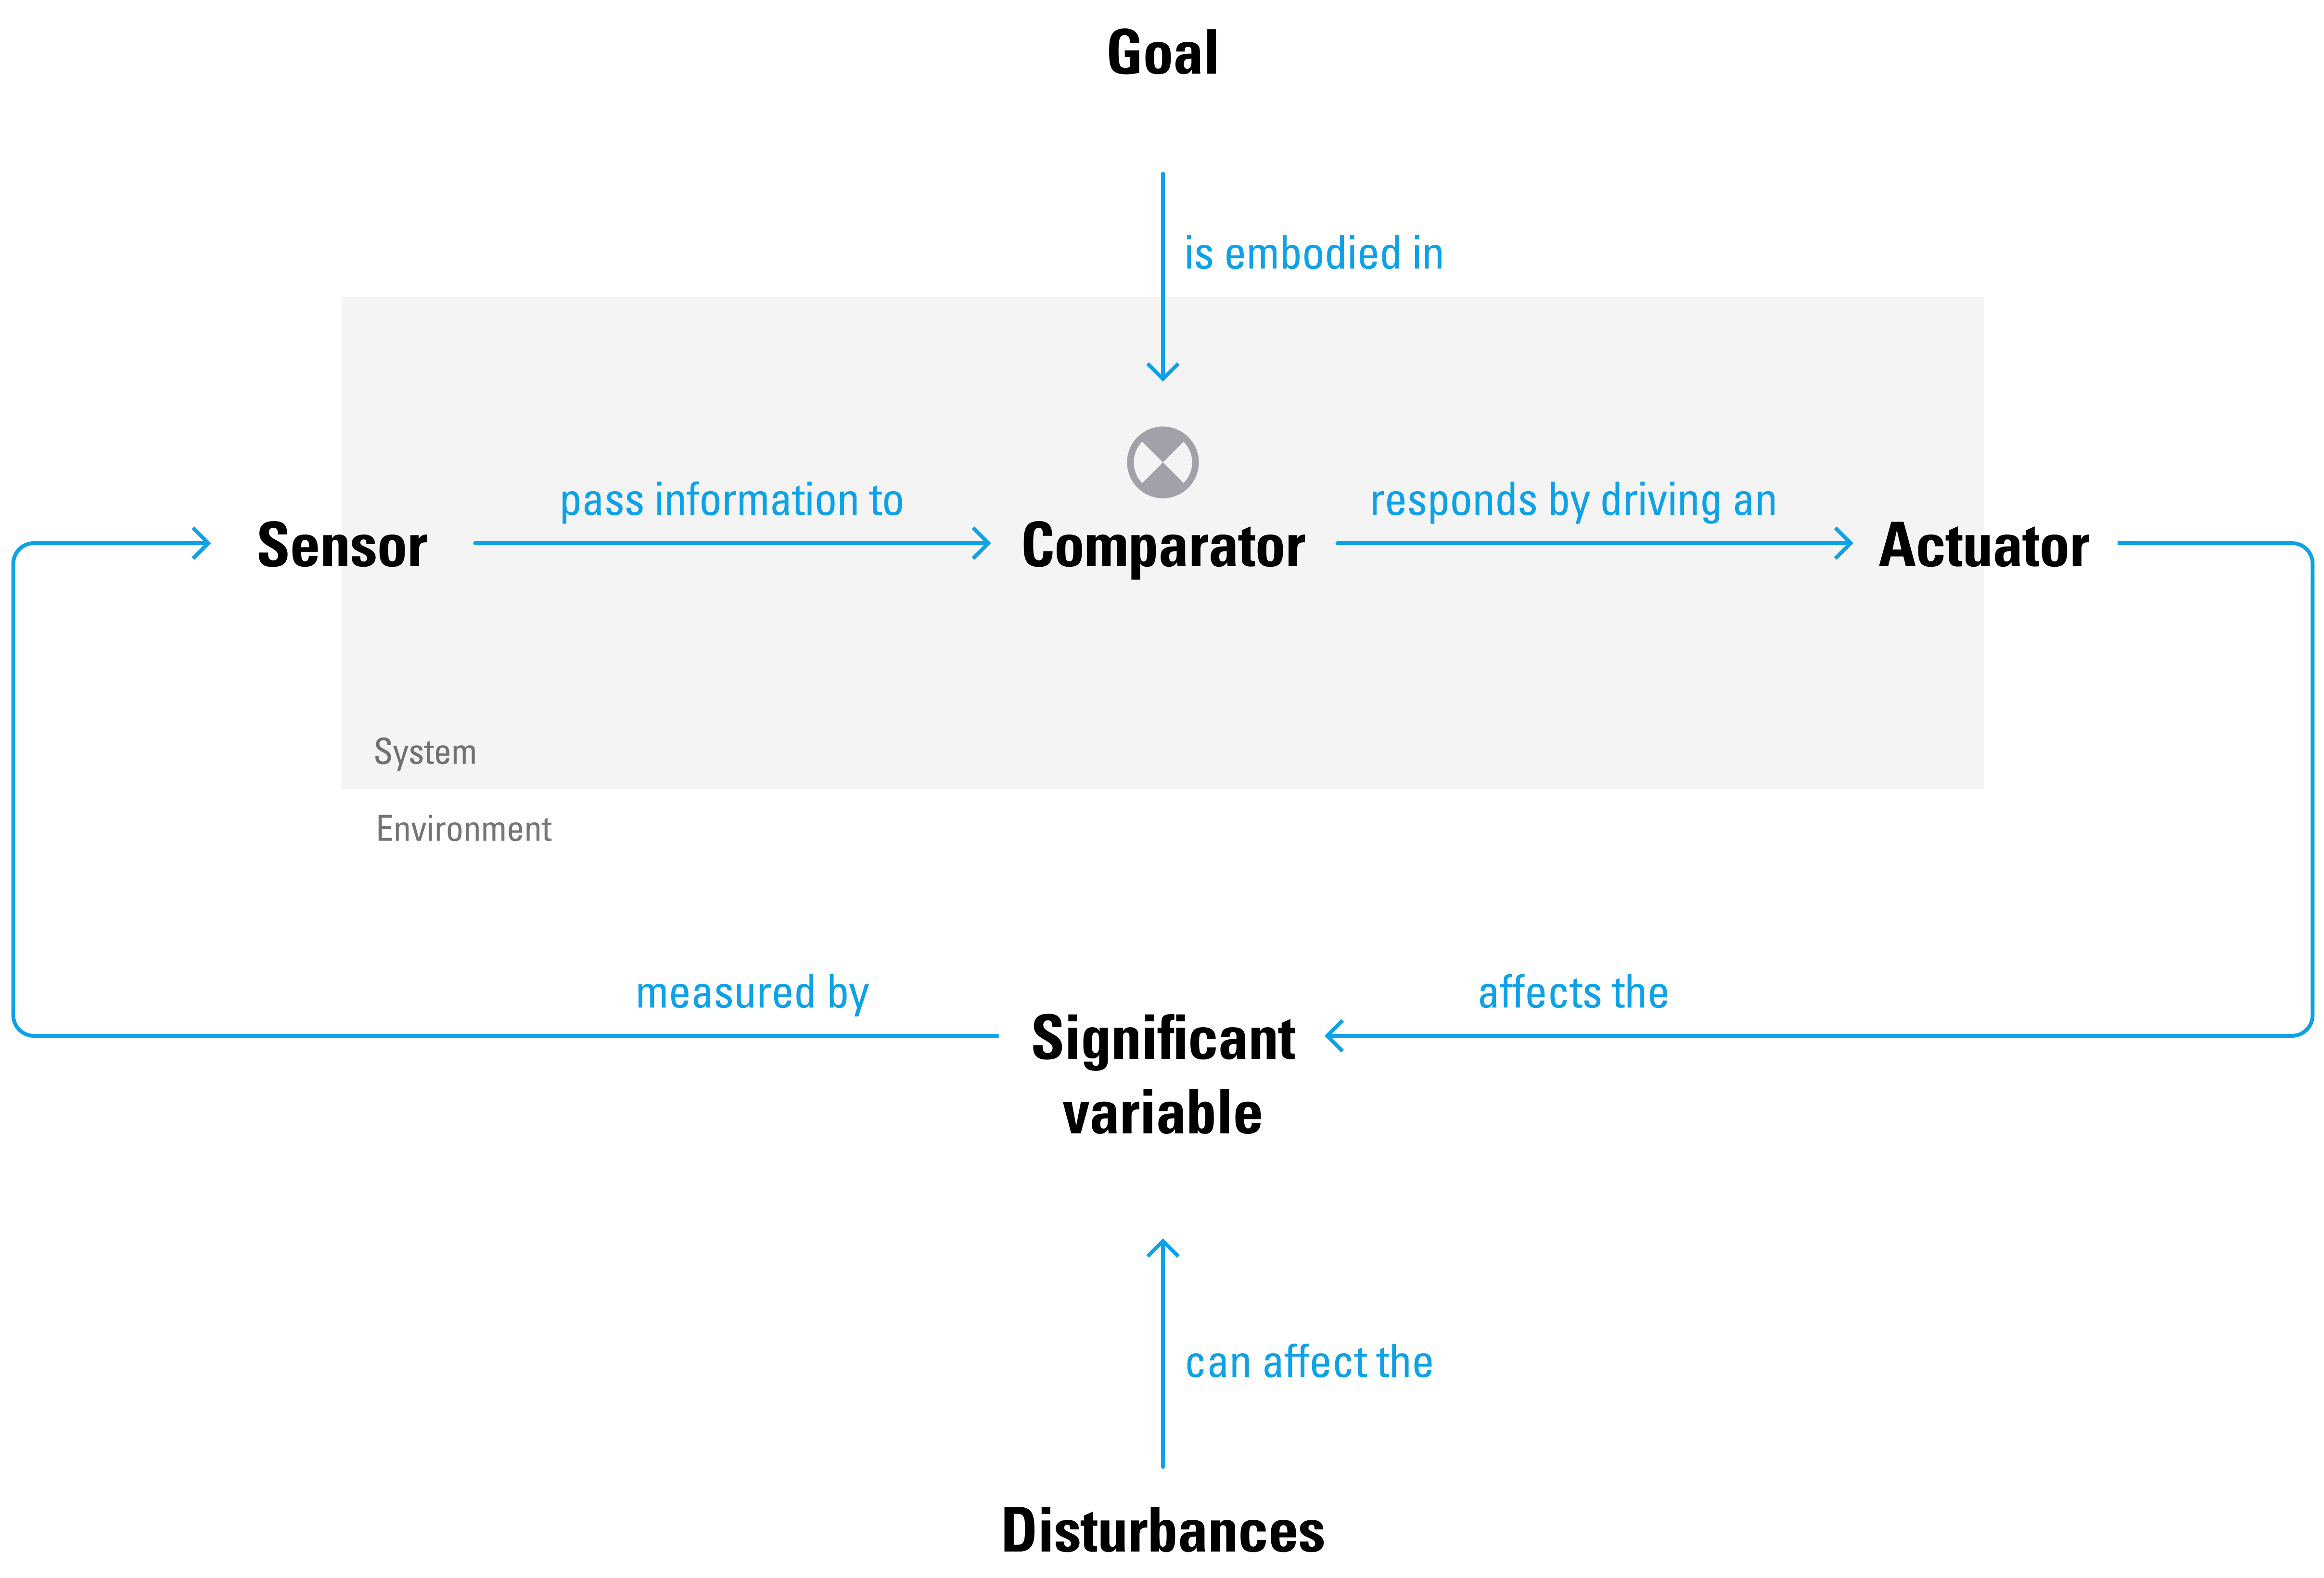

The intelligent agent’s ‘decision process’ is contained within the comparator,

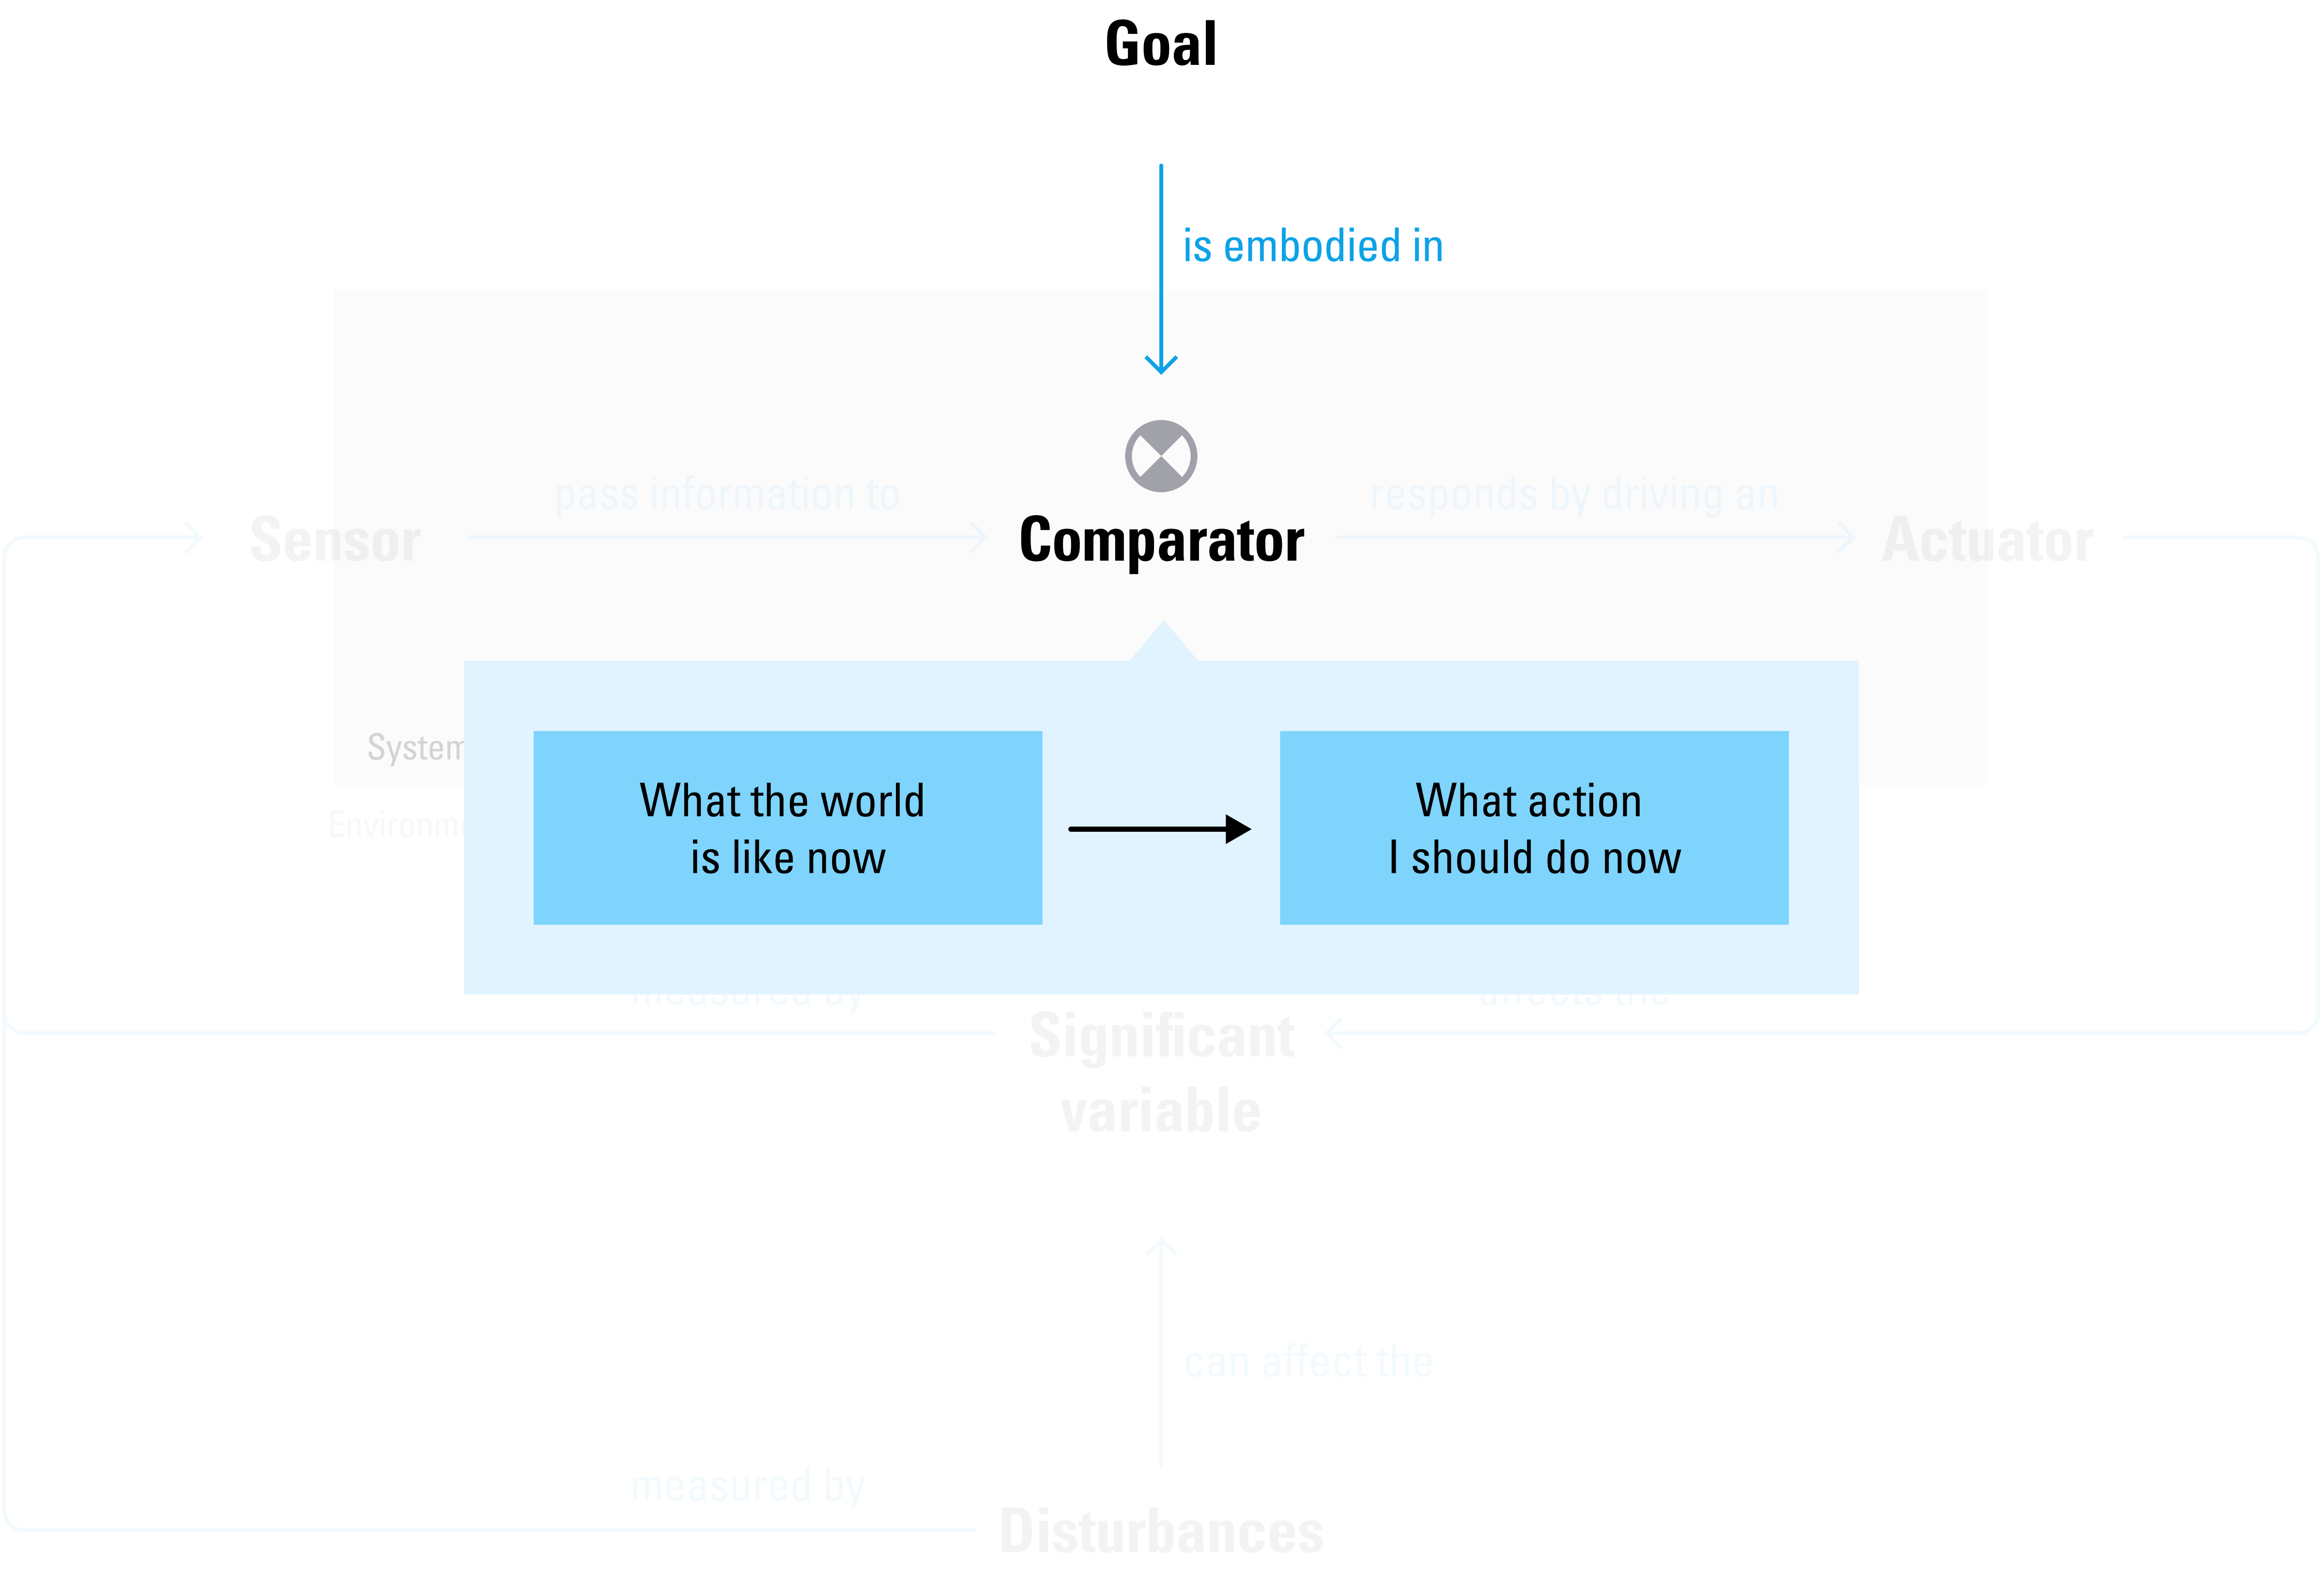

The intelligent agent's 'decision process' is contained within the comparator, with 'condition-action rules', 'goals', or 'utility' as the goal.

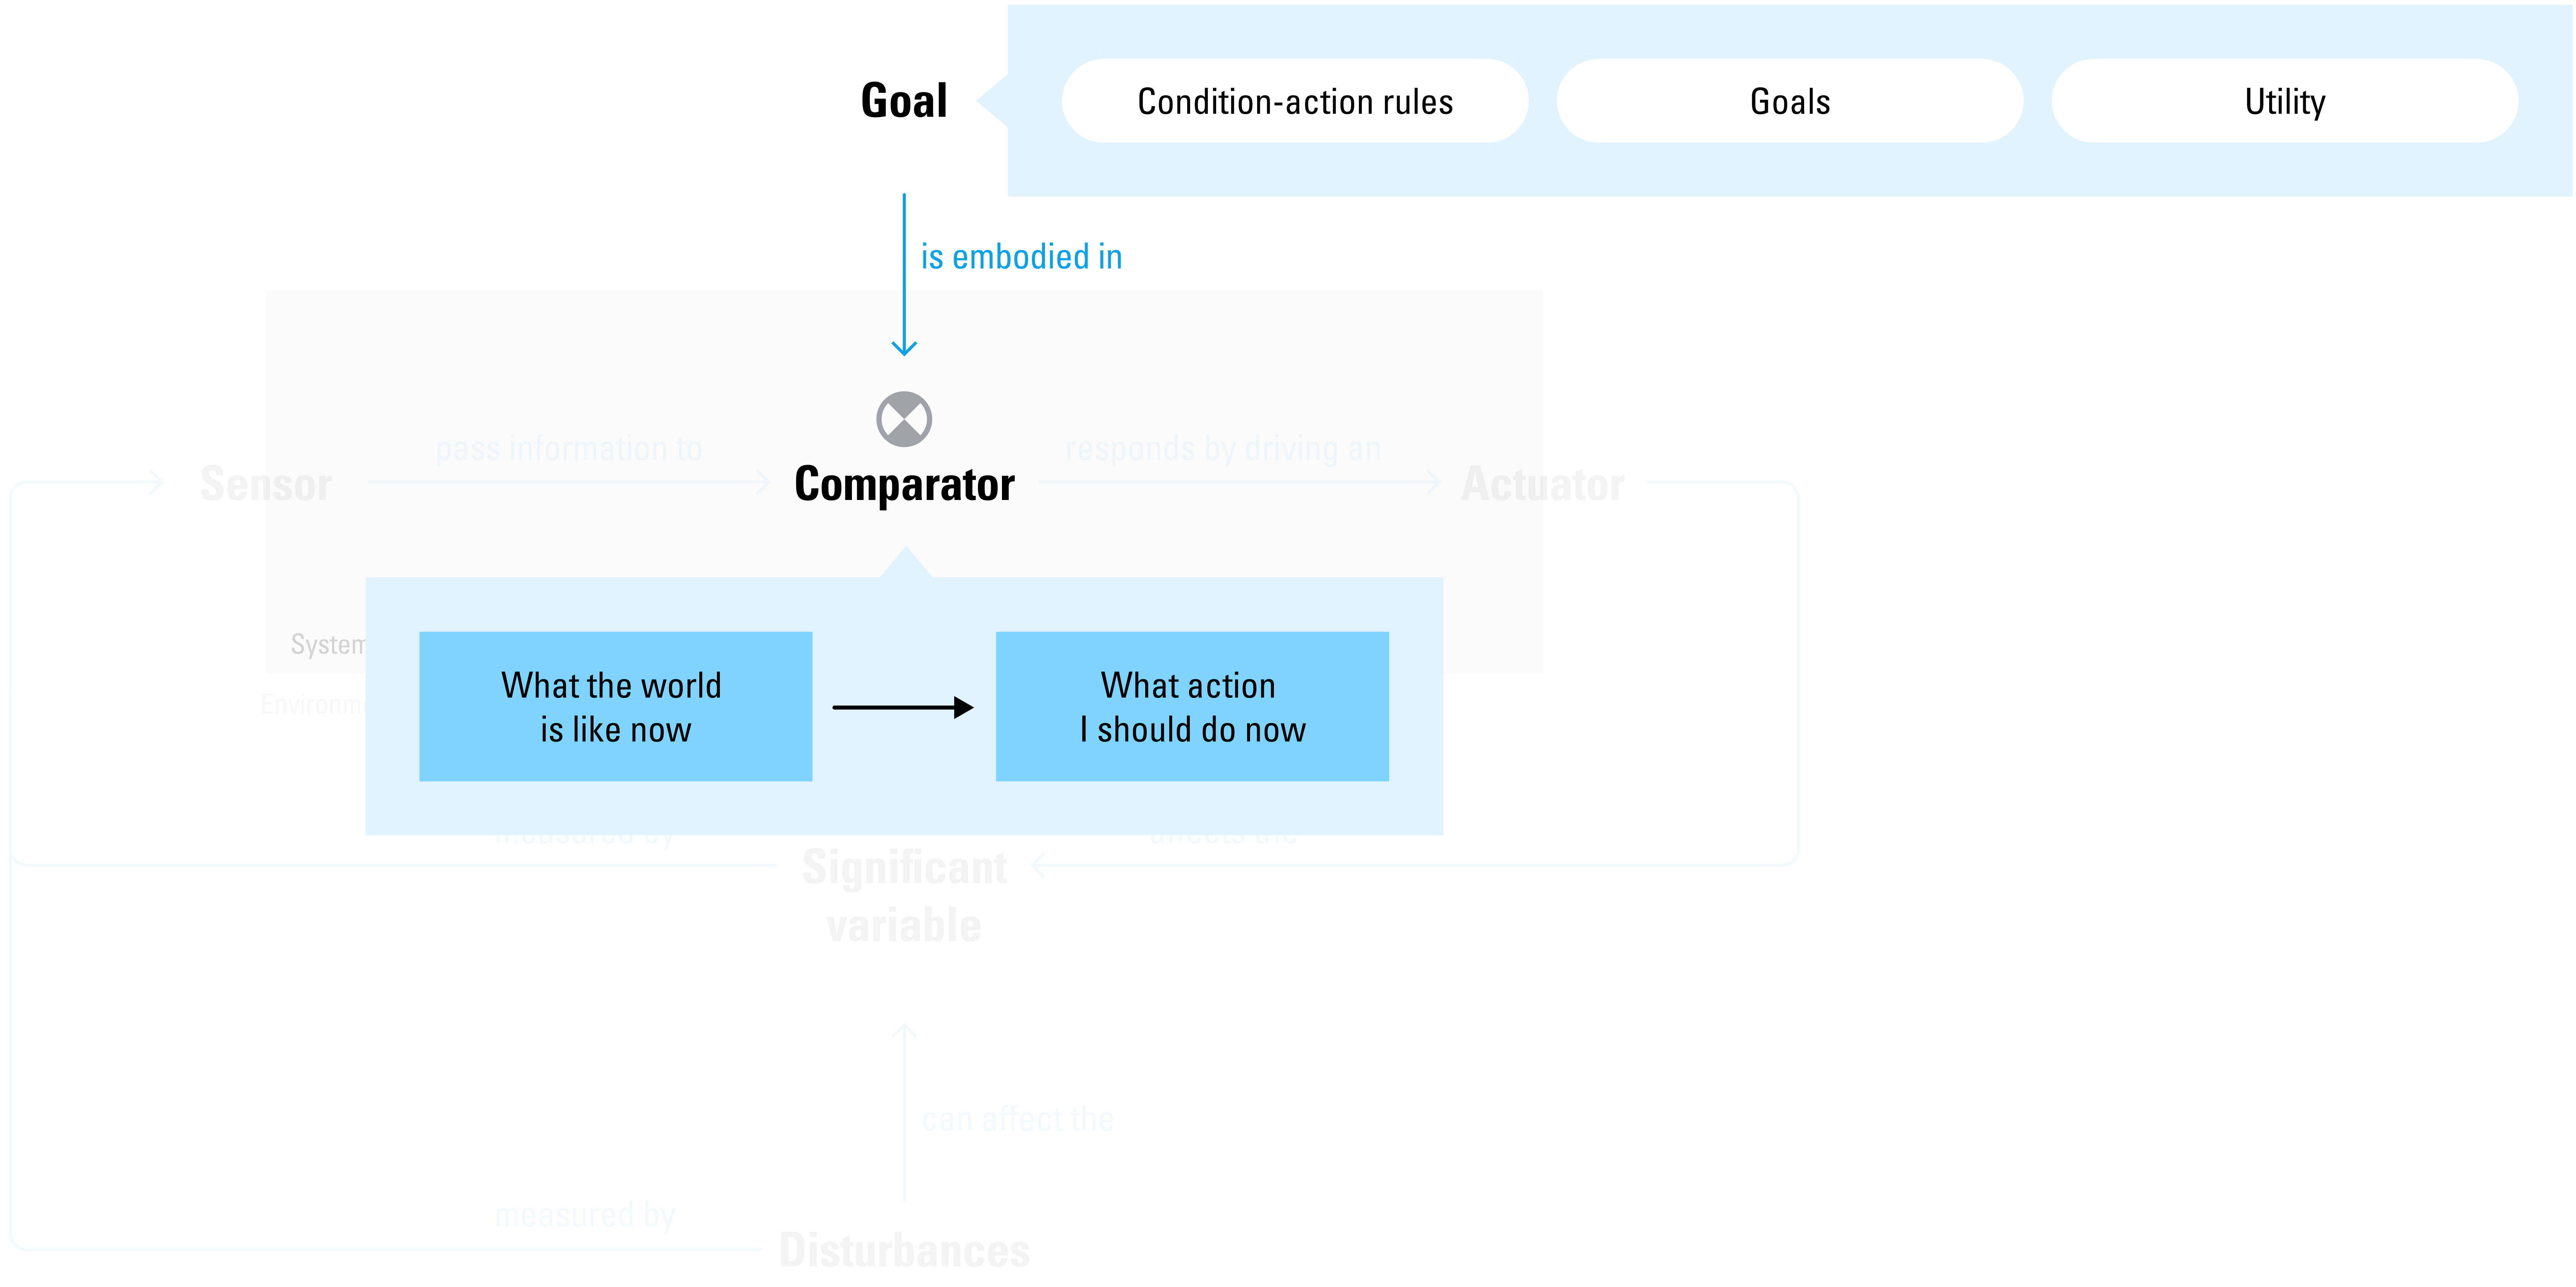

Digital twins as models in model-based agents

Using a model of ‘how the world works’, agents can predict future outcomes rather than only responding to current information.

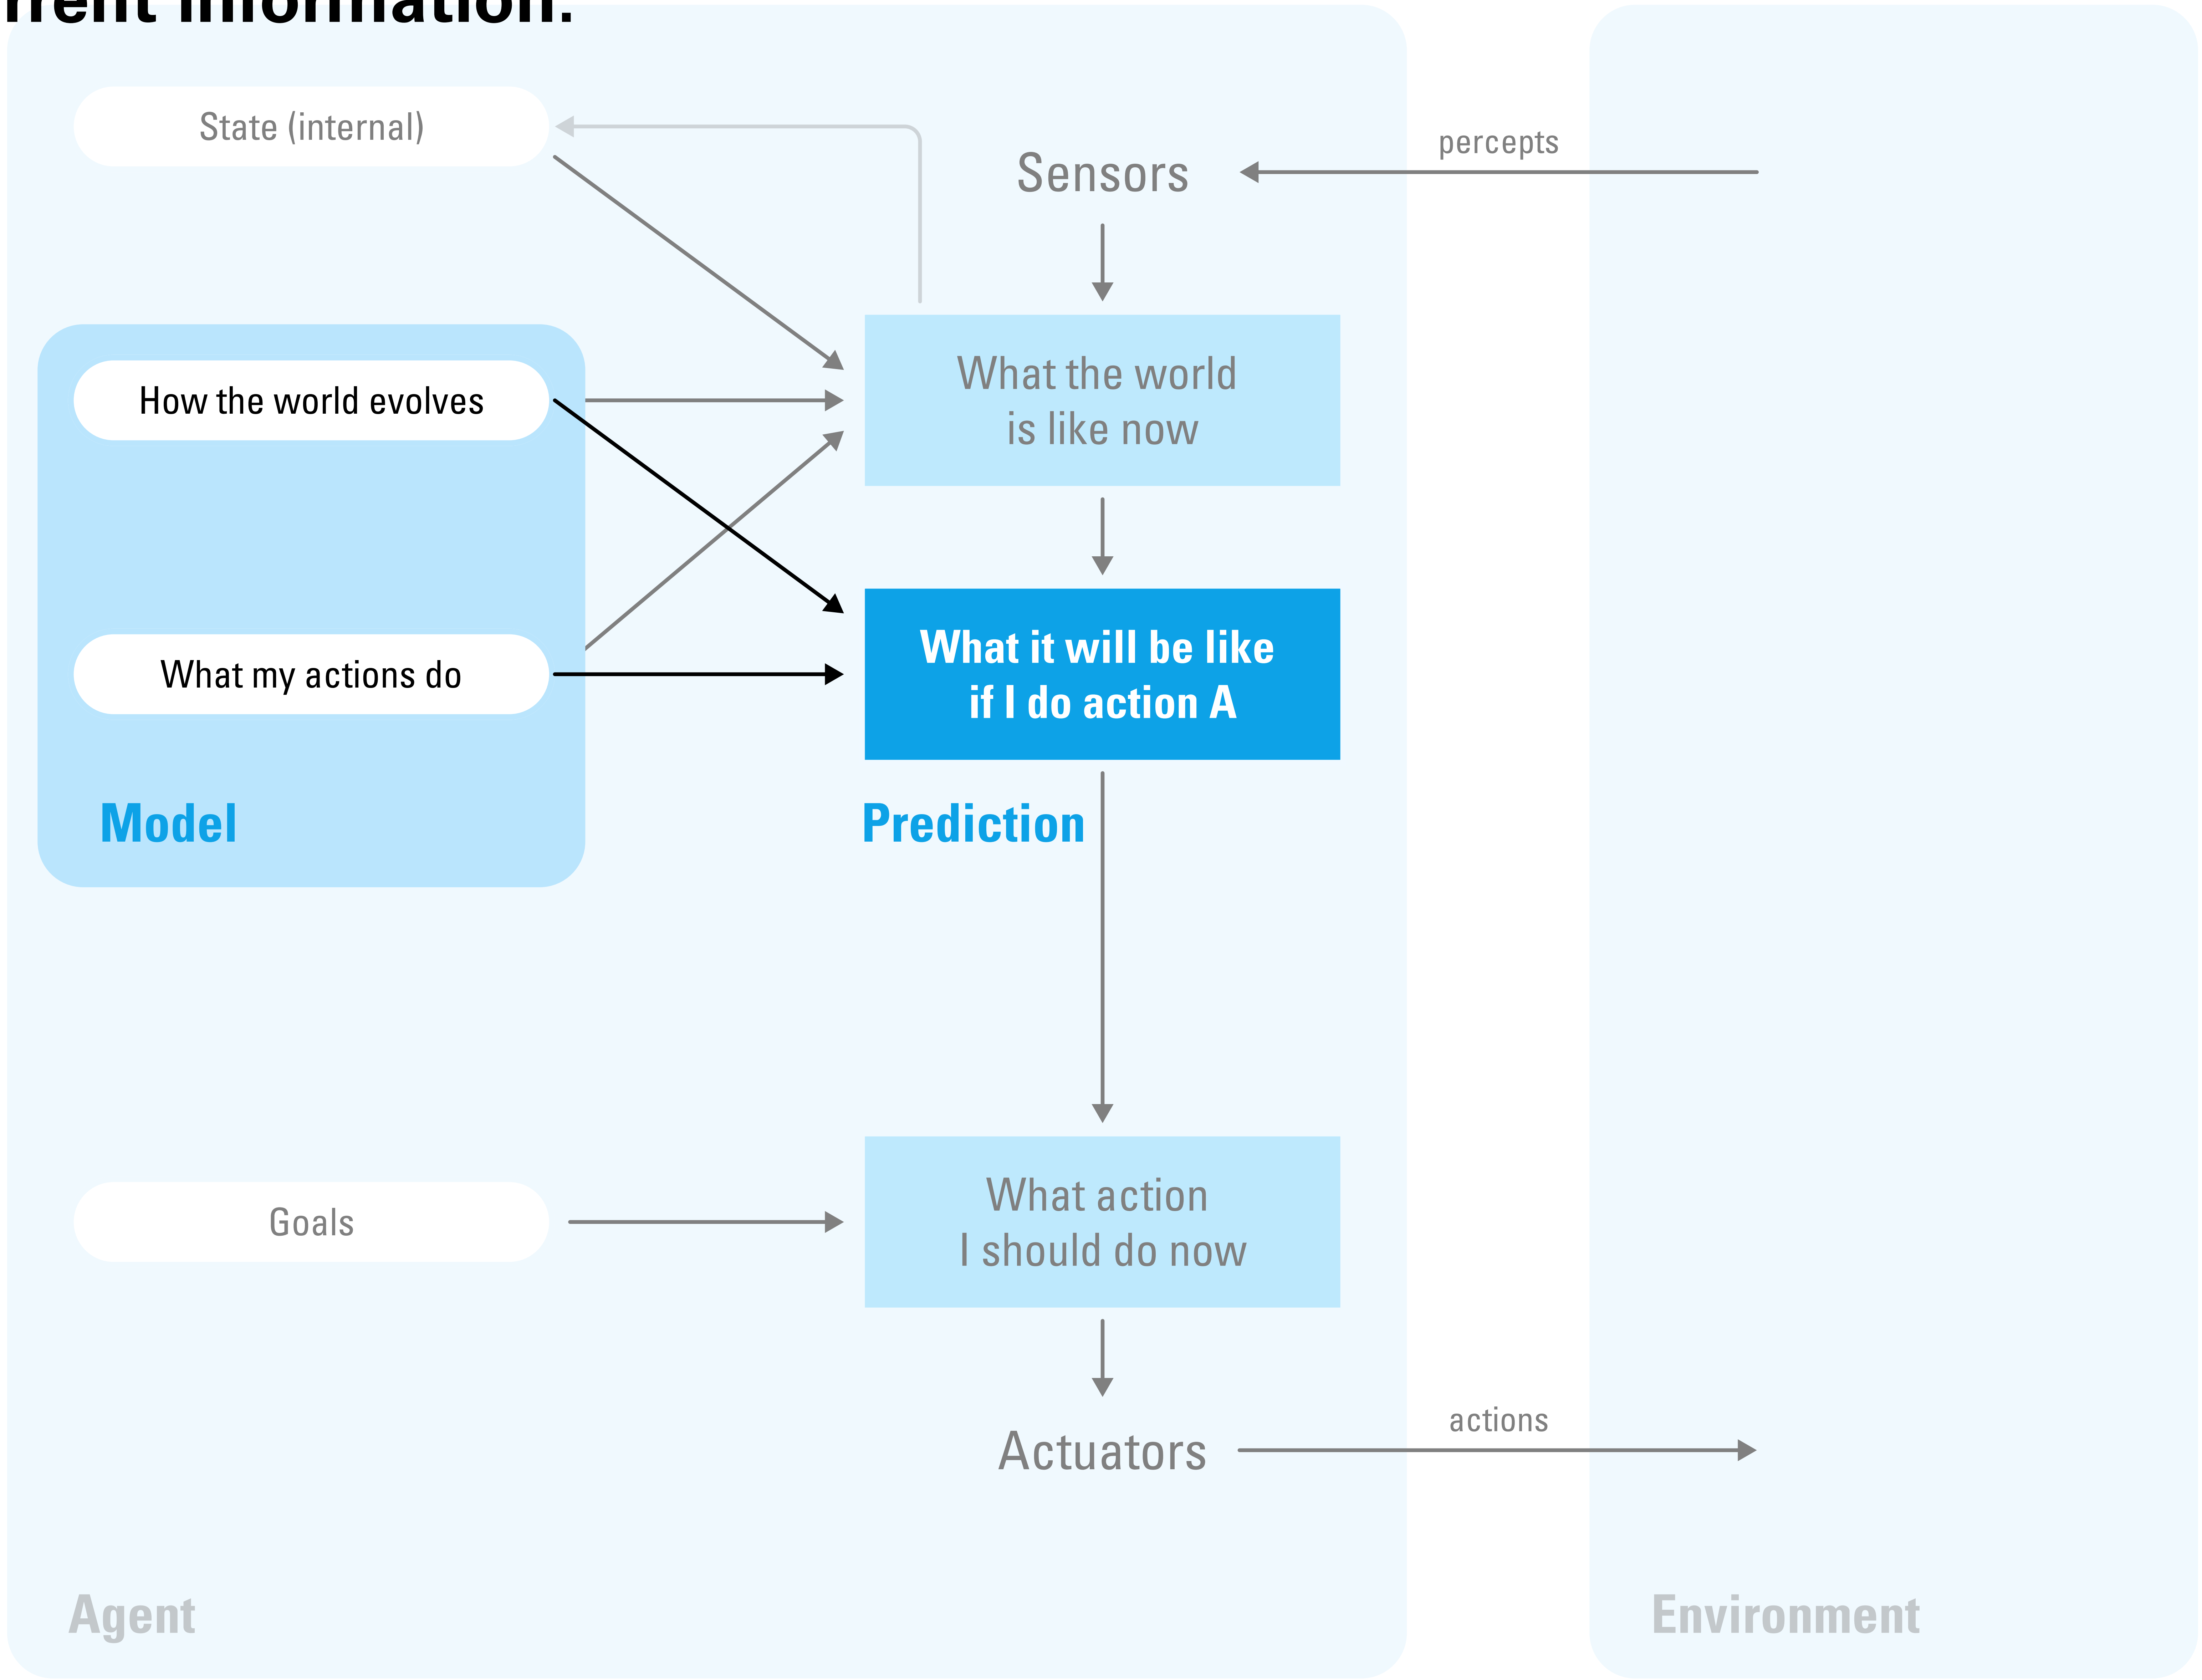

A ‘Digital Twin’ is created by storing the data collected by sensors, then using it to build models — which then can be used to predict future outcomes.

1. Gather histories
Sensors make a series of point-in-time measurements. As measurements accumulate, an historical record emerges.

2. Derive models
Sufficient historical data enable analysts to discover patterns and relationships — these are codified in models.

3. Predict futures
Once trained, new measurements are fed through the model to predict the future — enabling us to act today.

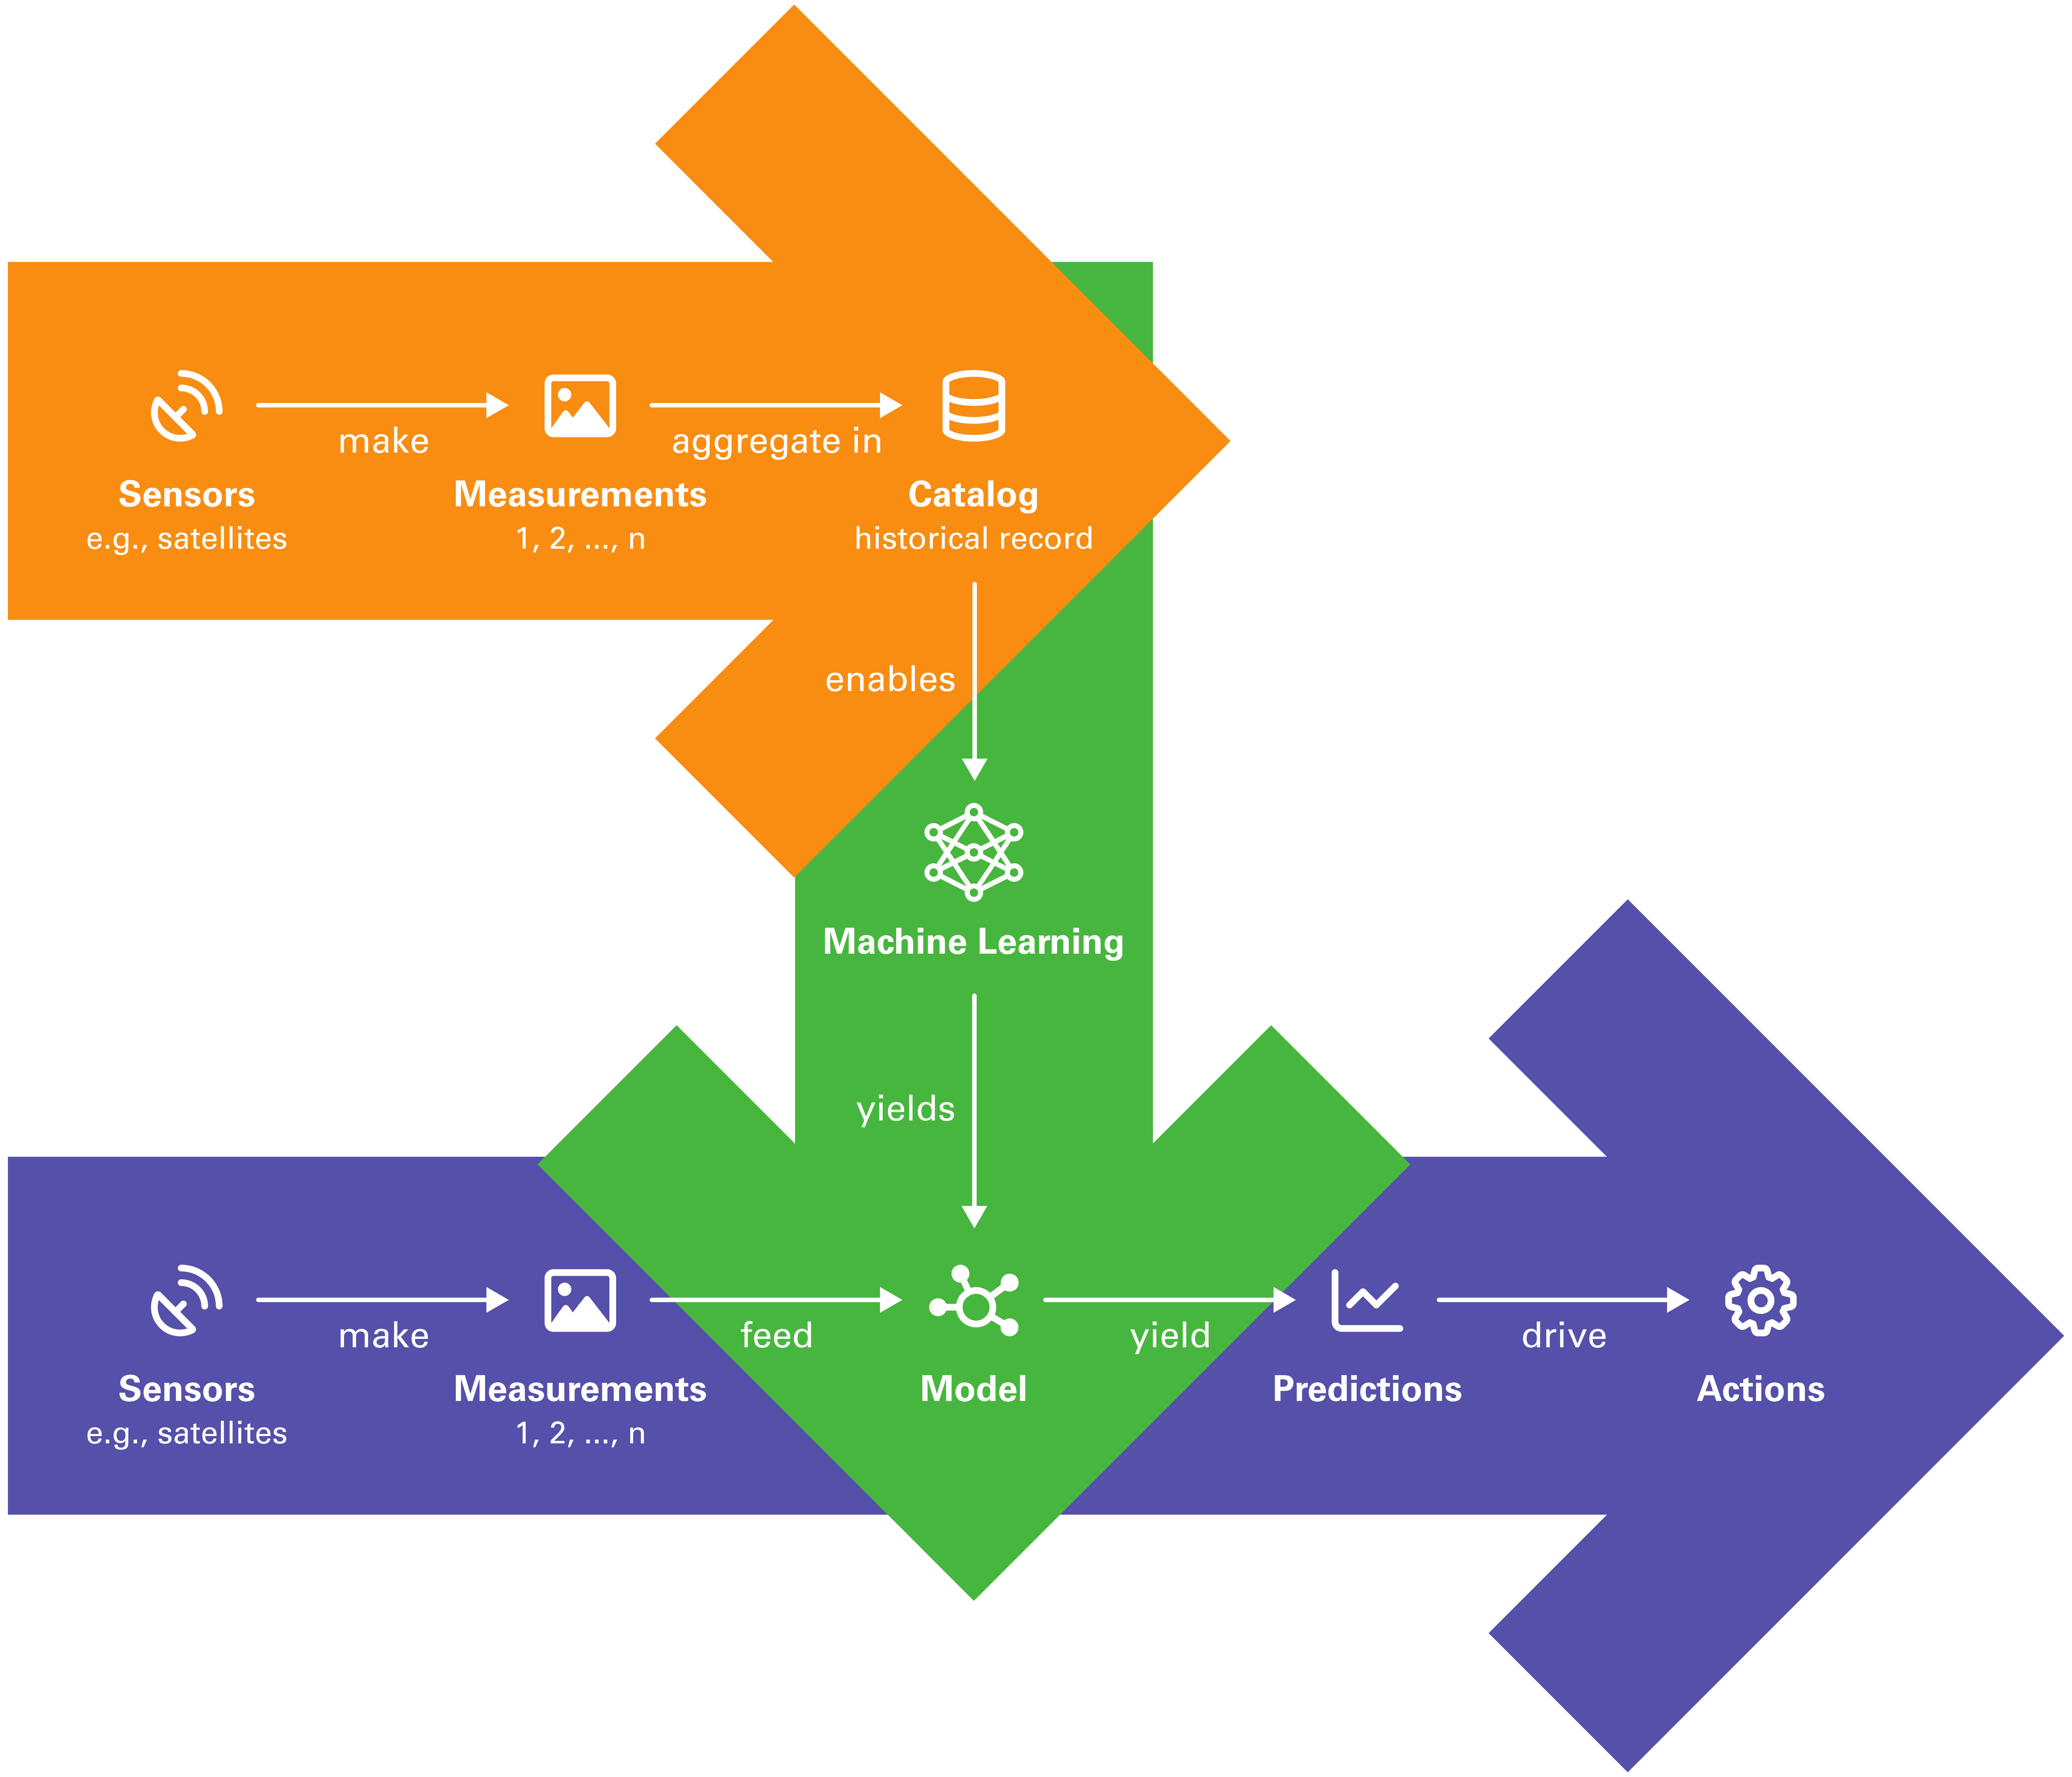

In a cybernetic system, the model contributes predictions to the decision process in the comparator.

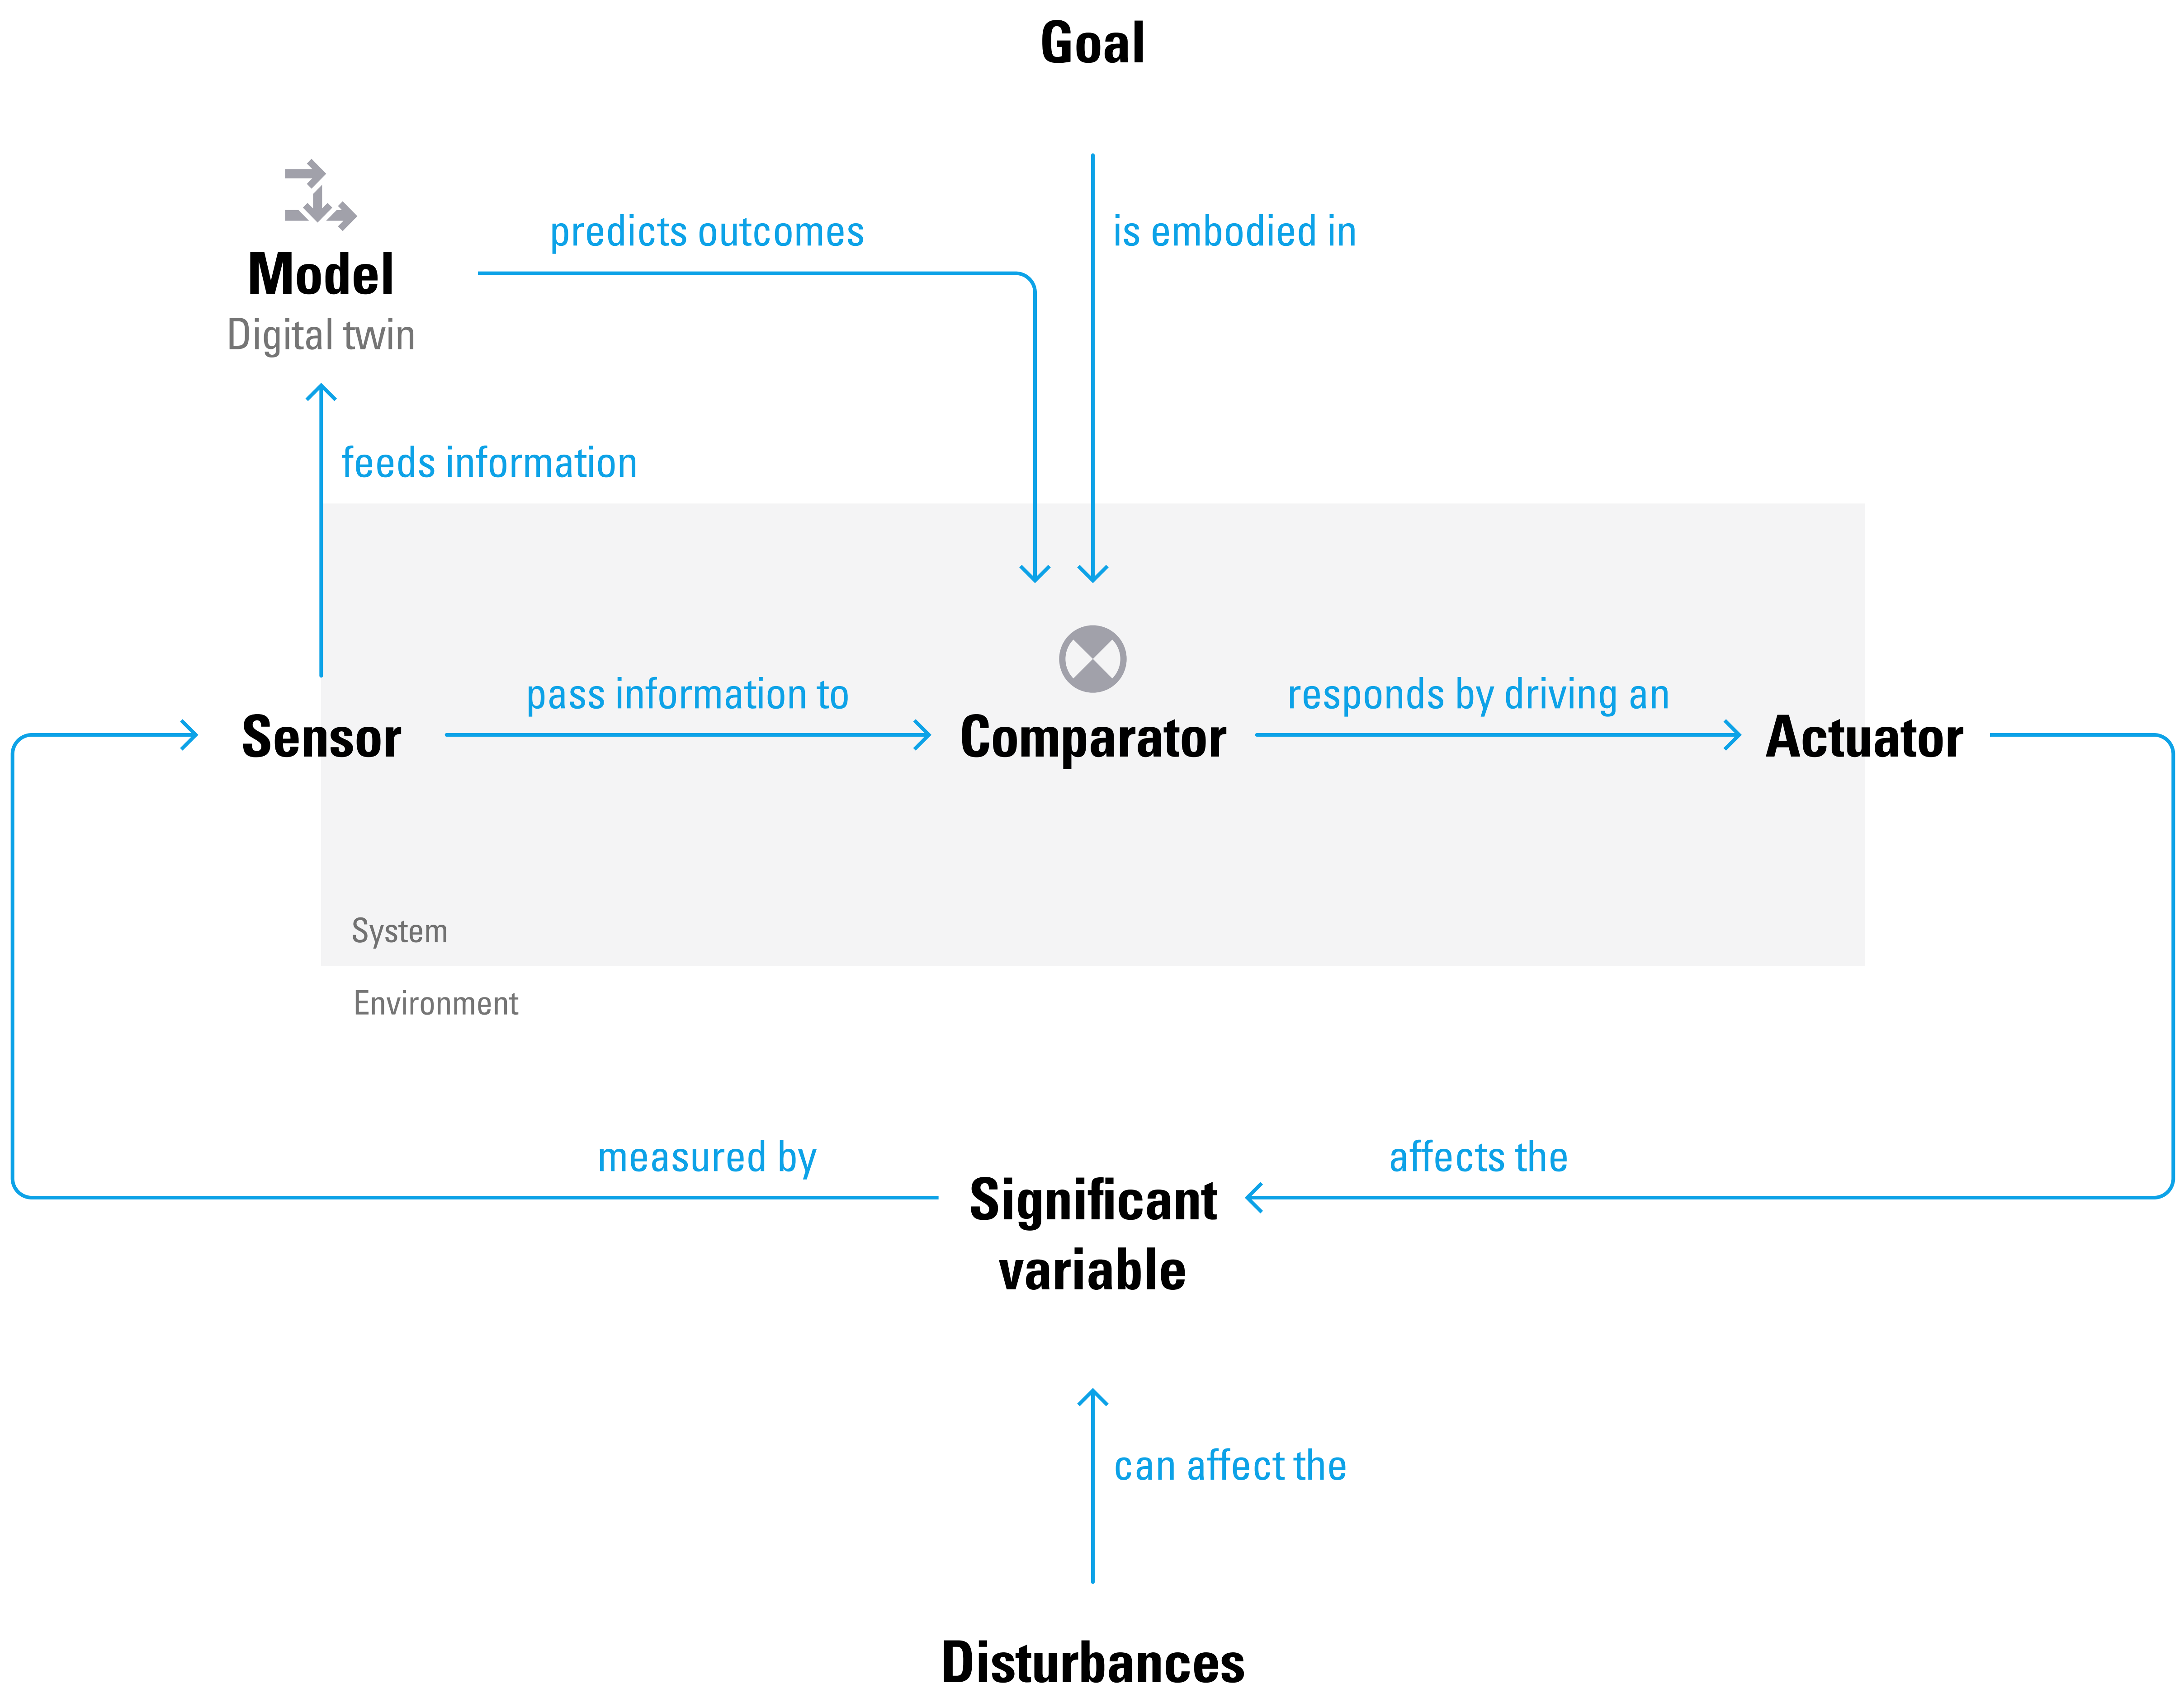

Feeding results back into the model can improve its performance by providing additional data — enabling iteration and 'learning'.

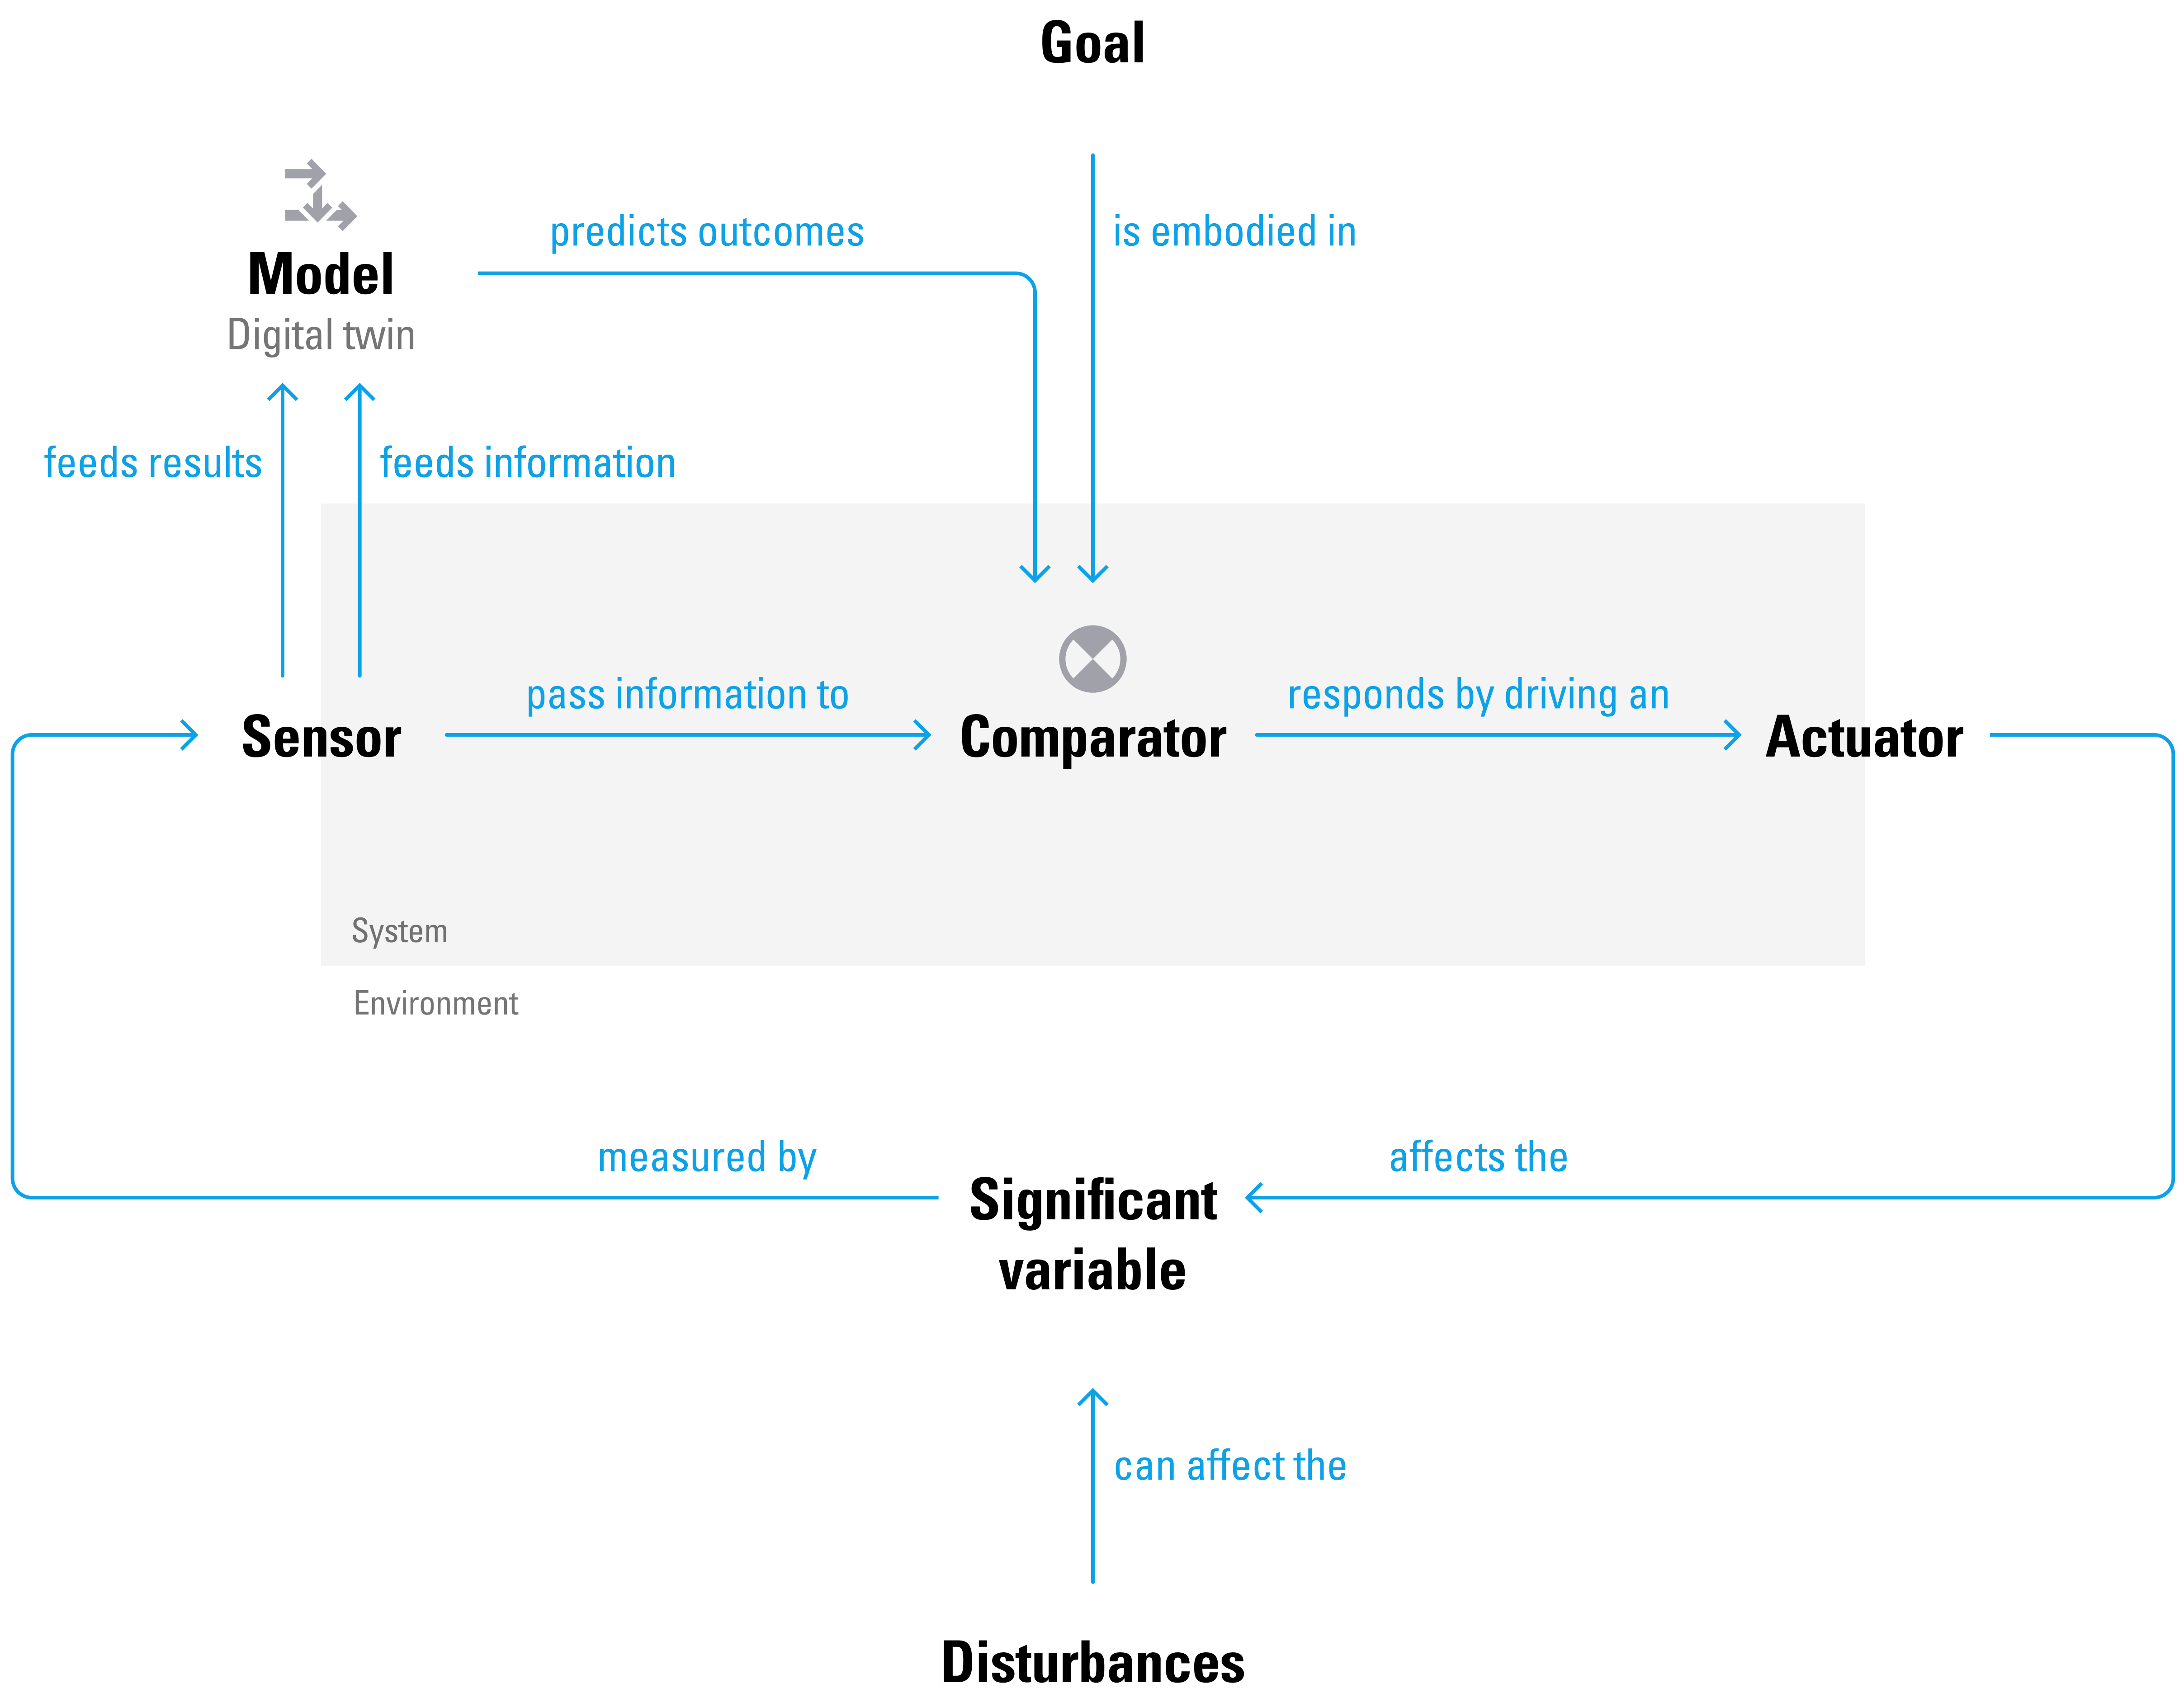

Incorporating feedforward by also measuring disturbances can enable the model to identify patterns in how disturbances affect the significant variable.

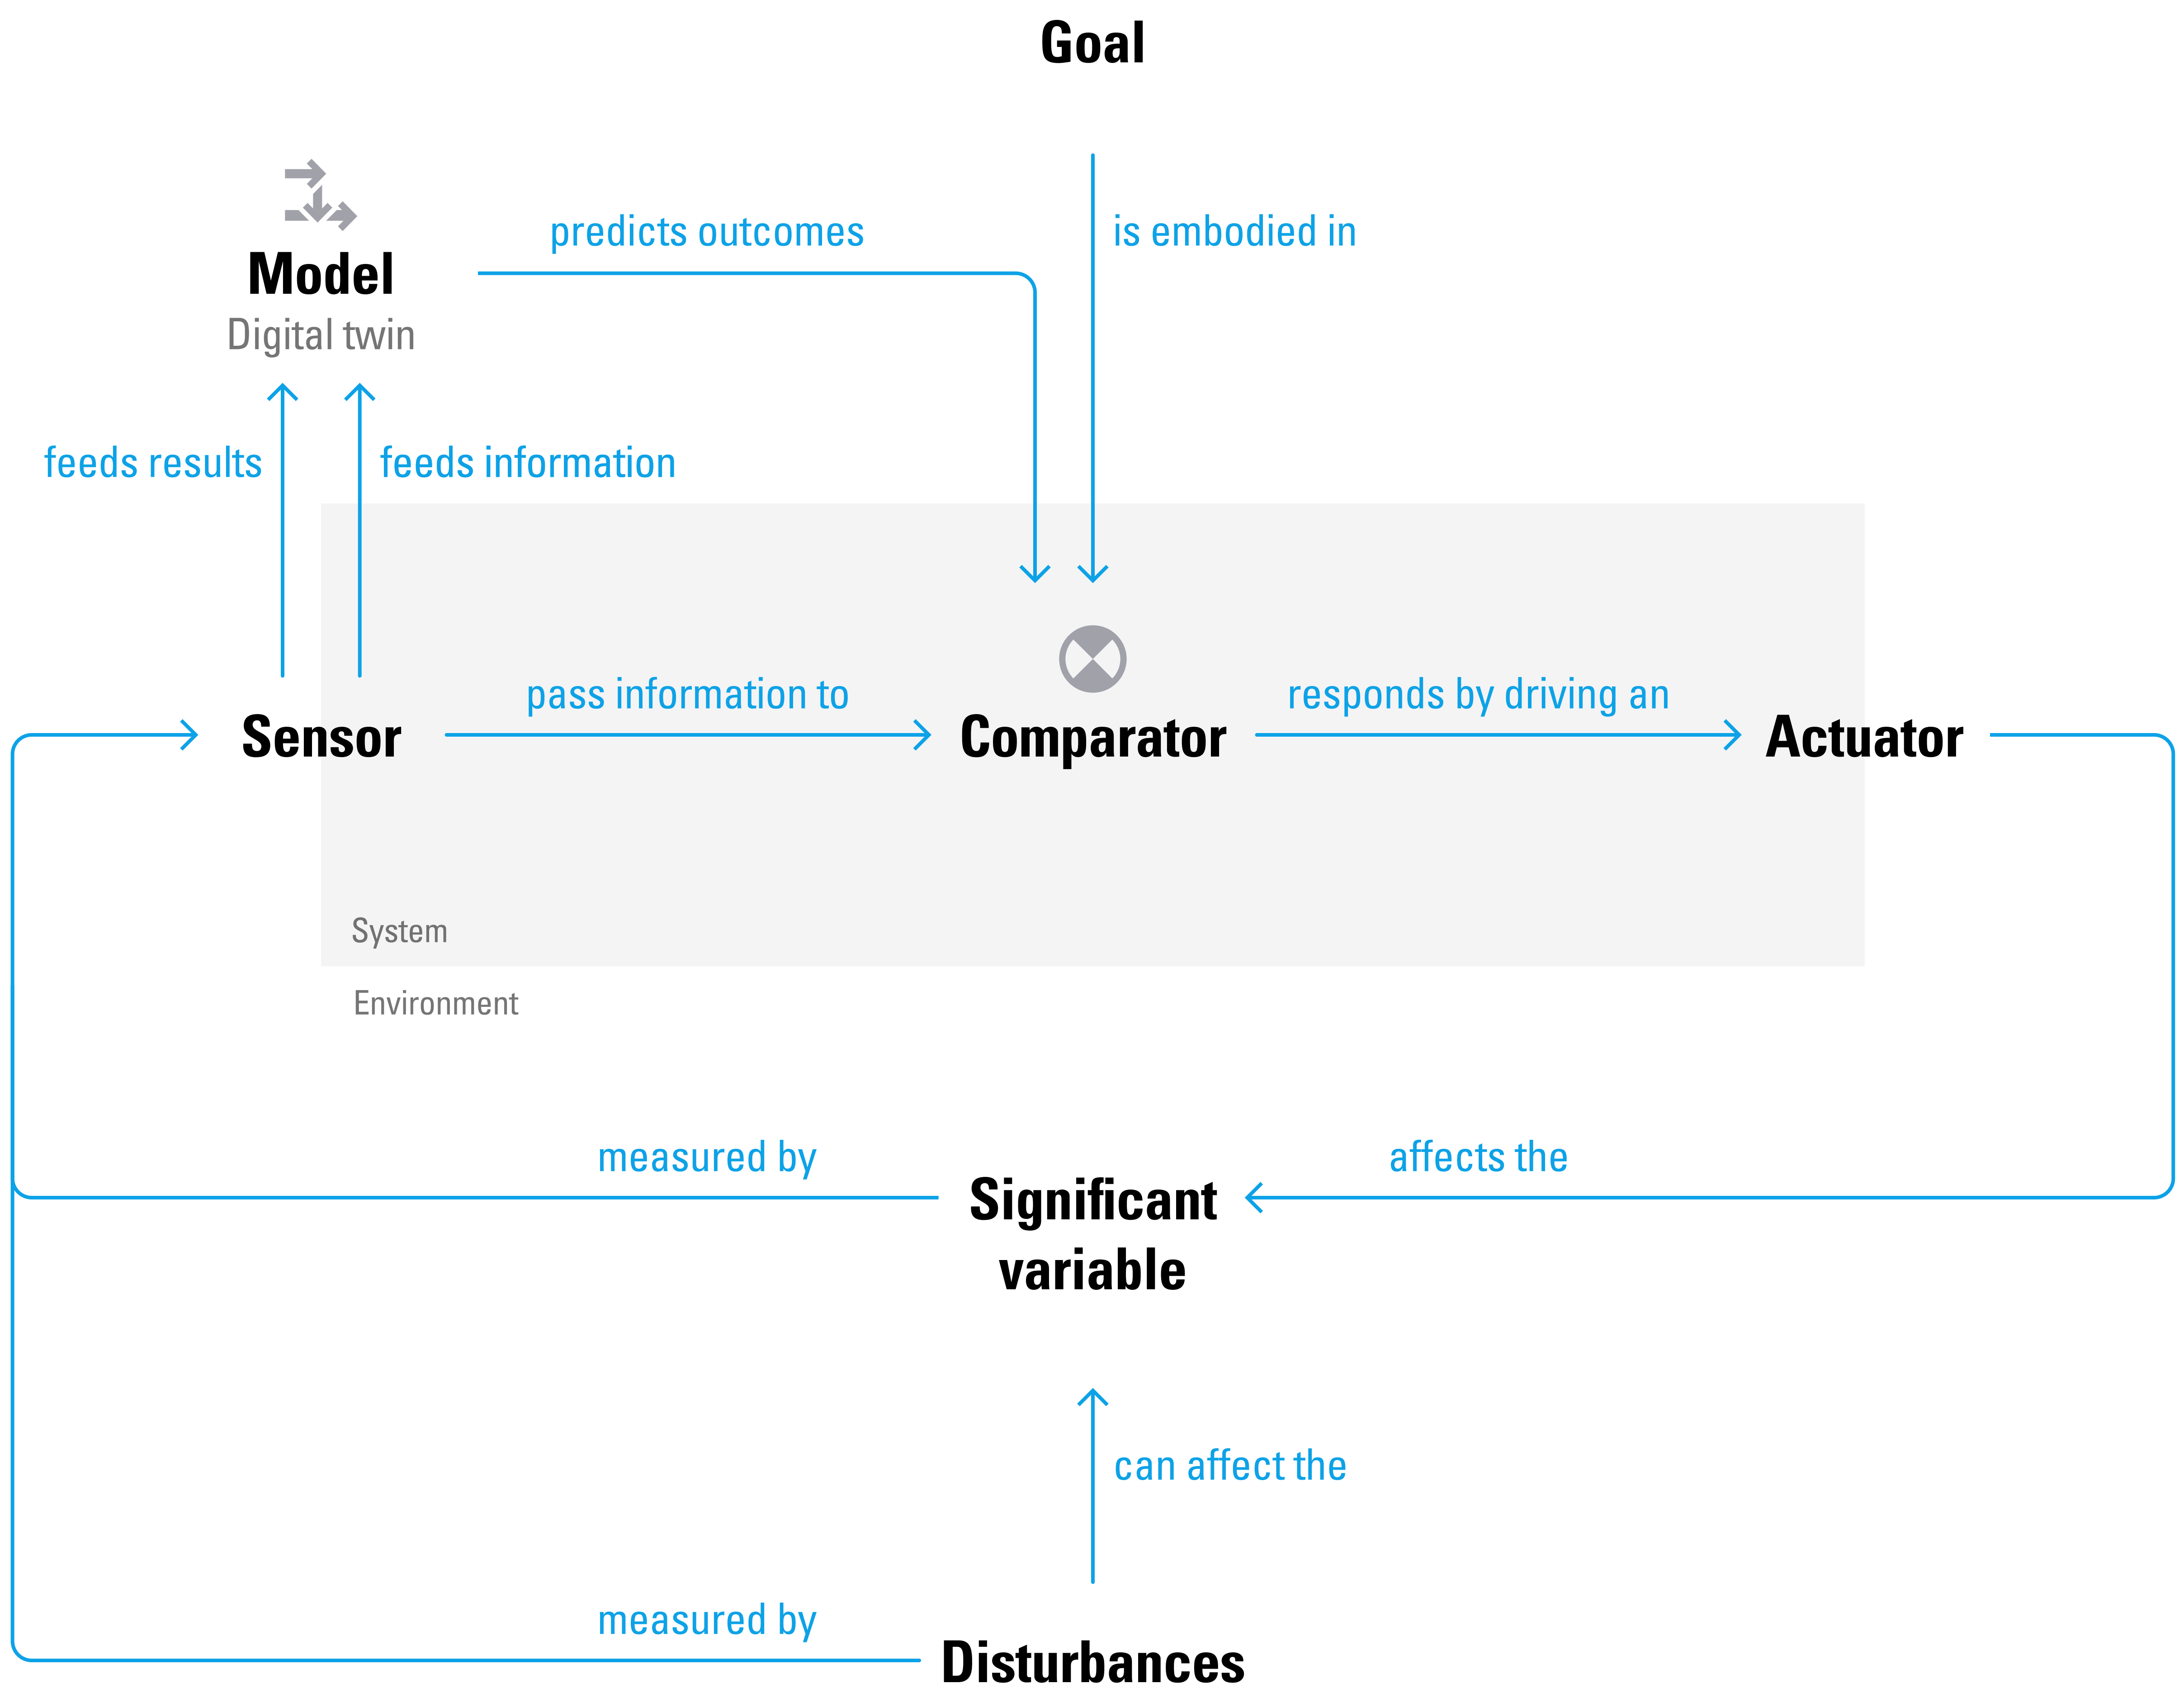

Second-order feedback in learning agents

A thermostat is a classic example of a cybernetic control loop.

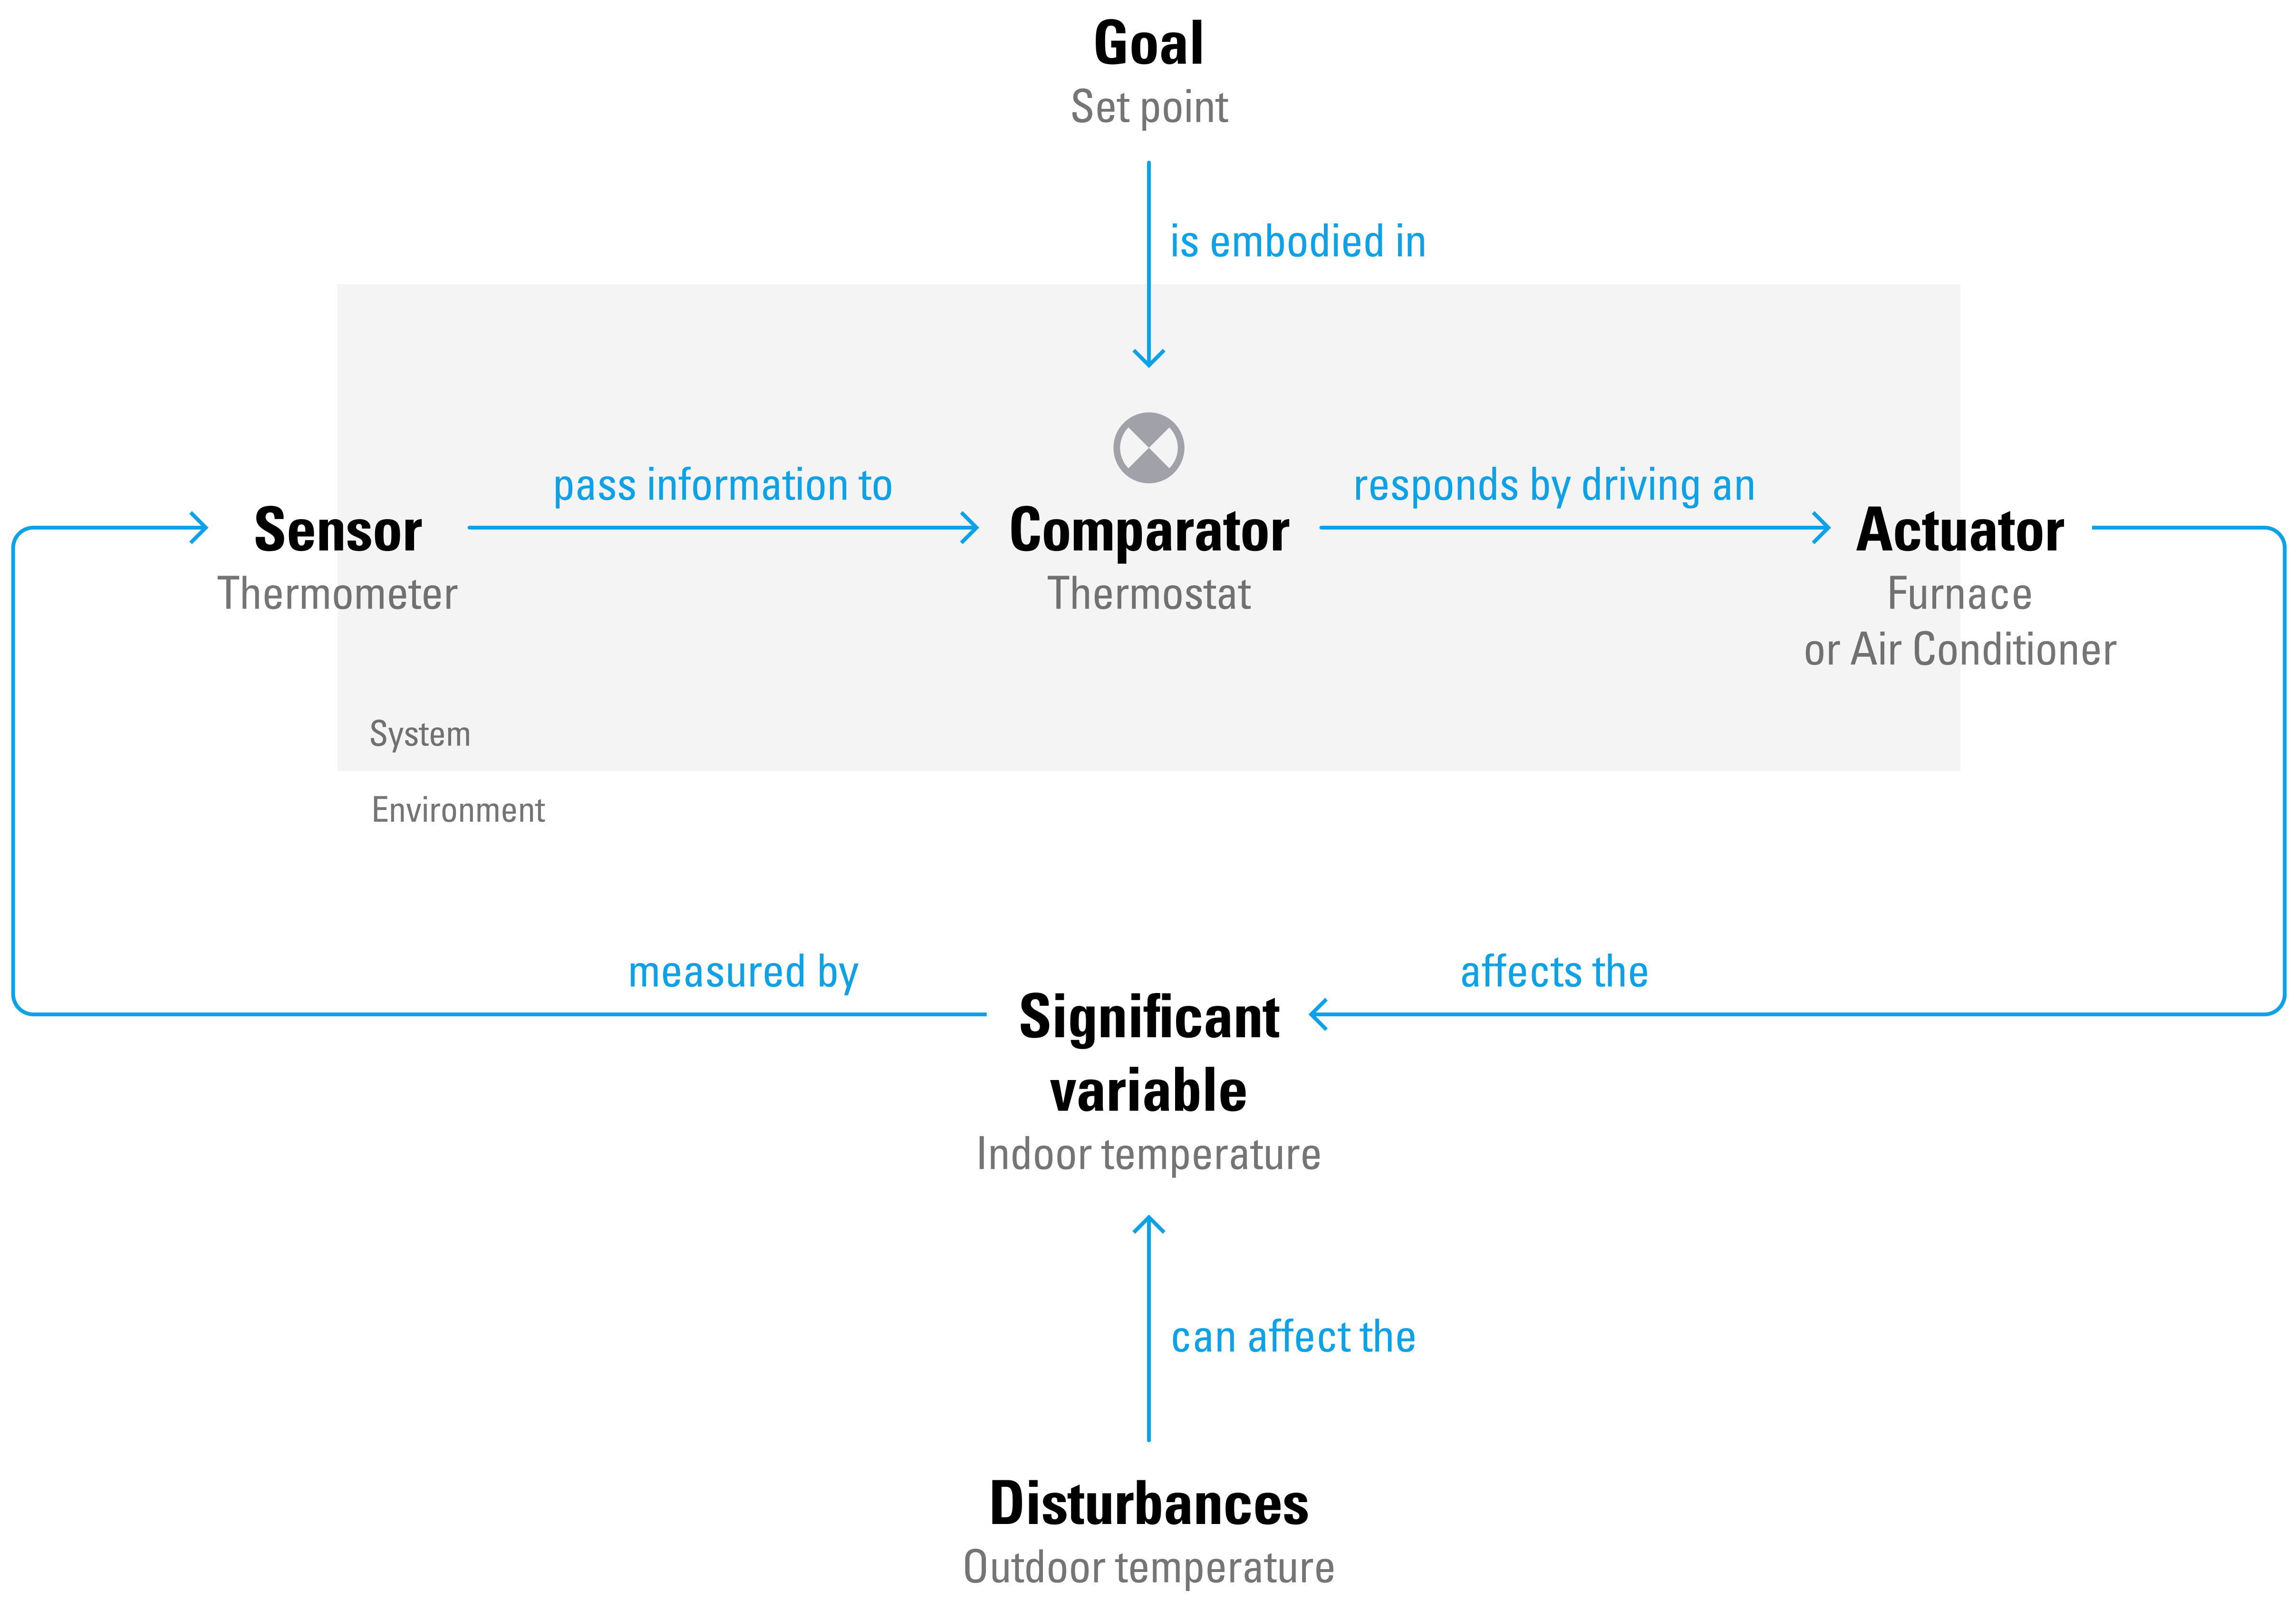

When depicted as a single-loop or first-order system, the goal or set point is assumed.

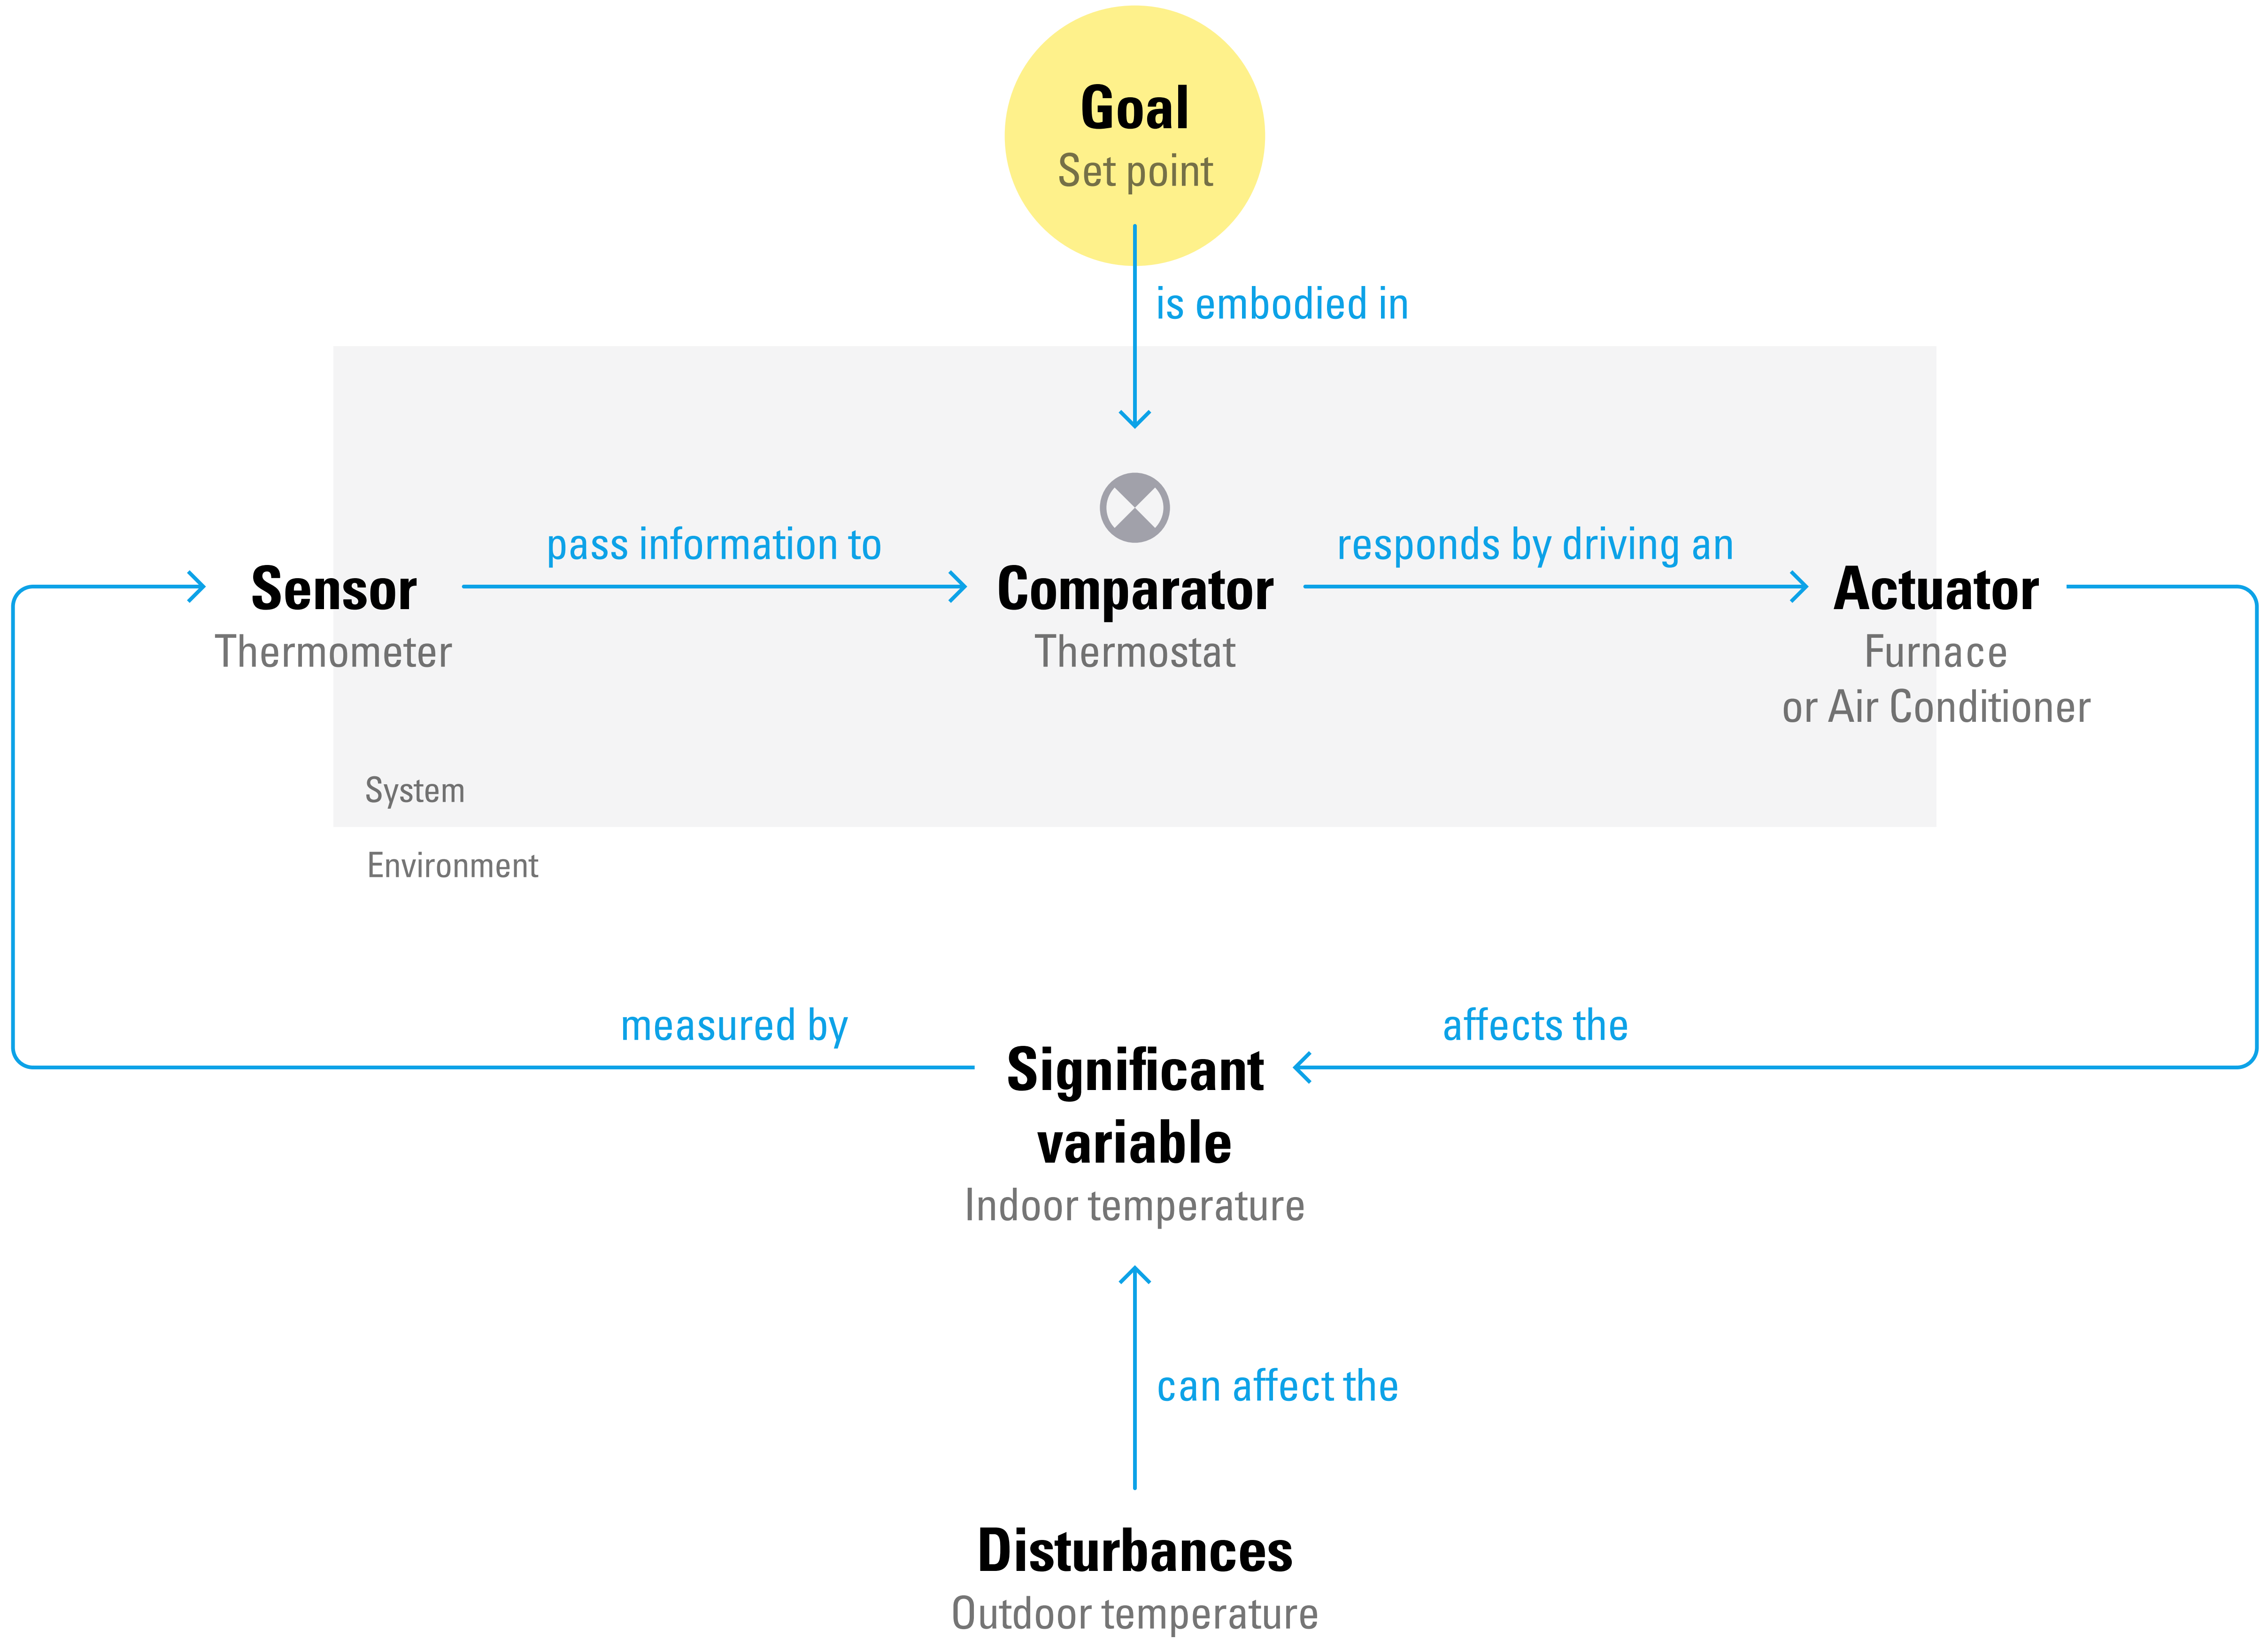

In practice, the situation is typically observed by a human 'user'.

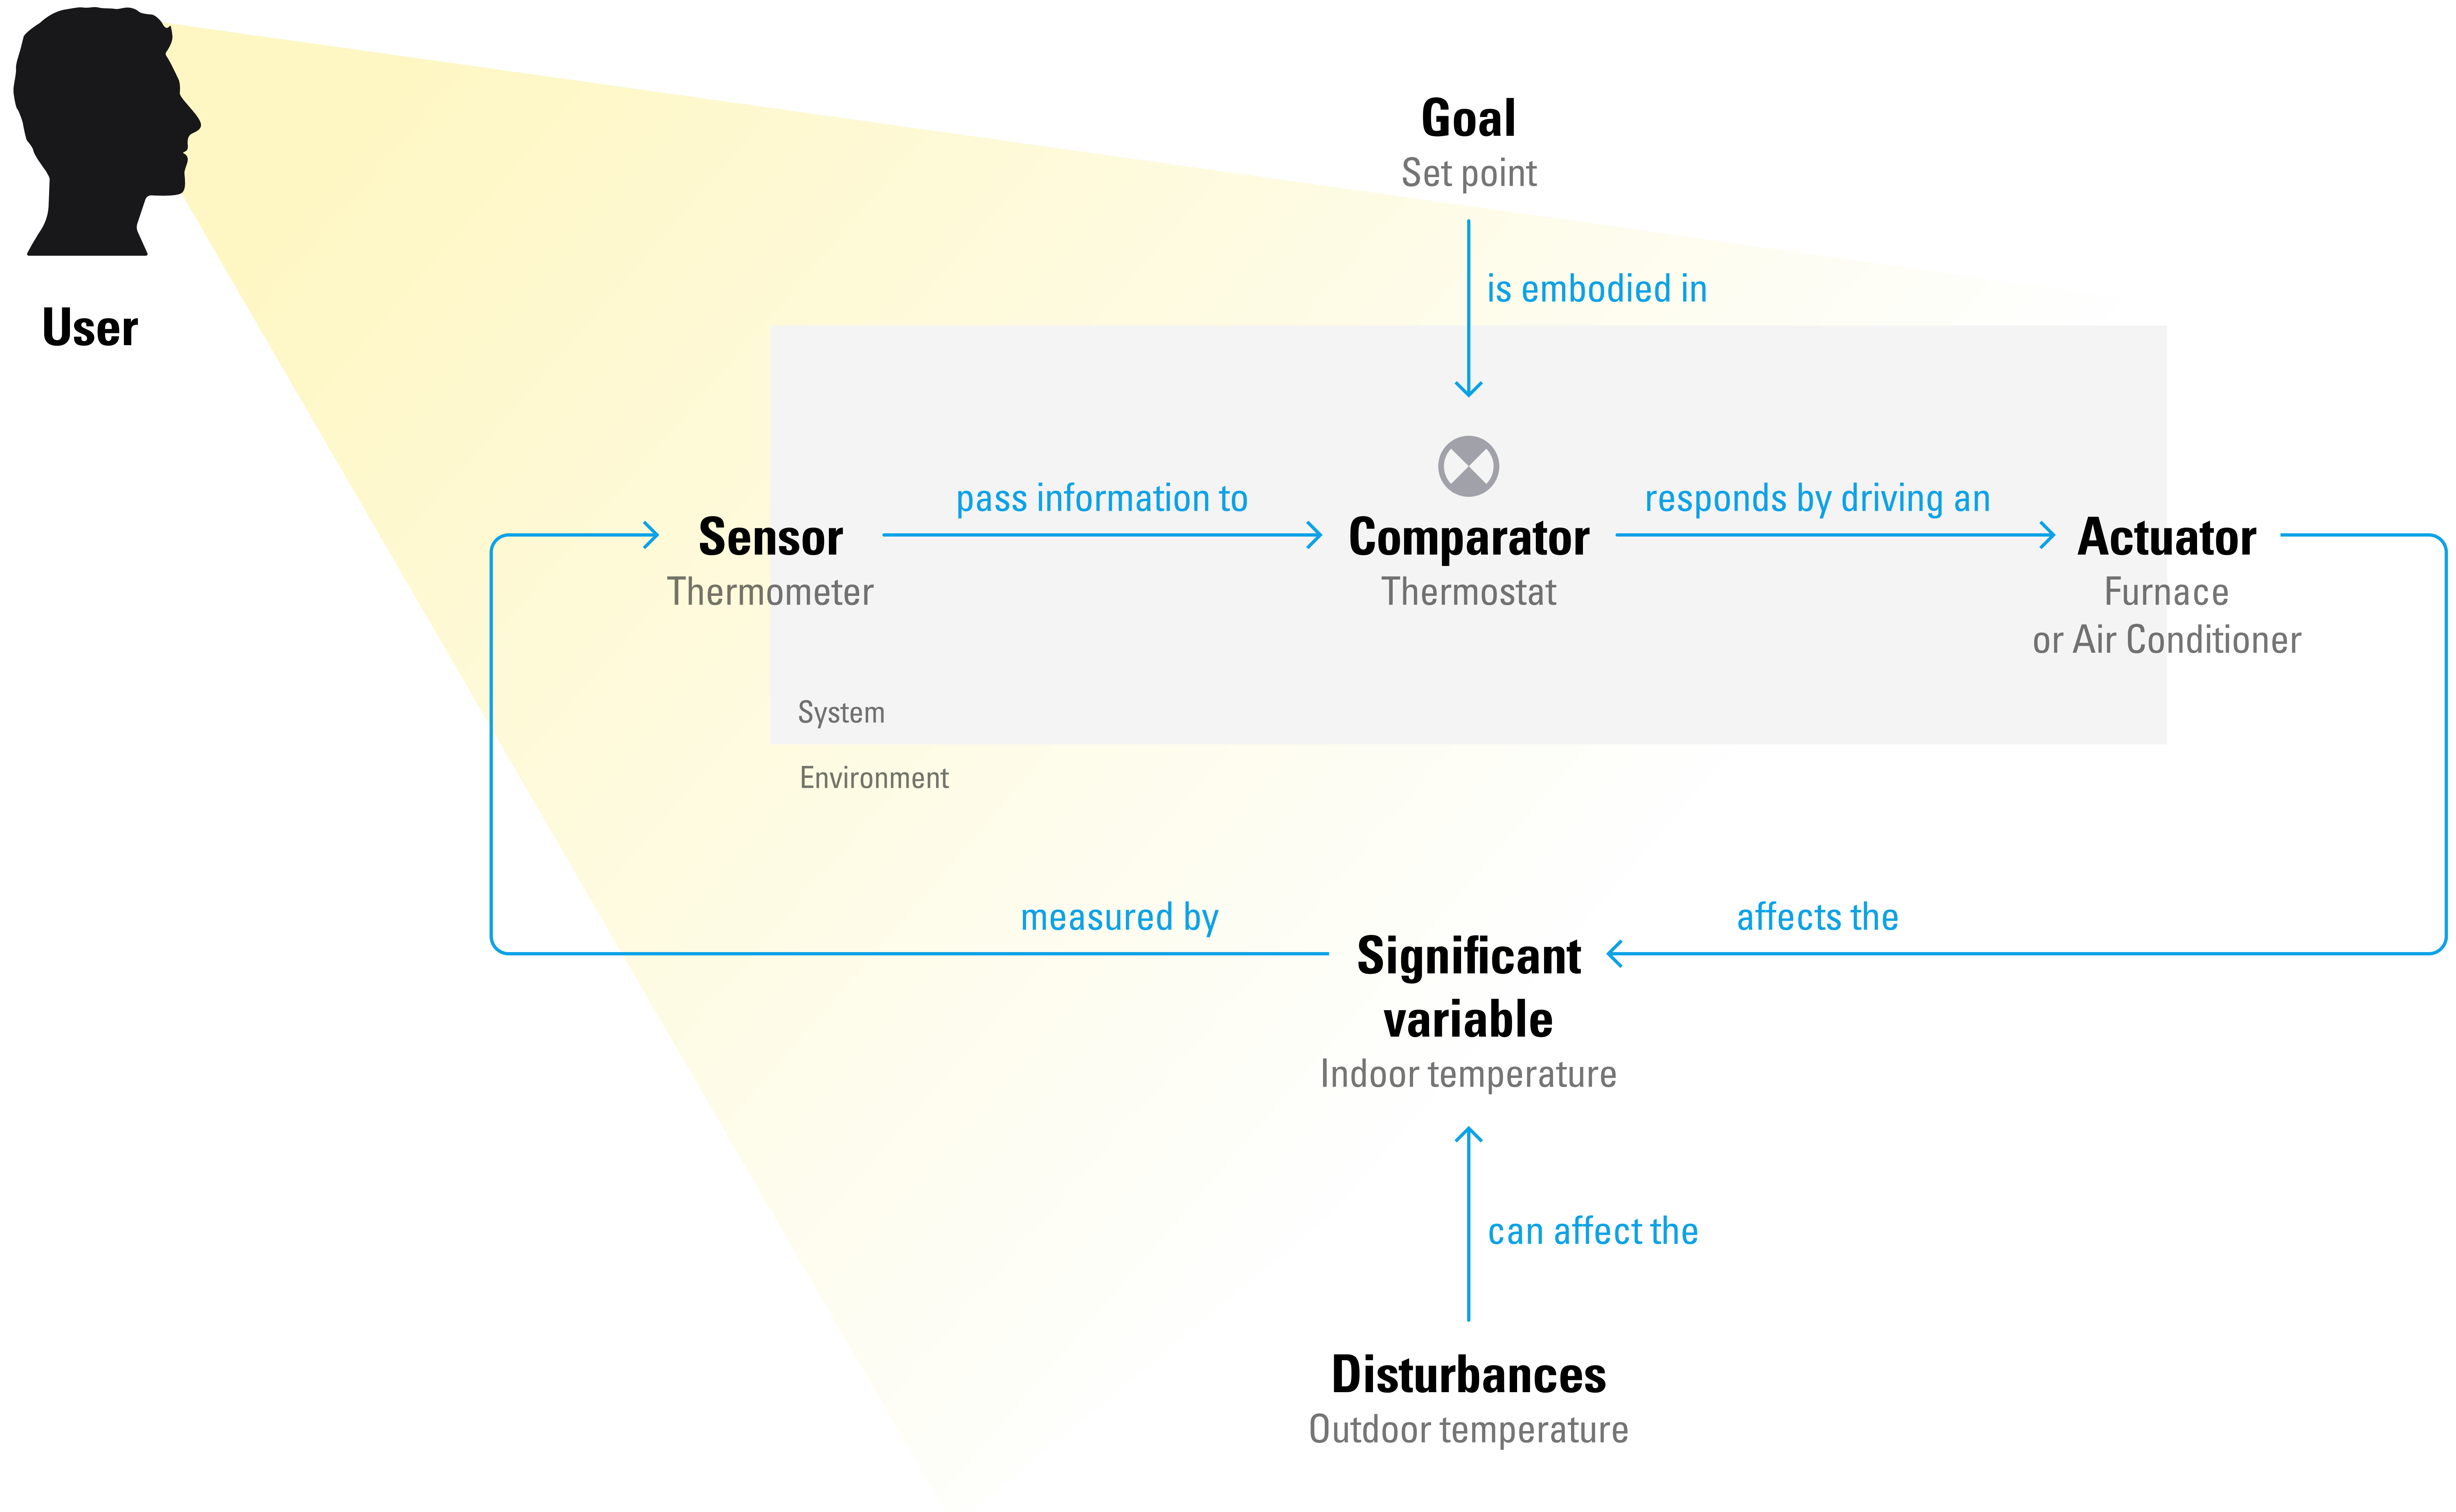

The user may set the thermostat’s set point.

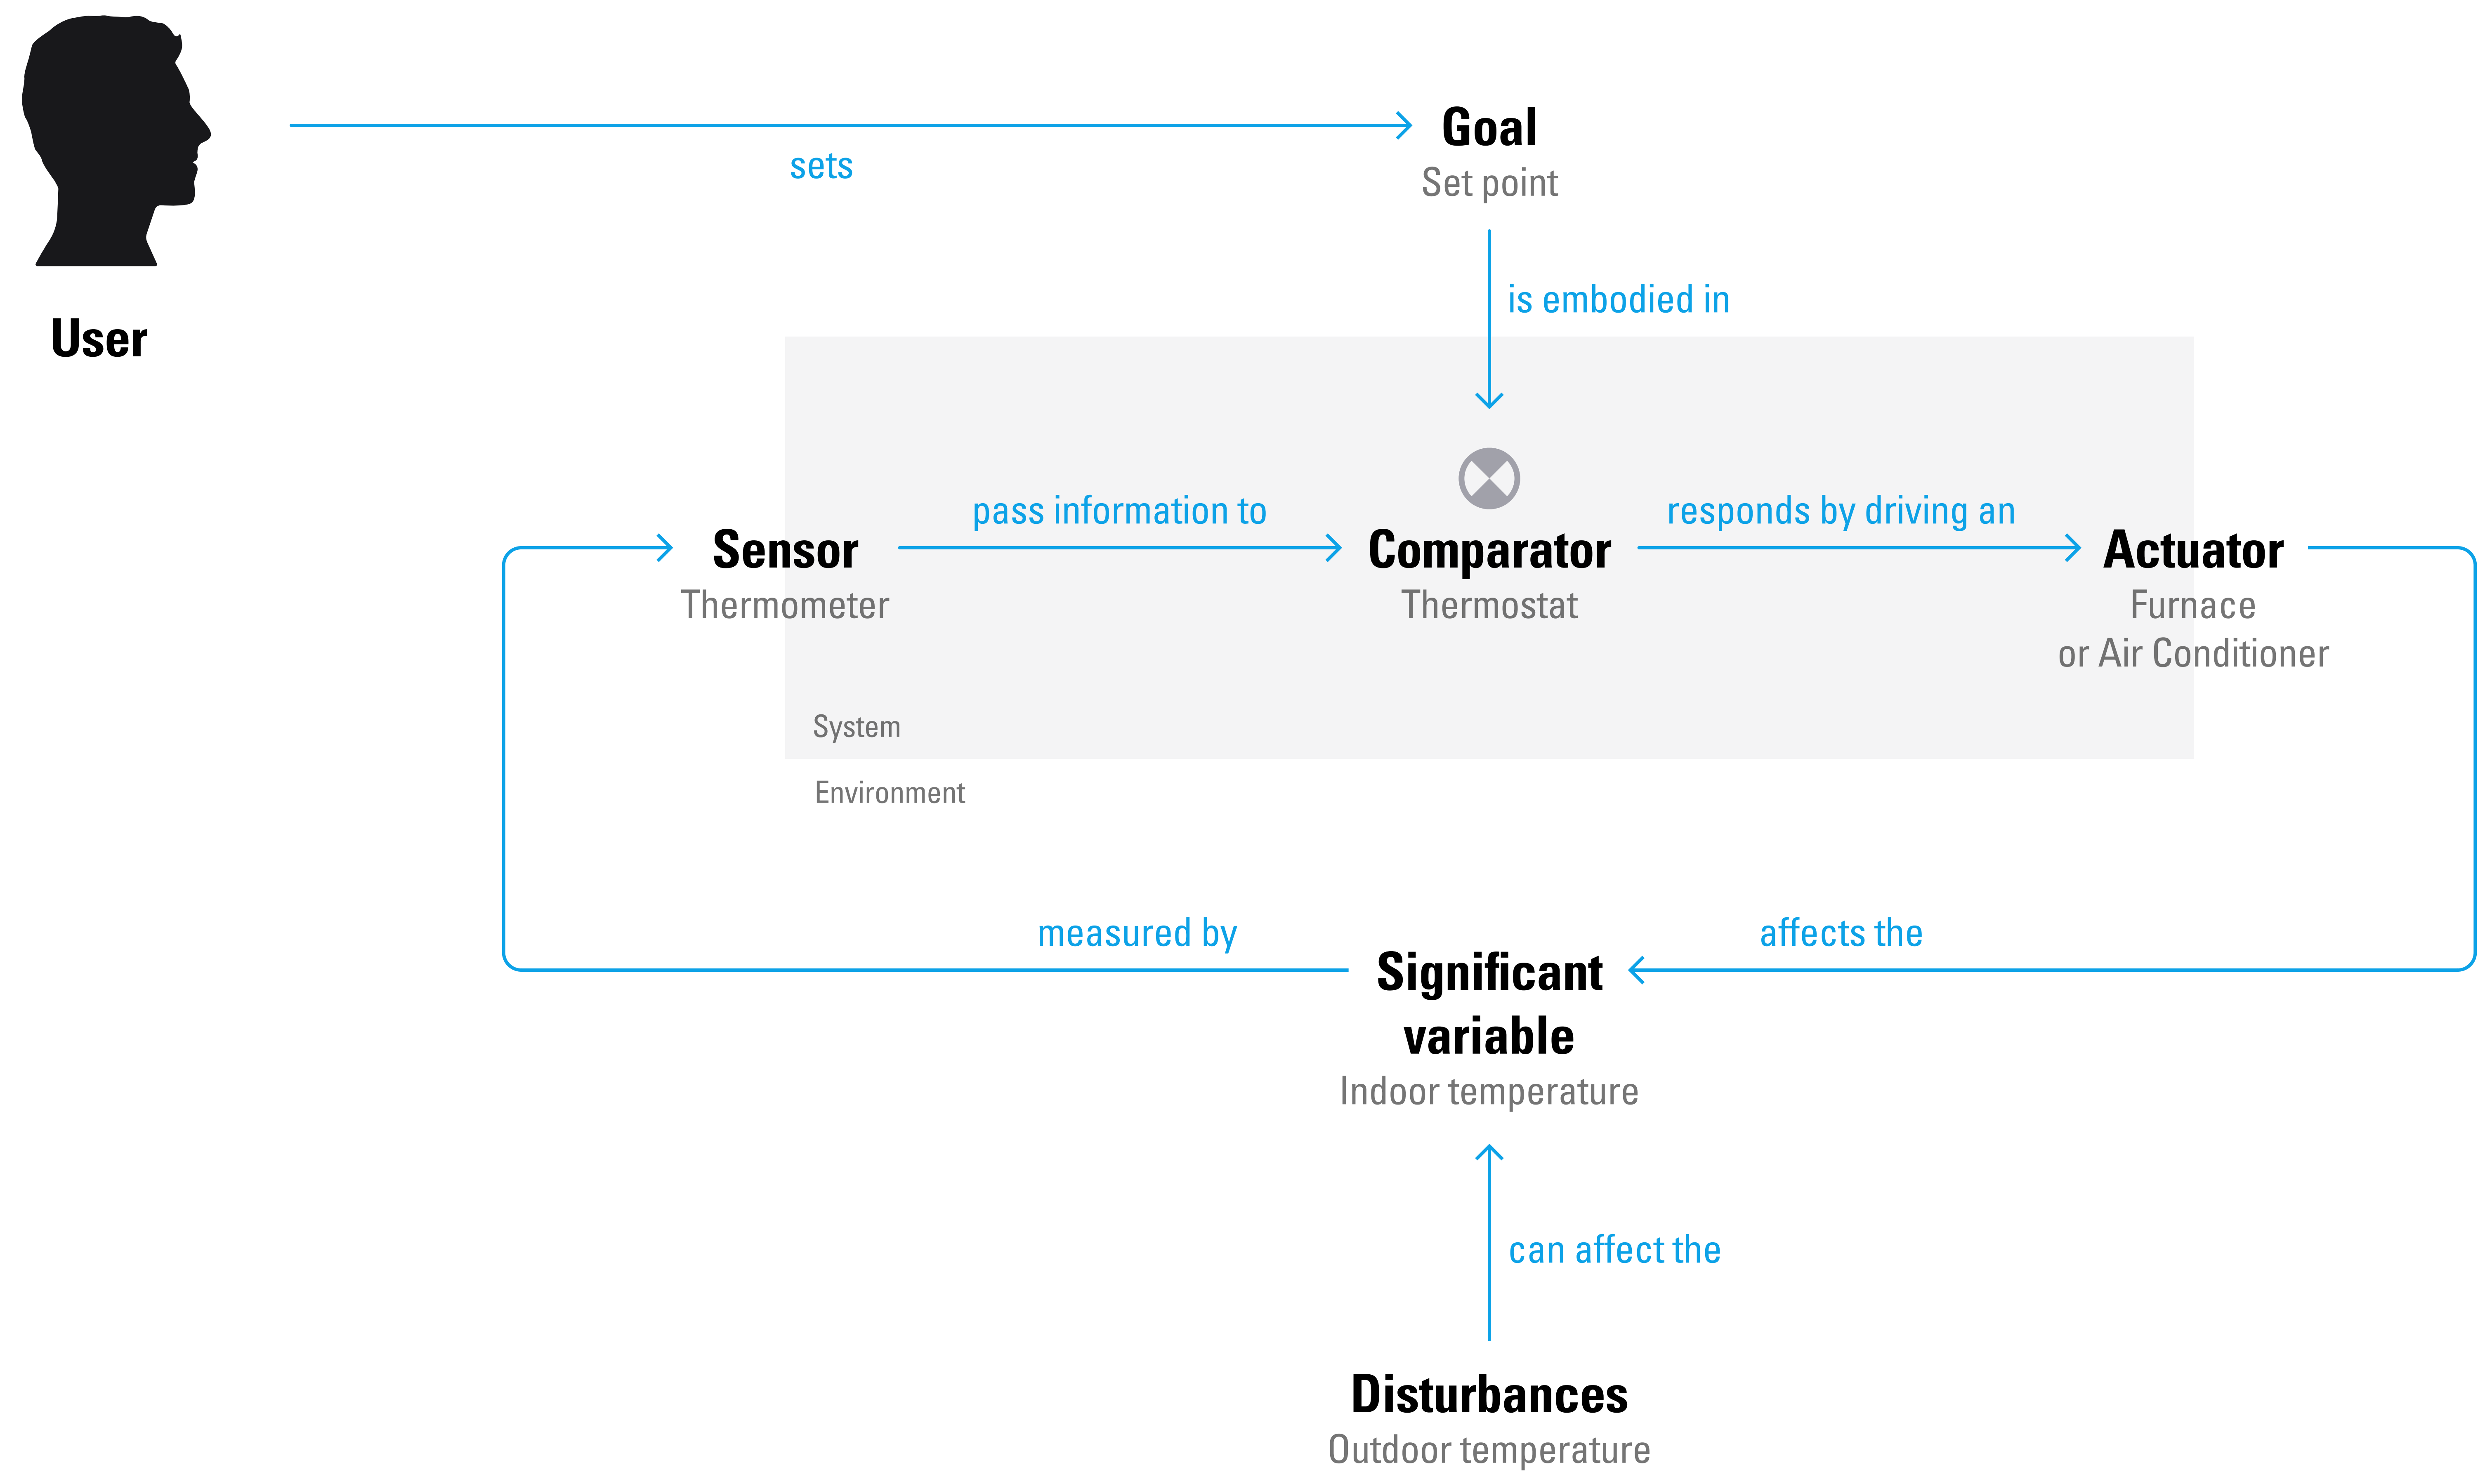

In addition, the user may monitor the significant variable.

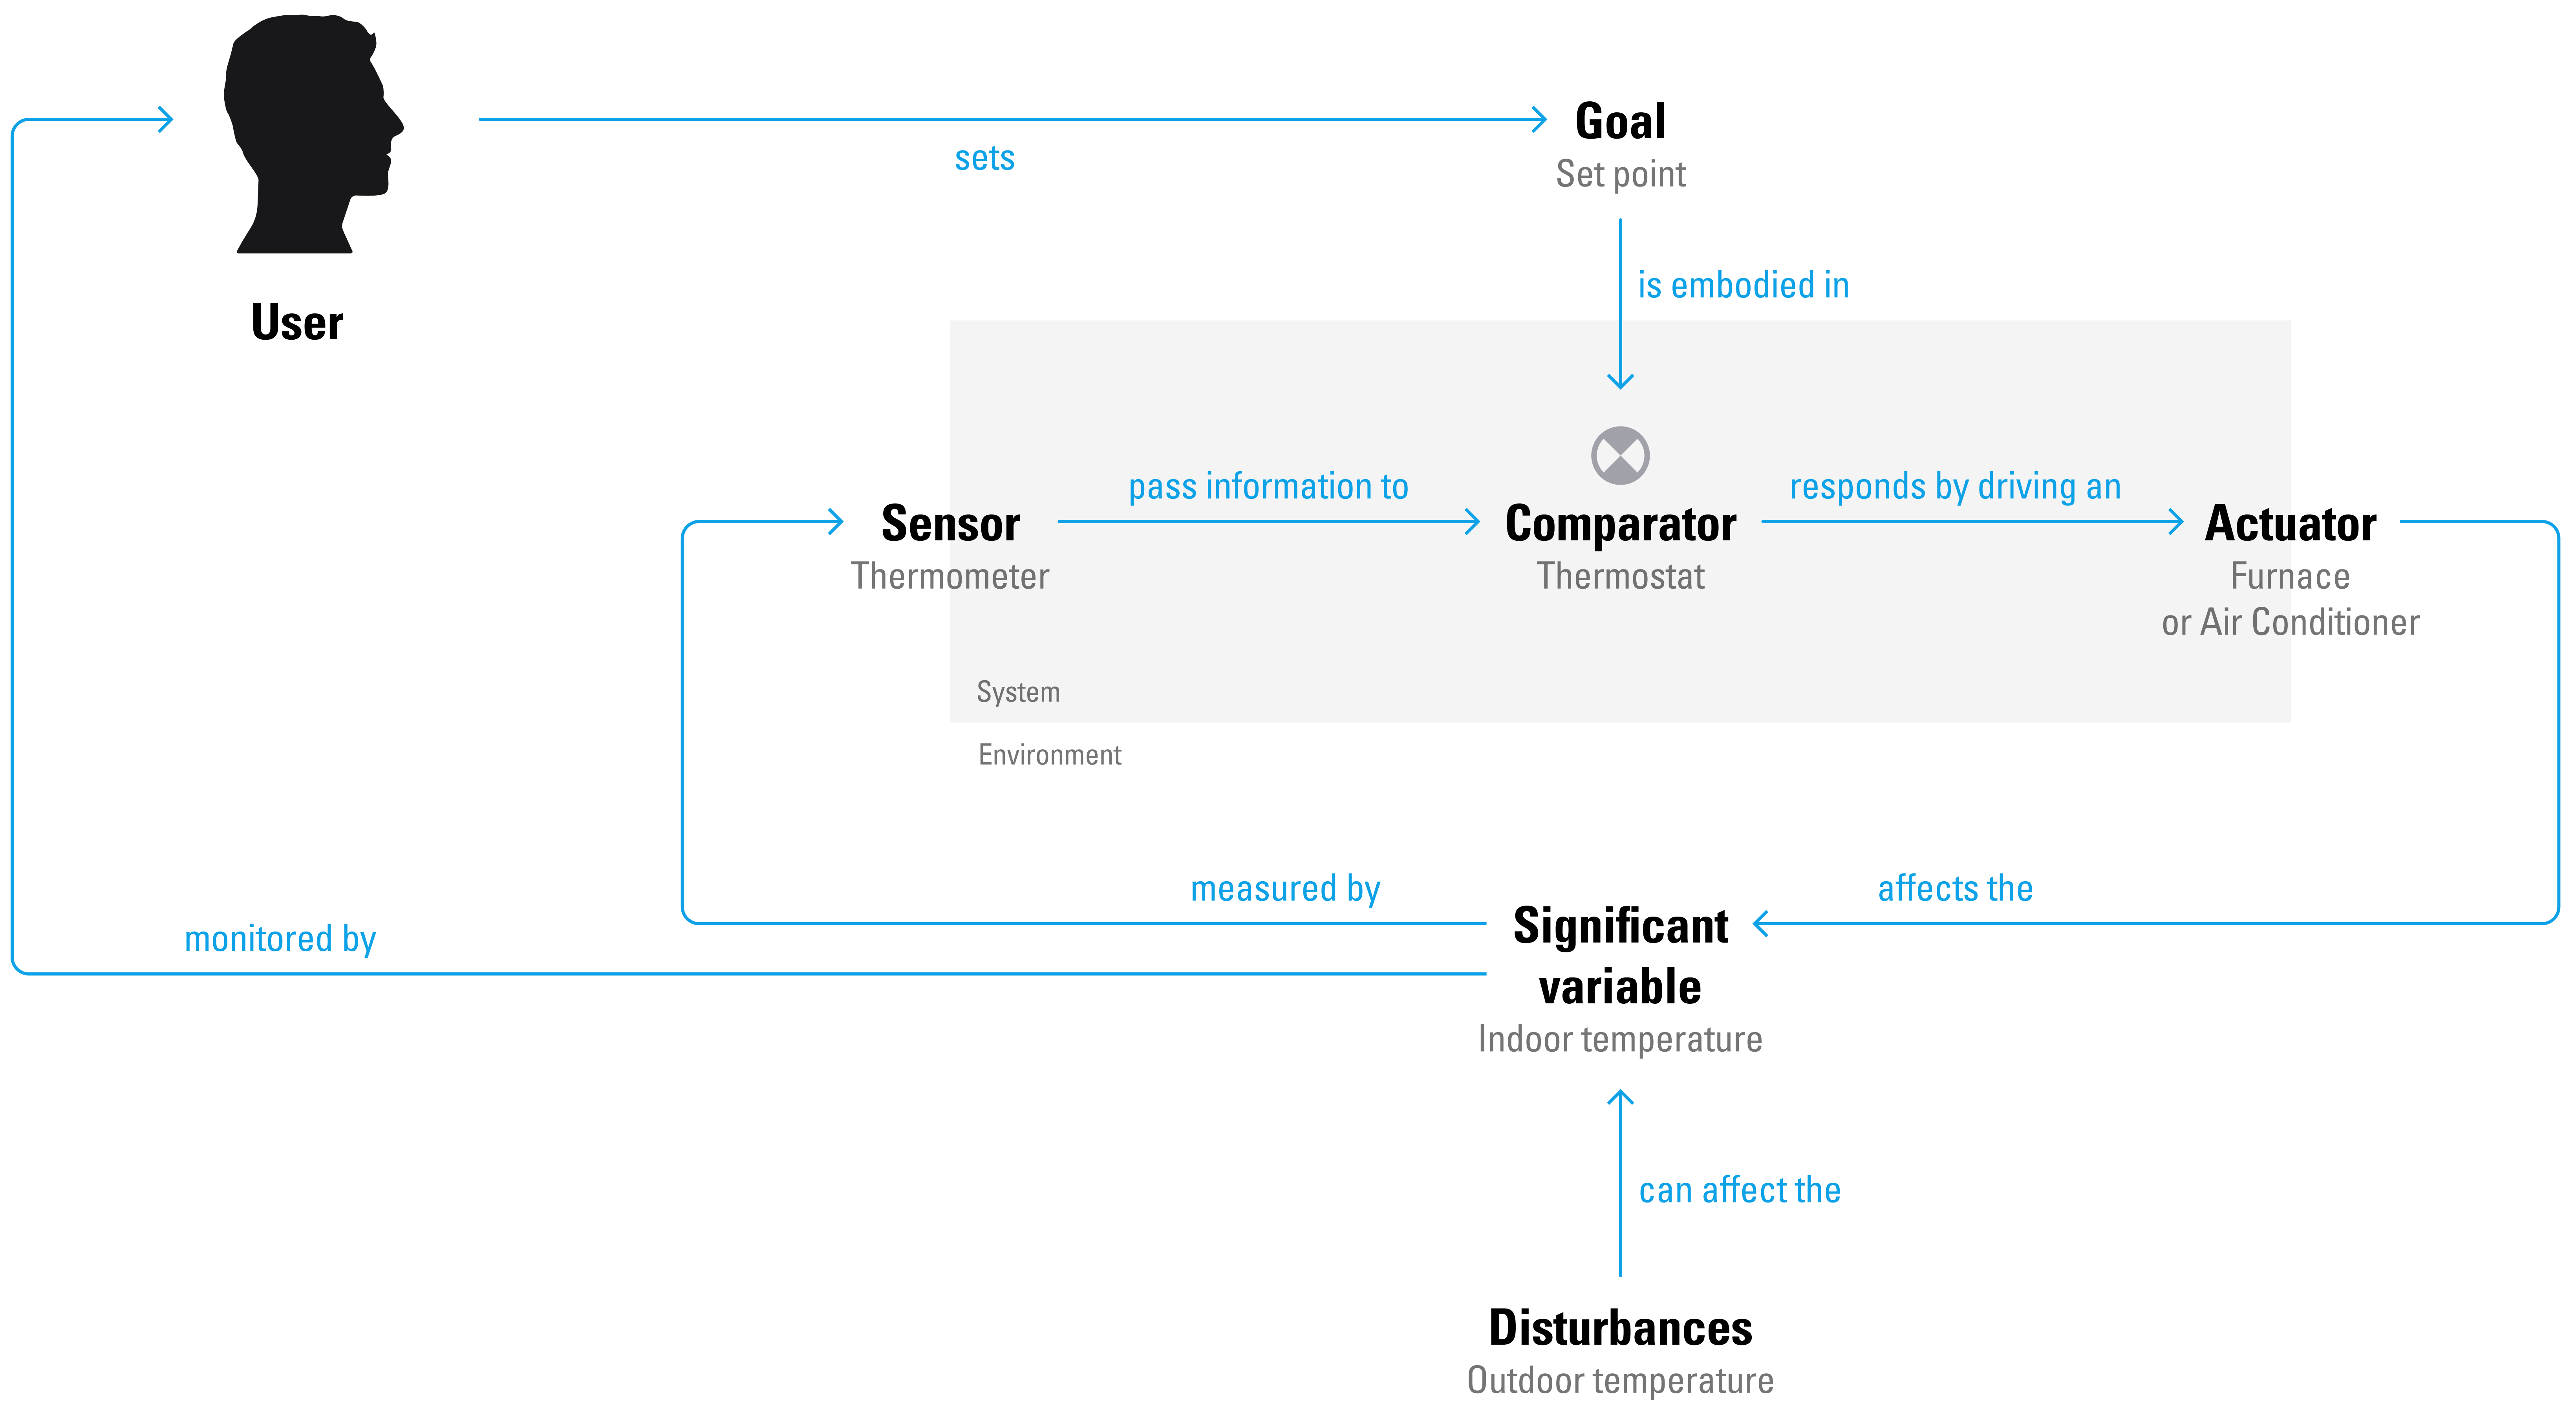

A second feedback loop appears above the first — creating a double-loop or second-order system.

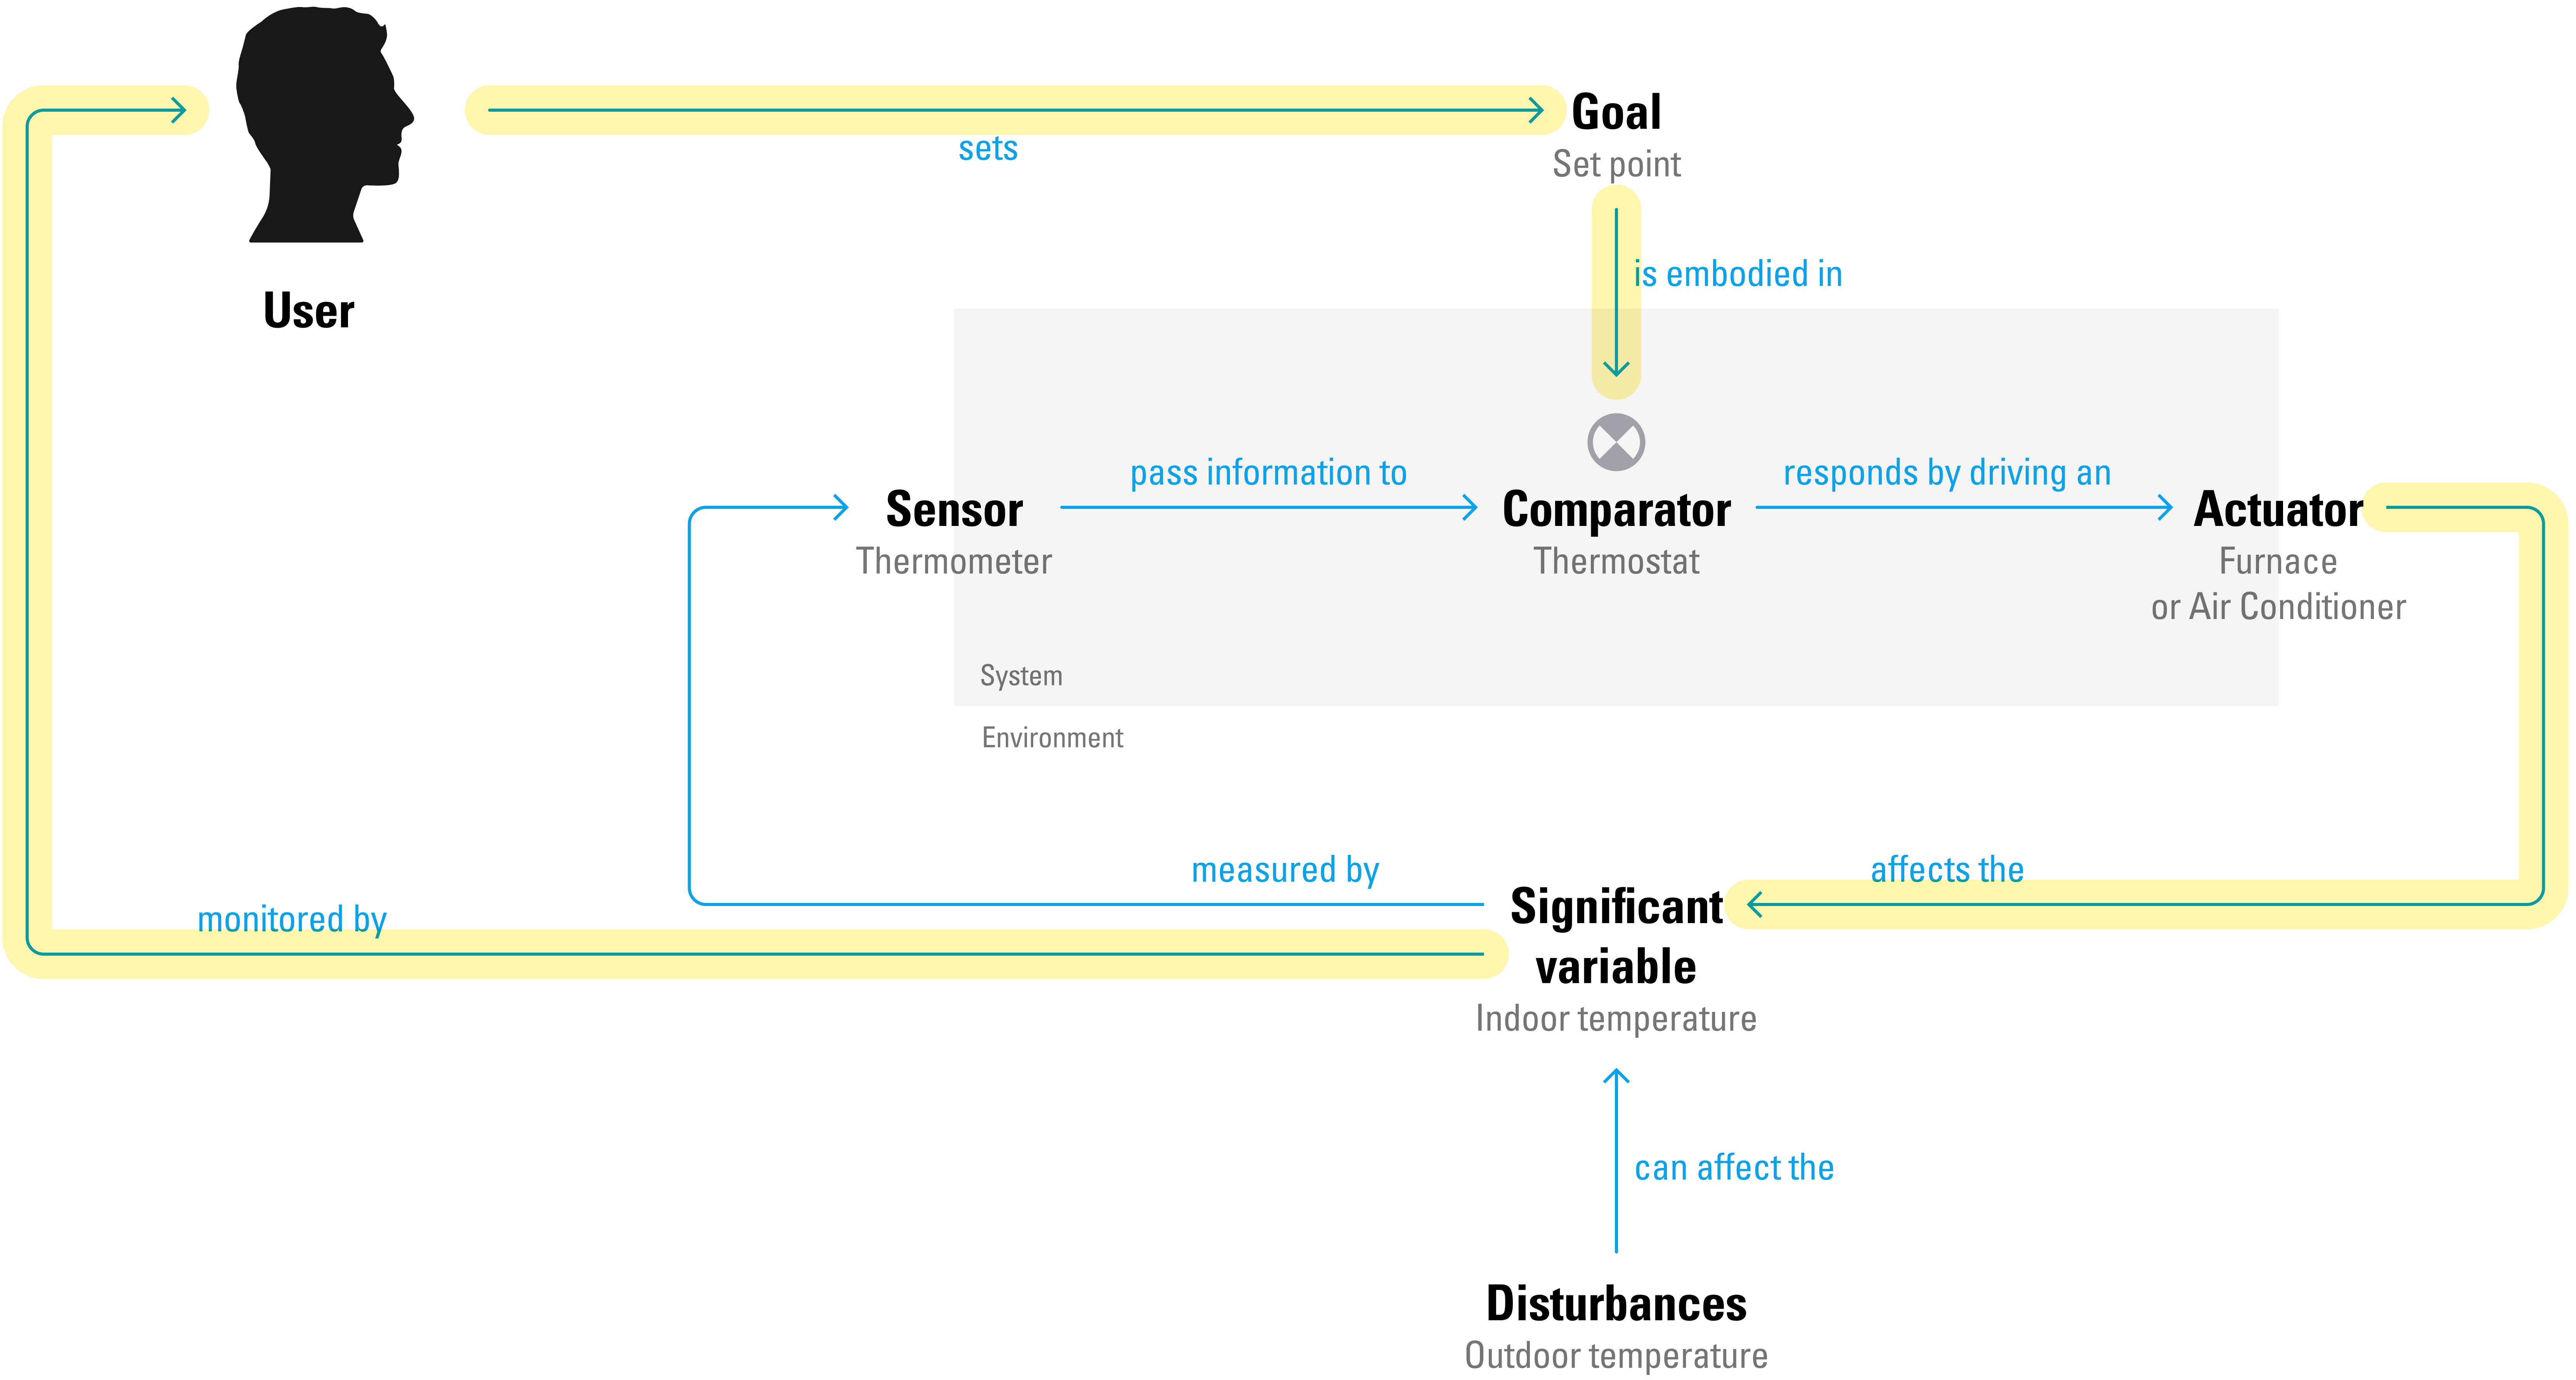

In our thermostat example, the user has the goal of ‘comfort‘ — let’s they set the thermostat to maintain 68ºF.

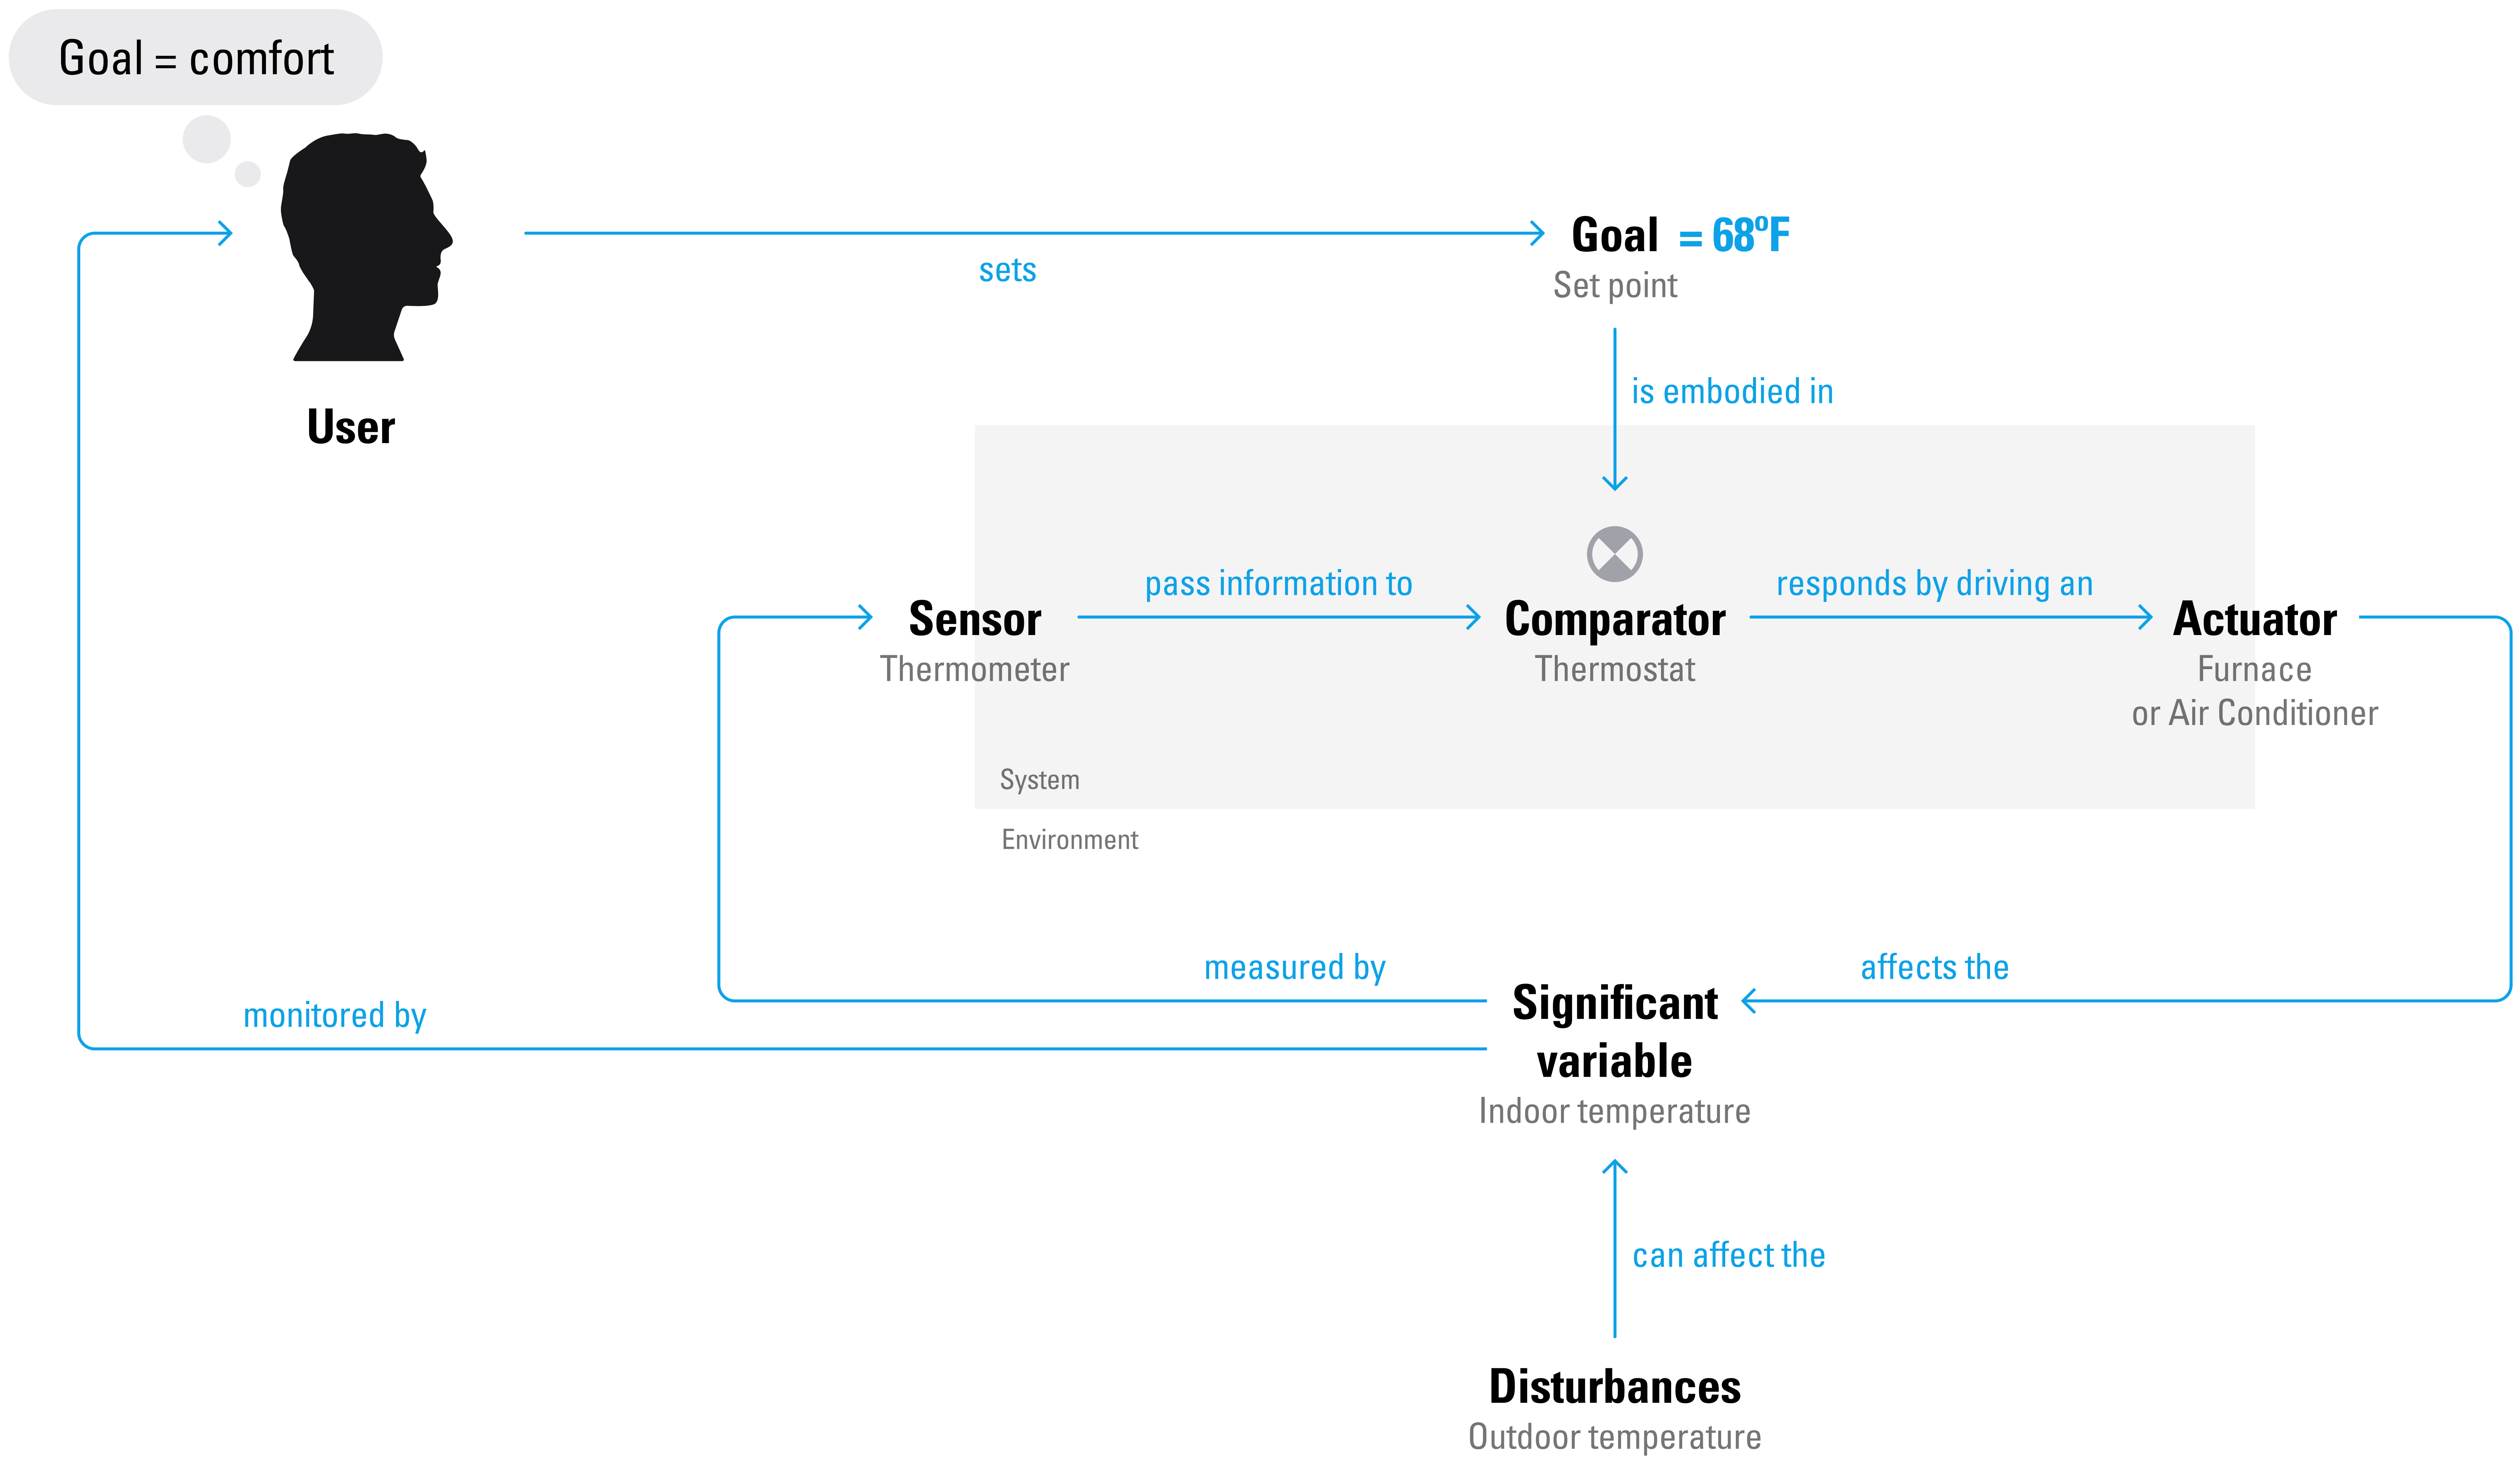

Later, they decide 68ºF is too cold — not comfortable.
In response, they may change the set point of the thermostat.

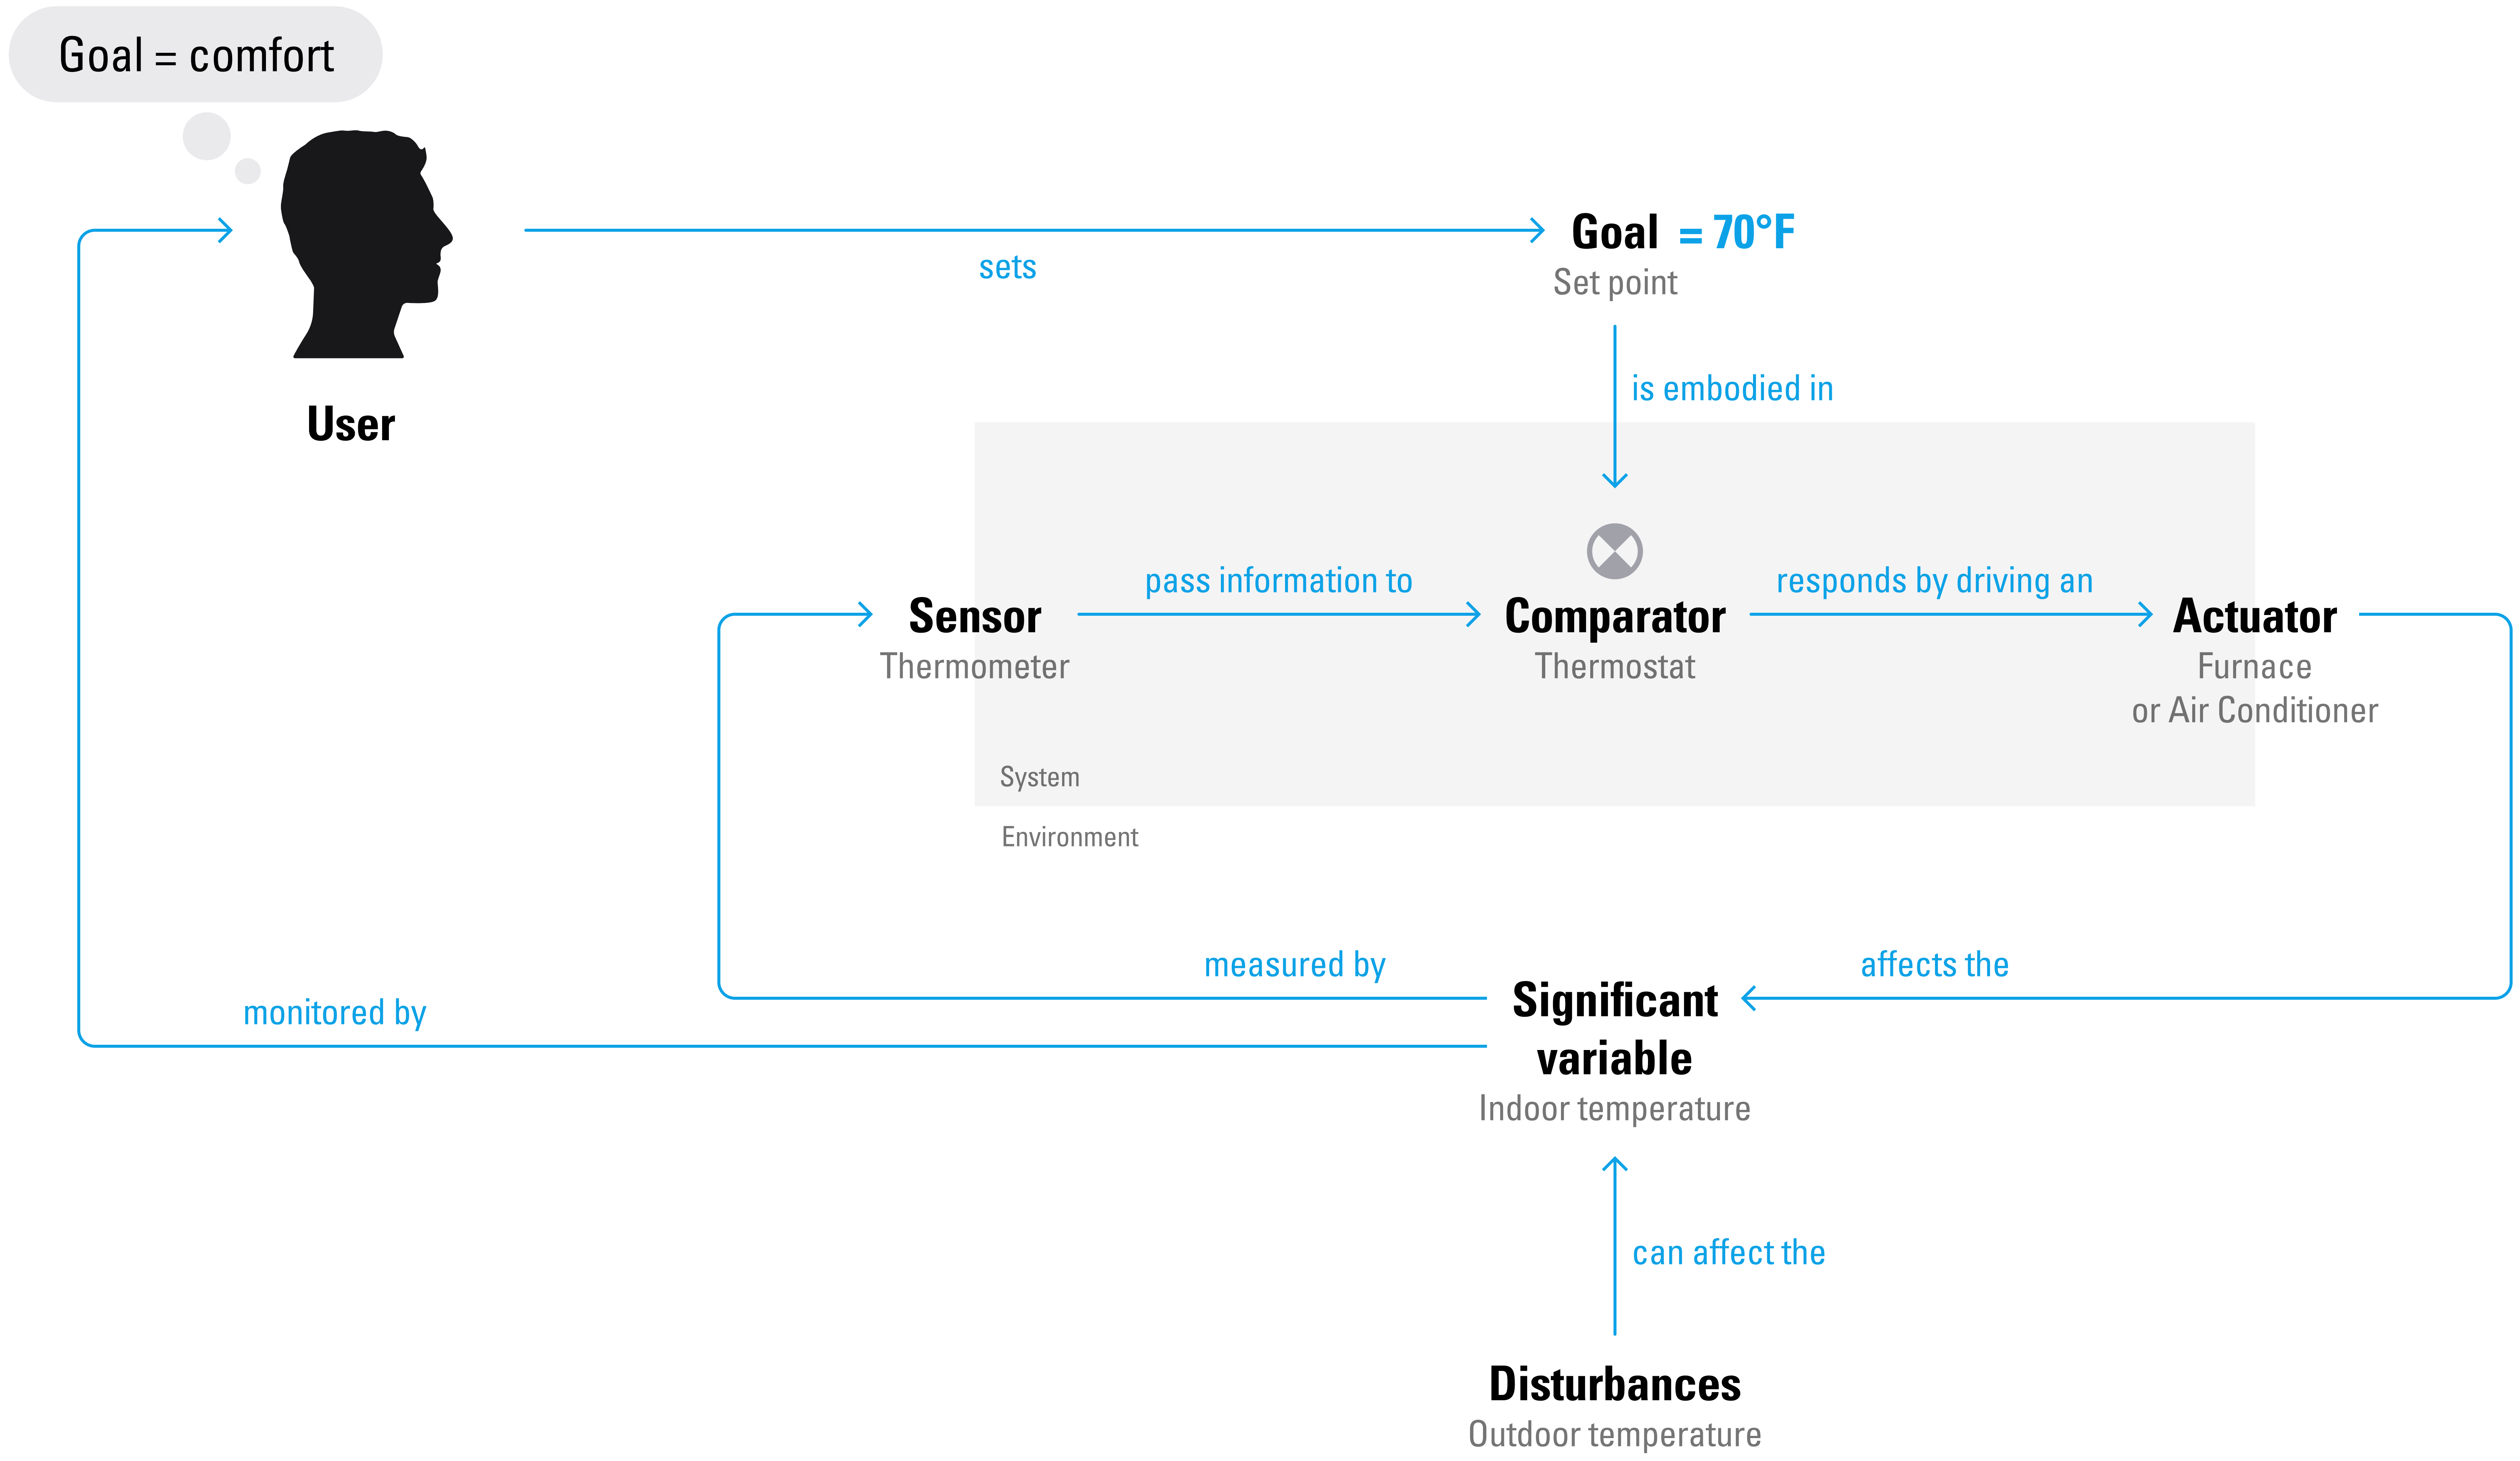

The second feedback loop is also a control system; the action of the second loop is to regulate the goal of the first loop.

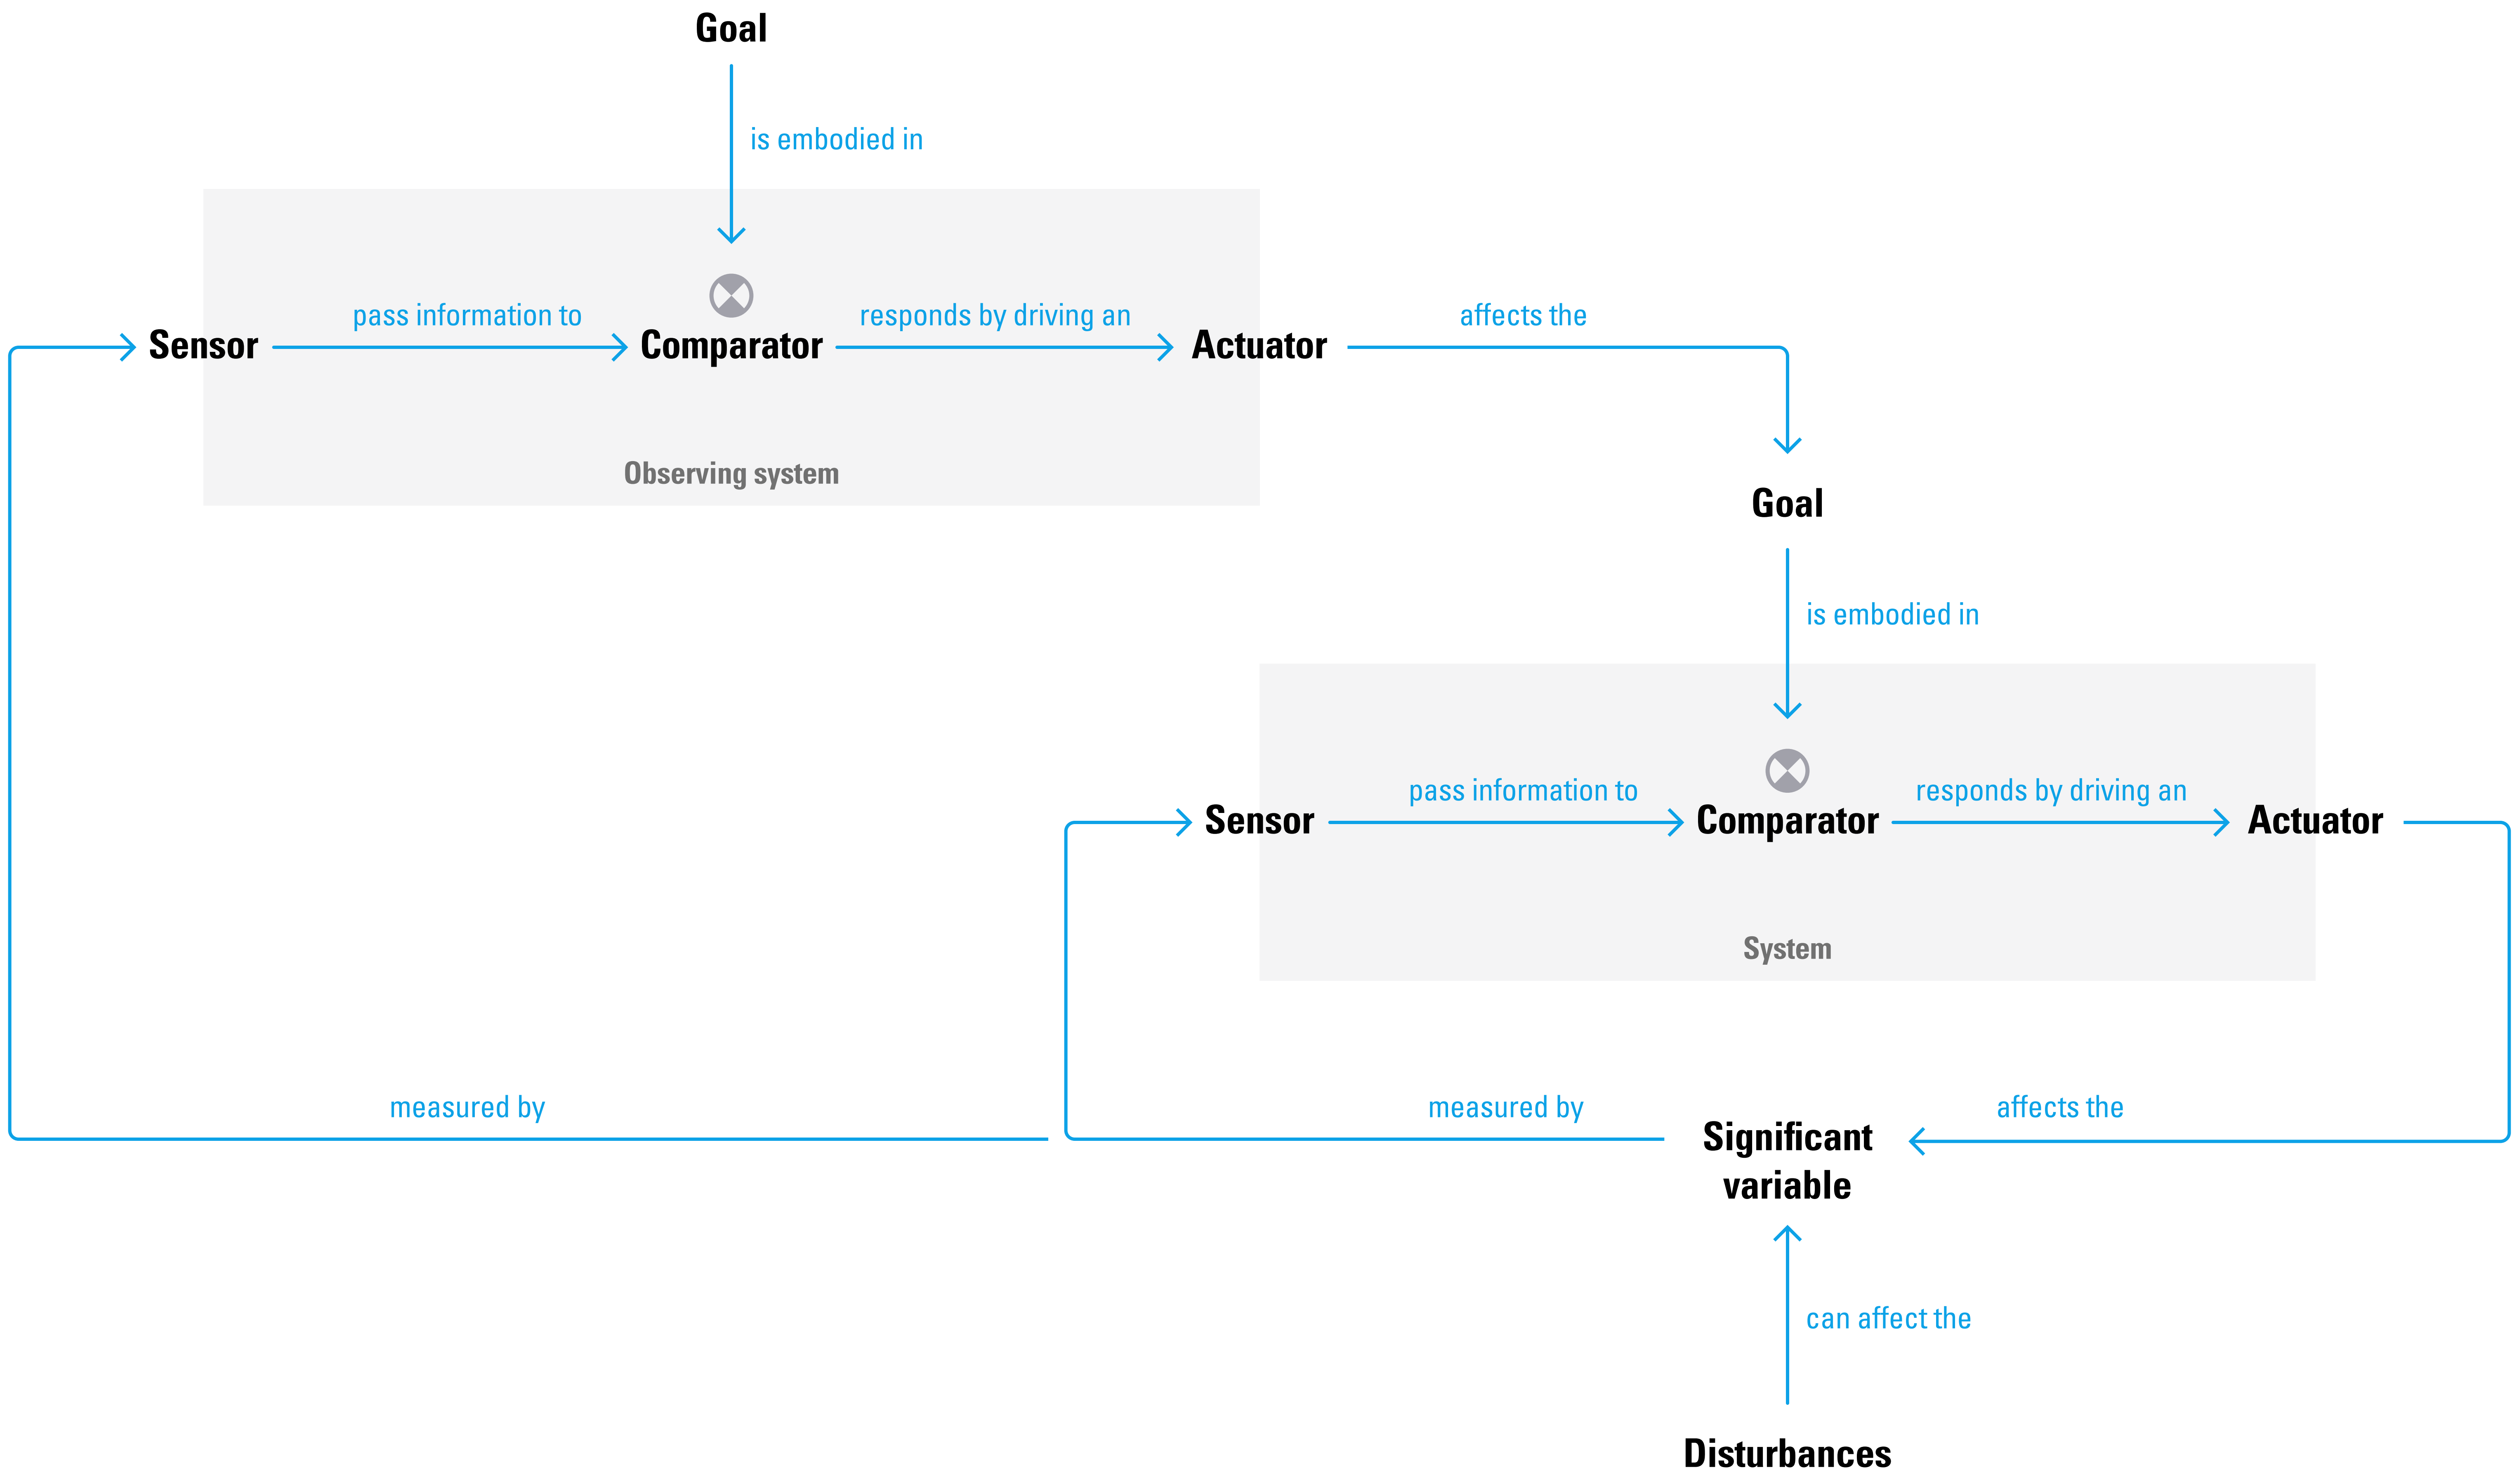

In Russell & Norvig's learning agent, the learning element modifies the agent's goals and models in order to improve its performance.

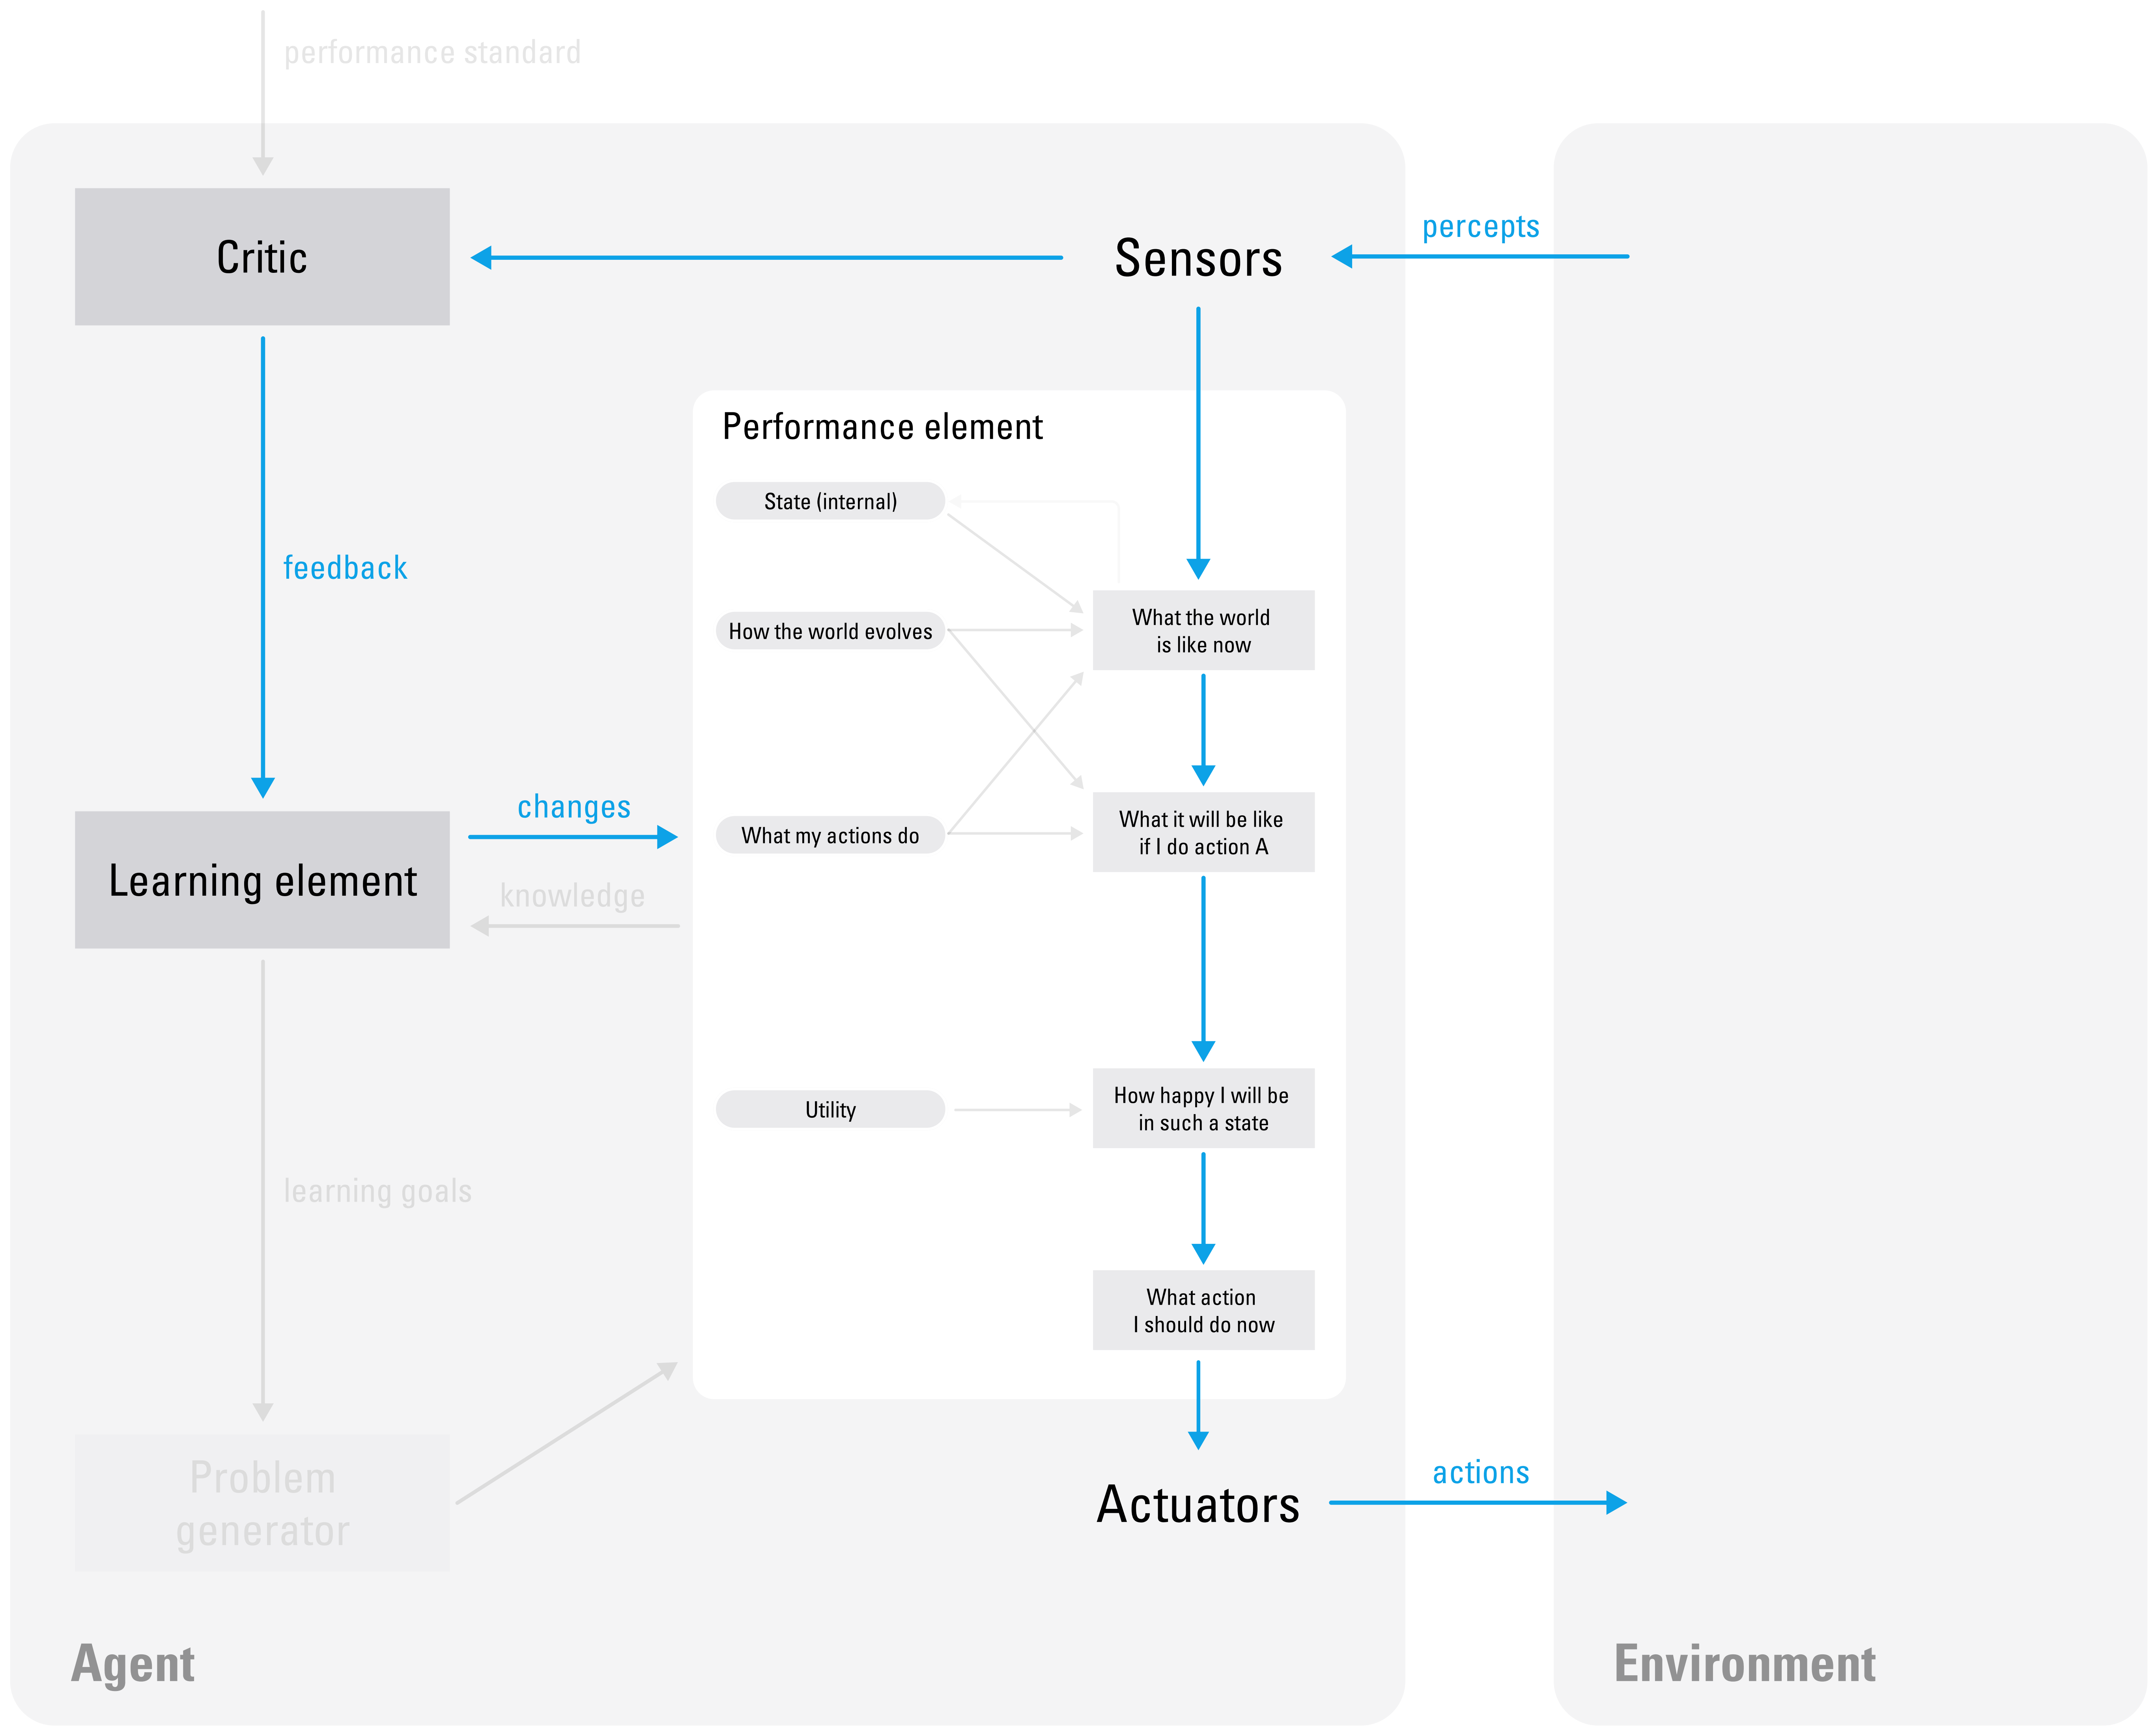

This learning process forms a second-order feedback loop in the agent.

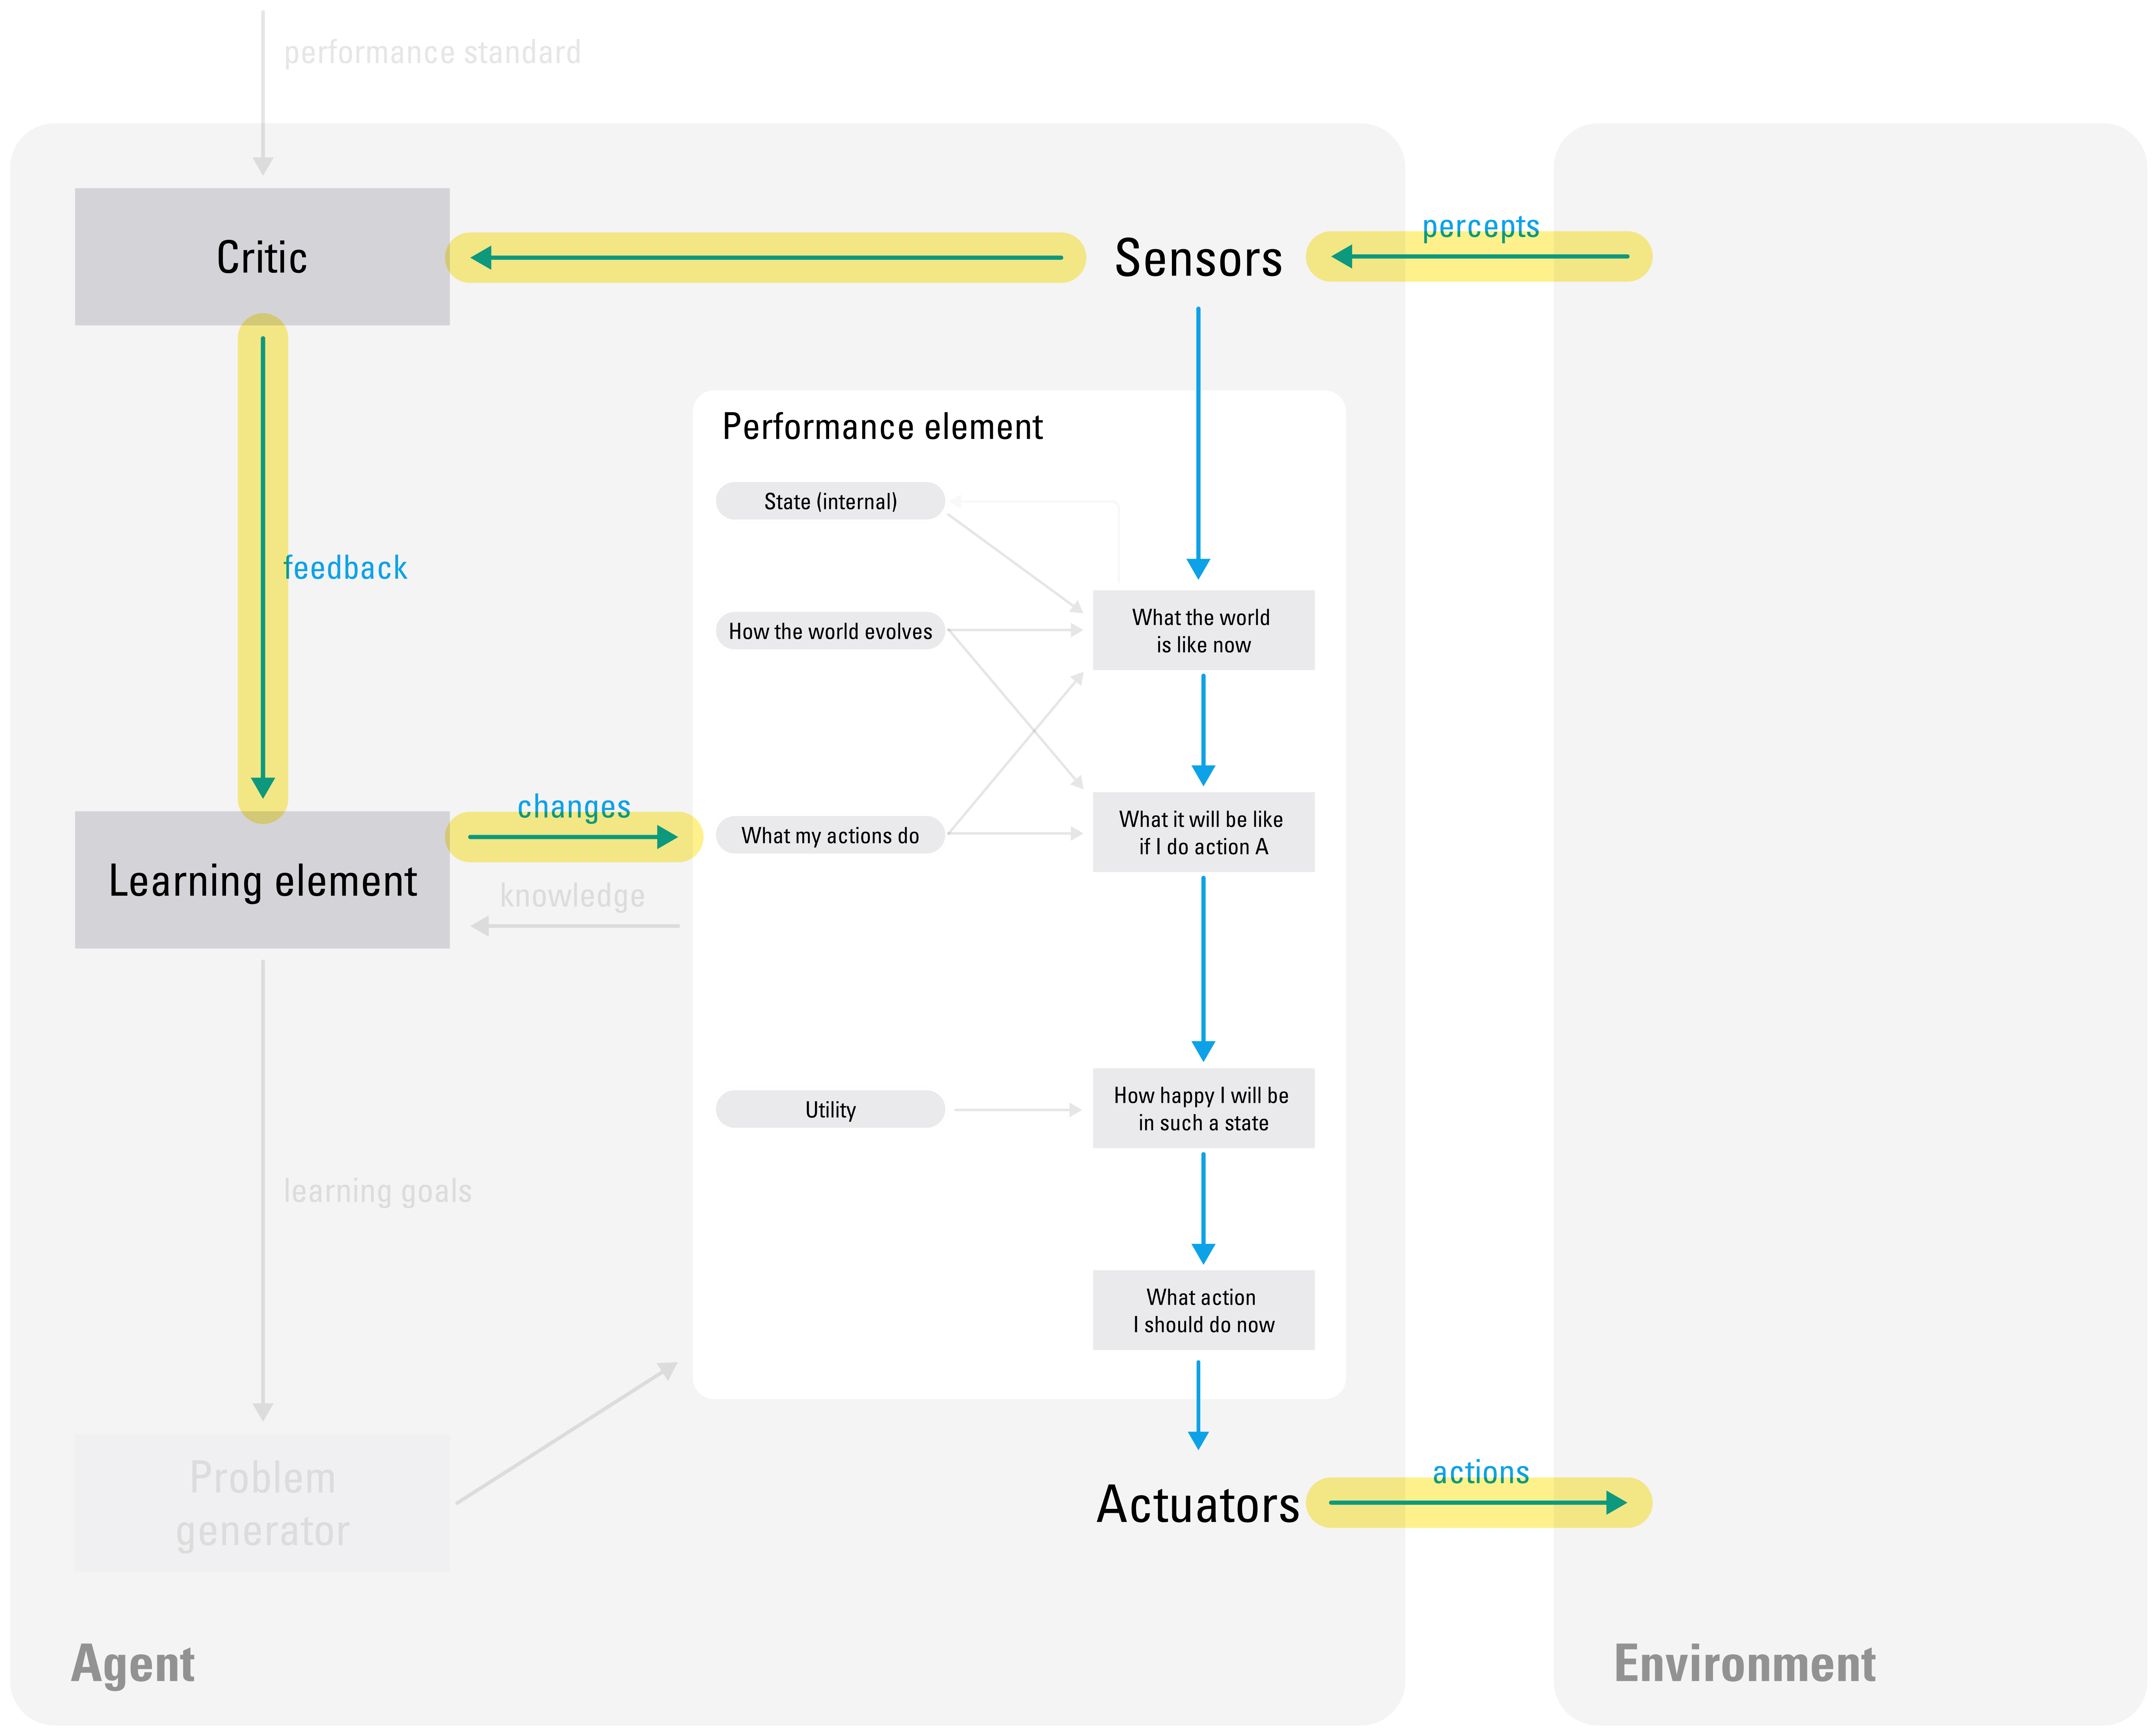

Intelligent agents as cybernetic systems

March 14, 2024
Dubberly Design Office