EECS 391: Introduction to AI

Soumya Ray

Website: http://vorlon.case.edu/~sray/eecs391_sp12/index.html

Email: sray@case.edu

Office: Olin 516

Office hours: Tue 2:30-4:00 or by appointment

Announcements

- HW1 will be posted online and on blackboard, due Jan 26
- PA1 will be posted, due Feb 7
- Reading assignment: Chapter 3 of textbook

Today

- What is AI?
- A brief history of AI
- Applications of AI today
- AI in the near future
- Intelligent Agents

What is AI?

- What is “intelligence”?
  - A human-centric view: “Can machines think?”
  - Intelligence is *not* some mystical special quality humans alone are endowed with
  - We don’t know exactly, but it is the result of a computational process

- Modern AI is about the *process* of intelligence, decoupled from the entity performing it

~~What is intelligence?~~

- What does intelligence *allow you to do?*
- Try to characterize intelligence "operationally"
  - I give you a sealed box and some rules for input/output. I claim, "This is an intelligent box."
  - How do you try to prove/disprove my claim?

What does intelligence do?

- Enables *Behavior*
  - That allows *flexible, rational problem solving*
  - That allows *learning*

Intelligence is a capability that enables flexible, adaptive, rational problem solving behavior.

What is AI?

- Intelligent behavior
- Embodied in an *artifact*
  - Something made by humans
  - Usually we take this to mean a digital computer

- Term “artificial intelligence” created by John McCarthy in 1956

Human Intelligence and modern AI

- Human intelligence can be used to identify a set of problems that could be a part of an intelligent system's repertoire of problems
- Several parts of AI are involved in simulating human intelligence

A Closer Look

- Person A tells person B, “Please pick up the pencil on the desk.” Person B does this.
  - What just happened?

A Closer Look (2)

- Person B hears a sequence of sounds, which have to be converted to *words* (speech recognition)

- The *sentence structure* must be recovered in order to know what needs to be done to whom (natural language processing)

A Closer Look (3)

- Person B must then *look around to identify* (a) the current state of the environment (desk, etc) and (b) the correct "pencil" (computer vision and classification)

- The facts about the pencil and its environment are *converted into a representation* that can be reasoned with (knowledge representation and reasoning)

A Closer Look (4)

- Given the current state of the environment and the "goal" (i.e. the pencil should be picked up), person B then *moves their hand to achieve the goal* (planning and reinforcement learning)
- The next time such a request (or a similar one) is made, we may expect person B to respond more quickly, having *learned to do this task* (machine learning)

A Closer Look (5)

- Finally, if we want to *build a physical system* to emulate this behavior, we also need robotics

AI subfields

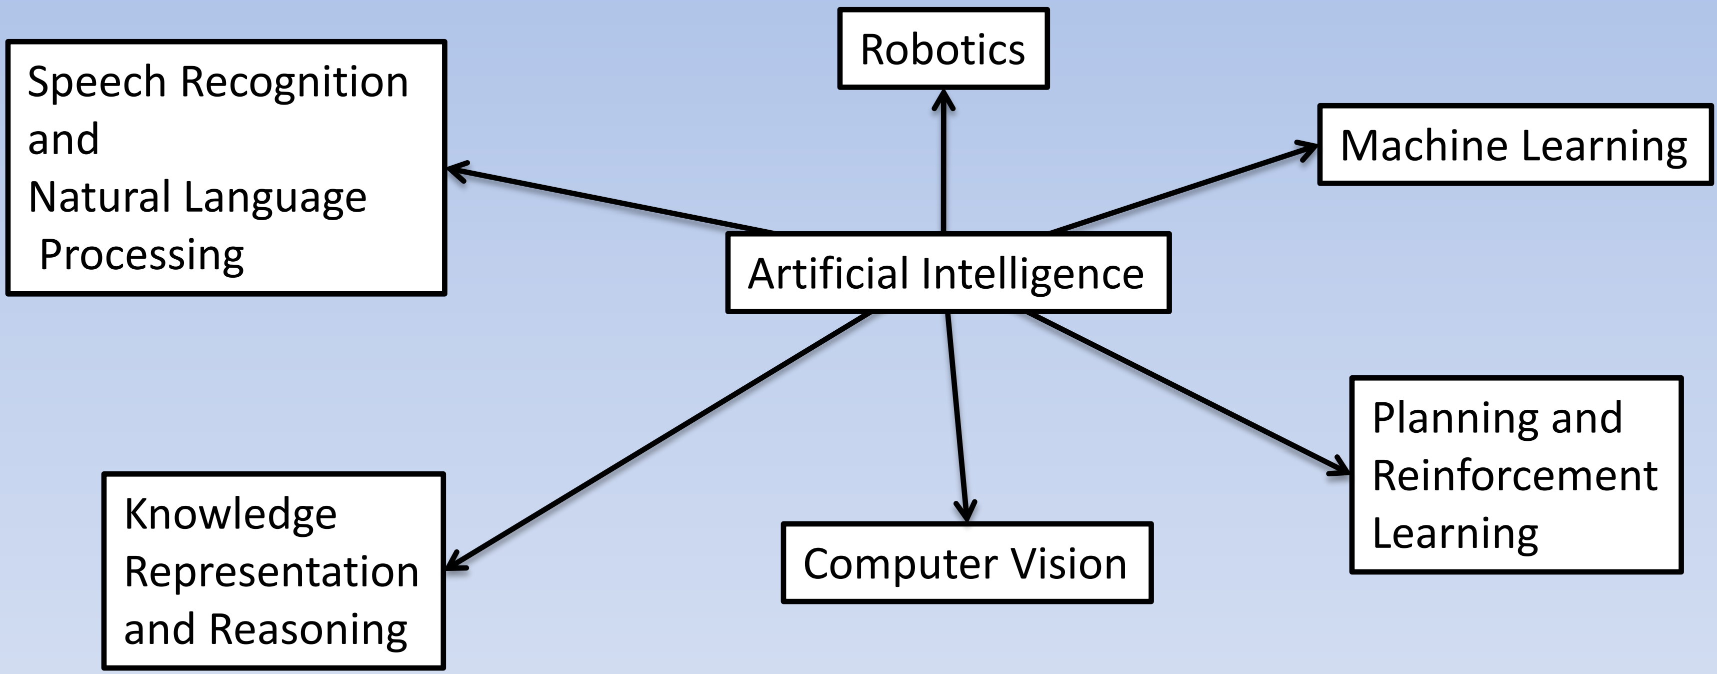

A brief history of AI

- Stage 1---Euphoria (1950-1965)
  - Computational models of the neuron proposed
  - Simple problem-solving demonstrated (checkers playing program, Logic Theorist)
- Stage 2---Reality (1965-1975)
  - Strong limits discovered on what problems could be solved with simple models
  - Ideas did not generalize to significantly different problems or scale to realistic problems

A brief history of AI (2)

- Stage 3---Resurgence (1975-1990s)
  - By combining human expertise with AI principles, rule-based systems started to solve complex problems
  - Neural networks made a comeback, with better understanding of underlying theory

A brief history of AI (3)

- Stage 4---Steady Progress (1990s-Now)
  - AI has matured
  - Concepts are mathematically grounded, ideas are scientifically evaluated
  - Very interdisciplinary
    - Uses tools from logic, probability and statistics, operations research along with computer science
    - AI technologies are used in numerous industries, from vacuum cleaners to Mars rovers
  - Many significant challenges remain, but progress is being made

"If you invent a breakthrough in artificial intelligence, so machines can learn, that is worth 10 Microsofts."

Bill Gates, at a speech in MIT in 2004


Poster by
AAAI
Association for the Advancement of Artificial Intelligence
AI magazine
Poster development supported in part by
Research
YAHOO! RESEARCH
Optimizing Paths & Flows
AI in Art
AI in Music
Autonomous Vehicles & Safety
Robot Guides & Assistants
Augmenting Cognition
Humanoid Robots
AI & Creative Expression
Planning
Plan Recognition
Intelligent Tutoring
Ubiquitous Computing
Social Computing
AI HISTORY TOUR
Gesture Recognition
Multimodal Interfaces
Mixed-Initiative Collaboration
Autonomous Space Exploration
AI and Preferences, Media & Entertainment
Home Robotics
Descartes
Aristotle
Leibniz
Whitehead
Lovelace
Russell
Turing
Robots For Education
Robotic Surgery
Diagnosis
Drug Design
Scientific Discovery
AI LAB
Assistive Technology
MEDICAL CENTER
Search & Retrieval
Security & Privacy
User Modeling
Vehicle Navigation
Handwriting & Sketch Recognition
Recommender Systems & Question Answering
BOOKS & COFFEE
Communications Triage
Machine Translation
Ecocomputing
See the AI timeline and more at
www.aaai.org/AILandscape
The AI Landscape
David Leake, Indiana University, Poster Development Committee Chair
Poster Design: Giacomo Marchesi, www.GiacomoMarchesi.com

Applications (1)

- Speech Recognition
  - "Please describe your problem in a few words"

- Board Games
  - Chess (Deep Blue beat Kasparov in 1997)
  - Checkers ("Checkers is Solved", Schaeffer et al, Science 2007)

Applications (2)

- Web search
- Recommender systems (Netflix prize)
- Medical diagnosis systems
- “Smart” user interfaces
- HSTS planner used by Hubble Space Telescope
- Chatbots (Loebner prize)

Applications (3)

- Robotics (http://www.youtube.com/watch?v=a9r4bvChWFc, http://video.google.com/videoplay?docid=-3227236507141963827&hl=en)

Applications of AI Today (4)

- http://video.google.com/videoplay?docid=-8274817955695344576&hl=en
- http://www.youtube.com/watch?v=lULl63ERek0
- urbanchallenge.case.edu

In the near future…

- Computer processing power is growing rapidly
- In a decade or two, we will have computers that are roughly hardware equivalent to the human brain

In the near future…

- Autonomous vehicles
- Artificial characters in video games/ online worlds

- But still some time away from a flexible, integrated intelligent system

Summary

- We learned about:
  - What AI is about and its subfields
  - A brief history of the field
  - Some current applications of AI

Intelligent Agents (Chapter 2)

Basic Agent Architecture

- Something that interacts with the world

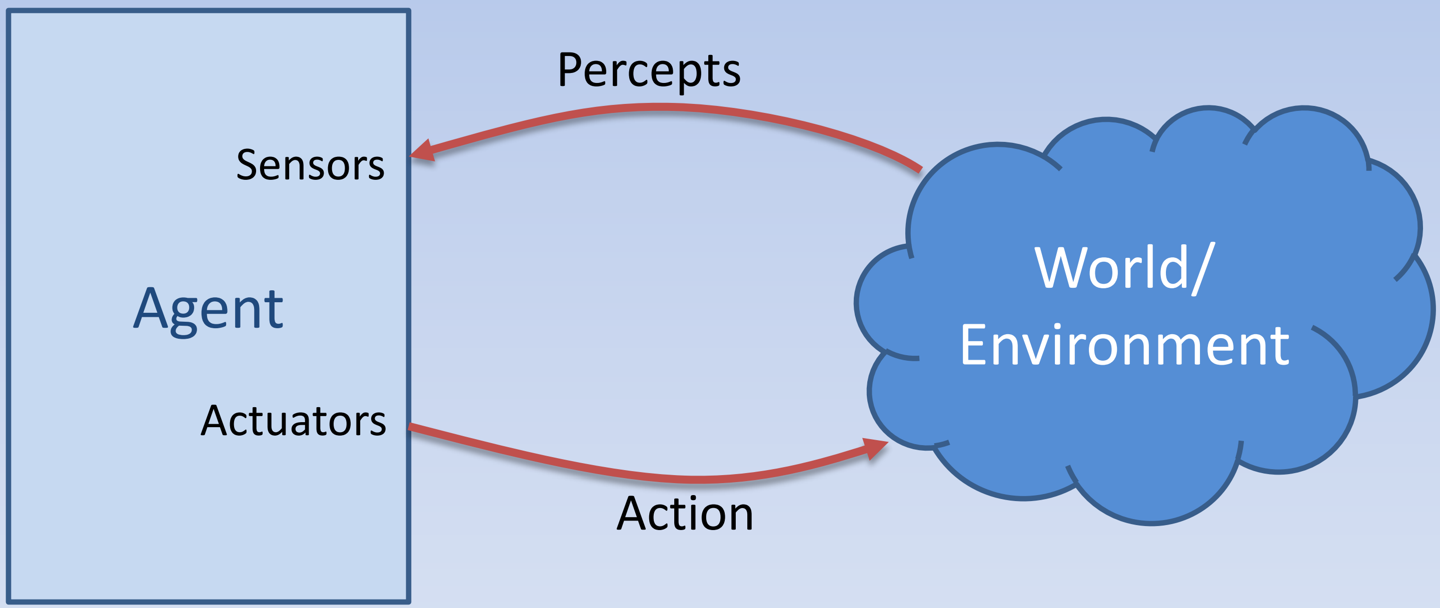

“Agent Function” A: Percept Sequences $\rightarrow$ Actions

Examples

- Chess-playing agent
  - Sensors?
  - Actuators?

- Autonomous vehicle agent
  - Sensors?
  - Actuators?

Performance Measures

- Agents are usually "goal-based", i.e. they are designed to achieve certain things in the world
  - Chess-playing agent?
  - Autonomous vehicle agent?
- These are often encoded using a "performance measure"
  - A function that maps a (percept, action) sequence to a real number
  - Externally imposed

Rational Agents

- A rational agent is one whose agent function always acts to maximize its performance measure, given its percept sequence until the current moment

Example

- Suppose the chess-playing agent is given the performance measure: every time it makes a move, it gets −1 point
  – What will a rational agent do?
- Suppose the autonomous vehicle agent is given the measure: +1 point every second without a collision
  – What might a rational agent do?

Information Gathering

- Agents may be used in unknown environments that we didn't anticipate when designing them

- In this case, it's a good idea to design the performance measure to reward the agent for collecting information about its environment along with achieving the goal

“PEAS” description

- To design the (rational) agent program we need four things:
  - The **P**erformance measure
  - A description of the **E**nvironment
  - A description of what **A**ctions the agent has
  - A description of what **S**ensors the agent has

Examples

- Chess-playing agent PEAS?
- Autonomous vehicle agent PEAS?

Types of Environments

Type	Definition
Fully observable (vs. *Partially Observable*)	Agent's sensors present complete, accurate picture of the world
Deterministic (vs. *Stochastic*)	The next state of the world is completely determined by current state and agent's action
Episodic (vs. *Sequential*)	Agent's current action does not affect future actions
Static (vs. *Dynamic*)	The world does not change until the agent takes an action
Discrete (vs. *Continuous*)	States, percepts and actions are discrete
Single Agent (vs. *Multiagent*)	The world has only one agent in it

Example

- Environment of chess-playing agent?

- Environment of autonomous vehicle agent?

Types of Agents

- Simple Reflex Agents
- Model-based Reflex Agents
- Goal-based Agents
- Utility-based Agents

Simple Reflex Agents

- Agent function maps current state directly to action

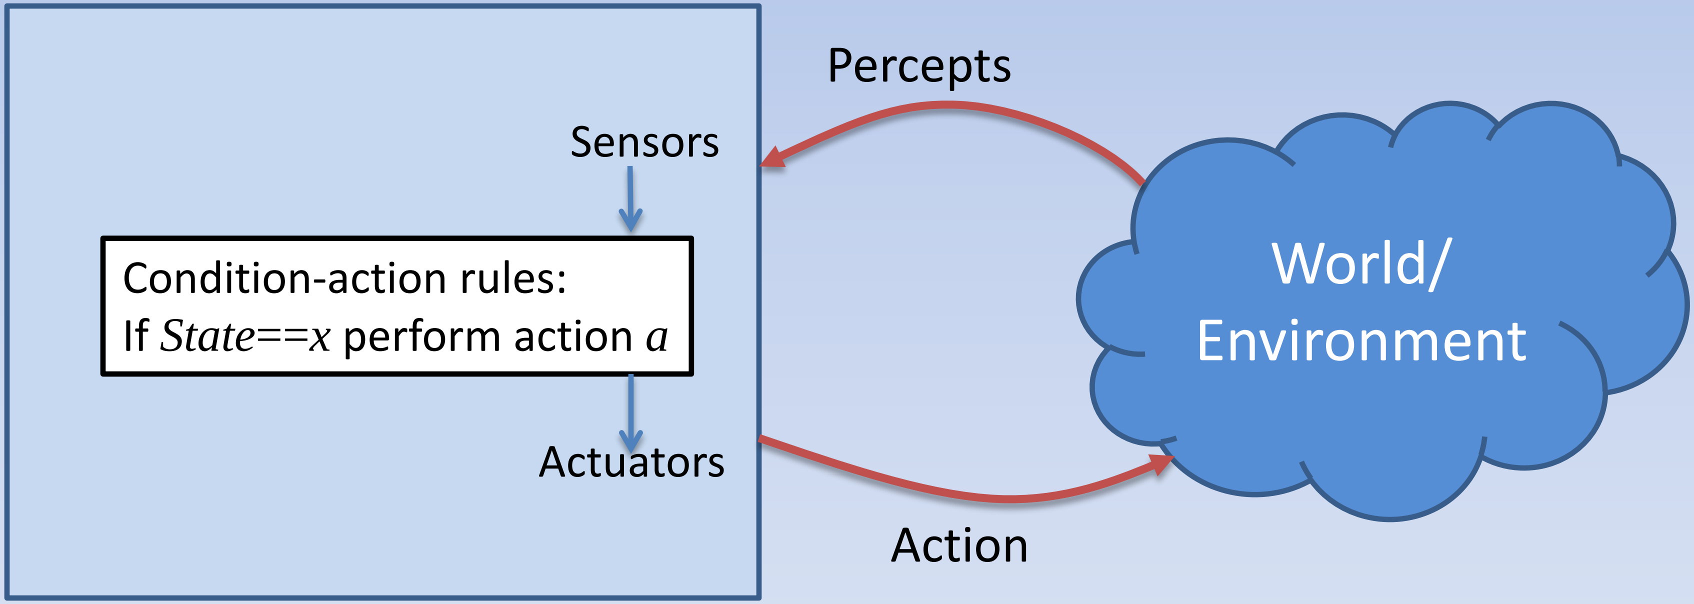

Would this work for the chess agent? For the vehicle agent?

Model-based Reflex Agent

- Maps current percept and *"world model"* to action

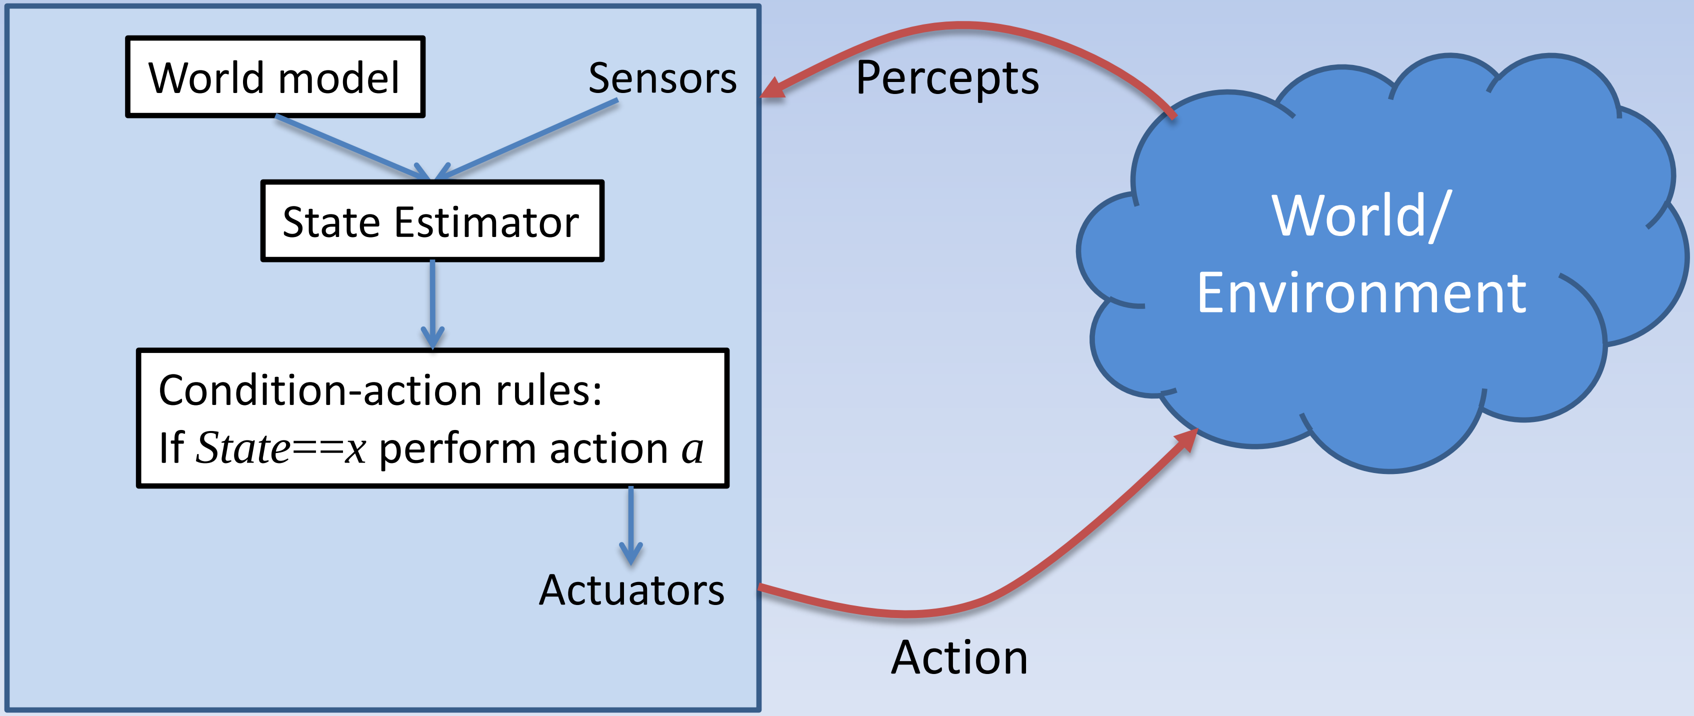

Goal-based Agent

- Maps current percept and knowledge of current goal to action

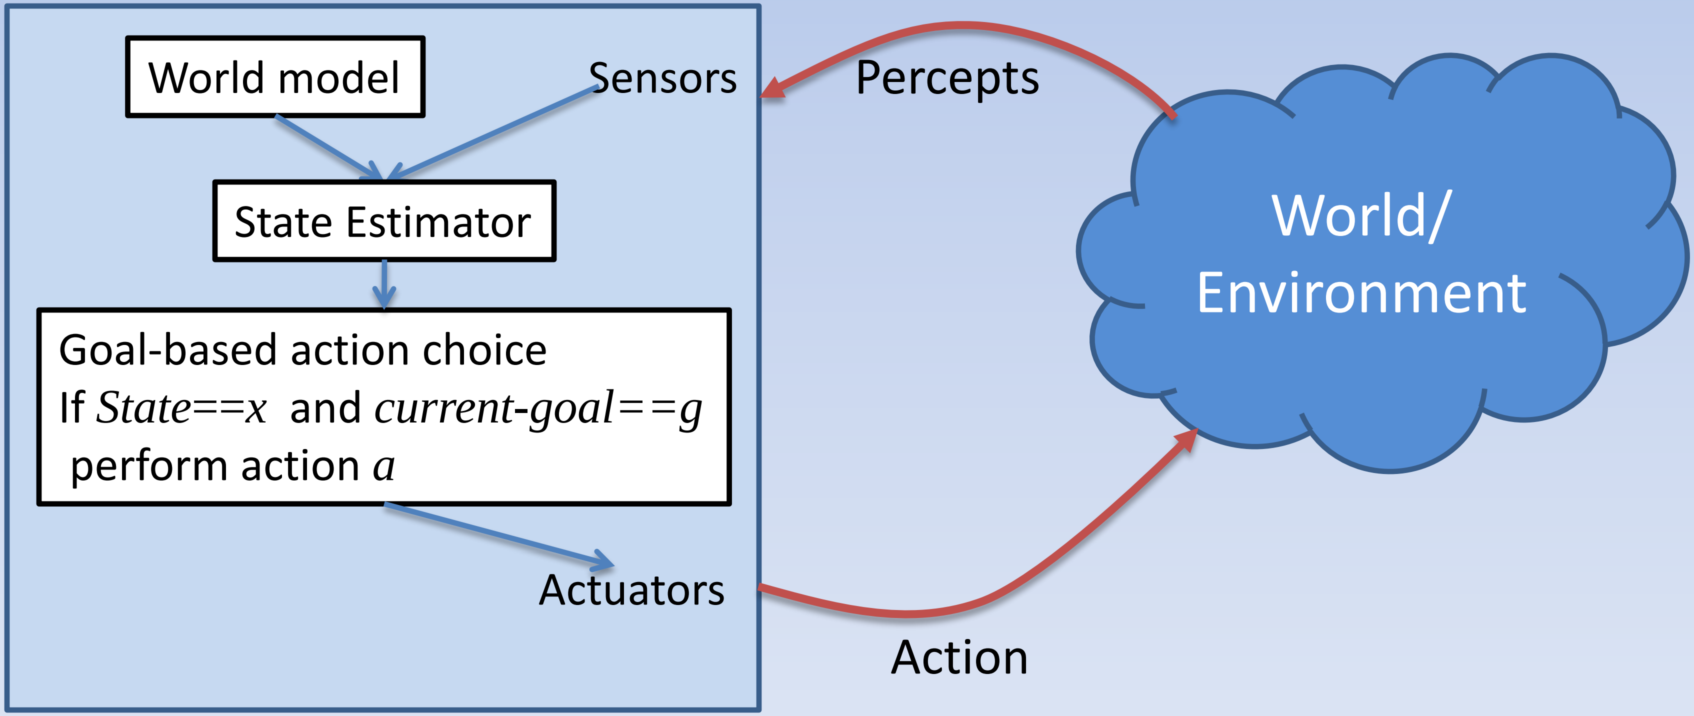

Utility-based Agent

- Instead of binary goal, an agent could have a fine-grained notion of how useful certain states are

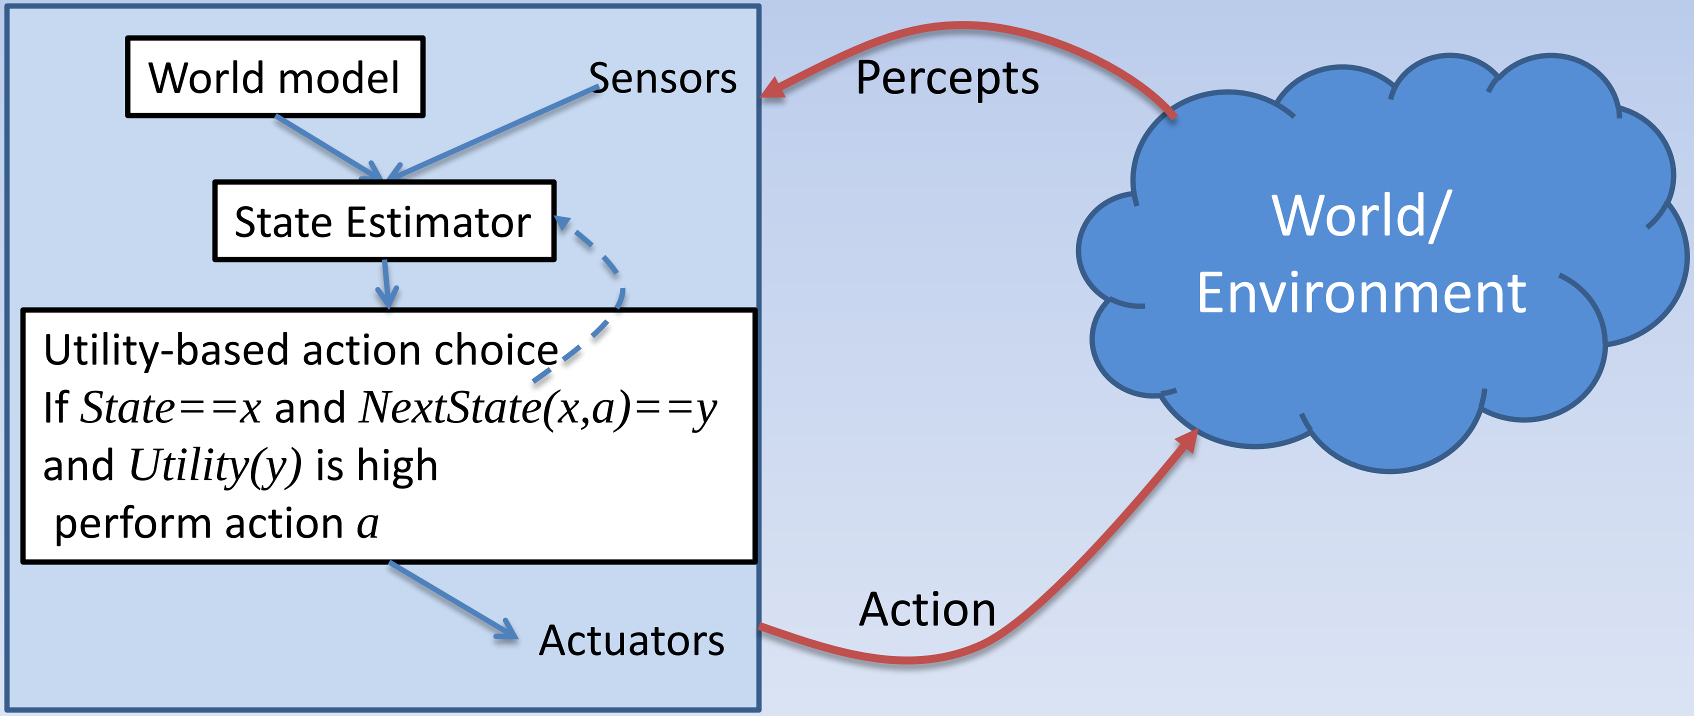

Learning Agents

- Who writes all these rules?
  - Or designs the agent function?
  - Could be quite complex!

- Could we have the agent *learn* the agent function on its own?
  - Would also help in unknown environments…

General Architecture

- Before, we had fixed rules (or functions)
  - We'll call this the "performance element"
- Now we need to add something that generates those functions, based on its (partial) knowledge of the environment and any feedback it receives
  - We'll call this the learning element

General Architecture

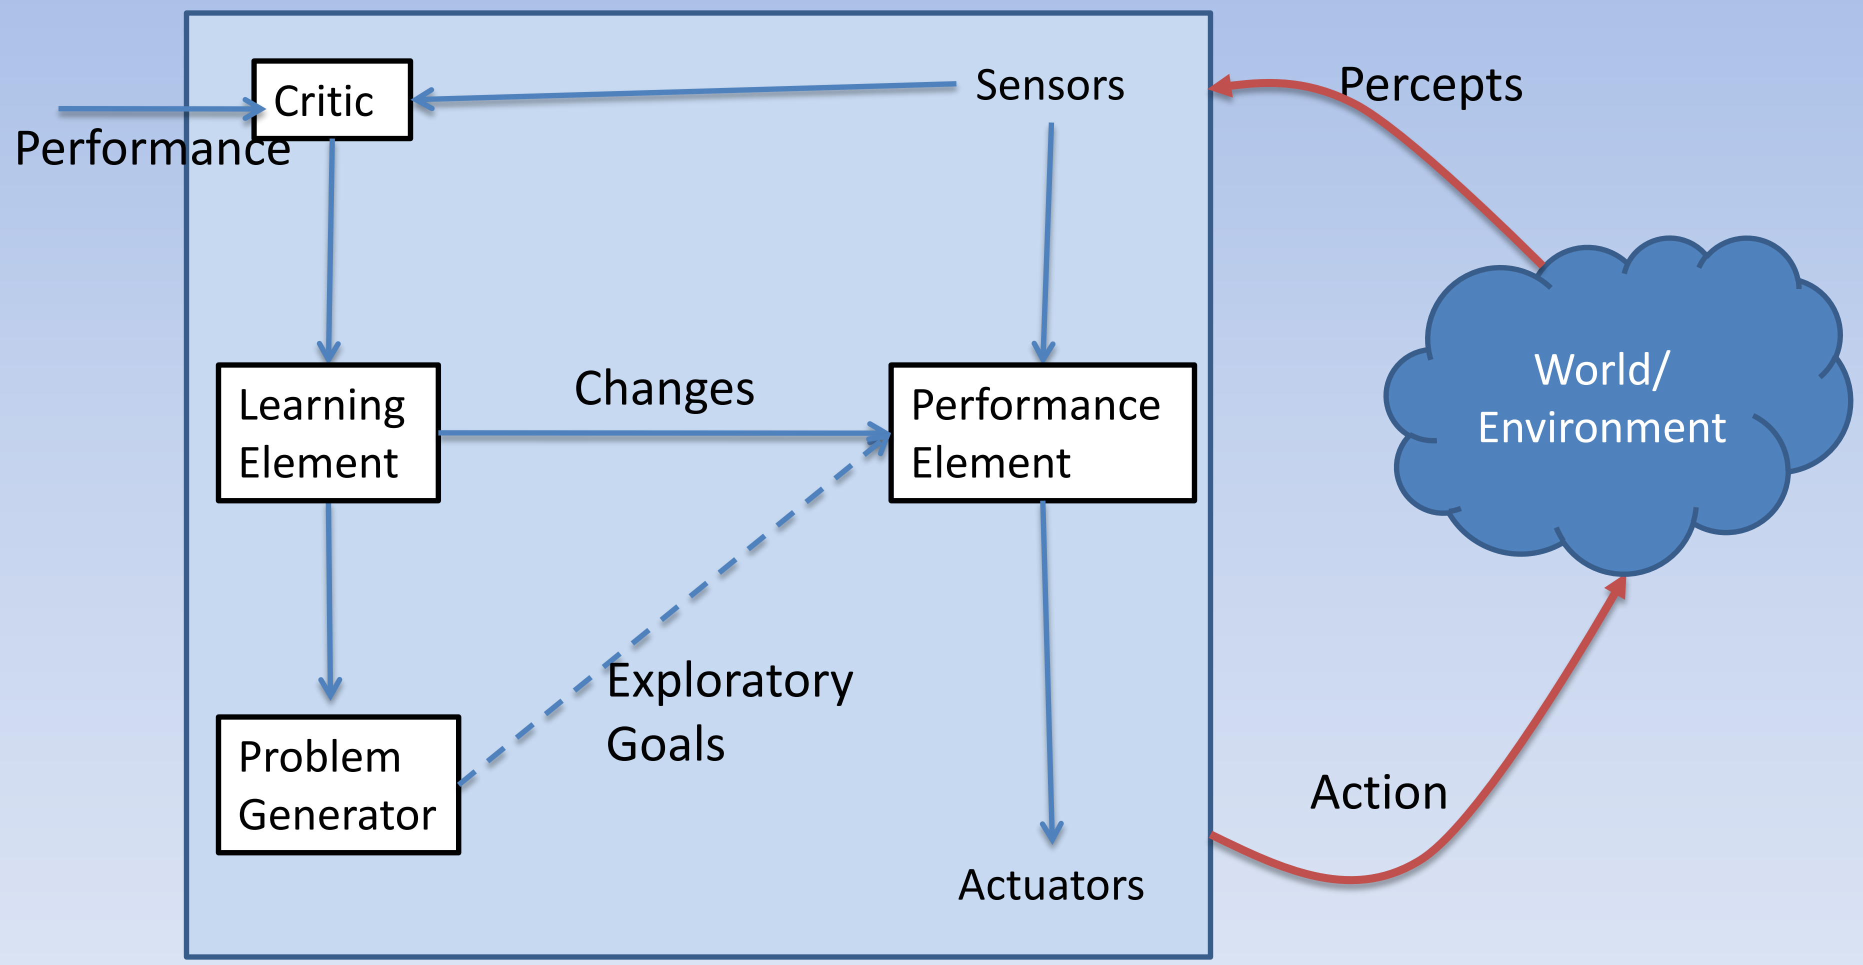

Summary

- We learned about:
  - Agent architecture
  - PEAS descriptions
  - Types of environments
  - Types of Agents