Artificial Intelligence

Mr. Prashant Atmakuri
Assistant Professor
Department of CSE
PVPSIT
prashant.a@pvpsiddhartha.ac.in
+91 8019709941

Textbook:

Artificial Intelligence
A Modern Approach
3rd Edition
Pearson Publications

AUTHORS: STUART RUSSELL
PETER NORVIG

STEM EDUCATION (SCIENCE-TECHNOLOGY-ENGINEERING-MATHEMATICS)

- Fostering Critical Thinking and Problem-Solving
- Promoting Innovation and Creativity
- Preparing for Future Careers
- Addressing Global Challenges

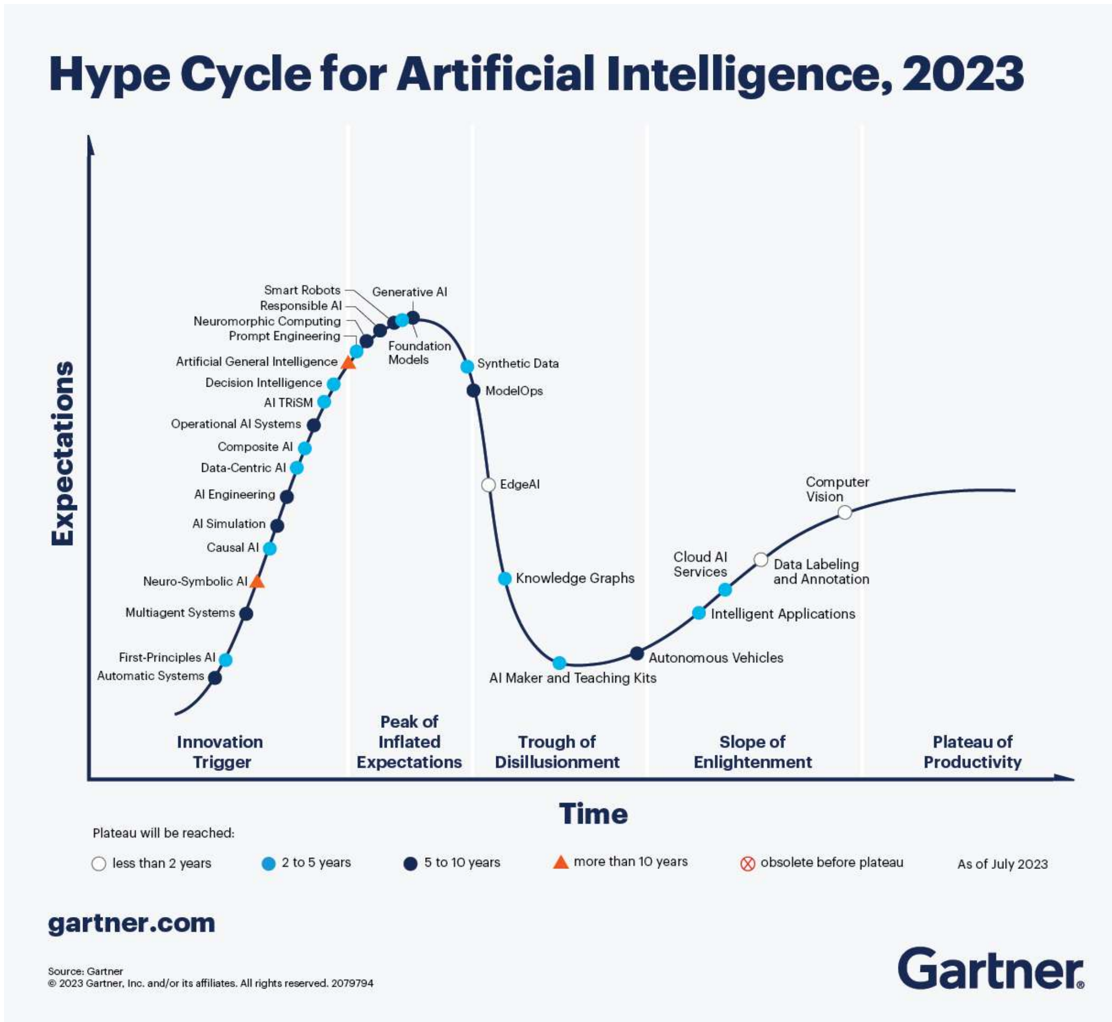
Hype Cycle for Artificial Intelligence, 2023
Expectations
Smart Robots
Responsible AI
Neuromorphic Computing
Prompt Engineering
Generative AI
Foundation Models
Artificial General Intelligence
Decision Intelligence
AI TRiSM
Operational AI Systems
Composite AI
Data-Centric AI
AI Engineering
AI Simulation
Causal AI
Neuro-Symbolic AI
Multiagent Systems
First-Principles AI
Automatic Systems
Synthetic Data
ModelOps
EdgeAI
Knowledge Graphs
AI Maker and Teaching Kits
Autonomous Vehicles
Intelligent Applications
Cloud AI Services
Data Labeling and Annotation
Computer Vision
Innovation Trigger
Peak of Inflated Expectations
Trough of Disillusionment
Slope of Enlightenment
Plateau of Productivity
Time
Plateau will be reached:
less than 2 years
2 to 5 years
5 to 10 years
more than 10 years
obsolete before plateau
As of July 2023
gartner.com
Source: Gartner
© 2023 Gartner, Inc. and/or its affiliates. All rights reserved. 2079794
Gartner

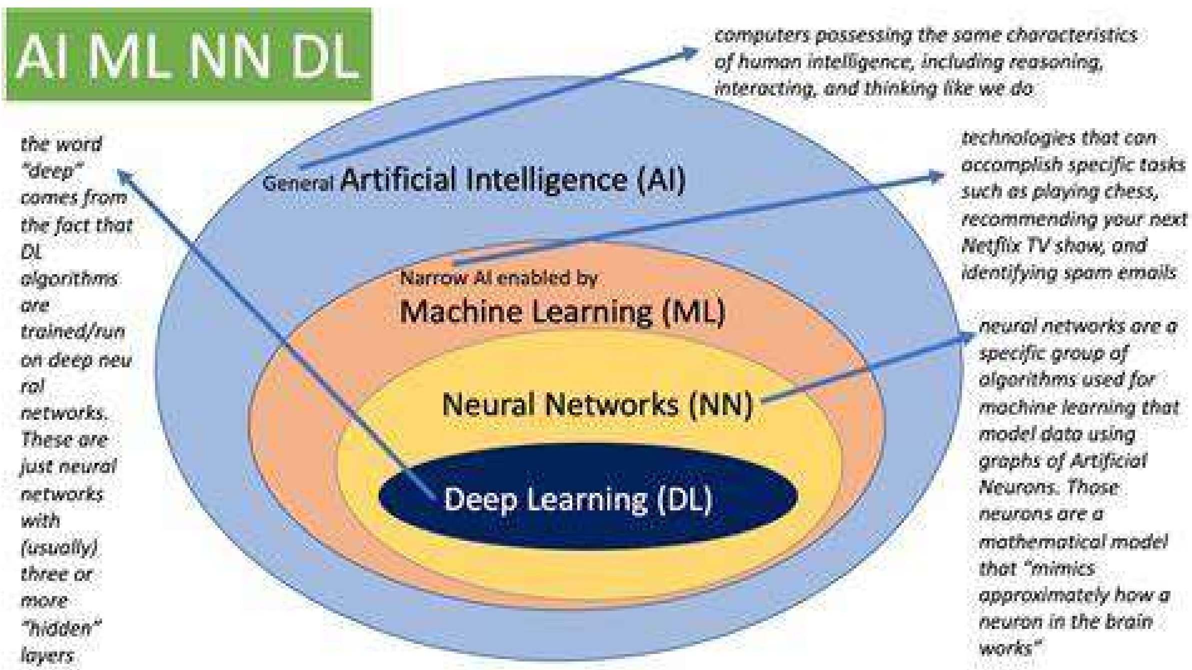
AI ML NN DL
computers possessing the same characteristics of human intelligence, including reasoning, interacting, and thinking like we do
technologies that can accomplish specific tasks such as playing chess, recommending your next Netflix TV show, and identifying spam emails
neural networks are a specific group of algorithms used for machine learning that model data using graphs of Artificial Neurons. Those neurons are a mathematical model that "mimics approximately how a neuron in the brain works"
the word "deep" comes from the fact that DL algorithms are trained/run on deep neu ral networks. These are just neural networks with (usually) three or more "hidden" layers
General Artificial Intelligence (AI)
Narrow AI enabled by
Machine Learning (ML)
Neural Networks (NN)
Deep Learning (DL)

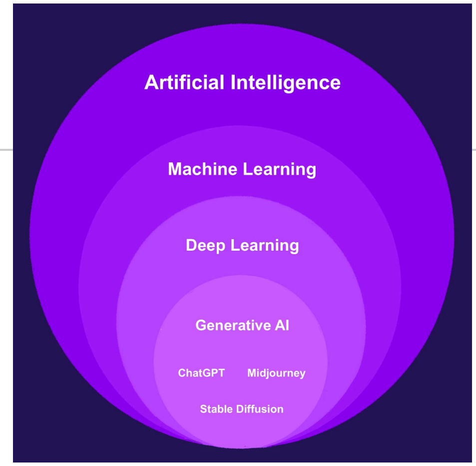
Artificial Intelligence
Machine Learning
Deep Learning
Generative AI
ChatGPT
Midjourney
Stable Diffusion

Topics – Unit I

- Defining Artificial Intelligence.
- Foundations of AI,
- Applications of AI.
- Intelligent agents: Agents and Environments.
- Structure of agents.

Introduction to AI

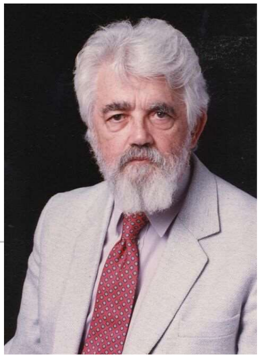

◉A branch of Computer Science named Artificial Intelligence (AI) pursues creating the computers / machines as intelligent as human beings.

John McCarthy the father of Artificial Intelligence described AI as, "***The science and engineering of making intelligent machines, especially intelligent computer programs***".

Artificial Intelligence (AI) is a branch of Science which deals with helping machines find solutions to complex problems in a more human-like fashion.

◉This generally involves borrowing characteristics from human intelligence, and applying them as algorithms in a computer friendly way.

◉ A more or less flexible or efficient approach can be taken depending on the requirements established, which influences how artificial the intelligent behaviour appears.

Defining AI techniques
APPROACHES TO AI

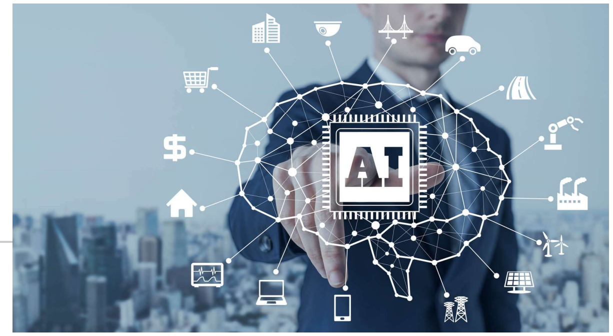

	HUMAN	IDEAL/RATIONAL
THOUGHT PROCESS & REASONING	Thinking Humanly **Systems that think like humans**	Thinking Rationally **Systems that think rationally**
BEHAVIOUR	Acting Humanly **Systems that act like humans**	Acting Rationally **Systems that act rationally**

A system is **rational** if it does the "right thing," given what it knows.

Thinking Humanly

- Humans as observed from 'inside'
- How do we know how humans think?
- Cognitive science: modeling the processes of human thought.
- Introspection: trying to catch our own thoughts as they go by.
- Psychological experiments: observing a person in action.
- Brain imaging: observing the brain in action
- Through a set of experiments and computational models, trying to build good explanations of what we do when we solve a particular task.
- Relevance to AI: to solve a problem that humans (or other living being) are capable of, it's good to know how we go about solving it.

Acting Humanly

- How do you distinguish intelligent behavior from intelligence?
- Turing test, by A. Turing, 1950: determining if a program qualifies as artificially intelligent by subjecting it to an interrogation along with a human counterpart.
- The program passes the test if a human judge cannot distinguish between the answers of the program and the answers of the human subject.
- Capabilities
  - NLP
  - Knowledge representation
  - Automated reasoning
  - Machine learning
- To pass Total Turing Test
  - Computer vision
  - robotics

Thinking Rationally

- Humans are not always 'rational'
- Rational thinking- defined in terms of logic?
- Systems capable of reasoning, capable of making logical deductions from a knowledge base.
- The theory of probability allows rigorous reasoning with uncertain information.
- This requires some capacity to make logical inferences, like "All humans are mortal; Socrates is a human; thus Socrates is mortal".

Acting Rationally

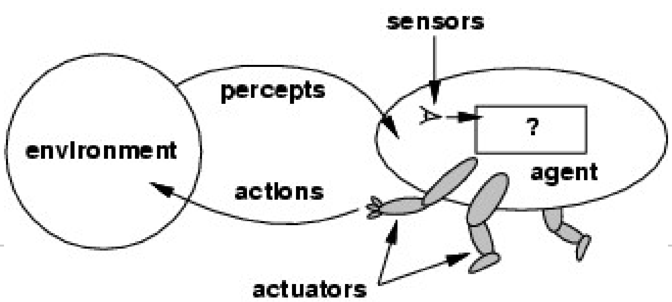

- Rational behavior: doing the right thing.
- Many AI applications adopt the intelligent agent approach.
- An agent is an entity capable of generating action.
- In AI a rational agent must be autonomous, capable of perceiving its environment, adaptable, with a given goal.
- Most often the agents are small pieces of code with a specific proficiency. The problem is solved by combining the skills of several agents.

Philosophy

- Can formal rules be used to draw valid conclusions?
- How does the mind arise from a physical brain?
- Where does knowledge come from?
- How does knowledge lead to action?

Foundations of AI

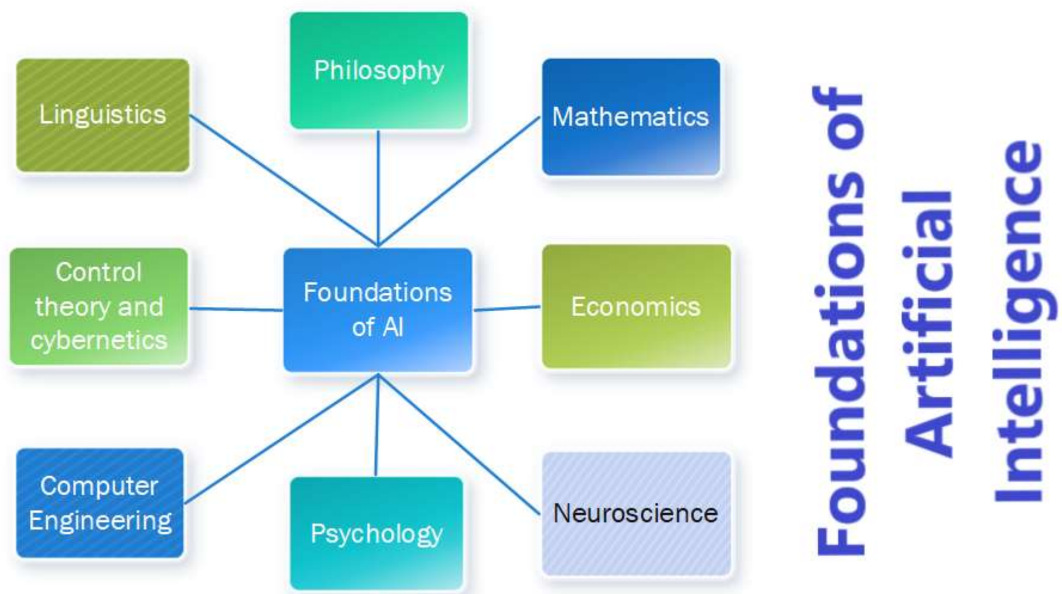

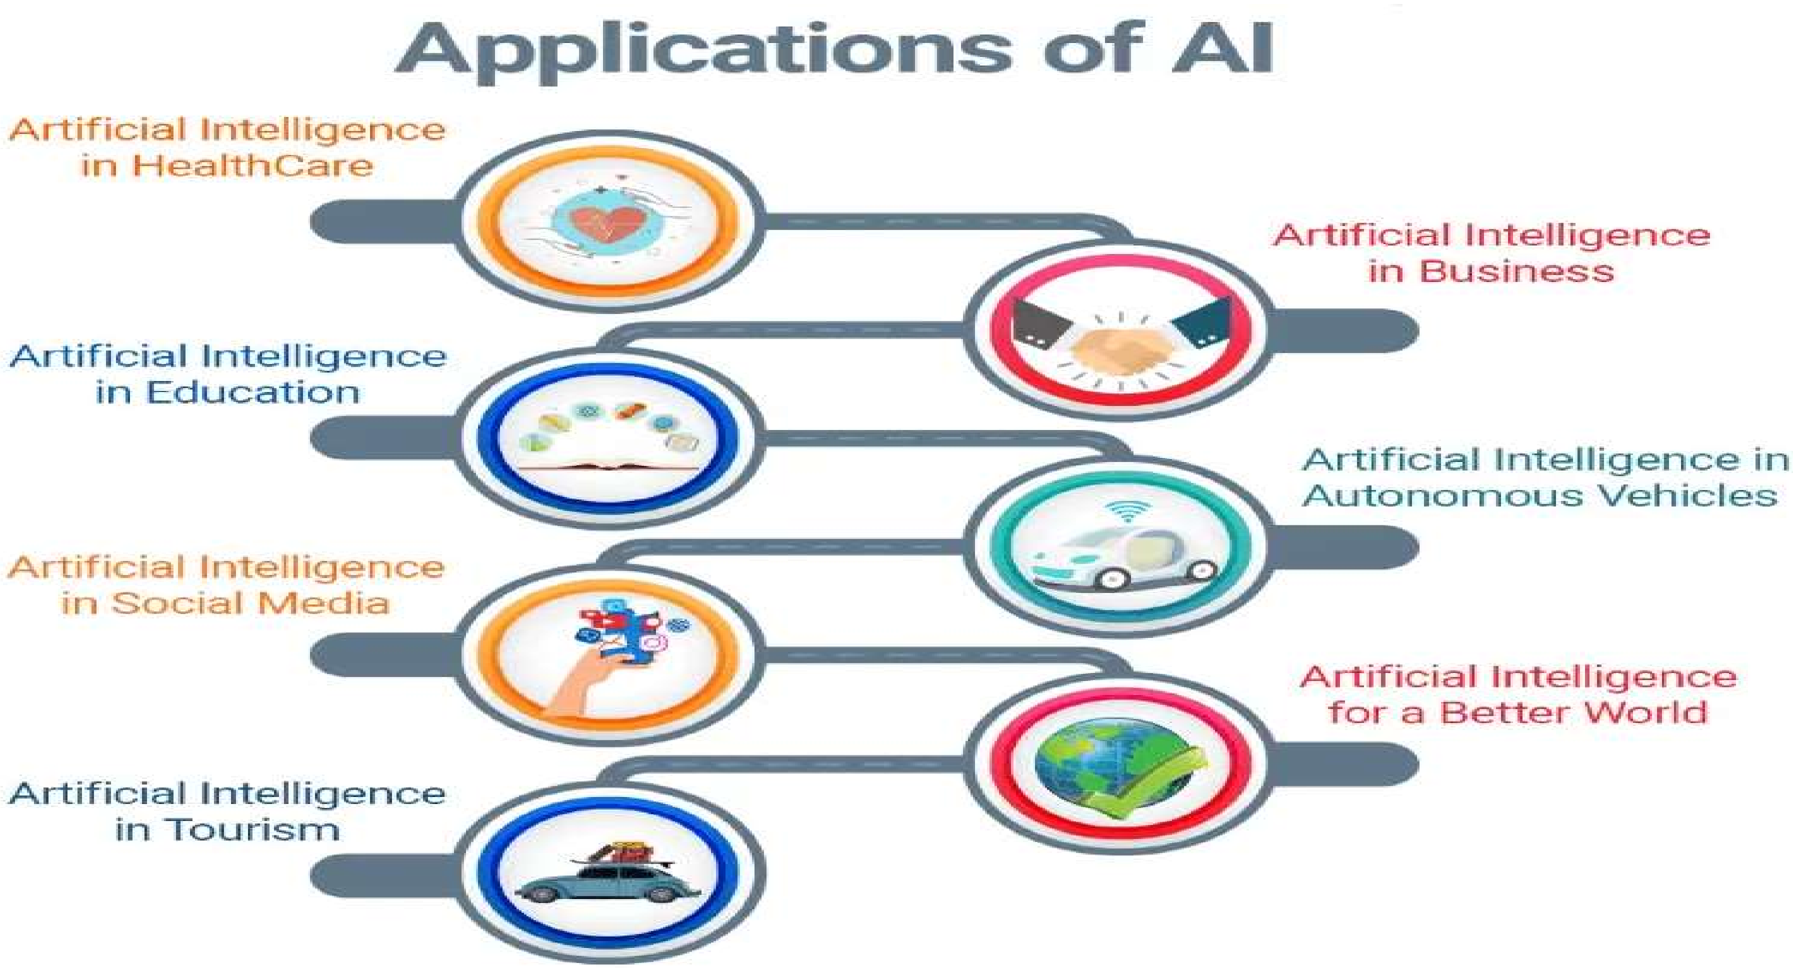
Applications of AI
Artificial Intelligence in HealthCare
Artificial Intelligence in Business
Artificial Intelligence in Education
Artificial Intelligence in Autonomous Vehicles
Artificial Intelligence in Social Media
Artificial Intelligence for a Better World
Artificial Intelligence in Tourism

Agents

- An agent is anything that can be viewed as perceiving its environment through sensors and acting upon that environment through actuators

- Human agent:
  - eyes, ears, and other organs for sensors;
  - hands, legs, mouth, and other body parts for actuators

- Robotic agent:
  - cameras and infrared range finders for sensors
  - various motors for actuators

Continued…

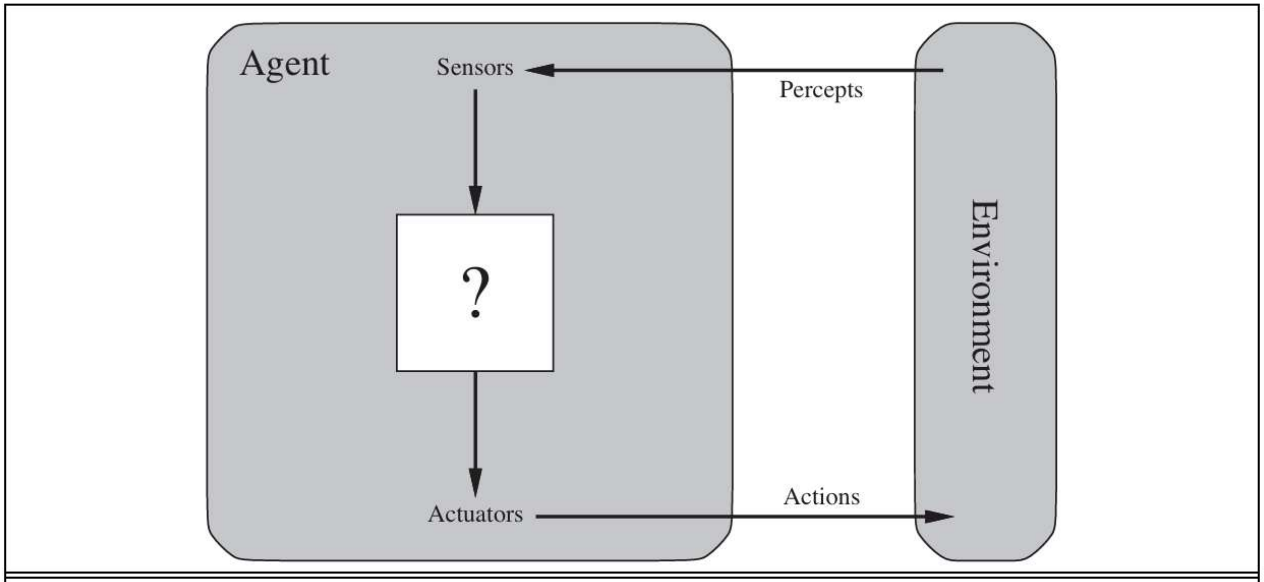

Figure 2.1 Agents interact with environments through sensors and actuators.

Continued…

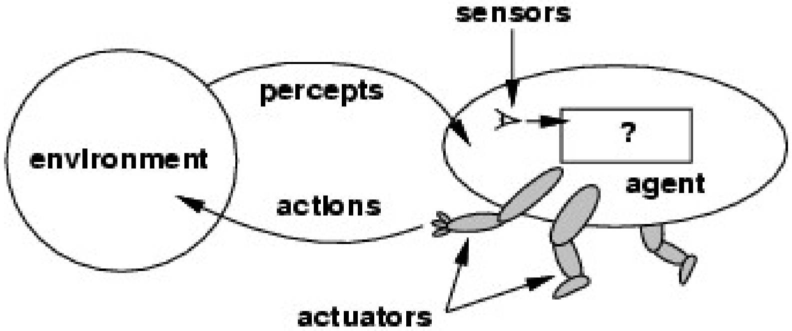

- The agent function maps from percept histories to actions:

$$[f: \mathcal{P}^{\star} \rightarrow \mathcal{A}]$$

- The agent program runs on the physical architecture to produce *f*
- agent = architecture + program

Vacuum-cleaner world

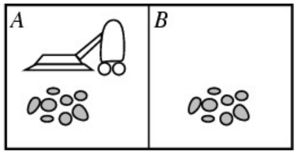

- Percepts: location and contents, e.g., [A,Dirty]
- Actions: *Left*, *Right*, *Suck*, *NoOp*
- *Agent's function → look-up table*

For many agents this is a very large table

Percept sequence	Action
$[A, Clean]$	$Right$
$[A, Dirty]$	$Suck$
$[B, Clean]$	$Left$
$[B, Dirty]$	$Suck$
$[A, Clean]$, $[A, Clean]$	$Right$
$[A, Clean]$, $[A, Dirty]$	$Suck$
$\vdots$	$\vdots$

Specifying the Task Environment - PEAS

PEAS*:* Performance measure, Environment, Actuators, Sensors

- Must first specify the setting for intelligent agent design
- Consider, e.g., *the task of designing an automated taxi driver*:
- ***Performance measure:*** Safe, fast, legal, comfortable trip, maximize profits
- ***Environment:*** Roads, other traffic, pedestrians, customers
- ***Actuators:*** Steering wheel, accelerator, brake, signal, horn
- ***Sensors:*** Cameras, sonar, speedometer, GPS, odometer, engine sensors, keyboard

PEAS

- **Performance Measure:** Performance measure is the unit used to define an agent's success. The performance of agents changes according to their distinct principles.
- **Actuator:** An actuator is a component of the agent that provides the action's output to the environment.
- **Sensor:** Sensors are the receptive components of an agent that receive input.

PEAS

- **Environment:** The environment is an agent's immediate surroundings. If the agent is set in motion, it changes over time. There are five primary types of environments:

Fully Observable & Partially Observable

Episodic & Sequential

Static & Dynamic

Discrete & Continuous

Deterministic & Stochastic

Continued…

- **Agent:** Interactive English tutor
- **Performance measure:** Maximize student's score on test
- **Environment:** Set of students
- **Actuators:** Screen display (exercises, suggestions, corrections)
- **Sensors:** Keyboard

Environment types

- Fully observable (vs. partially observable)
- Deterministic (vs. stochastic)
- Episodic (vs. sequential)
- Static (vs. dynamic)
- Discrete (vs. continuous)
- Single agent (vs. multiagent):

Continued…

◉ **1) Fully observable / Partially observable**

If an agent's sensors give it access to the complete state of the environment needed to choose an action, the environment is fully observable.

(e.g. chess)

Continued…

Deterministic / Stochastic

- An environment is **deterministic** if the next state of the environment is completely determined by the current state of the environment and the action of the agent;
- In a **stochastic** environment, there are multiple, unpredictable outcomes. (If the environment is deterministic except for the actions of other agents, then the environment is **strategic**).

Continued…

- We say an environment is uncertain if it is not fully observable or not deterministic.
- In a fully observable, deterministic environment, the agent need not deal with uncertainty.
- “stochastic” generally implies that uncertainty about outcomes is quantified in terms of probabilities.

Continued…

Episodic / Sequential

In an episodic environment, the agent's experience is divided into atomic episodes. Each episode consists of the agent perceiving and then performing a single action.

Subsequent episodes do not depend on what actions occurred in previous episodes. Choice of action in each episode depends only on the episode itself.

(E.g., classifying images.)

In a sequential environment, the agent engages in a series of connected episodes. Current decision can affect future decisions. (E.g., chess and driving)

Continued…

Static / Dynamic

A static environment does not change while the agent is thinking.

The passage of time as an agent deliberates is irrelevant.

The environment is semidynamic if the environment itself does not change with the passage of time but the agent's performance score does.

Continued…

Discrete / Continuous

If the number of distinct percepts and actions is limited, the environment is discrete, otherwise it is continuous.

Single agent / Multi-agent

If the environment contains other intelligent agents, the agent needs to be concerned about strategic, game-theoretic aspects of the environment (for either cooperative *or* competitive agents).

Most engineering environments don't have multi-agent properties, whereas most social and economic systems get their complexity from the interactions of (more or less) rational agents.

Continued…

- **Known vs. unknown**: Strictly speaking, this distinction refers not to the environment itself but to the agent's (or designer's) state of knowledge about the "laws of physics" of the environment.

Continued…

	Chess with a clock	Chess without a clock	Taxi driving
Fully observable	Yes	Yes	No
Deterministic	Strategic	Strategic	No
Episodic	No	No	No
Static	Semi	Yes	No
Discrete	Yes	Yes	No
Single agent	No	No	No

The environment type largely determines the agent design

The real world is (of course) partially observable, stochastic, sequential, dynamic, continuous, multi-agent

Activity

- Crossword -1
- Crossword -2
- Team Activity
- Quiz -1
- Quiz -2

Agent Examples for Team Activity

- drone for package delivery
- self-driving taxi
- tracking of a person in an organization
- dynamic traffic control system
- child care taker
- self-driving train
- drone to rescue people in danger at sea
- Tutor to teach AI
- elderly care taker
- receptionist
- paper valuation
- Shopping for used AI books on the Internet.
- Playing a tennis match.
- Practicing tennis against a wall.
- Performing a high jump.
- Knitting a sweater.
- Bidding on an item at an auction.
- bank cashier.

Types of Agents

- Simple reflex agents
- Reflex agents with state/model
- Goal-based agents
- Utility-based agents

Simple reflex agents

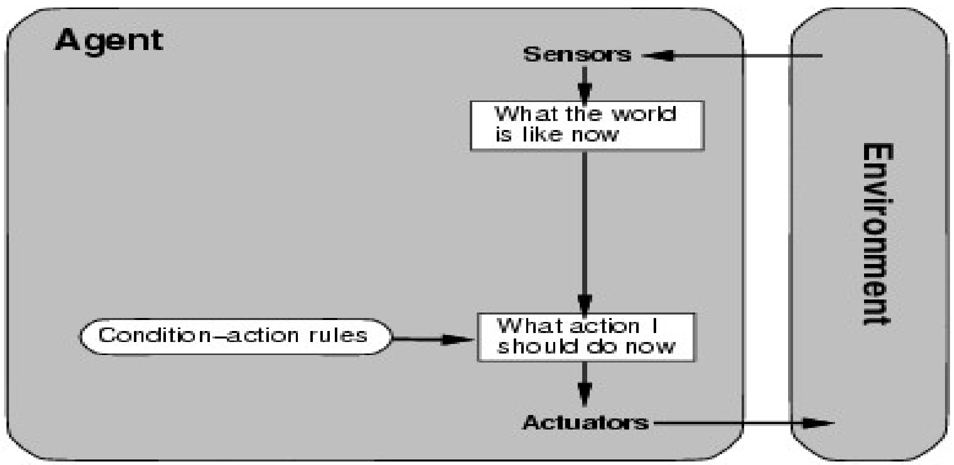

- Simple but very limited intelligence.
- **Action does not depend on percept history, only on current percept.**
- Therefore no memory requirements.

```
function REFLEX-VACUUM-AGENT( [location,status]) returns an action
    if status = Dirty then return Suck
    else if location = A then return Right
    else if location = B then return Left
```

Simple reflex agents

- A Simple Reflex Agent is typically employed when all the information of the current game state is directly observable, (eg: **Chess, Checkers, Tic Tac Toe, Connect-Four**) and the decision regarding the move only depends on the current state.
- That is, when the agent does not need to remember any information of the past state to make a decision.

Model-based reflex agents

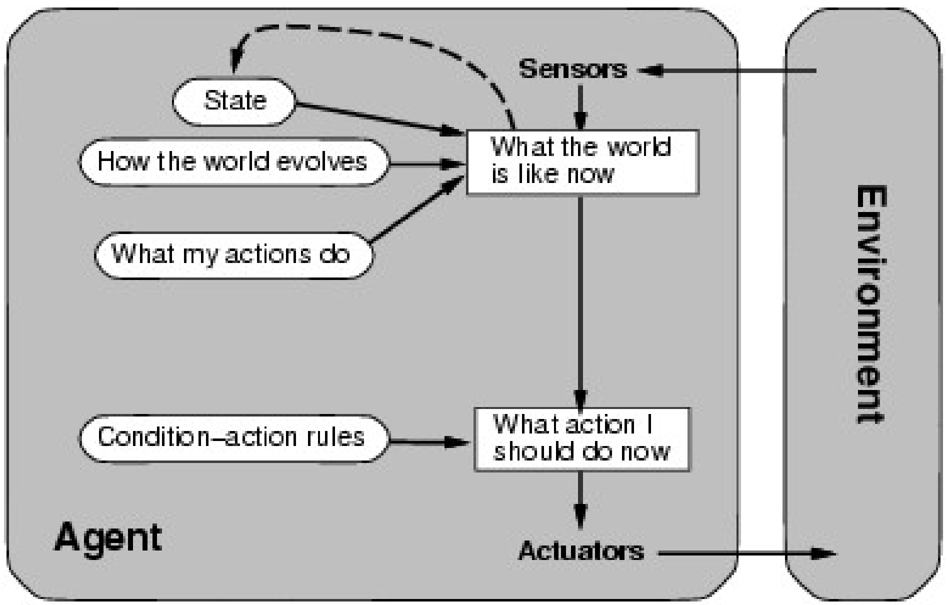

- Know how world evolves
  - Overtaking car gets closer from behind
- How agents actions affect the world
  - Wheel turned clockwise takes you right

- Model base agents update their state

```
function REFLEX-AGENT-WITH-STATE(percept) returns action
   static: state, a description of the current world state
           rules, a set of condition-action rules

   state ← UPDATE-STATE(state, percept)
   rule ← RULE-MATCH(state, rules)
   action ← RULE-ACTION[rule]
   state ← UPDATE-STATE(state, action)
   return action
```

Model-based reflex agents

- Model-based reflex agents are AI agents that **make decisions based on their current perception** of the environment and an **internal model that represents their knowledge and predictions** about the world. They combine the reactive nature of simple reflex agents with the ability to reason about the environment using their internal models.

Model-based reflex agents

- While simple reflex agents make decisions solely based on their current perception, model-based reflex agents go a step further by considering their internal models.
- These models allow them to make more informed decisions by taking into account not only the immediate sensory information but also their knowledge and predictions about the environment.

Model-based reflex agents

- Internal models play a crucial role in the decision-making process of model-based reflex agents. These models represent the agent's understanding of the environment, including the current state, possible future states, and the effects of actions. By leveraging these models, agents can make predictions, plan ahead, and adapt their behavior based on their knowledge and expectations.

Architecture of Model-Based Reflex Agents

- The **perception module** is responsible for capturing and processing sensory information from the environment.
- The **internal model** is a core component of model-based reflex agents.
- The **decision-making module** uses the information from the perception module and the internal model to determine the appropriate action to take.
- Once the decision-making module selects an action, the **action execution** component carries out the chosen action in the environment.

Model-based reflex agents- Types of Internal Models

- **State-based** models represent the environment as a set of discrete states and the transitions between them.
- **Predictive models** focus on forecasting future states or outcomes based on the current state and the agent's actions.
- **Hybrid models** combine elements of state-based and predictive models to create a more comprehensive representation of the environment.

Advantages of Model-Based Reflex Agents

- More Informed Decision-Making
- Ability to Handle Uncertainty
- Adaptability to Changing Environments

Model-based reflex agents- Challenges and Limitations

- Computational Complexity
- Model Accuracy and Maintenance
- Handling Incomplete or Noisy Information

Applications of Model-Based Reflex Agents

- Robotics and Autonomous Systems
- Gaming and Simulation
- Intelligent Control Systems

Future Developments:

- Integration with Machine Learning
- Scalability and Real-Time Performance
- Ethical Considerations

Goal-based agents

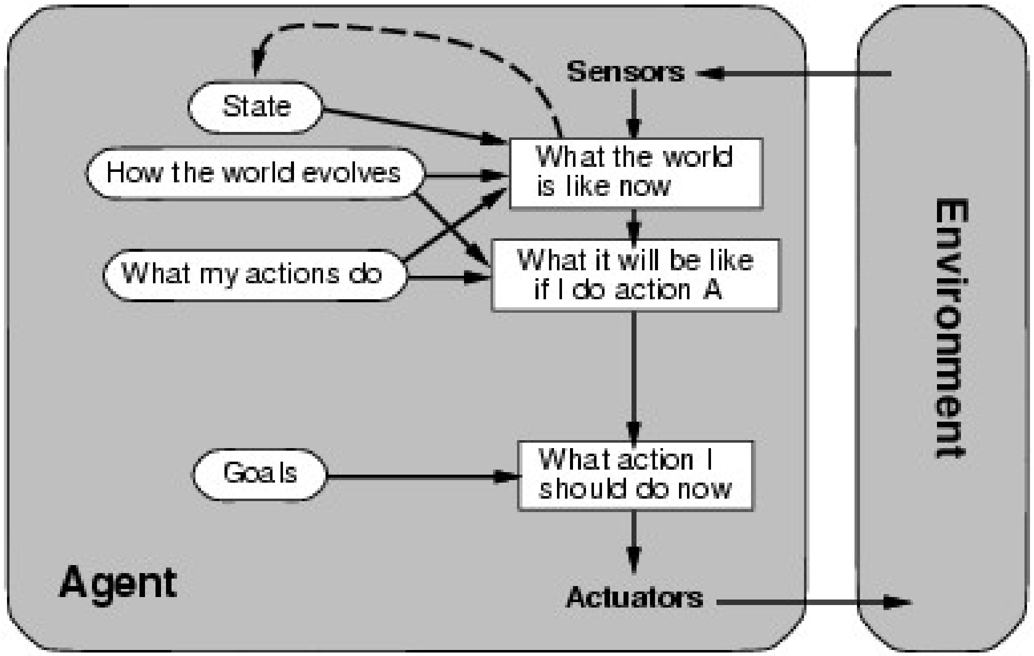

•knowing state and environment? Enough?

–Taxi can go left, right, straight

•Have a goal

A destination to get to

Uses knowledge about a goal to guide its actions

E.g., Search, planning

Goal-based agents

- The goal-based agent is also known as a planning or goal-seeking agent.
- **We sometimes call it a rule-based agent as it follows a set of rules to achieve its goal.**
- It uses search algorithms to find the most efficient path to the goal. Additionally, it also utilizes heuristics and AI techniques to improve its performance.

Utility-based agents

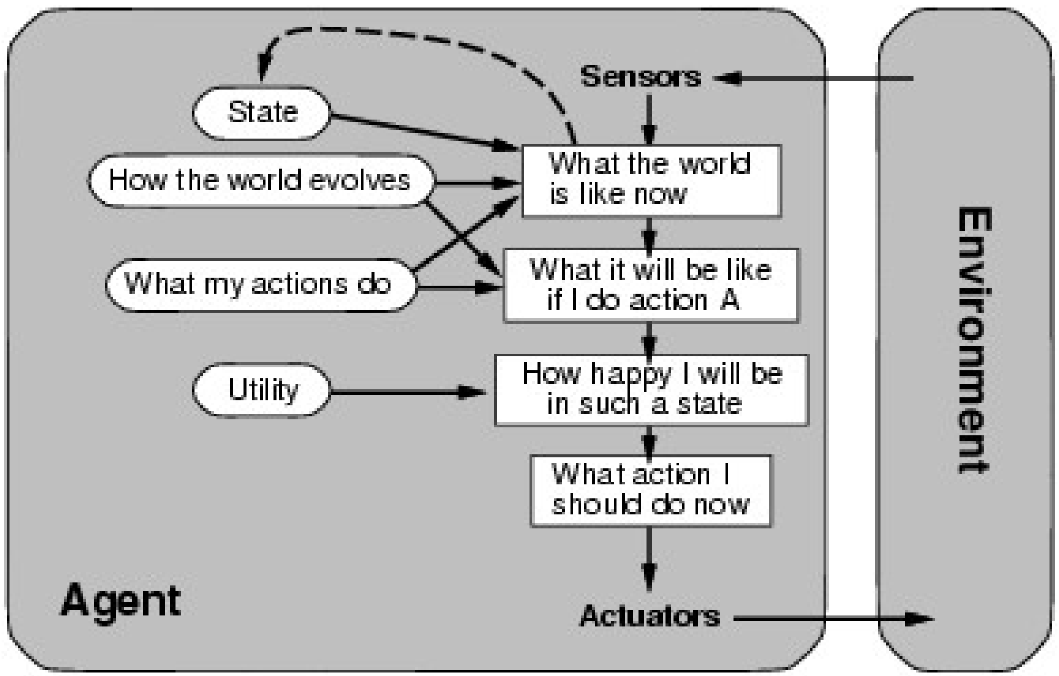

Goals are not always enough

- Many action sequences get taxi to destination
- Consider other things. How fast, how safe…..

A utility function maps a state onto a real number which describes the associated degree of "happiness", "goodness", "success".

Utility-based agents

- **It makes decisions based on a set of predetermined criteria or a utility function.**
- These criteria represent the goals or objectives of the agent.
- **The goal of the agent is to maximize the utility function.**

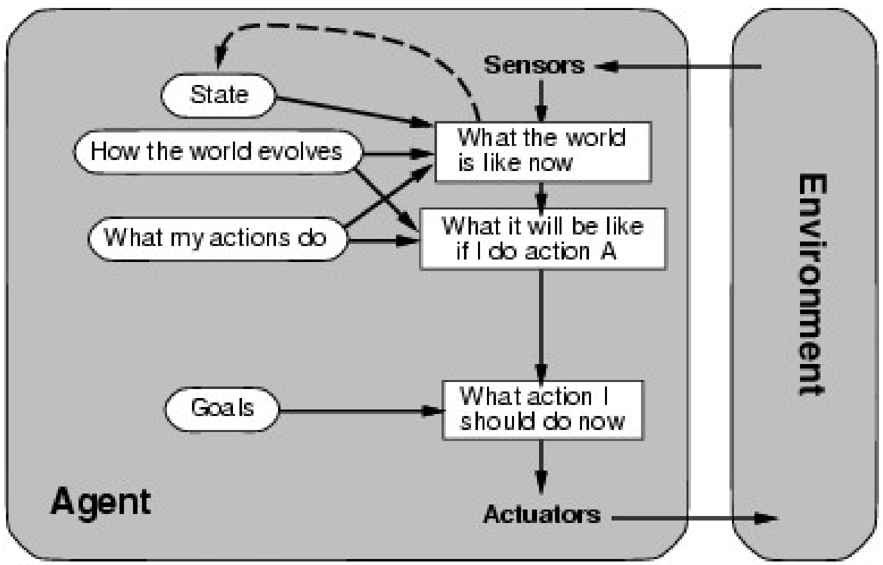

Goal-based agents

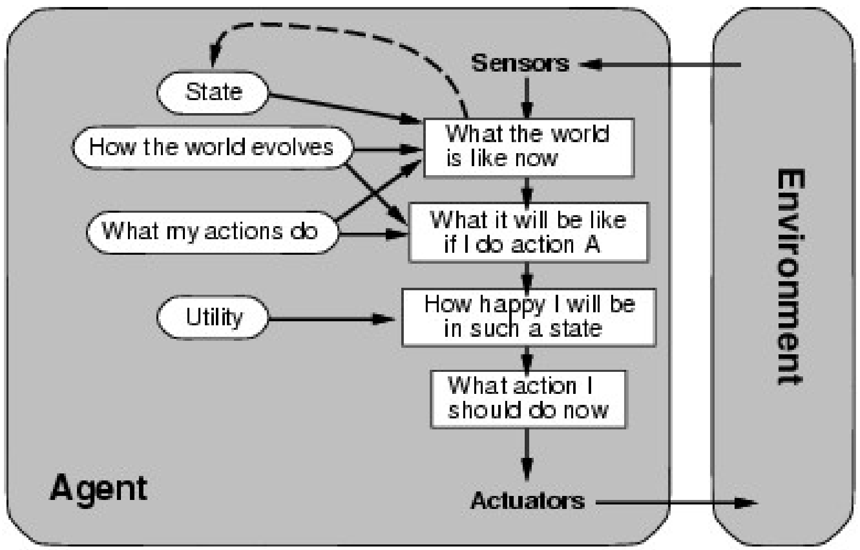

Utility-based agents

Goal-Based Agents	Utility-Based Agents
Goal-based agents may perform in a way that produces an unexpected outcome because their search space is limited	Utility-based agents are more reliable because they can learn from their environment and perform most efficiently
Makes decisions based on the goal and the available information	Makes decisions based on the utility and general information
Goal-based agents are easier to program	Implementing utility-based agents can be a complex task
Considers a set of possible actions before deciding whether the goal is achieved or not	Maps each state to an actual number to check how efficiently each step achieves its goals
Utilized in computer vision, robotics, and NLP	Used in GPS and tracking systems

Learning agent

learning element -> to make improvements

performance element (agent)-> selecting external actions

The learning element uses feedback from the **critic** on how the agent is doing and determines how the performance element should be modified to do better in the future.

problem generator ->suggesting actions that will lead to new and informative experiences.

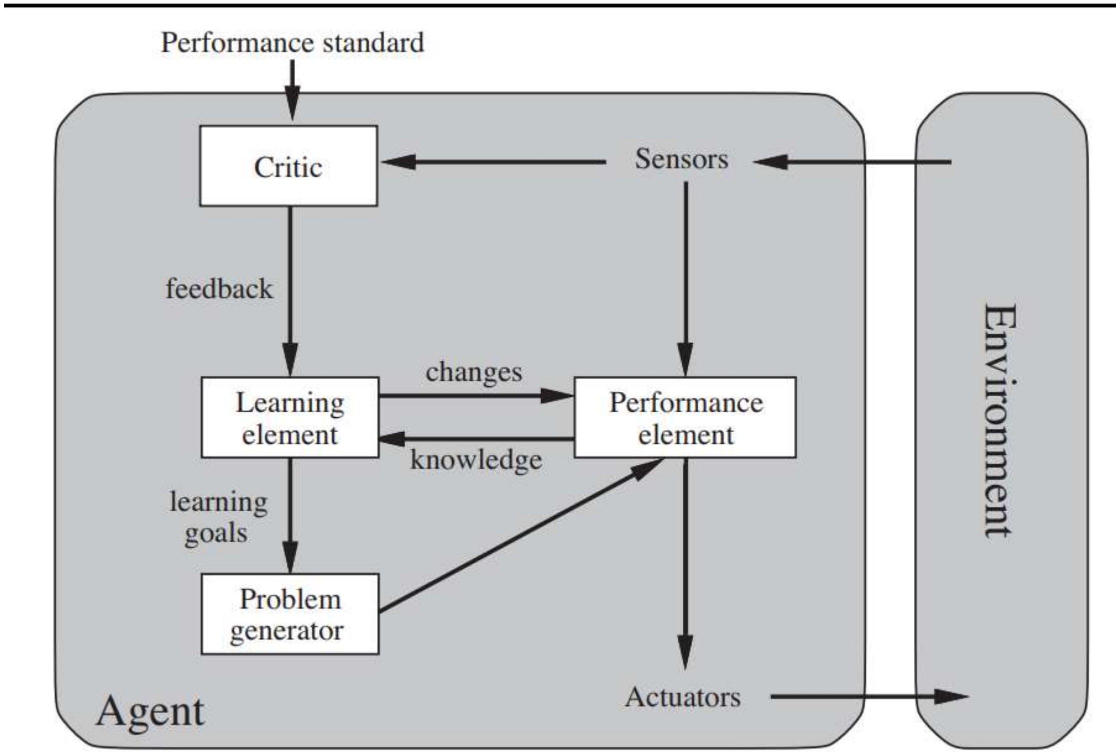

How the components of agent programs work

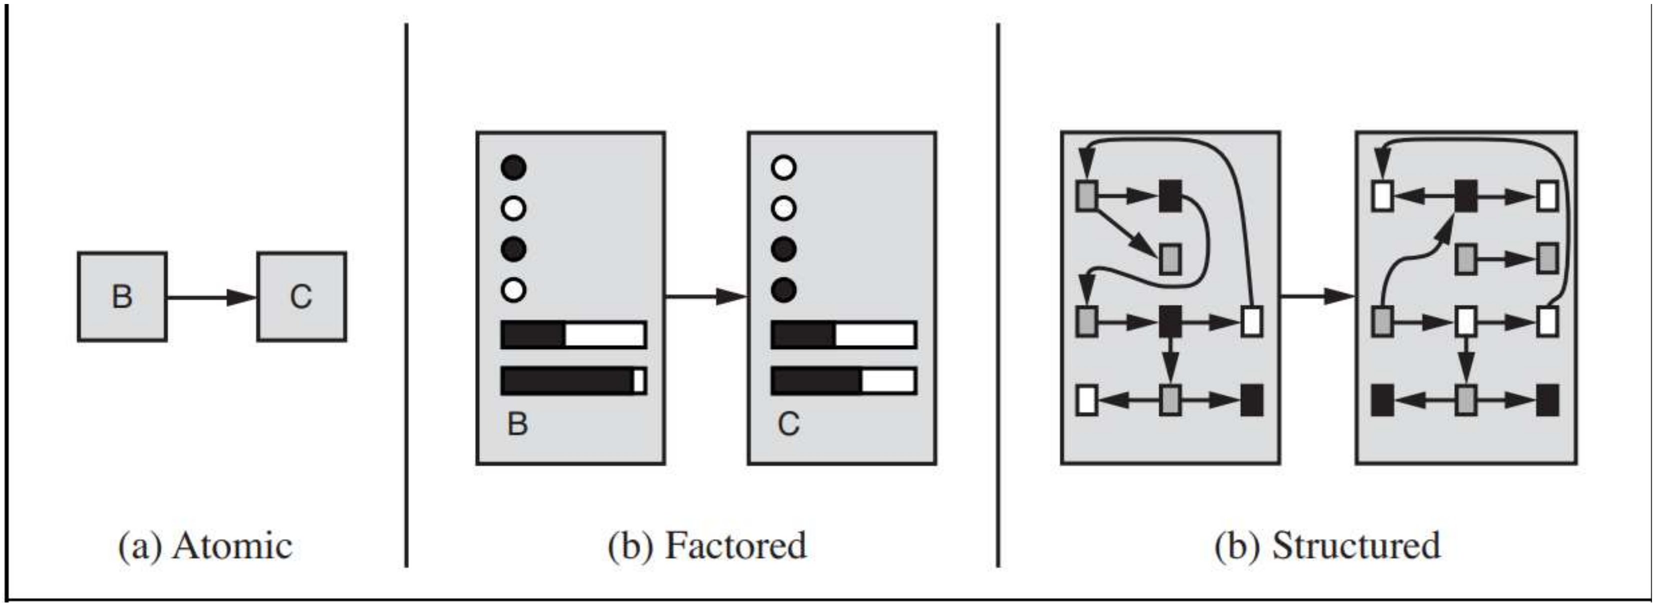

(a) Atomic (b) Factored (b) Structured

Continue….

- (a) Atomic representation: a state (such as B or C) is a black box with no internal structure;
- (b) Factored representation: a state consists of a vector of attribute values; values can be Boolean, real valued, or one of a fixed set of symbols.
- (c) Structured representation: a state includes objects, each of which may have attributes of its own as well as relationships to other objects.
- a more expressive representation can capture, at least as concisely, everything a less expressive one can capture, plus some more

- THANK YOU