Smart and Adaptive Cyber-Physical Systems

Chapters 1,2

Cyber-Physical Systems

- Smart mobility
- Smart factory
- Smart grid
- Smart health care
- Smart city

Smart XX

But what does it mean to be smart ?

Being Smart

We call ourselves Home Sapiens = Wise Man

Understanding human intelligence

- How do we perceive, understand, predict and manipulate
- A world far larger and more complicated than ourselves?

Building intelligent systems

- Originally: the aim of artificial intelligence
- Nowadays: control and decision theory, CPS theory, etc

Being Smart (Artificial Intelligence)

Thinking Humanly	Thinking Rationally
Make computers think, make machines with mind Automation of all activities we associate with human thinking	Study of mental faculties through computational models Computations allowing to perceive, act and reason

Being Smart (Artificial Intelligence)

Thinking Humanly	**Thinking Rationally**
Make computers think, make machines with mind Automation of all activities we associate with human thinking	Study of mental faculties through computational models Computations allowing to perceive, act and reason
Acting Humanly	**Acting Rationally**
Machines that perform actions requiring intelligence Make computers do things at which we are currently better	The study of the design of intelligent agents Concerned with intelligent behavior in artifacts

Acting Humanly

Turing Test (Proposed in 1950)

- After posing some written questions a human interrogator
- Cannot tell if the written responses came from a computer

A Computer Needs the Following Capabilities

- Natural language processing: To communicate
- Knowledge representation: To store what it knows
- Automated reasoning: To draw new conclusions
- Machine learning: To adapt to new circumstances

Acting Humanly

Total Turing Test Contains in Addition

- A video signal: To test subject's perceptual abilities
- A hatch: To pass physical objects to the subject

A Computer Needs Additional Capabilities

- Computer vision: To perceive objects
- Robotics: To manipulate objects and move about

Thinking Humanly

Need to get inside the actual working of human minds

- Introspection: Try to catch our own thoughts
- Psychological experiments: Observing a person in action
- Brain imaging: Observing the brain in action

Cognitive science brings together

- Computer models: From artificial intelligence
- Experimental techniques: From psychology
- In order to: Construct precise and testable theories

What About the Brain?

- Brains (human minds) are very good at making rational decisions (but not perfect)
- "Brains are to intelligence as wings are to flight"
- Brains aren't as modular as software
- Lessons learned: **prediction** and **simulation** are key to decision making

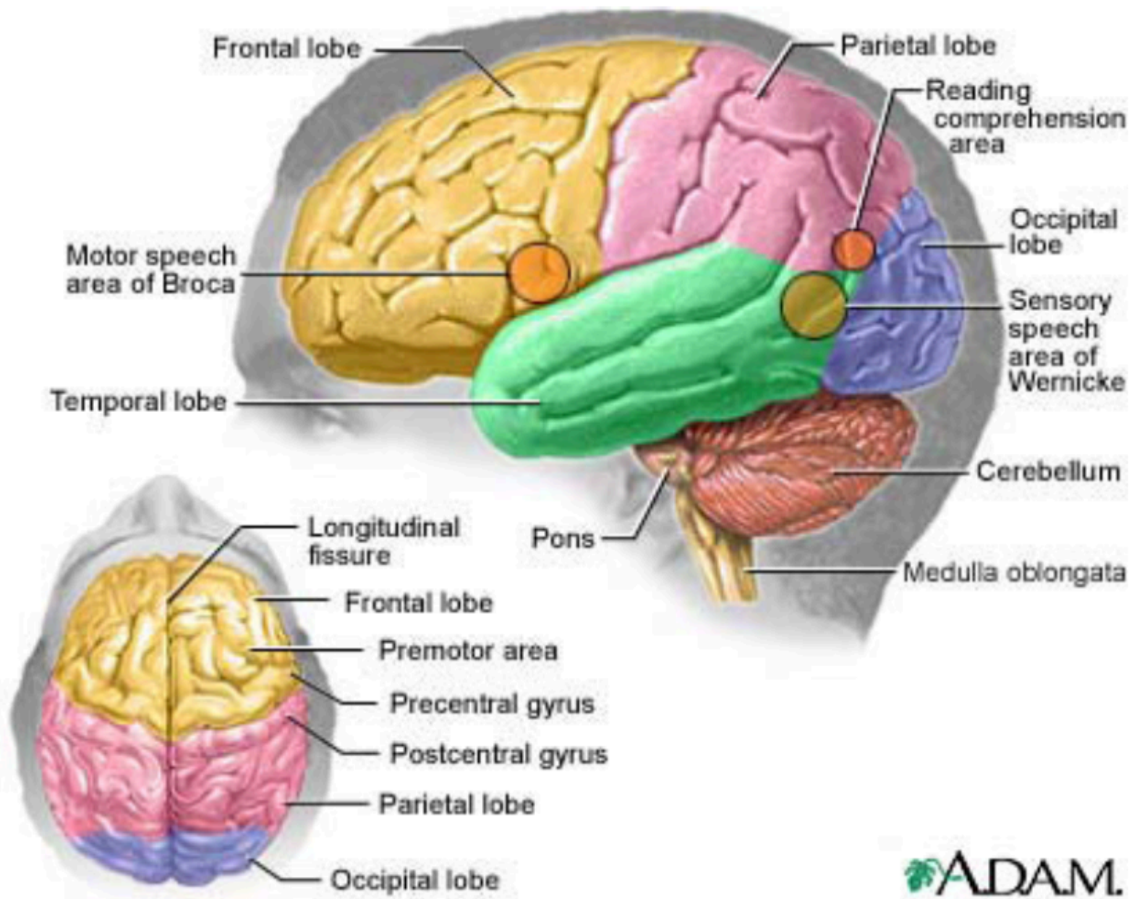

Thinking Rationally

The syllogism of Greek philosopher Aristotle

- Pattern for right thinking: Always yield correct conclusions
- Main pattern: $A \wedge (A \rightarrow B) = A \wedge B = B \wedge (B \rightarrow A)$
- Problem: World is not black and white (qualitative)

The extension of syllogisms to a logic of science

- Main pattern: $P(A)\,P(B \mid A) = P(A \wedge B) = P(B)\,P(A \mid B)$
- Advantage: Shades of gray (quantitative)

Acting Rationally

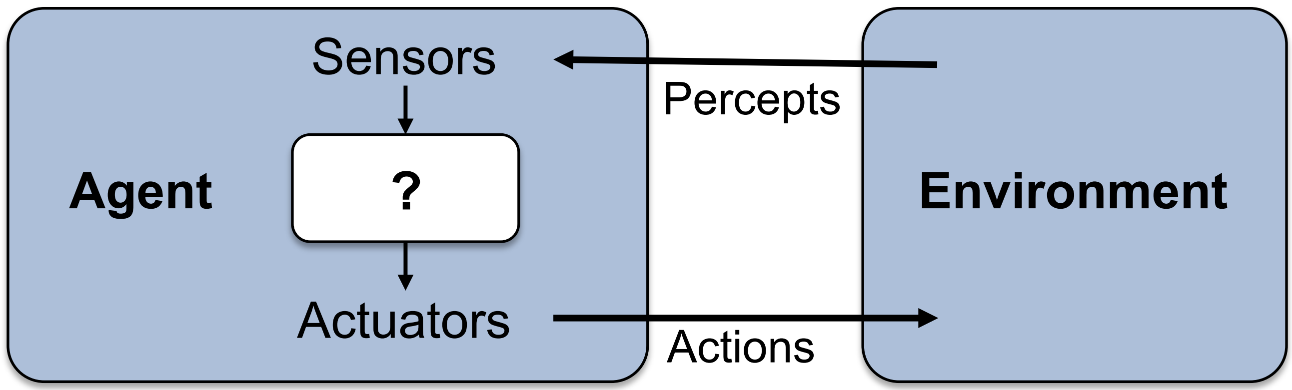

Computer Agent Latin agere = doing

- Operates autonomously and persists over long time
- Perceives, acts upon and adapts to its environment
- Creates and pursues its own goals

Pacman as an Agent

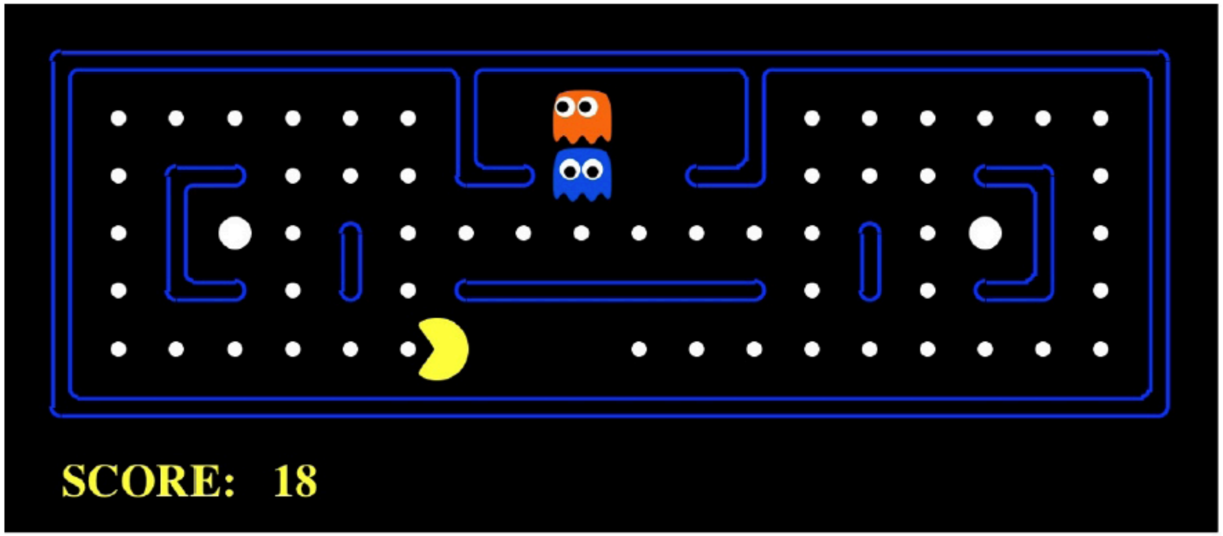

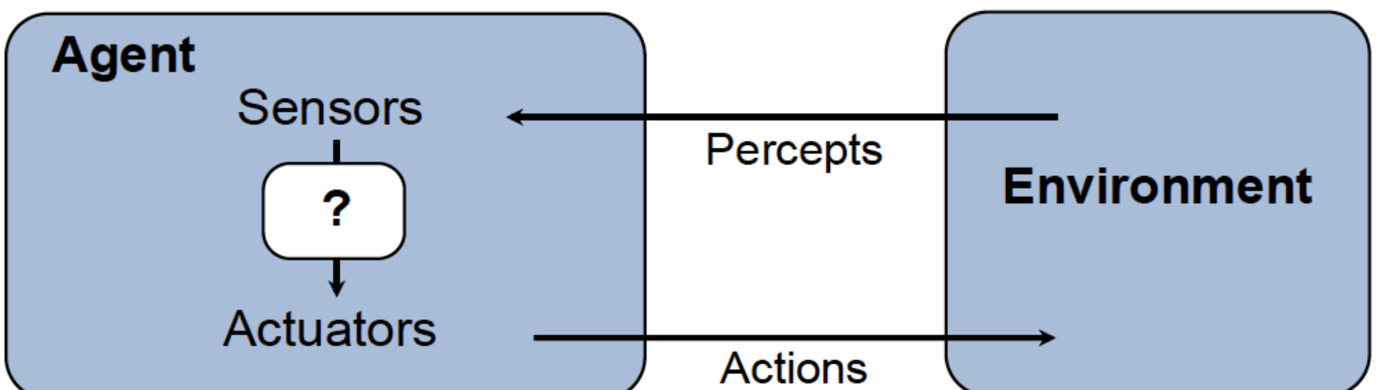

Acting Rationally

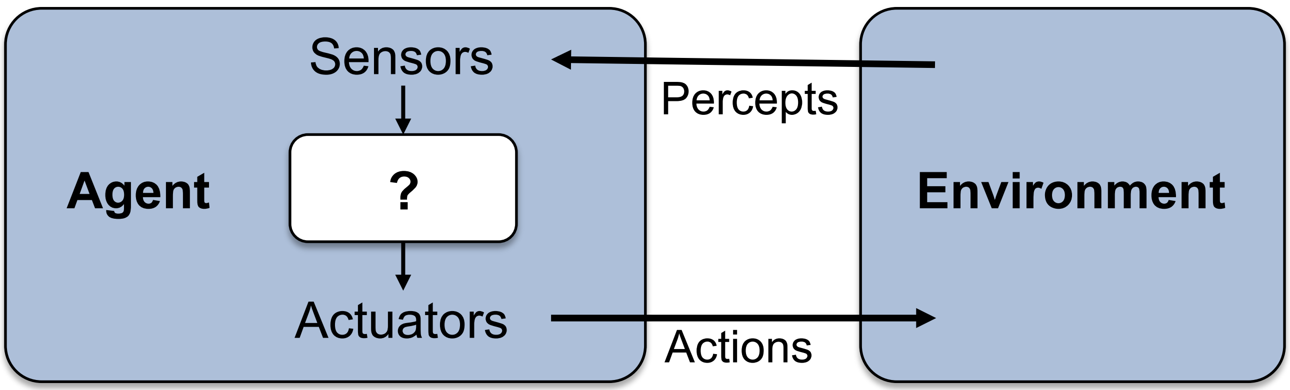

Rational Agent Extension of Computer Agent

- Acts so as to achieve the best outcome, and when
- There is uncertainty, the best expected outcome

Advantages of Rational Agents

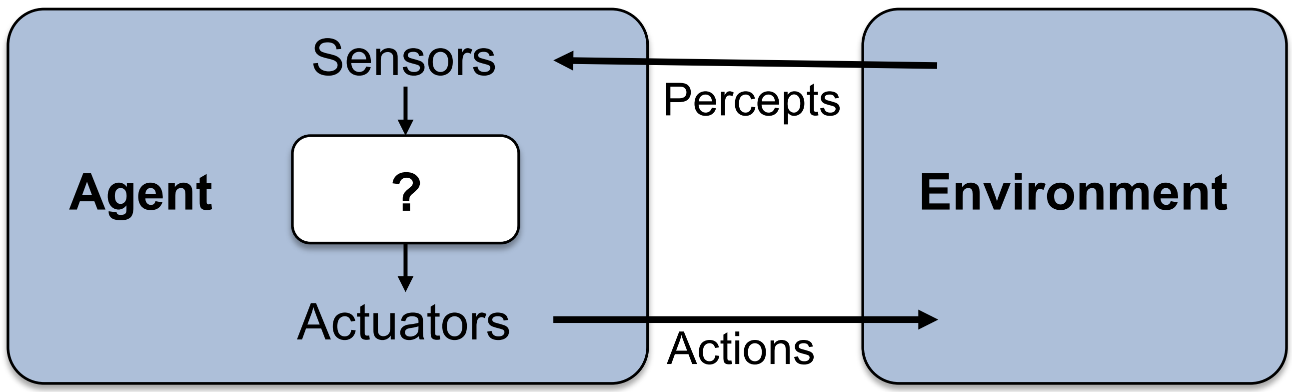

The rational-agent approach has two advantages

- It is more general than the laws-of-thought approach
  - Correct inference is just one way of achieving rationality
- It is more amenable to scientific development
  - Rationality is mathematically well defined and very general

In this class: Smart = Rational !

Rational Agent (RA)

What is rational at any given time depends on

- The performance measure that defines success
- The agent’s prior knowledge of the environment
- The actions that the agent can perform
- The agent’s percept sequence to date

For each percept sequence a RA should

- Select an action that is expected to
- Maximize its performance measure, given
- The evidence provided by the percept sequence and
- Whatever built-in knowledge the agent has

Task Environment (TE)

Agent Type	Performance Measure	Environment Type	Actuators Type	Sensors Type
Taxi Driver	Safe, Fast, Legal, Comfortable trip, Maximize profits	Roads, Other Traffic, Pedestrians, Customers	Steering, Accelerator, Brake, Signal, Horn, Display	Cameras, Sonar, Speedometer, GPS, Accelerometer, Engine Sensors

PEAS=(Performance, Environment, Actuators, Sensors)

Task Environment (TE)

Agent Type	Performance Measure	Environment Type	Actuators Type	Sensors Type
Taxi Driver	Safe, Fast, Legal, Comfortable trip, Maximize profits	Roads, Other Traffic, Pedestrians, Customers	Steering, Accelerator, Brake, Signal, Horn, Display	Cameras, Sonar, Speedometer, GPS, Accelerometer, Engine Sensors
Medical Diagnosis	Healthy Patient	Patient, Hospital, Physician, Staff	Diagnoses, Treatments	Patient Answers, Monitors

PEAS=(Performance, Environment, Actuators, Sensors)

Task Environment (TE)

Agent Type	Performance Measure	Environment Type	Actuators Type	Sensors Type
Taxi Driver	Safe, Fast, Legal, Comfortable trip, Maximize profits	Roads, Other Traffic, Pedestrians, Customers	Steering, Accelerator, Brake, Signal, Horn, Display	Cameras, Sonar, Speedometer, GPS, Accelerometer, Engine Sensors
Medical Diagnosis	Healthy Patient	Patient, Hospital, Physician, Staff	Diagnoses, Treatments	Patient Answers, Monitors
Refinery Controller	Purity, Yield, Safety	Refinery, Operators	Valves, Pumps, Heaters	Temperature, Pressure

PEAS=(Performance, Environment, Actuators, Sensors)

Task Environment (TE)

Agent Type	Performance Measure	Environment Type	Actuators Type	Sensors Type
Taxi Driver	Safe, Fast, Legal, Comfortable trip, Maximize profits	Roads, Other Traffic, Pedestrians, Customers	Steering, Accelerator, Brake, Signal, Horn, Display	Cameras, Sonar, Speedometer, GPS, Accelerometer, Engine Sensors
Medical Diagnosis	Healthy Patient	Patient, Hospital, Physician, Staff	Diagnoses, Treatments	Patient Answers, Monitors
Refinery Controller	Purity, Yield, Safety	Refinery, Operators	Valves, Pumps, Heaters	Temperature, Pressure
Interactive English Tutor	Student's Score on a Test	Set of Students, Testing Agency	Display Exercises Suggestions	Keyboard Entry

PEAS=(Performance, Environment, Actuators, Sensors)

TE: Observable-Unobservable

Fully observable task environments (crossword)

- Sensors always capture complete environment's state
- Effective if all aspects relevant to action choices detected
- Relevance depends on the performance measure
- Agent does not need to maintain the world's state

Partially observable task environments (taxi driving)

- Noisy and inaccurate sensors
- Parts of the state are simply missing from sensor data
- E.g. a taxi driver cannot see what other drivers think

Unobservable task environments

- The agent has no sensors

TE: Single-Multi Agent

Single agent task environment

- E.g. an agent solving a crossword puzzle by itself

Multi agent task environment

- E.g. a game of chess is a two agent environment
- Agent’s B performance depends on agent’s A performance
- Cooperative or competitive multi agents (taxi, chess)

Multi-agent design is quite different

- Communication is rational as it receives hidden state
- Randomization is rational as it avoids predictability

TE: Certain-Uncertain

Deterministic task environments (certain)

- Environment's next state completely determined by
  - Current state of the environment
  - The action executed by the agent (e.g. crossword)
- Uncertainty ignored in fully observable deterministic TE

Stochastic task environments (uncertain)

- Partially observable environments appear as stochastic
- In most real situations this is the case, e.g. in taxi driving

Nondeterministic task environments (uncertain)

- Uncertain but no probabilities attached (e.g. chess)
- The agent needs to succeed for all possible outcomes

TE: Episodic-Sequential

Episodic task environments

- Agent's experience divided in episodes. In each:
  - It first receives a single percept from its sensors
  - It than performs a single action with its actuators
- Next episode does not depend on actions in previous ones
- Many classification tasks are episodic
  - Spotting defective parts on an assembly line

Sequential task environments

- Current decision could affect all future decisions
  - Crossword, chess game and taxi driving are sequential
- In sequential environments one needs to think ahead

TE: Static-Dynamic

Static task environments

- Environment doesn't change while agent deliberates
  - Need not look at the world while deciding on an action
  - Need not worry about the passage of time
  - E.g. crossword puzzles

Dynamic task environments

- Environment can change while agent deliberates
  - Are continuously asking the agent what it wants to do
  - If it hasn't decided yet, that counts as deciding to do nothing
  - E.g. taxi driving

Semi-dynamic task environments

- Environment cannot change in time but agent's score can
  - E.g. chess playing with a clock

TE: Discrete-Continuous

Discrete task environments

- State of environment is discrete
- Time of environment is discrete
- Percepts and/or actions are discrete
  - E.g. chess has discrete state, percepts and actions

Continuous task environments

- State of environment is continuous
- Time of environment is continuous
- Percepts and/or actions are continuous
  - E.g. taxi driving is continuous state, time, percepts and actions

Types of Task Environments

Task Env	Observ	Agents	Stochastic	Episodic	Static	Discrete
Crossword	Fully	Single	Deterministic	Sequential	Static	Discrete
Chess (clock)	Fully	Multi	Deterministic	Sequential	Semi	Discrete

Types of Task Environments

Task Env	Observ	Agents	Stochastic	Episodic	Static	Discrete
Crossword Chess (clock)	Fully Fully	Single Multi	Deterministic Deterministic	Sequential Sequential	Static Semi	Discrete Discrete
Poker Backgammon	Partially Fully	Multi Multi	Stochastic Stochastic	Sequential Sequential	Static Static	Discrete Discrete

Types of Task Environments

Task Env	Observ	Agents	Stochastic	Episodic	Static	Discrete
Crossword Chess (clock)	Fully Fully	Single Multi	Deterministic Deterministic	Sequential Sequential	Static Semi	Discrete Discrete
Poker Backgammon	Partially Fully	Multi Multi	Stochastic Stochastic	Sequential Sequential	Static Static	Discrete Discrete
Taxi-driving Medical-diag	Partially Partially	Multi Single	Stochastic Stochastic	Sequential Sequential	Dynamic Dynamic	Continuous Continuous

Types of Task Environments

Task Env	Observ	Agents	Stochastic	Episodic	Static	Discrete
Crossword Chess (clock)	Fully Fully	Single Multi	Deterministic Deterministic	Sequential Sequential	Static Semi	Discrete Discrete
Poker Backgammon	Partially Fully	Multi Multi	Stochastic Stochastic	Sequential Sequential	Static Static	Discrete Discrete
Taxi-driving Medical-diag	Partially Partially	Multi Single	Stochastic Stochastic	Sequential Sequential	Dynamic Dynamic	Continuous Continuous
Image-analysis Part-pick-robot	Fully Partially	Single Single	Deterministic Stochastic	Episodic Episodic	Semi Dynamic	Continuous Continuous

Types of Task Environments

Task Env	Observ	Agents	Stochastic	Episodic	Static	Discrete
Crossword Chess (clock)	Fully Fully	Single Multi	Deterministic Deterministic	Sequential Sequential	Static Semi	Discrete Discrete
Poker Backgammon	Partially Fully	Multi Multi	Stochastic Stochastic	Sequential Sequential	Static Static	Discrete Discrete
Taxi-driving Medical-diag	Partially Partially	Multi Single	Stochastic Stochastic	Sequential Sequential	Dynamic Dynamic	Continuous Continuous
Image-analysis Part-pick-robot	Fully Partially	Single Single	Deterministic Stochastic	Episodic Episodic	Semi Dynamic	Continuous Continuous
Refinery ctrl English tutor	Partially Partially	Single Multi	Stochastic Stochastic	Sequential Sequential	Dynamic Dynamic	Continuous Discrete

Simple Reflex Agent (Episodic)

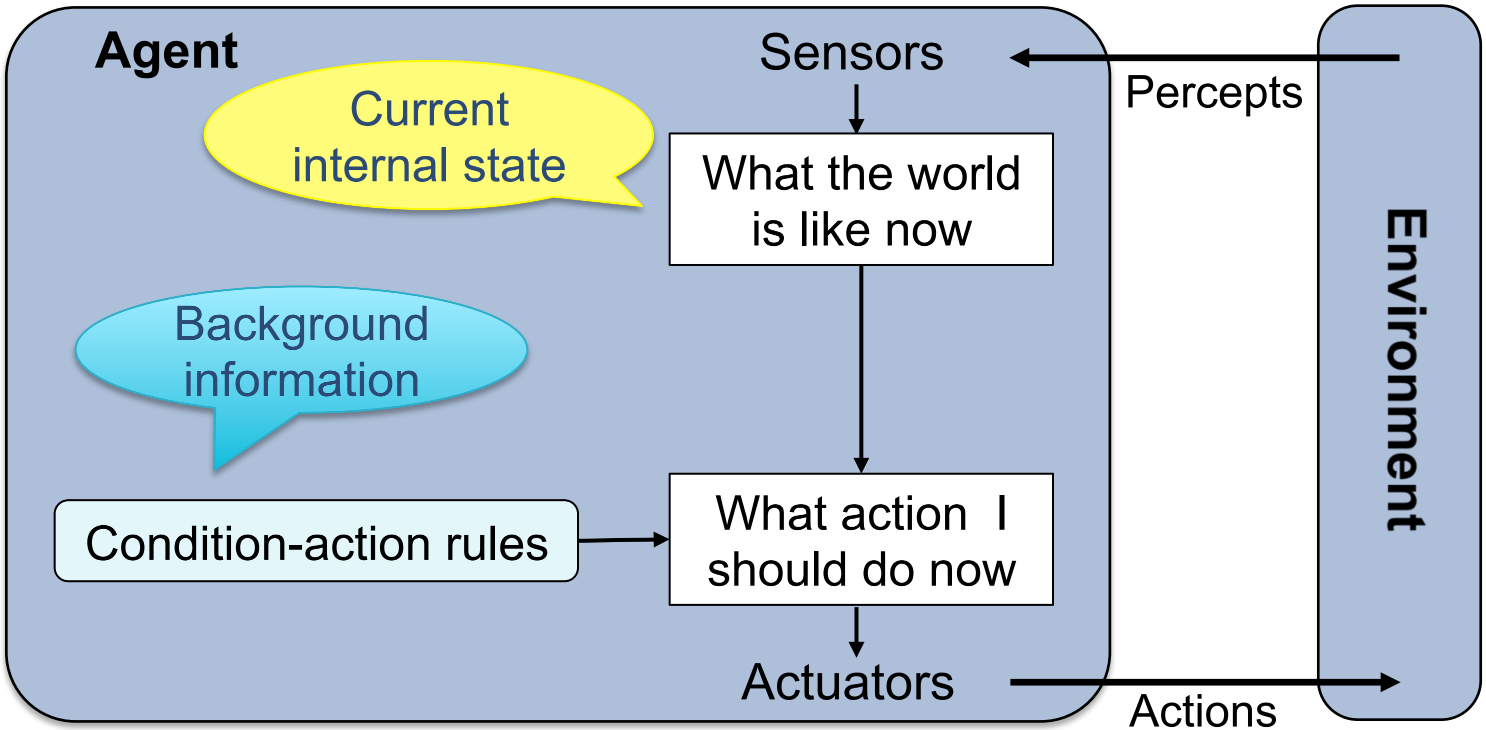

function Simple-Reflex-Agent (percept) **returns** action
persistent set-of-condition-action rules

state = Interpret-Input(percept)
rule = Rule-Match(state, rules)
return rule.Action

Simple Reflex Agent (Episodic)

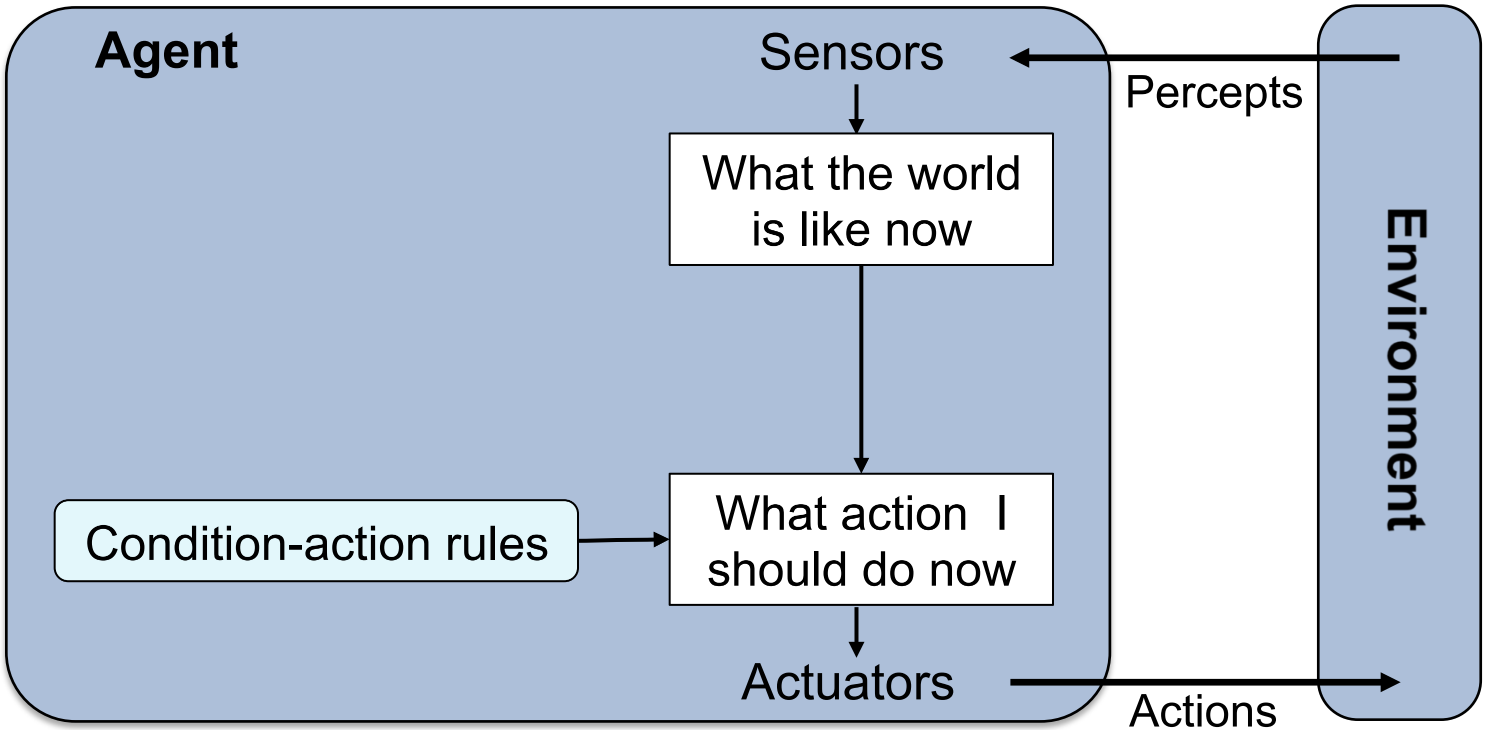

Example: The car in front is braking

- Very simple and very fast as a consequence
- However, one can do better by learning front-car behavior

Model-Based Reflex Agent (Seq)

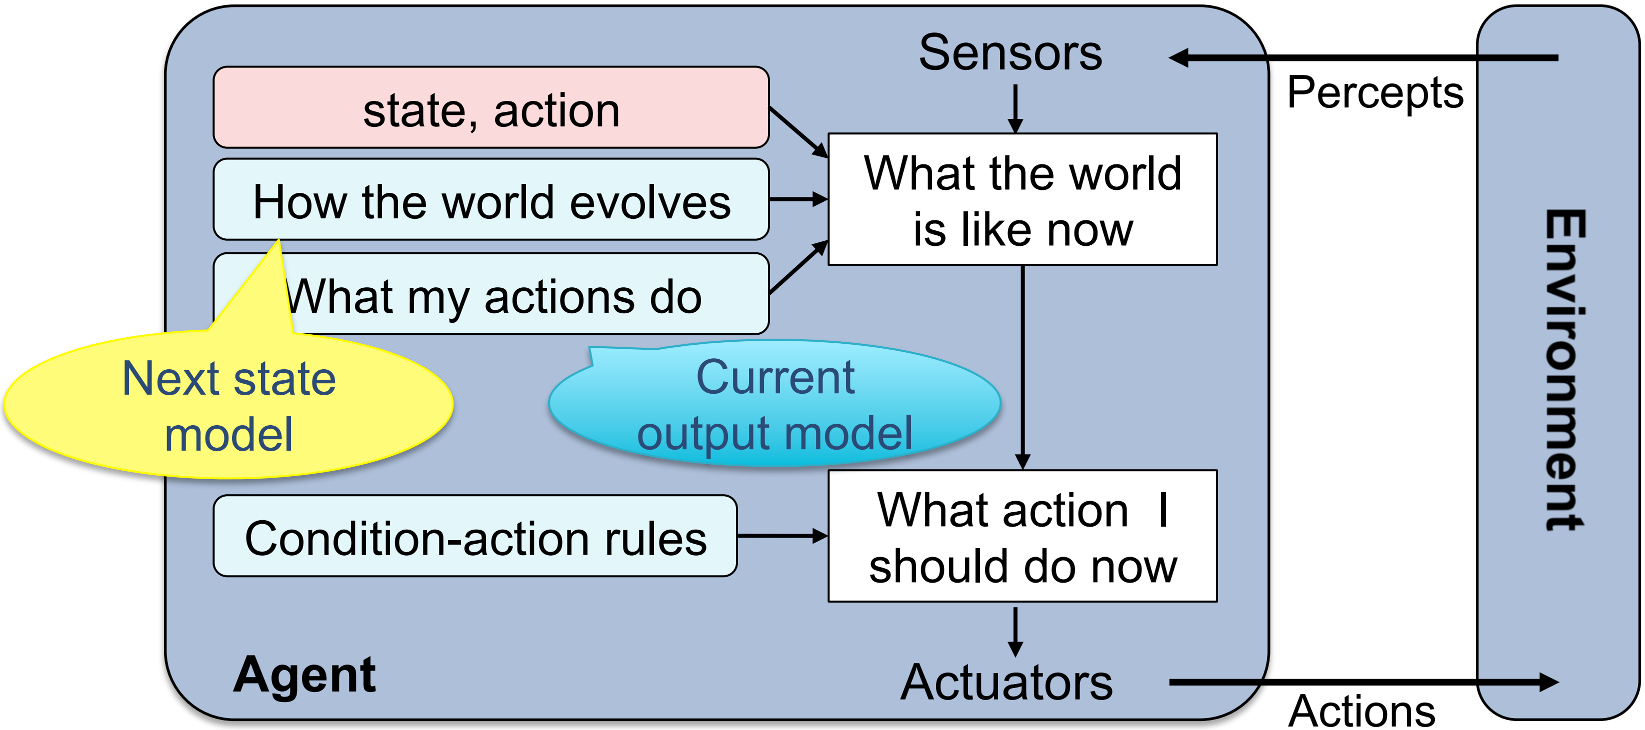

The most effective way to handle partial observability

- Keep track of the part of the world it can't see now
- The agent should maintain the previous state and action
- Which depend on the percept history. Braking car: 1-2 frames

Model-Based Reflex Agent (Seq)

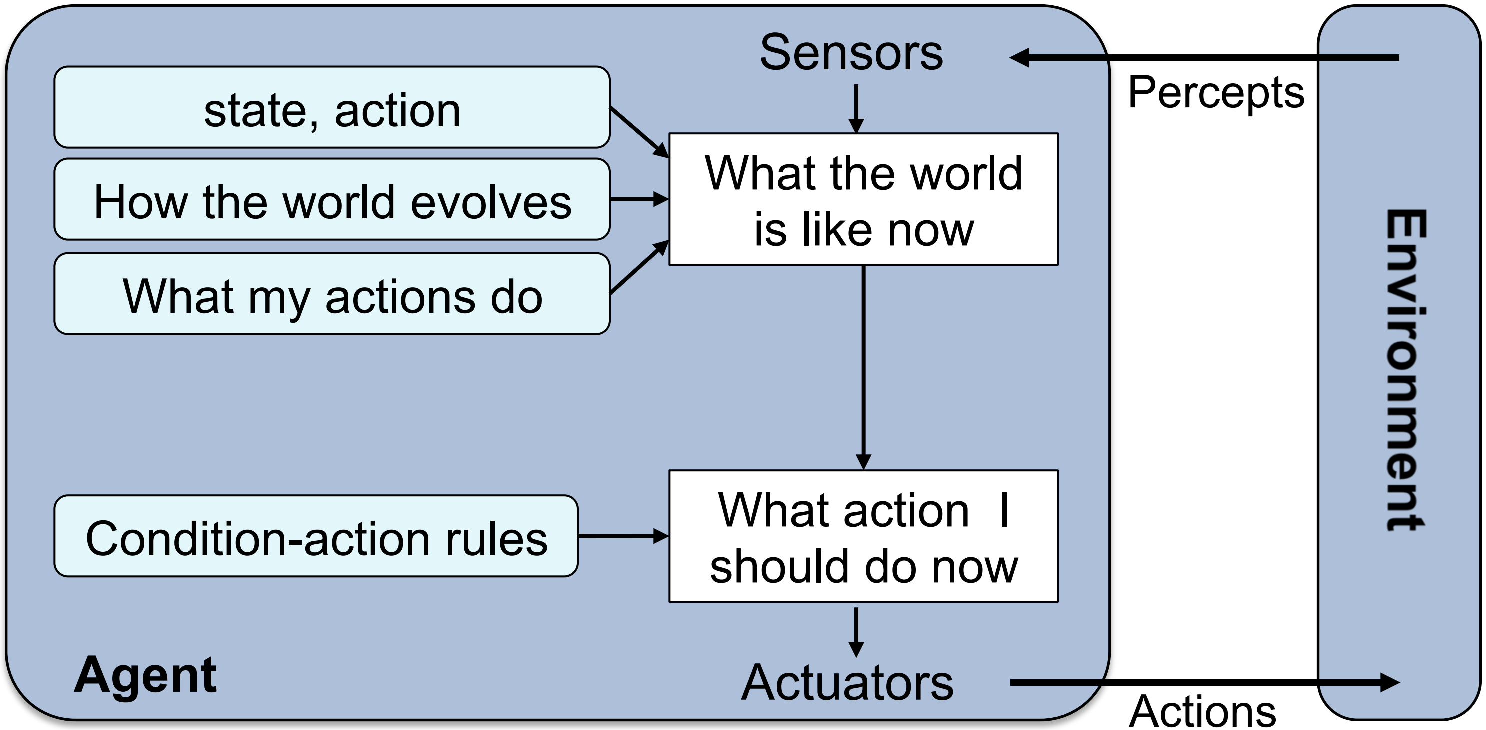

function Model-Based-Reflex-Agent (percept) **returns** action
  persistent state, action, model, rules

  state = Update-State(state, action, percept, model)
  rule = Rule-Match(state, rules)
  action = rule.Action
  return action

State-based description of an agent (white box)

Model-Based Reflex Agent (Seq)

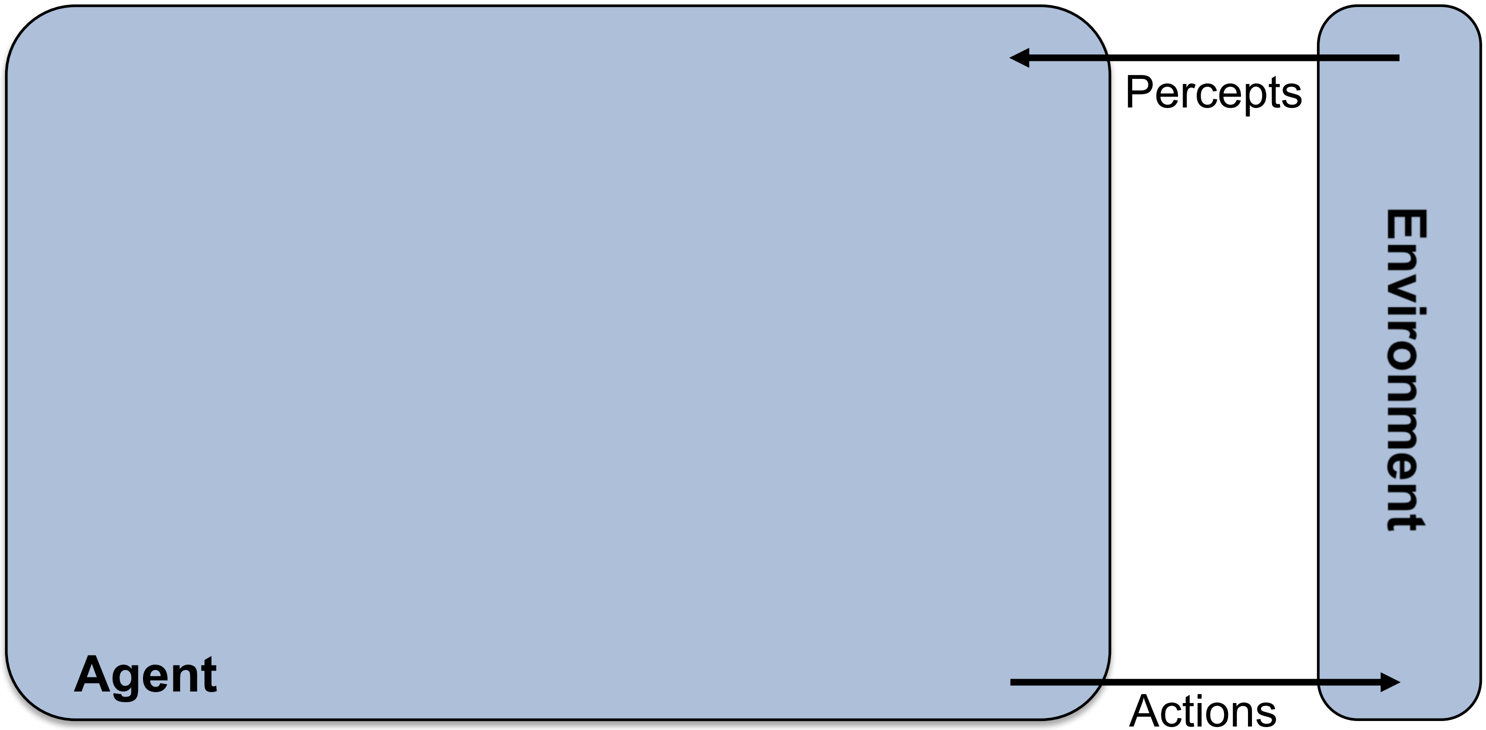

function Model-Based-Reflex-Agent (percepts) **returns** actions

Input-output description of an agent (black box)
This is what is observable and on what the performance is measured!

Model-Based Goal-Based Agent

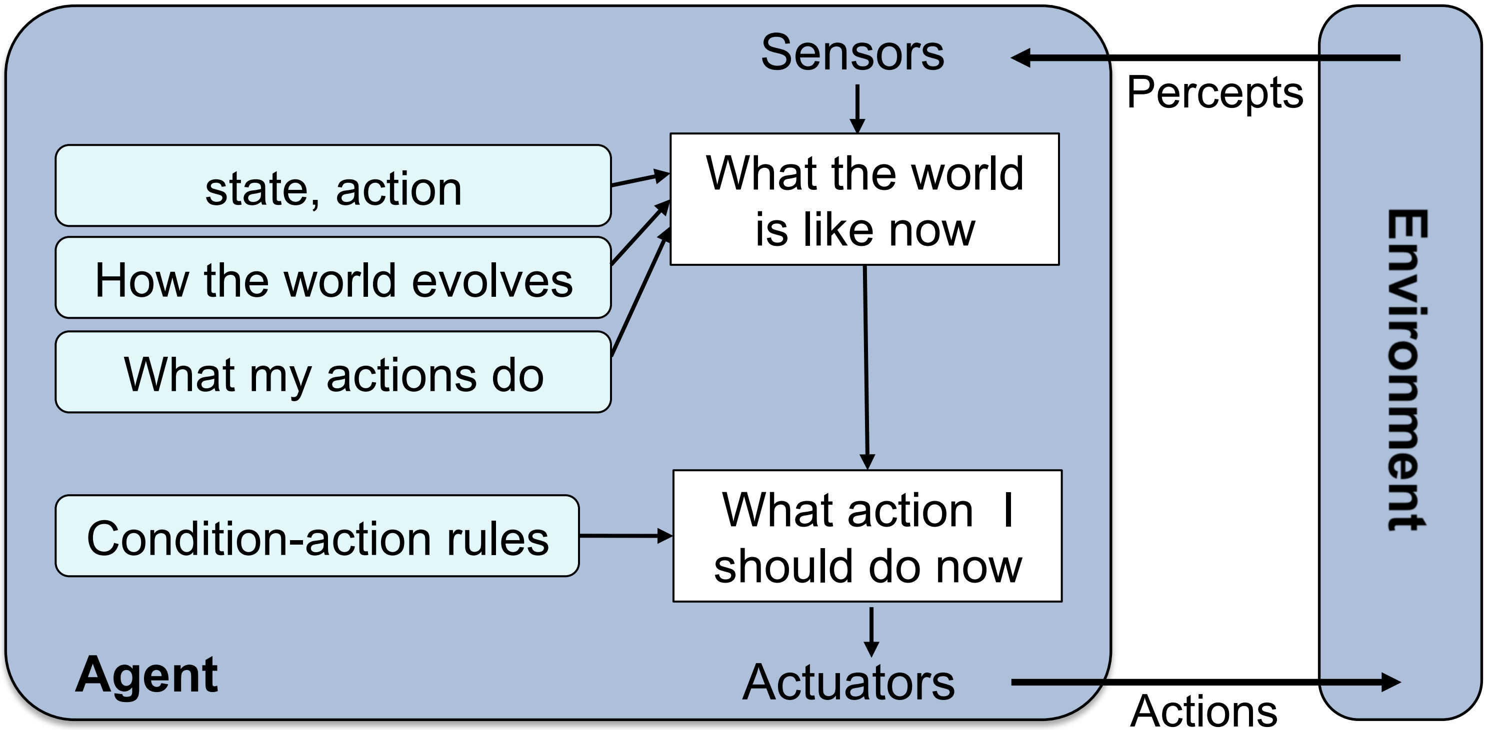

Knowing internal state not always enough to decide what to do

- At a road junction a car can turn left, right or go straight

Model-Based Goal-Based Agent

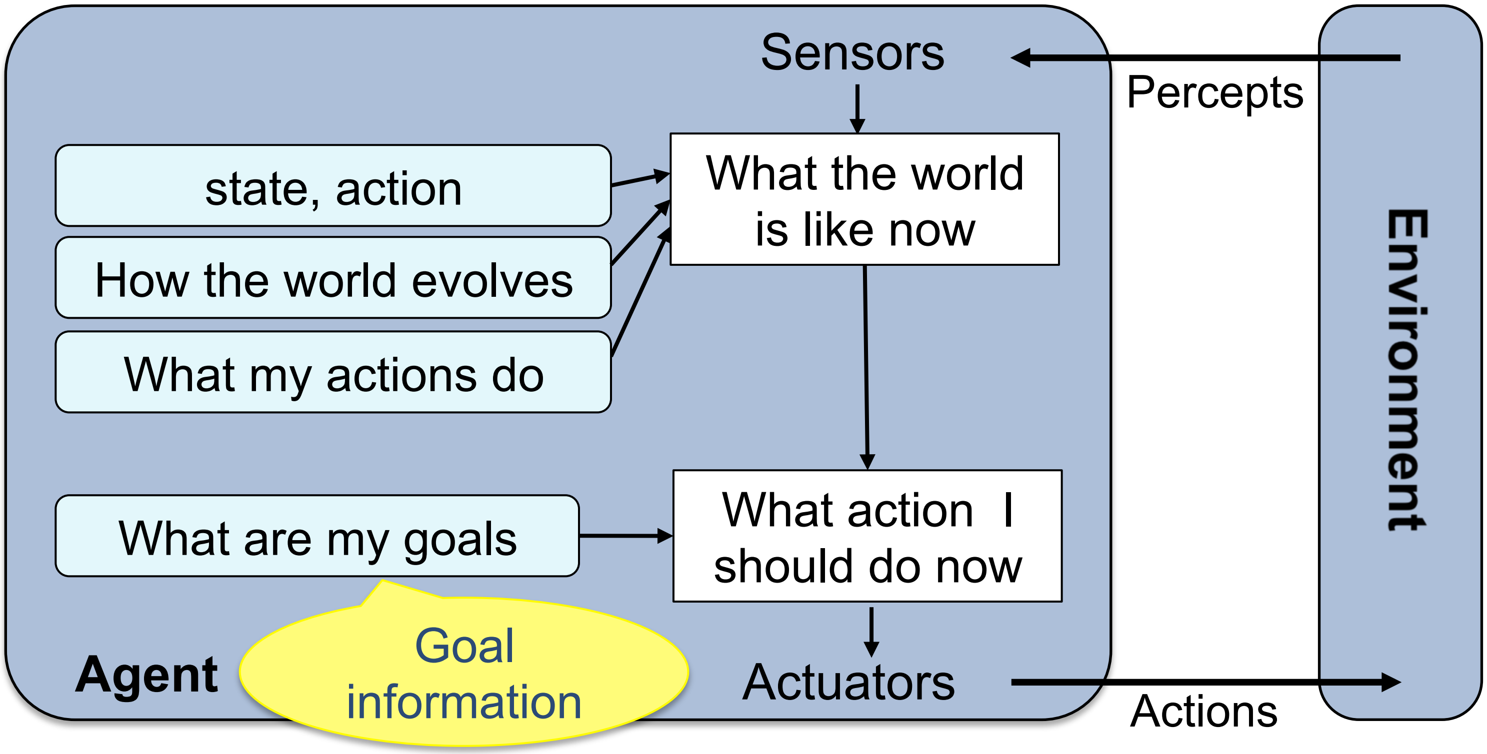

Knowing internal state not always enough to decide what to do

- At a road junction a car can turn left, right or go straight
- Correct decision depends on where the car wants to go

Model-Based Goal-Based Agent

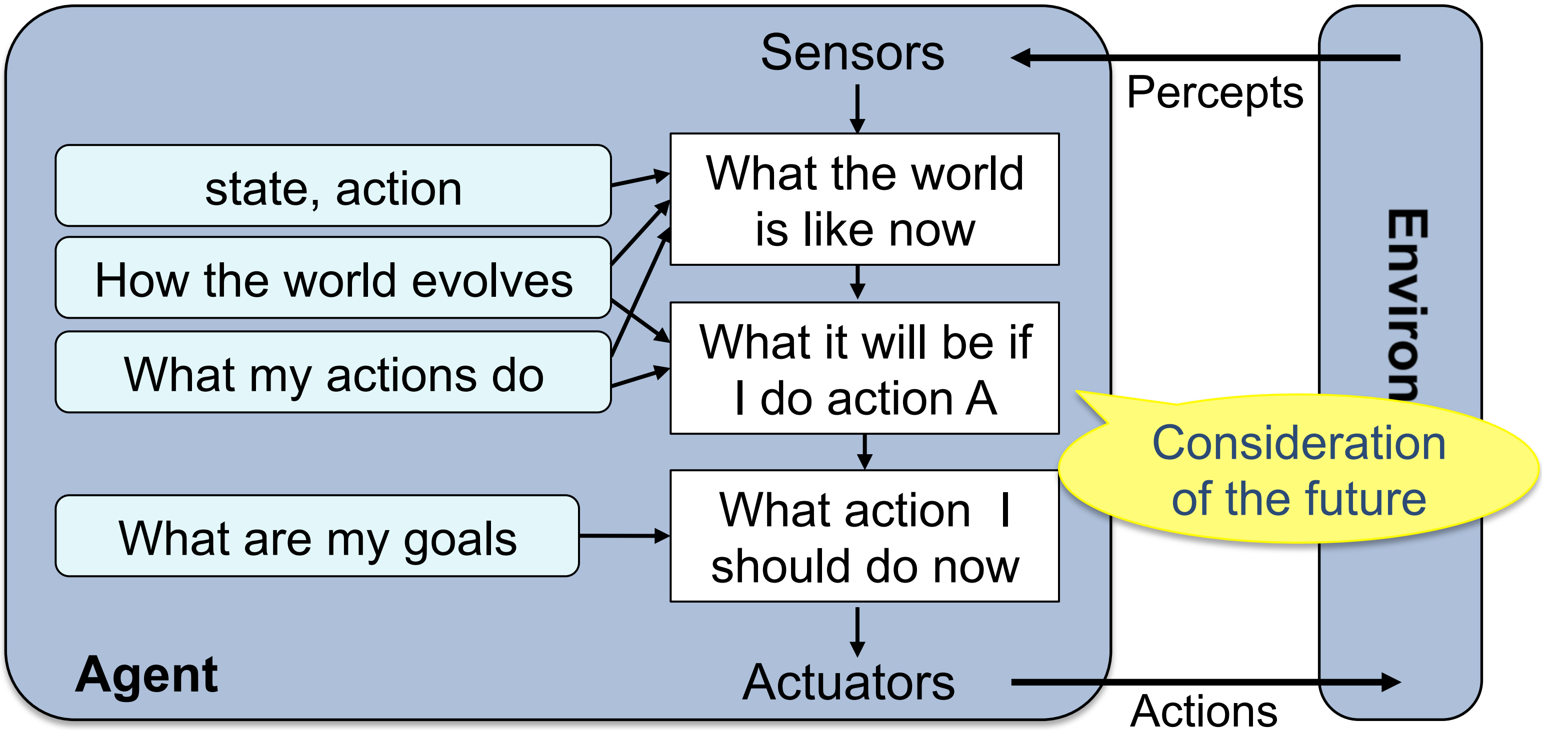

Knowing internal state not always enough to decide what to do

- At a road junction a car can turn left, right or go straight
- Correct decision depends on where the car wants to go
- Search and planning involves consideration of the future

Model-Based Goal-Based Agent

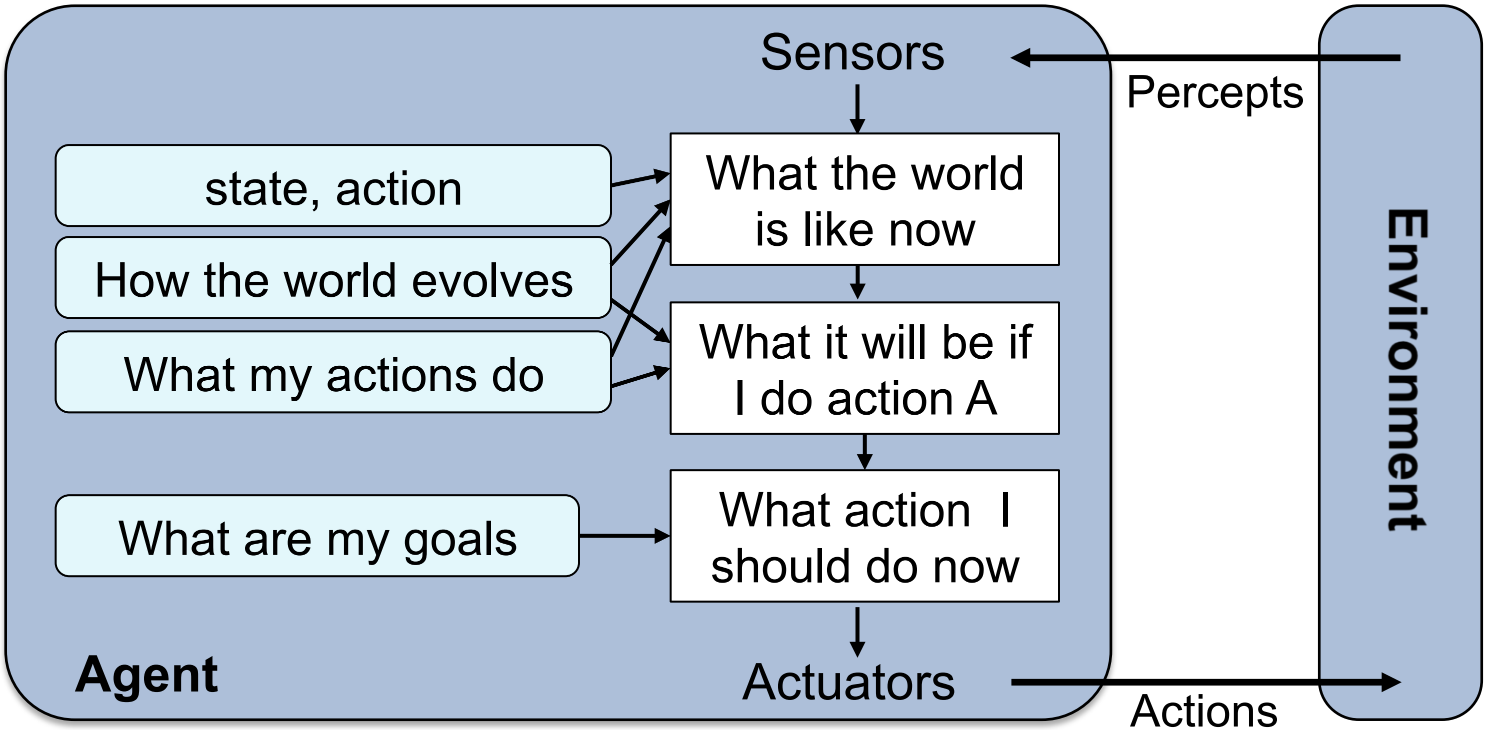

Goal-based agent versus model-based reflex agent

- Less efficient but more flexible as knowledge is explicitly represented
- Goals alone are not sufficient as they do not consider performance
- Taxi driving: faster, cheaper, more reliable, safer.

Model-Based Utility-Based Agent

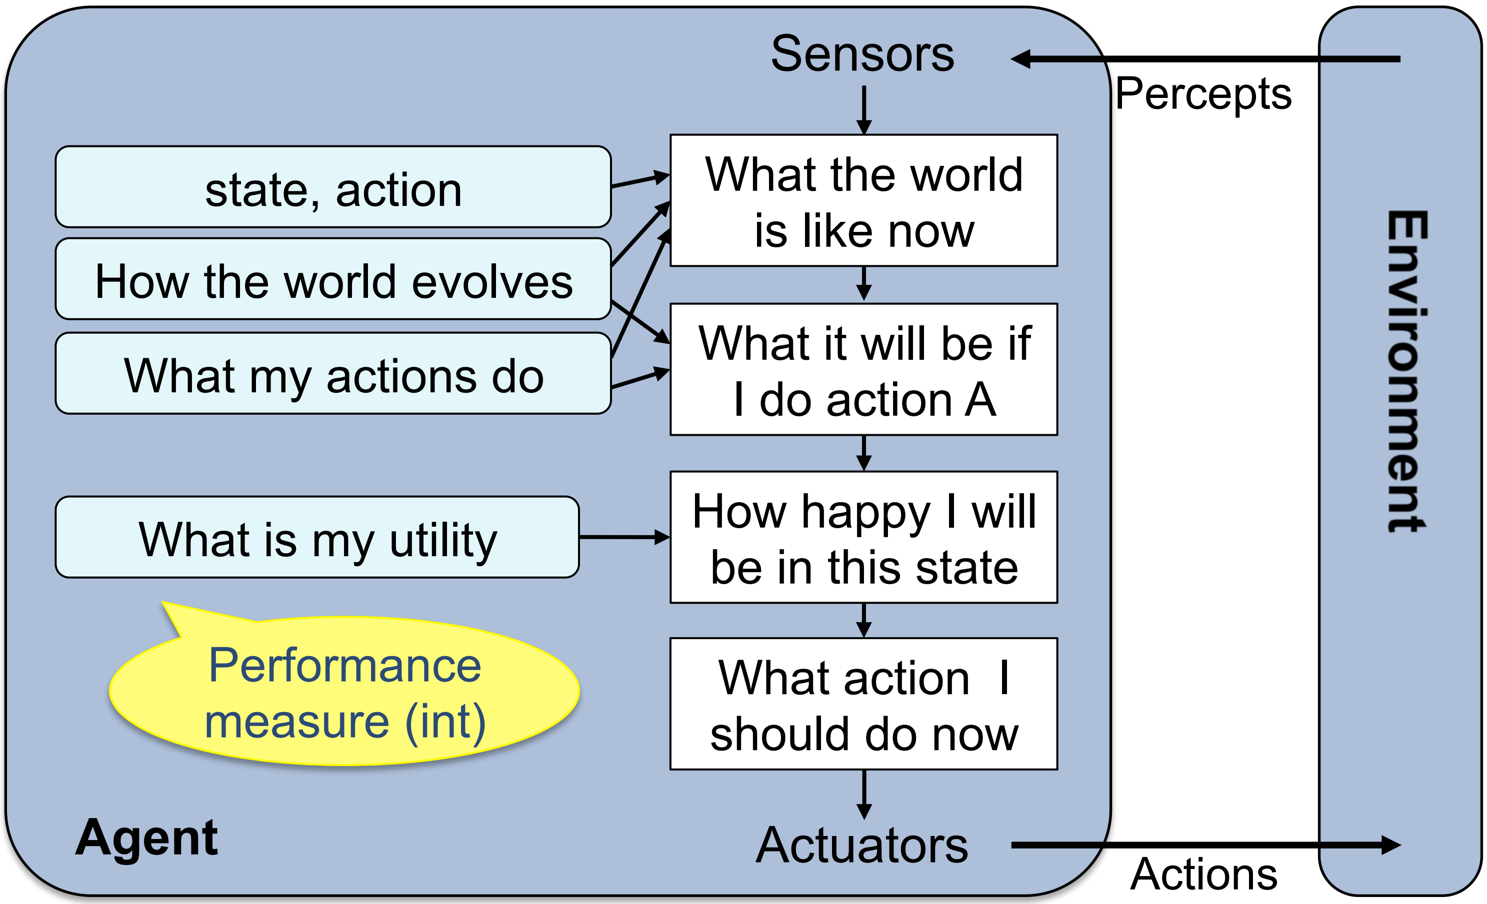

- Goals provide only a crude binary distinction: happy, unhappy
- Utilities provide a more general internalization of performance measure

Model-Based Utility-Based Agent

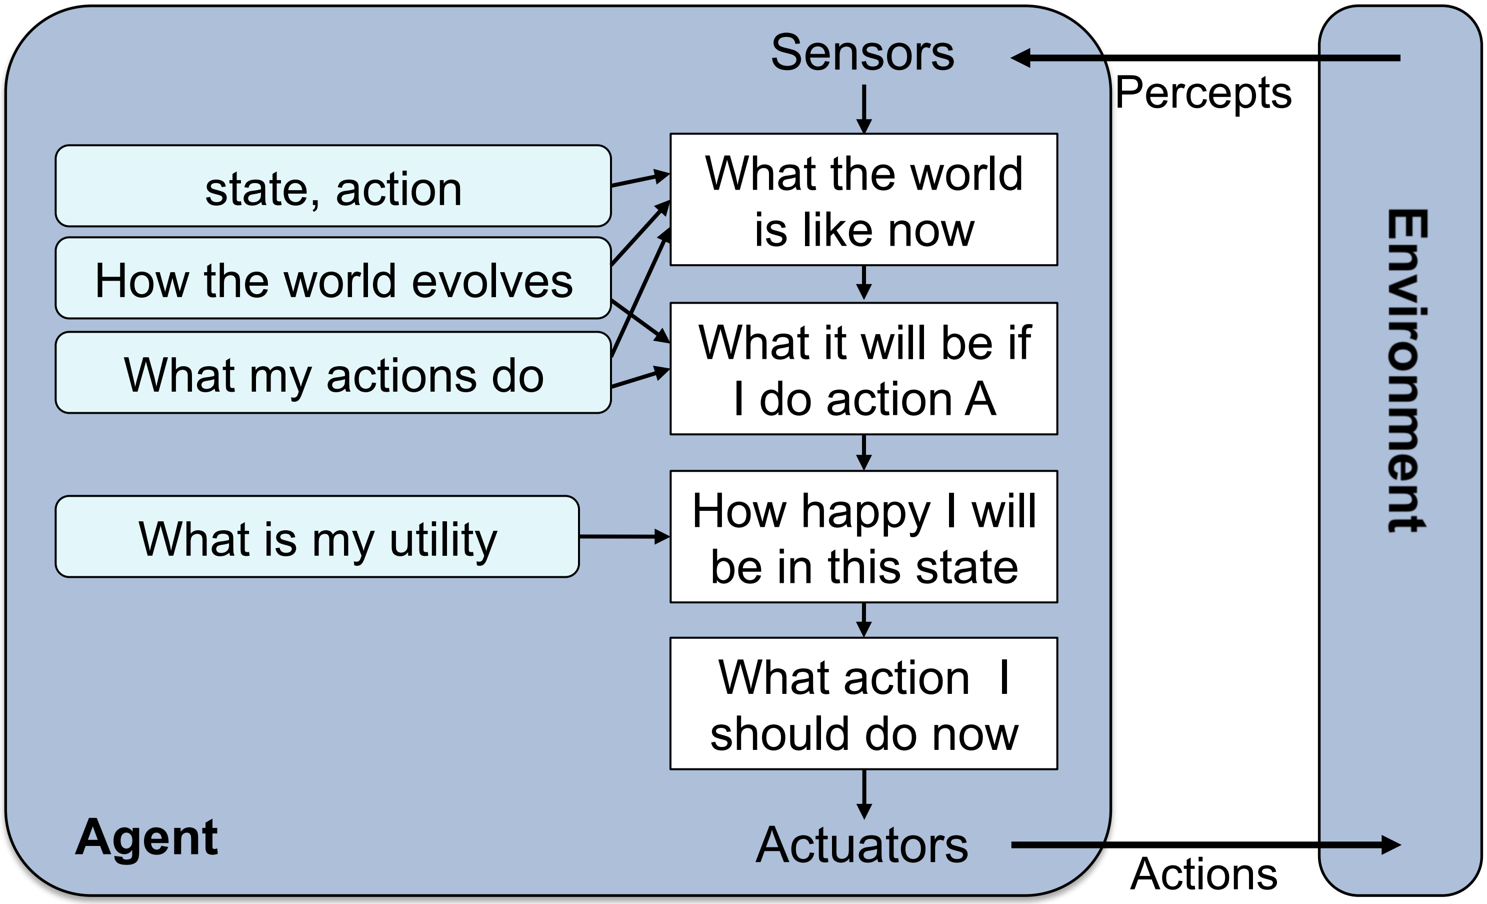

Is it that simple? Just build agents maximizing expected utility?

- Keep track of environment: perception, modeling, reasoning, learning

Summary: Rational Agent

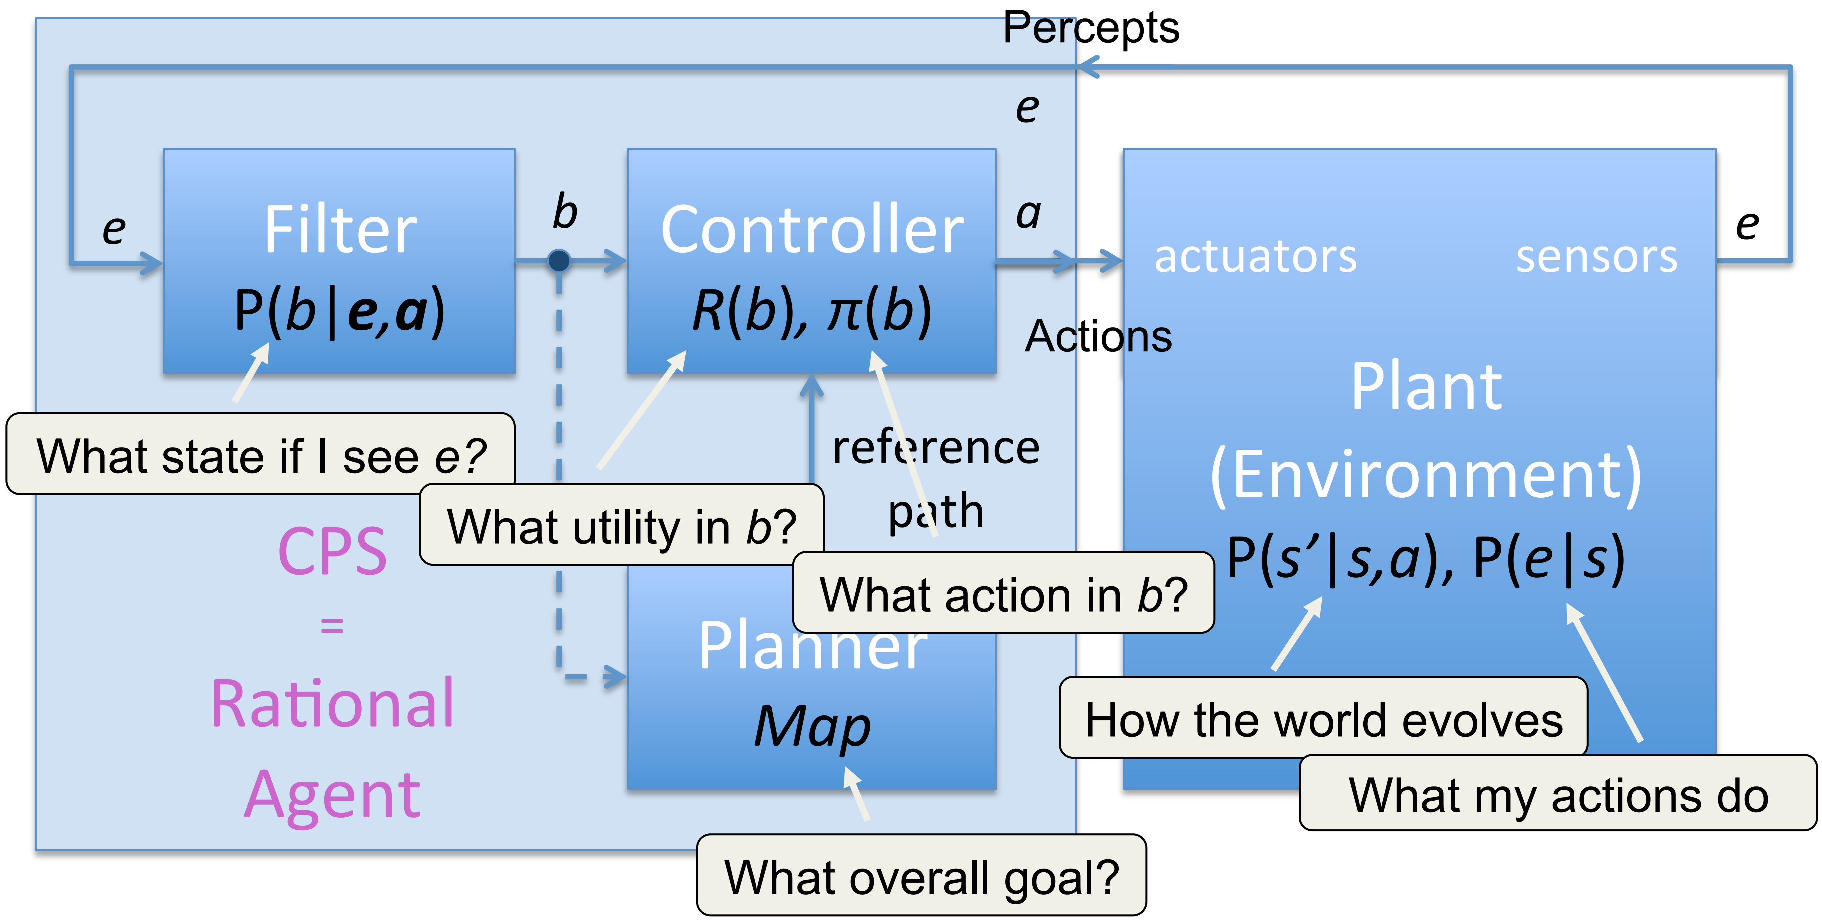

Model-Based Utility-Based Agent

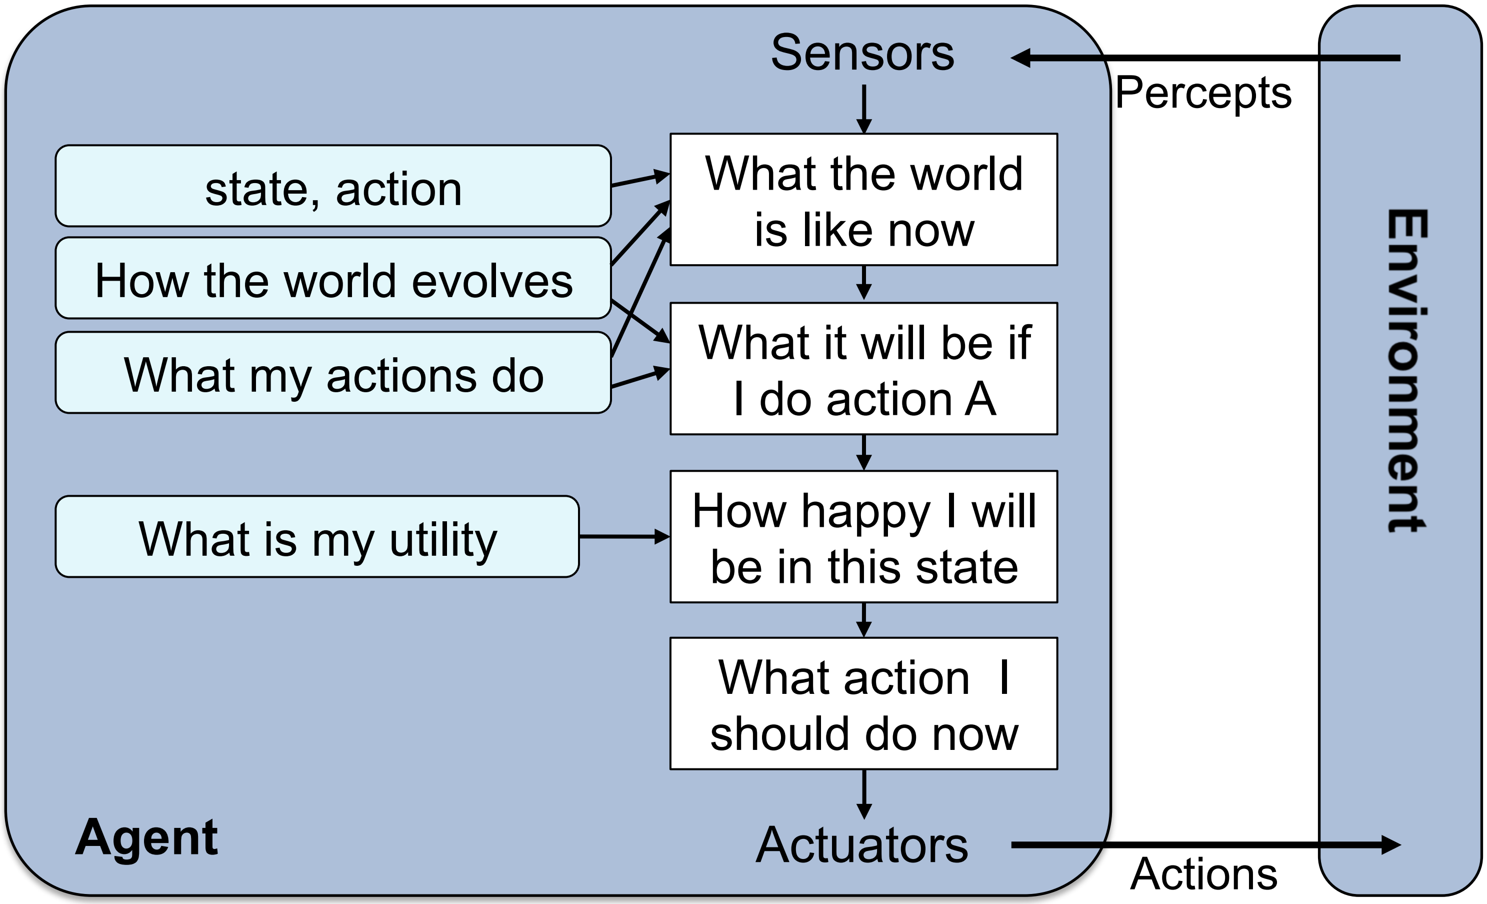

Hoiw does one develop such agents?

- Turing: Manually is too tedious. One should learn them

General Learning Agent

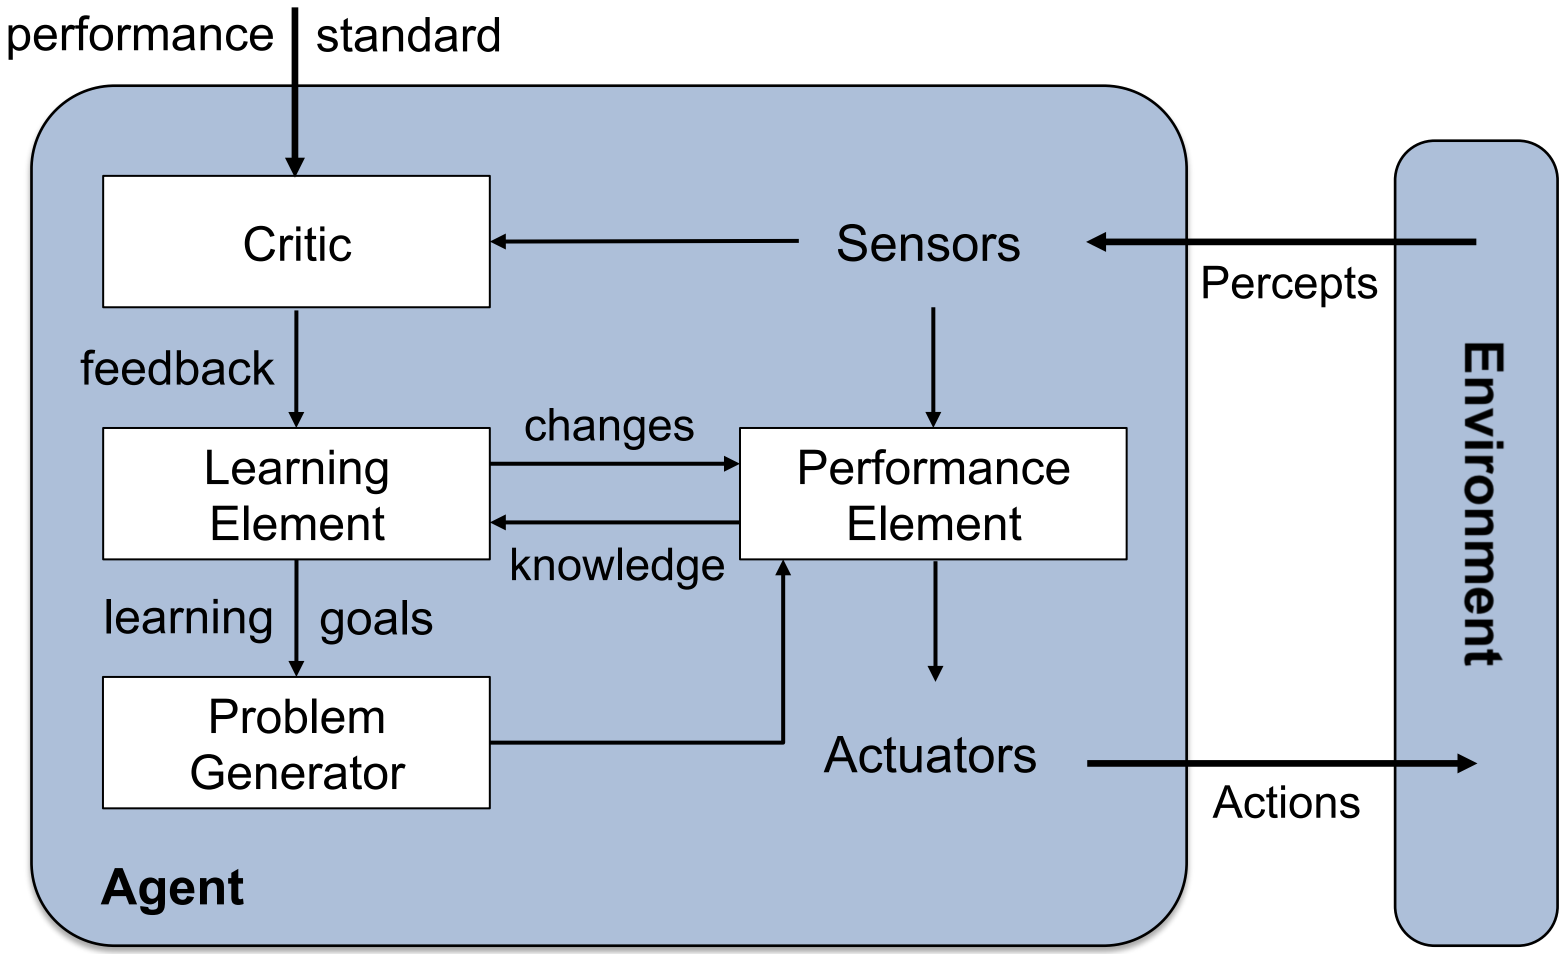

Turing proposes to build learning machines and teach them

- 4 components: Learning, performance, and critic elements, problem gen.

General Learning Agent

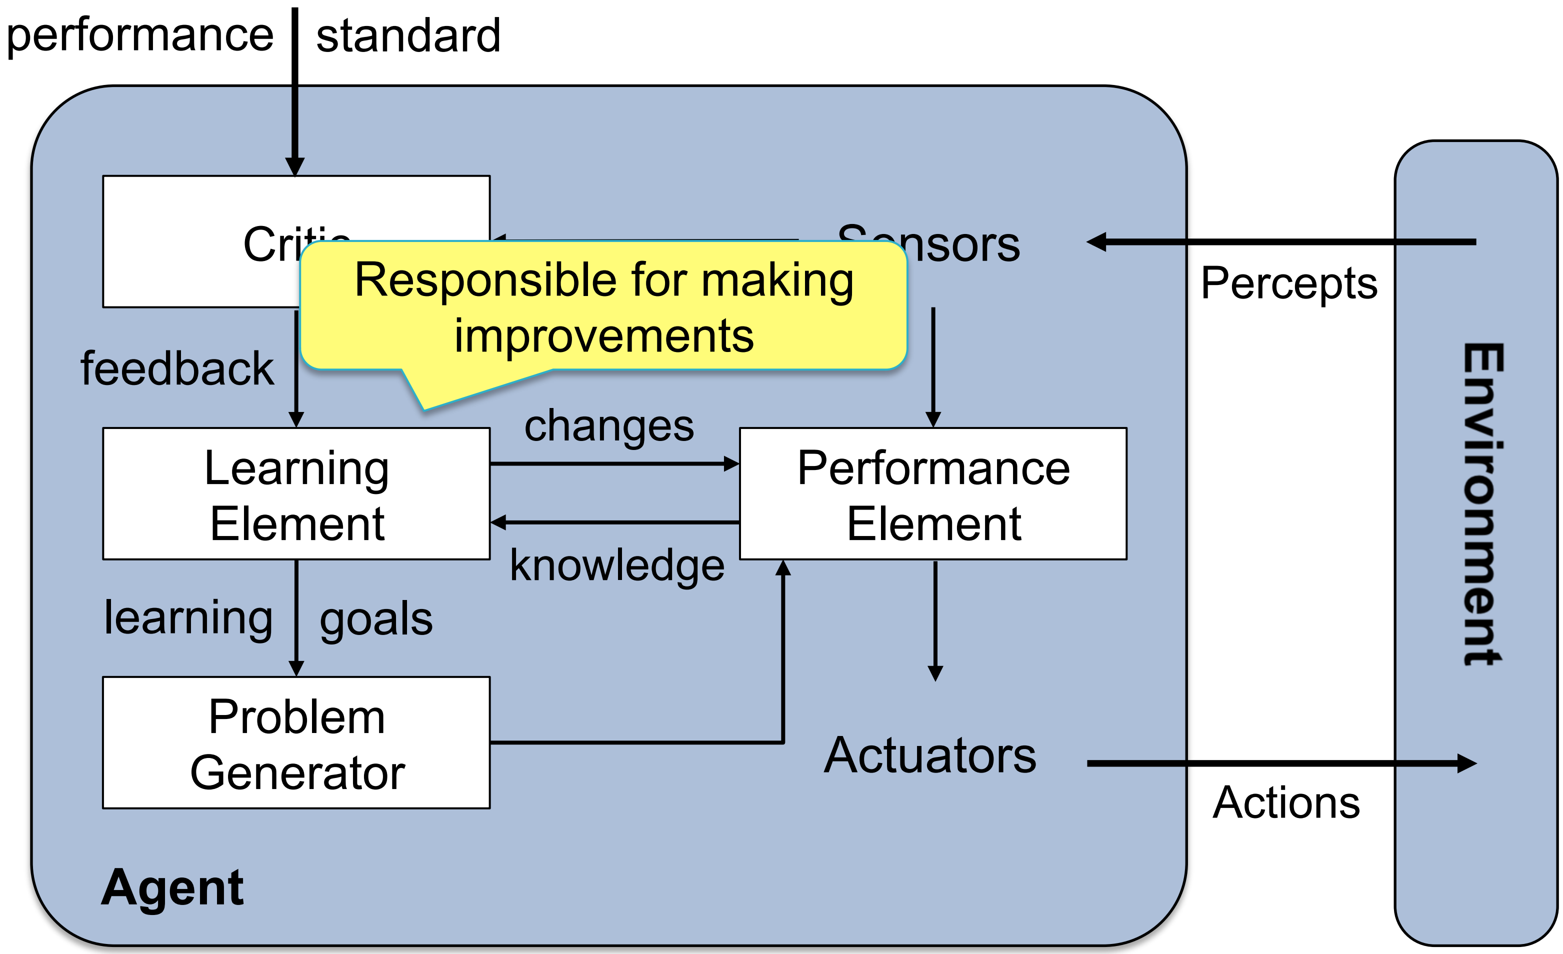

General Learning Agent

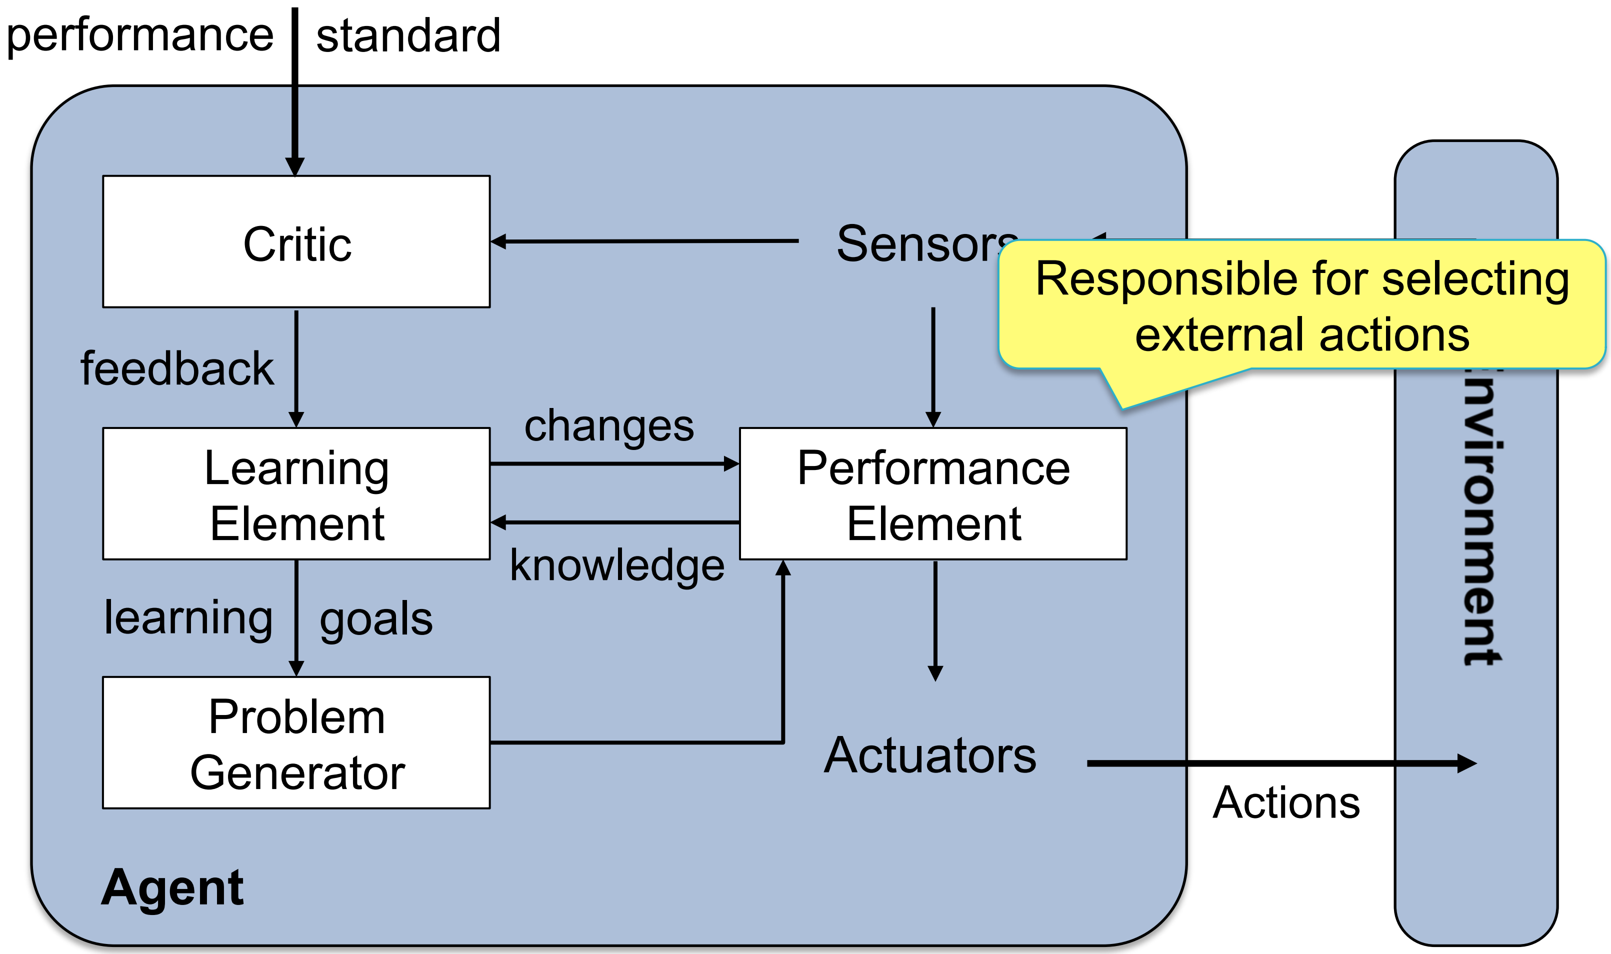

General Learning Agent

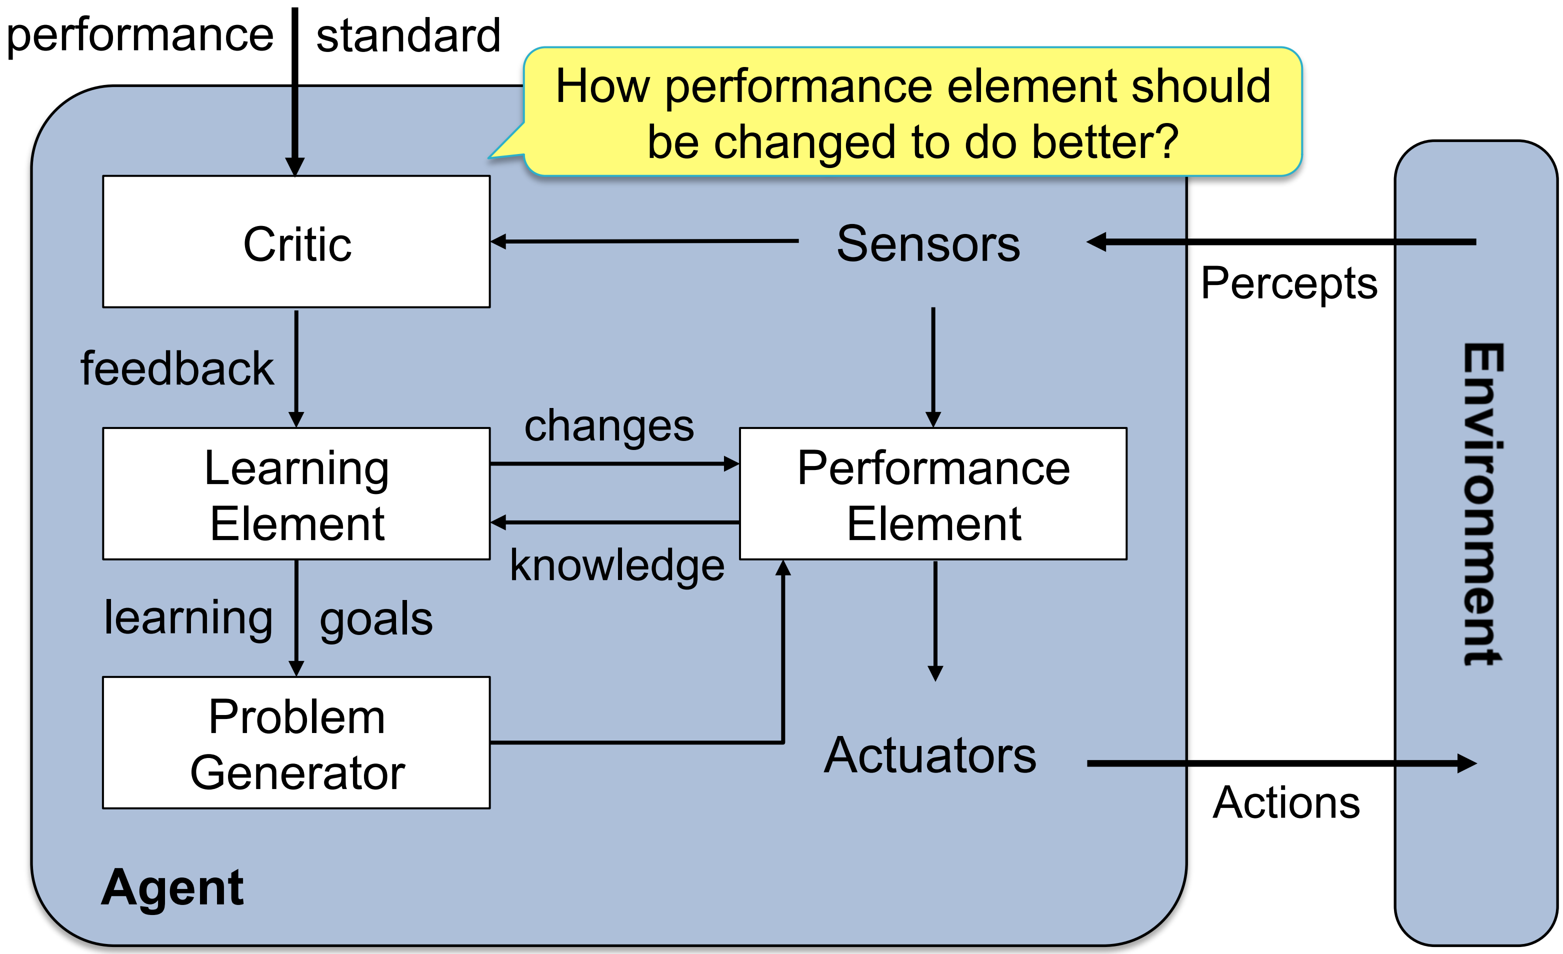

General Learning Agent

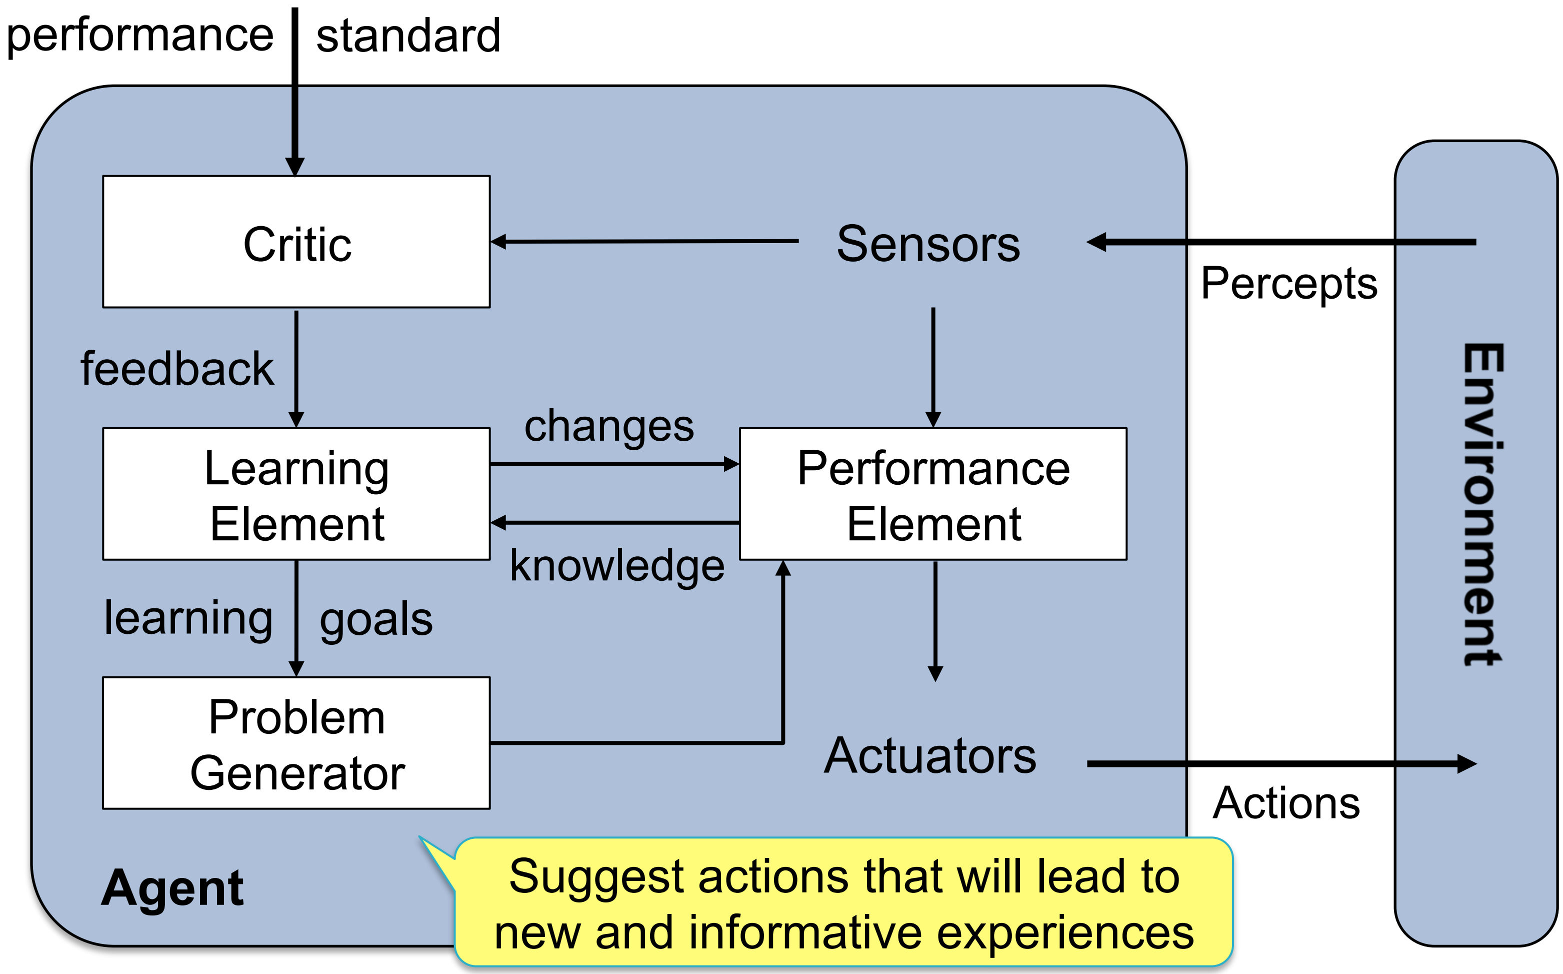

General Learning Agent

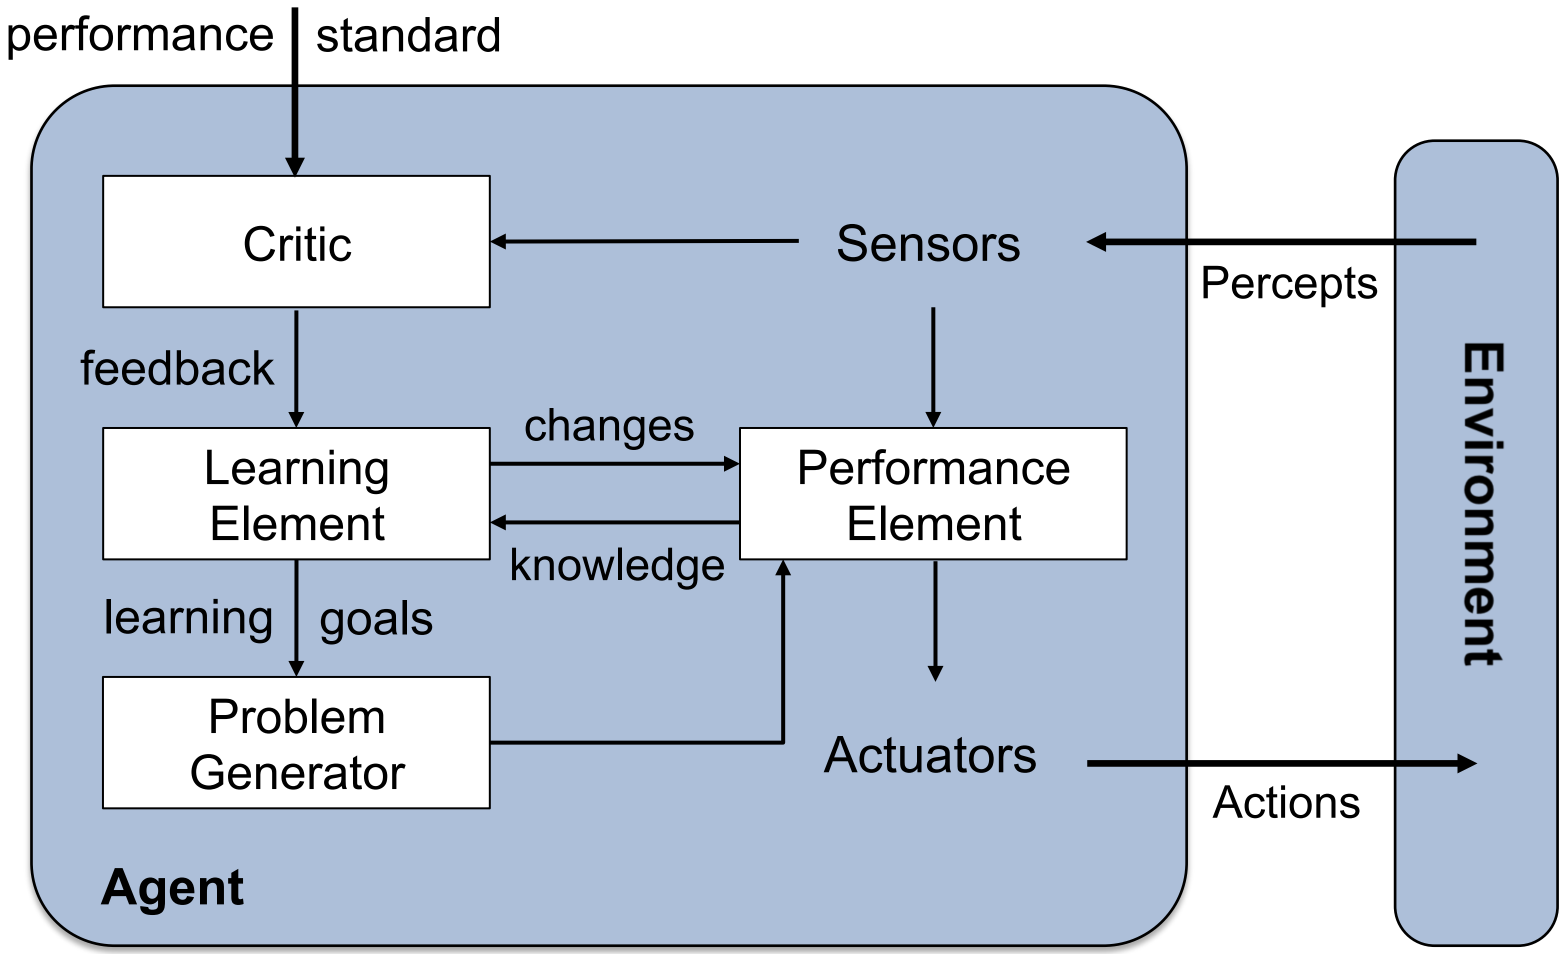

Preferred method of creating agents in many AI areas

- Advantage: Allows the agent to operate in initially unknown environments

A (Short) History of AI

- 1940-1950: Early days
  - 1943: McCulloch & Pitts: Boolean circuit model of brain
  - 1950: Turing's "Computing Machinery and Intelligence"

A (Short) History of AI

- 1940-1950: Early days
  - 1943: McCulloch & Pitts: Boolean circuit model of brain
  - 1950: Turing's "Computing Machinery and Intelligence"
- 1950—70: Excitement: Look, Ma, no hands!
  - 1950s: Early AI programs, including
    - Samuel's checkers program,
    - Newell & Simon's Logic Theorist,
    - Gelernter's Geometry Engine
  - 1956: Dartmouth meeting: "Artificial Intelligence" adopted
  - 1965: Robinson's complete algorithm for logical reasoning
    - E.g., generate plan for driving to the airport
  - 1966: Weizenbaum's Eliza / Turing test

Herb Simon, 1957

It is not my aim to surprise or shock you---but the simplest way I can summarize is to say that there are now in the world machines that think, that learn and that create. Moreover, their ability to do these things is going to increase rapidly until---in a visible future---the range of problems they can handle will be coextensive with the range to which human mind has been applied.

More precisely: within 10 years a computer would be chess champion, and an important new mathematical theorem would be proved by a computer.

Harder than originally thought

- Herb Simon's prediction came true, but after roughly 40 years instead of after 10

Harder than originally thought

- Herb Simon's prediction came true, but after roughly 40 years instead of after 10
- Eliza:
  - " … mother …" → "Tell me more about your family"
  - "I wanted to adopt a puppy, but it's too young to be separated from its mother." → ???

Harder than originally thought

- Herb Simon's prediction came true, but after roughly 40 years instead of after 10
- Eliza:
  - " … mother …" → "Tell me more about your family"
  - "I wanted to adopt a puppy, but it's too young to be separated from its mother." → ???
- 1957: Sputnik
  - Automatic Russian → English translation
  - Famous example:
    - "The spirit is willing but the flesh is weak."
    - E → R → E: "The vodka is strong but the meat is rotten."

Observations

- Need some understanding about the world

Observations

- Need some understanding about the world
- Computational tractability, NP-completeness, exponential scaling.

A (Short) History of AI (ctd)

- 1970—88: Knowledge-based approaches
  - 1969—79: Early development of knowledge-based systems
  - 1980—88: Expert systems industry booms
  - 1988—93: Expert systems industry busts: “AI Winter”

A (Short) History of AI (ctd)

- 1970—88: Knowledge-based approaches
  - 1969—79: Early development of knowledge-based systems
  - 1980—88: Expert systems industry booms
  - 1988—93: Expert systems industry busts: "AI Winter"

- 1988—: Statistical approaches
  - Resurgence of probability, focus on uncertainty
  - General increase in technical depth
  - Agents and learning systems… "AI Spring"?

A (Short) History of AI (ctd)

- 1970—88: Knowledge-based approaches
  - 1969—79: Early development of knowledge-based systems
  - 1980—88: Expert systems industry booms
  - 1988—93: Expert systems industry busts: “AI Winter”
- 1988—: Statistical approaches
  - Resurgence of probability, focus on uncertainty
  - General increase in technical depth
  - Agents and learning systems… “AI Spring”?
- 2000—: Where are we now?

What Can AI Do?

Quiz: Which of the following can be done at present?

Play a decent game of table tennis?

What Can AI Do?

Quiz: Which of the following can be done at present?

- ✓ Play a decent game of table tennis?

What Can AI Do?

Quiz: Which of the following can be done at present?

- ✓ Play a decent game of table tennis?
- Drive safely along a curving mountain road?

What Can AI Do?

Quiz: Which of the following can be done at present?

- ✓ Play a decent game of table tennis?
- ✓ Drive safely along a curving mountain road?

What Can AI Do?

Quiz: Which of the following can be done at present?

- ✓ Play a decent game of table tennis?
- ✓ Drive safely along a curving mountain road?
- Drive safely along Telegraph Avenue?

What Can AI Do?

Quiz: Which of the following can be done at present?

- ✓ Play a decent game of table tennis?
- ✓ Drive safely along a curving mountain road?
- ? Drive safely along Telegraph Avenue?

What Can AI Do?

Quiz: Which of the following can be done at present?

- ✓ Play a decent game of table tennis?
- ✓ Drive safely along a curving mountain road?
- ? Drive safely along Telegraph Avenue?
- Buy a week's worth of groceries on the web?

What Can AI Do?

Quiz: Which of the following can be done at present?

- ✓ Play a decent game of table tennis?
- ✓ Drive safely along a curving mountain road?
- ? Drive safely along Telegraph Avenue?
- ✓ Buy a week's worth of groceries on the web?

What Can AI Do?

Quiz: Which of the following can be done at present?

- ✓ Play a decent game of table tennis?
- ✓ Drive safely along a curving mountain road?
- ? Drive safely along Telegraph Avenue?
- ✓ Buy a week's worth of groceries on the web?
- Buy a week's worth of groceries at Berkeley Bowl?

What Can AI Do?

Quiz: Which of the following can be done at present?

- ✓ Play a decent game of table tennis?
- ✓ Drive safely along a curving mountain road?
- ? Drive safely along Telegraph Avenue?
- ✓ Buy a week's worth of groceries on the web?
- ✗ Buy a week's worth of groceries at Berkeley Bowl?

What Can AI Do?

Quiz: Which of the following can be done at present?

- ✓ Play a decent game of table tennis?
- ✓ Drive safely along a curving mountain road?
- ? Drive safely along Telegraph Avenue?
- ✓ Buy a week's worth of groceries on the web?
- ✗ Buy a week's worth of groceries at Berkeley Bowl?
- Discover and prove a new mathematical theorem?

What Can AI Do?

Quiz: Which of the following can be done at present?

- Play a decent game of table tennis?
- Drive safely along a curving mountain road?
- Drive safely along Telegraph Avenue?
- Buy a week's worth of groceries on the web?
- Buy a week's worth of groceries at Berkeley Bowl?
- Discover and prove a new mathematical theorem?

What Can AI Do?

Quiz: Which of the following can be done at present?

- Play a decent game of table tennis?
- Drive safely along a curving mountain road?
- Drive safely along Telegraph Avenue?
- Buy a week's worth of groceries on the web?
- Buy a week's worth of groceries at Berkeley Bowl?
- Discover and prove a new mathematical theorem?
- Converse successfully with another person for an hour?

What Can AI Do?

Quiz: Which of the following can be done at present?

- ✓ Play a decent game of table tennis?
- ✓ Drive safely along a curving mountain road?
- ? Drive safely along Telegraph Avenue?
- ✓ Buy a week's worth of groceries on the web?
- ✗ Buy a week's worth of groceries at Berkeley Bowl?
- ? Discover and prove a new mathematical theorem?
- ✗ Converse successfully with another person for an hour?

What Can AI Do?

Quiz: Which of the following can be done at present?

- ✓ Play a decent game of table tennis?
- ✓ Drive safely along a curving mountain road?
- ? Drive safely along Telegraph Avenue?
- ✓ Buy a week's worth of groceries on the web?
- ✗ Buy a week's worth of groceries at Berkeley Bowl?
- ? Discover and prove a new mathematical theorem?
- ✗ Converse successfully with another person for an hour?
- Perform a complex surgical operation?

What Can AI Do?

Quiz: Which of the following can be done at present?

- ✓ Play a decent game of table tennis?
- ✓ Drive safely along a curving mountain road?
- ? Drive safely along Telegraph Avenue?
- ✓ Buy a week's worth of groceries on the web?
- ✗ Buy a week's worth of groceries at Berkeley Bowl?
- ? Discover and prove a new mathematical theorem?
- ✗ Converse successfully with another person for an hour?
- ? Perform a complex surgical operation?

What Can AI Do?

Quiz: Which of the following can be done at present?

- Play a decent game of table tennis?
- Drive safely along a curving mountain road?
- Drive safely along Telegraph Avenue?
- Buy a week's worth of groceries on the web?
- Buy a week's worth of groceries at Berkeley Bowl?
- Discover and prove a new mathematical theorem?
- Converse successfully with another person for an hour?
- Perform a complex surgical operation?
- Unload a dishwasher and put everything away?

What Can AI Do?

Quiz: Which of the following can be done at present?

- ✓ Play a decent game of table tennis?
- ✓ Drive safely along a curving mountain road?
- ? Drive safely along Telegraph Avenue?
- ✓ Buy a week's worth of groceries on the web?
- ✗ Buy a week's worth of groceries at Berkeley Bowl?
- ? Discover and prove a new mathematical theorem?
- ✗ Converse successfully with another person for an hour?
- ? Perform a complex surgical operation?
- ? Unload a dishwasher and put everything away?

What Can AI Do?

Quiz: Which of the following can be done at present?

- ✓ Play a decent game of table tennis?
- ✓ Drive safely along a curving mountain road?
- ? Drive safely along Telegraph Avenue?
- ✓ Buy a week's worth of groceries on the web?
- ✗ Buy a week's worth of groceries at Berkeley Bowl?
- ? Discover and prove a new mathematical theorem?
- ✗ Converse successfully with another person for an hour?
- ? Perform a complex surgical operation?
- ? Unload a dishwasher and put everything away?
- Translate spoken Chinese into spoken English in real time?

What Can AI Do?

Quiz: Which of the following can be done at present?

- ✓ Play a decent game of table tennis?
- ✓ Drive safely along a curving mountain road?
- ? Drive safely along Telegraph Avenue?
- ✓ Buy a week's worth of groceries on the web?
- ✗ Buy a week's worth of groceries at Berkeley Bowl?
- ? Discover and prove a new mathematical theorem?
- ✗ Converse successfully with another person for an hour?
- ? Perform a complex surgical operation?
- ? Unload a dishwasher and put everything away?
- ✓ Translate spoken Chinese into spoken English in real time?

What Can AI Do?

Quiz: Which of the following can be done at present?

- Play a decent game of table tennis?
- Drive safely along a curving mountain road?
- Drive safely along Telegraph Avenue?
- Buy a week's worth of groceries on the web?
- Buy a week's worth of groceries at Berkeley Bowl?
- Discover and prove a new mathematical theorem?
- Converse successfully with another person for an hour?
- Perform a complex surgical operation?
- Unload a dishwasher and put everything away?
- Translate spoken Chinese into spoken English in real time?
- Write an intentionally funny story?

What Can AI Do?

Quiz: Which of the following can be done at present?

- ✓ Play a decent game of table tennis?
- ✓ Drive safely along a curving mountain road?
- ? Drive safely along Telegraph Avenue?
- ✓ Buy a week's worth of groceries on the web?
- ✗ Buy a week's worth of groceries at Berkeley Bowl?
- ? Discover and prove a new mathematical theorem?
- ✗ Converse successfully with another person for an hour?
- ? Perform a complex surgical operation?
- ? Unload a dishwasher and put everything away?
- ✓ Translate spoken Chinese into spoken English in real time?
- ✗ Write an intentionally funny story?

Natural Language

- Speech technologies
  - Automatic speech recognition (ASR)
  - Text-to-speech synthesis (TTS)
  - Dialog systems

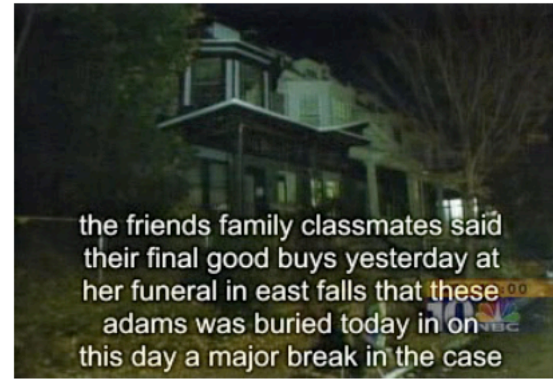

Natural Language

- Speech technologies
  - Automatic speech recognition (ASR)
  - Text-to-speech synthesis (TTS)
  - Dialog systems

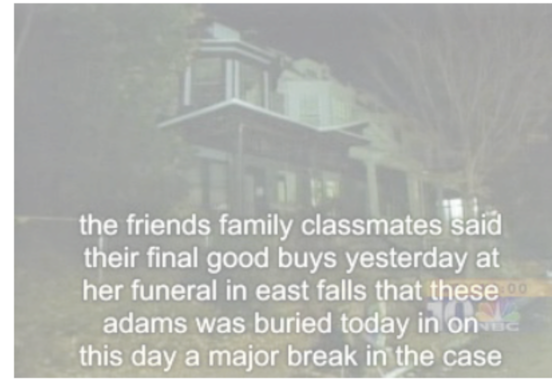

- Language processing technologies
  - Machine translation

Compétitivité : pourquoi l'écart se creuse entre la France et l'Allemagne

La question de la compétitivité de l'économie française mobilise, tout particulièrement par rapport au voisin allemand. Ce sujet est au cœur d'un rapport qui sera remis, jeudi , au ministre de l'industrie, Eric Besson.

- Les Français travaillent moins et sont plus productifs
- Christine Lagarde prévoit au moins 1,5 % de croissance en 2010
- L'inflation a fait son retour en 2010
- La croissance depuis 2007

Competitiveness: why the gap between France and Germany

The issue of competitiveness of the French economy mobilizes, especially compared to neighboring Germany. This topic is in the heart of a report to be presented Thursday to the Minister of Industry, Eric Besson.

- The French work less and are more productive
- Christine Lagarde provides at least 1.5% growth in 2010
- Inflation has made a comeback in 2010
- Growth since 2007

Natural Language

- Speech technologies
  - Automatic speech recognition (ASR)
  - Text-to-speech synthesis (TTS)
  - Dialog systems

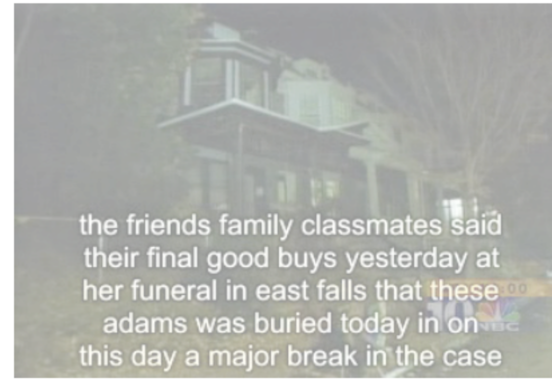

- Language processing technologies
  - Machine translation

Compétitivité : pourquoi l'écart se creuse entre la France et l'Allemagne

La question de la compétitivité de l'économie française mobilise, tout particulièrement par rapport au voisin allemand. Ce sujet est au cœur d'un rapport qui sera remis, jeudi , au ministre de l'industrie, Eric Besson.

- Les Français travaillent moins et sont plus productifs
- Christine Lagarde prévoit au moins 1,5 % de croissance en 2010
- L'inflation a fait son retour en 2010
- La croissance depuis 2007

Competitiveness: why the gap between France and Germany

The issue of competitiveness of the French economy mobilizes, especially compared to neighboring Germany. This topic is in the heart of a report to be presented Thursday to the Minister of Industry, Eric Besson.

- The French work less and are more productive
- Christine Lagarde provides at least 1.5% growth in 2010
- Inflation has made a comeback in 2010
- Growth since 2007

  - Information extraction
  - Information retrieval, question answering
  - Text classification, spam filtering, etc…

Vision (Perception)

- Object and character recognition
- Scene segmentation
- 3D reconstruction
- Image classification

Robotics

- Robotics
  - Part mech. eng.
  - Part AI
  - Reality much harder than simulations!

- Technologies
  - Vehicles
  - Rescue
  - Soccer!
  - Lots of automation…
- In this class:
  - We ignore mechanical aspects
  - Methods for planning
  - Methods for control

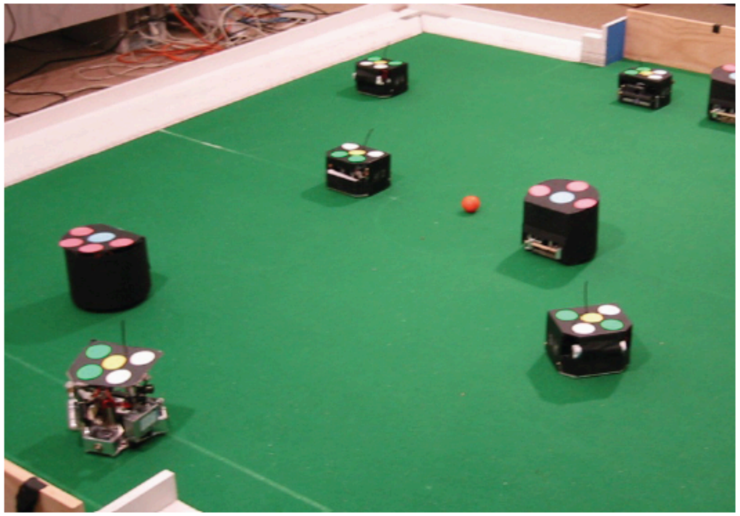

Images from stanfordracing.org, CMU RoboCup, Honda ASIMO sites

Logic

- Logical systems
  - Theorem provers
  - NASA fault diagnosis
  - Question answering

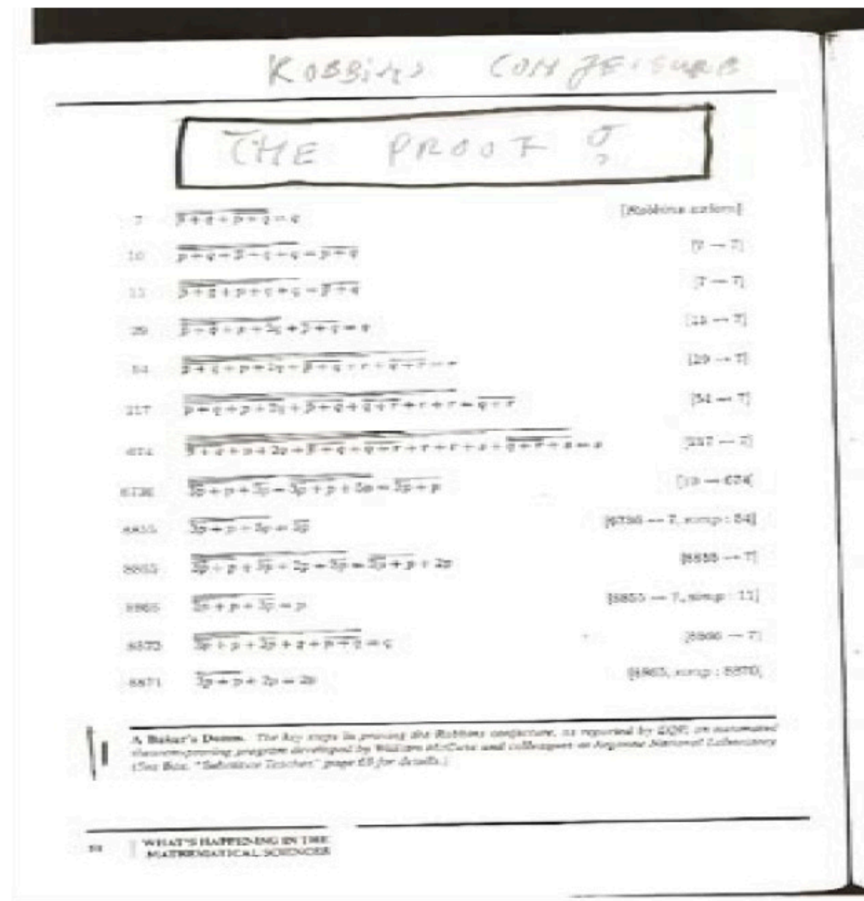

Image from Bart Selman

Logic

- Logical systems
  - Theorem provers
  - NASA fault diagnosis
  - Question answering
- Methods:
  - Deduction systems
  - Constraint satisfaction
  - Satisfiability solvers (huge advances here!)

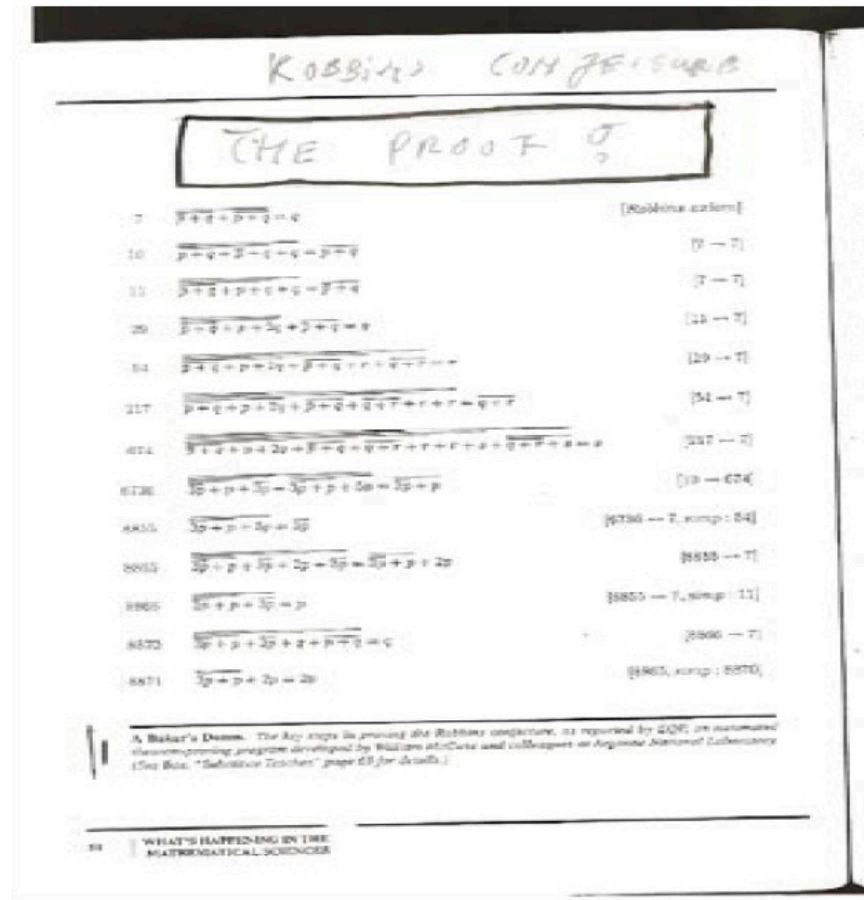

Image from Bart Selman

Game Playing

- May, '97: Deep Blue vs. Kasparov
  - First match won against world-champion
  - “Intelligent creative” play
  - 200 million board positions per second!
  - Humans understood 99.9 of Deep Blue's moves
  - Can do about the same now with a big PC cluster

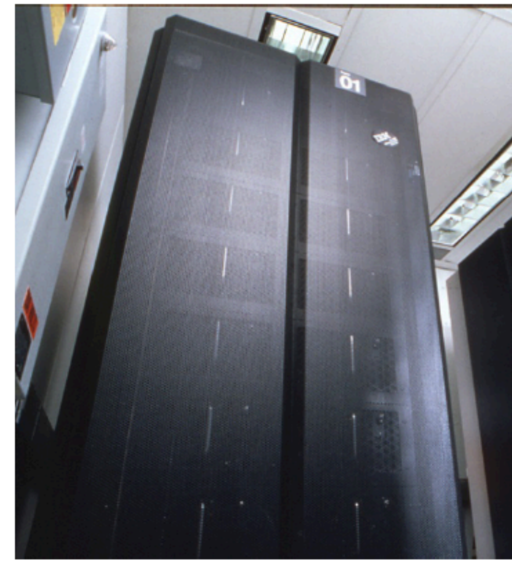

Text from Bart Selman, image from IBM’s Deep Blue pages

Game Playing

- May, '97: Deep Blue vs. Kasparov
  - First match won against world-champion
  - “Intelligent creative” play
  - 200 million board positions per second!
  - Humans understood 99.9 of Deep Blue's moves
  - Can do about the same now with a big PC cluster
- Open question:
  - How does human cognition deal with the search space explosion of chess?
  - Or: how can humans compete with computers at all??

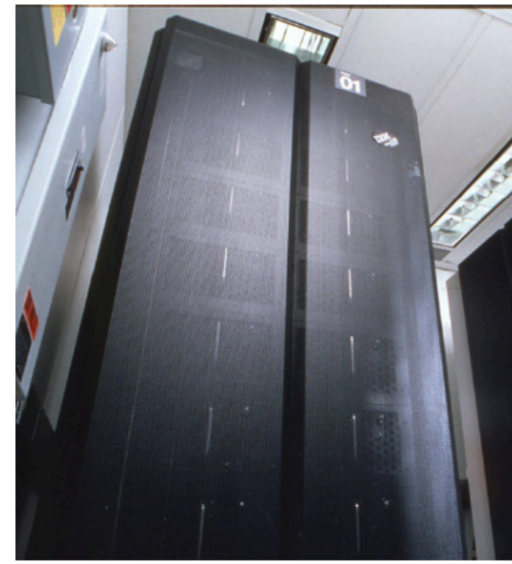

Text from Bart Selman, image from IBM’s Deep Blue pages

Game Playing

- May, '97: Deep Blue vs. Kasparov
  - First match won against world-champion
  - “Intelligent creative” play
  - 200 million board positions per second!
  - Humans understood 99.9 of Deep Blue's moves
  - Can do about the same now with a big PC cluster

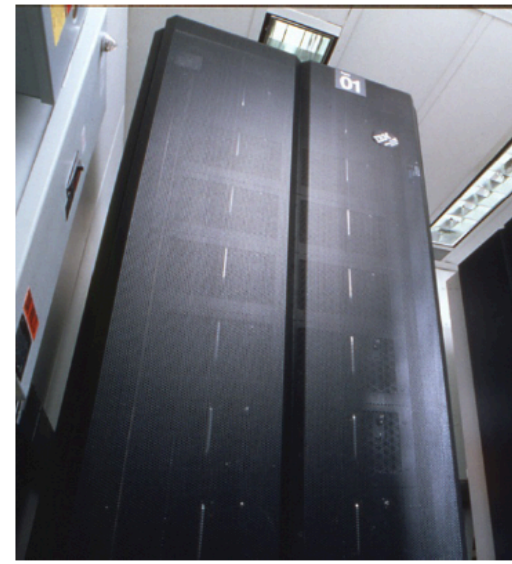

- Open question:
  - How does human cognition deal with the search space explosion of chess?
  - Or: how can humans compete with computers at all??

- 1996: Kasparov Beats Deep Blue

  “I could feel --- I could smell --- a new kind of intelligence across the table.”

Text from Bart Selman, image from IBM’s Deep Blue pages

Game Playing

- May, '97: Deep Blue vs. Kasparov
  - First match won against world-champion
  - “Intelligent creative” play
  - 200 million board positions per second!
  - Humans understood 99.9 of Deep Blue's moves
  - Can do about the same now with a big PC cluster
- Open question:
  - How does human cognition deal with the search space explosion of chess?
  - Or: how can humans compete with computers at all??

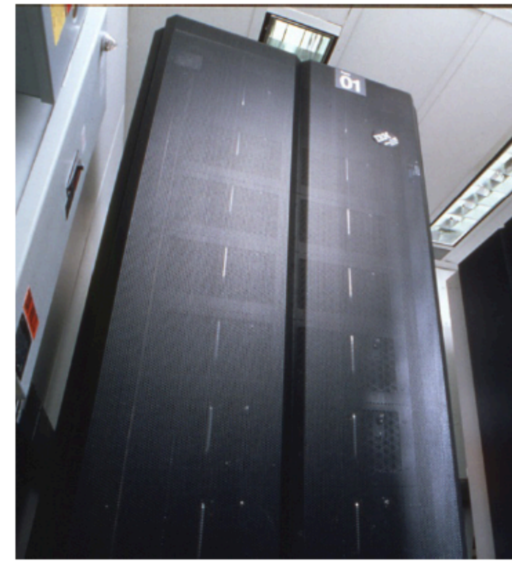

- 1996: Kasparov Beats Deep Blue
  “I could feel --- I could smell --- a new kind of intelligence across the table.”
- 1997: Deep Blue Beats Kasparov
  “Deep Blue hasn't proven anything.”

Text from Bart Selman, image from IBM’s Deep Blue pages

Decision Making

- Scheduling, e.g. airline routing, military

Decision Making

- Scheduling, e.g. airline routing, military
- Route planning, e.g. google maps

Decision Making

- Scheduling, e.g. airline routing, military
- Route planning, e.g. google maps
- Medical diagnosis

Decision Making

- Scheduling, e.g. airline routing, military
- Route planning, e.g. google maps
- Medical diagnosis
- Automated help desks

Decision Making

- Scheduling, e.g. airline routing, military
- Route planning, e.g. google maps
- Medical diagnosis
- Automated help desks
- Fraud detection

Decision Making

- Scheduling, e.g. airline routing, military
- Route planning, e.g. google maps
- Medical diagnosis
- Automated help desks
- Fraud detection
- Spam classifiers

Decision Making

- Scheduling, e.g. airline routing, military
- Route planning, e.g. google maps
- Medical diagnosis
- Automated help desks
- Fraud detection
- Spam classifiers
- Web search engines

Decision Making

- Scheduling, e.g. airline routing, military
- Route planning, e.g. google maps
- Medical diagnosis
- Automated help desks
- Fraud detection
- Spam classifiers
- Web search engines
- Movie and book recommendations

Decision Making

- Scheduling, e.g. airline routing, military
- Route planning, e.g. google maps
- Medical diagnosis
- Automated help desks
- Fraud detection
- Spam classifiers
- Web search engines
- Movie and book recommendations

- … Lots more!