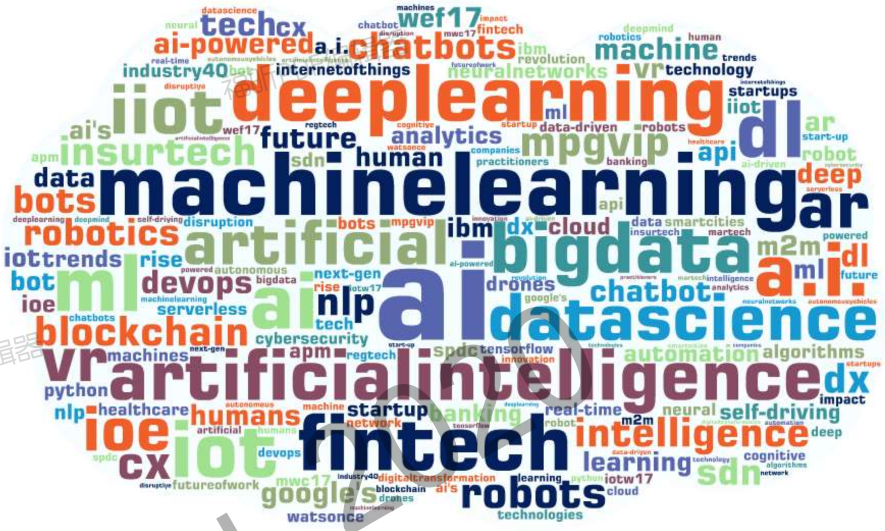

Agent

Pengju Ren

Institute of Artificial Intelligence and Robotics

pengjuren@xjtu.edu.cn

Recap: What is AI ?

Artificial intelligence (AI) is intelligence exhibited by machines. In computer science, the field of AI research defines itself as the study of "intelligent agents": any device that perceives its environment and takes actions that maximize its chance of success at some goal.

--- from Wikipedia

Agents

- An **agent(Software/Hardware)** is anything that can be viewed as perceiving its **environment** through **Sensors** and acting upon that environment though **actuators.**

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

- **Robotic agents:** cameras and infrared range finders for sensors;
  Various motors for actuators

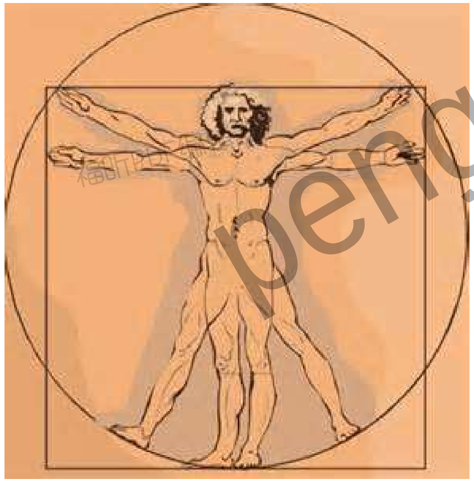

Xianer Robot monk

Rethink Robotics

Beyond the Human Senses

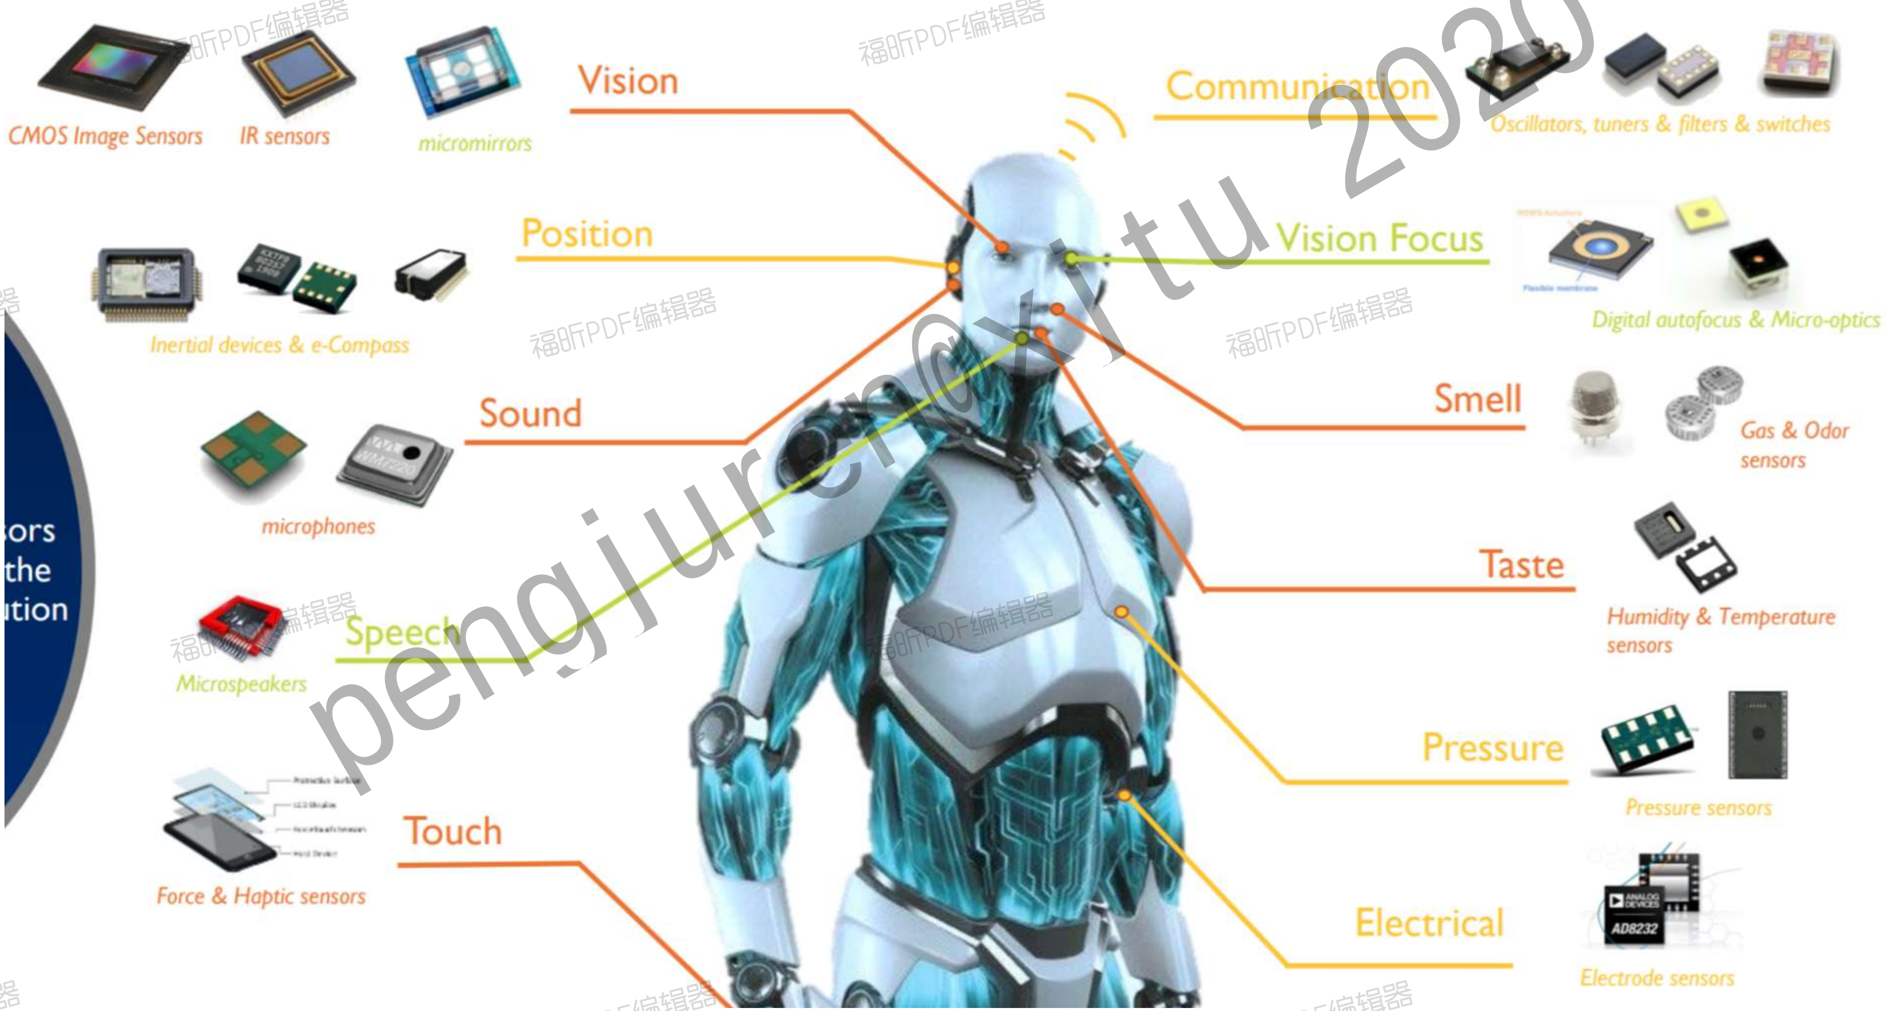

Sensors for Robotics and Drones

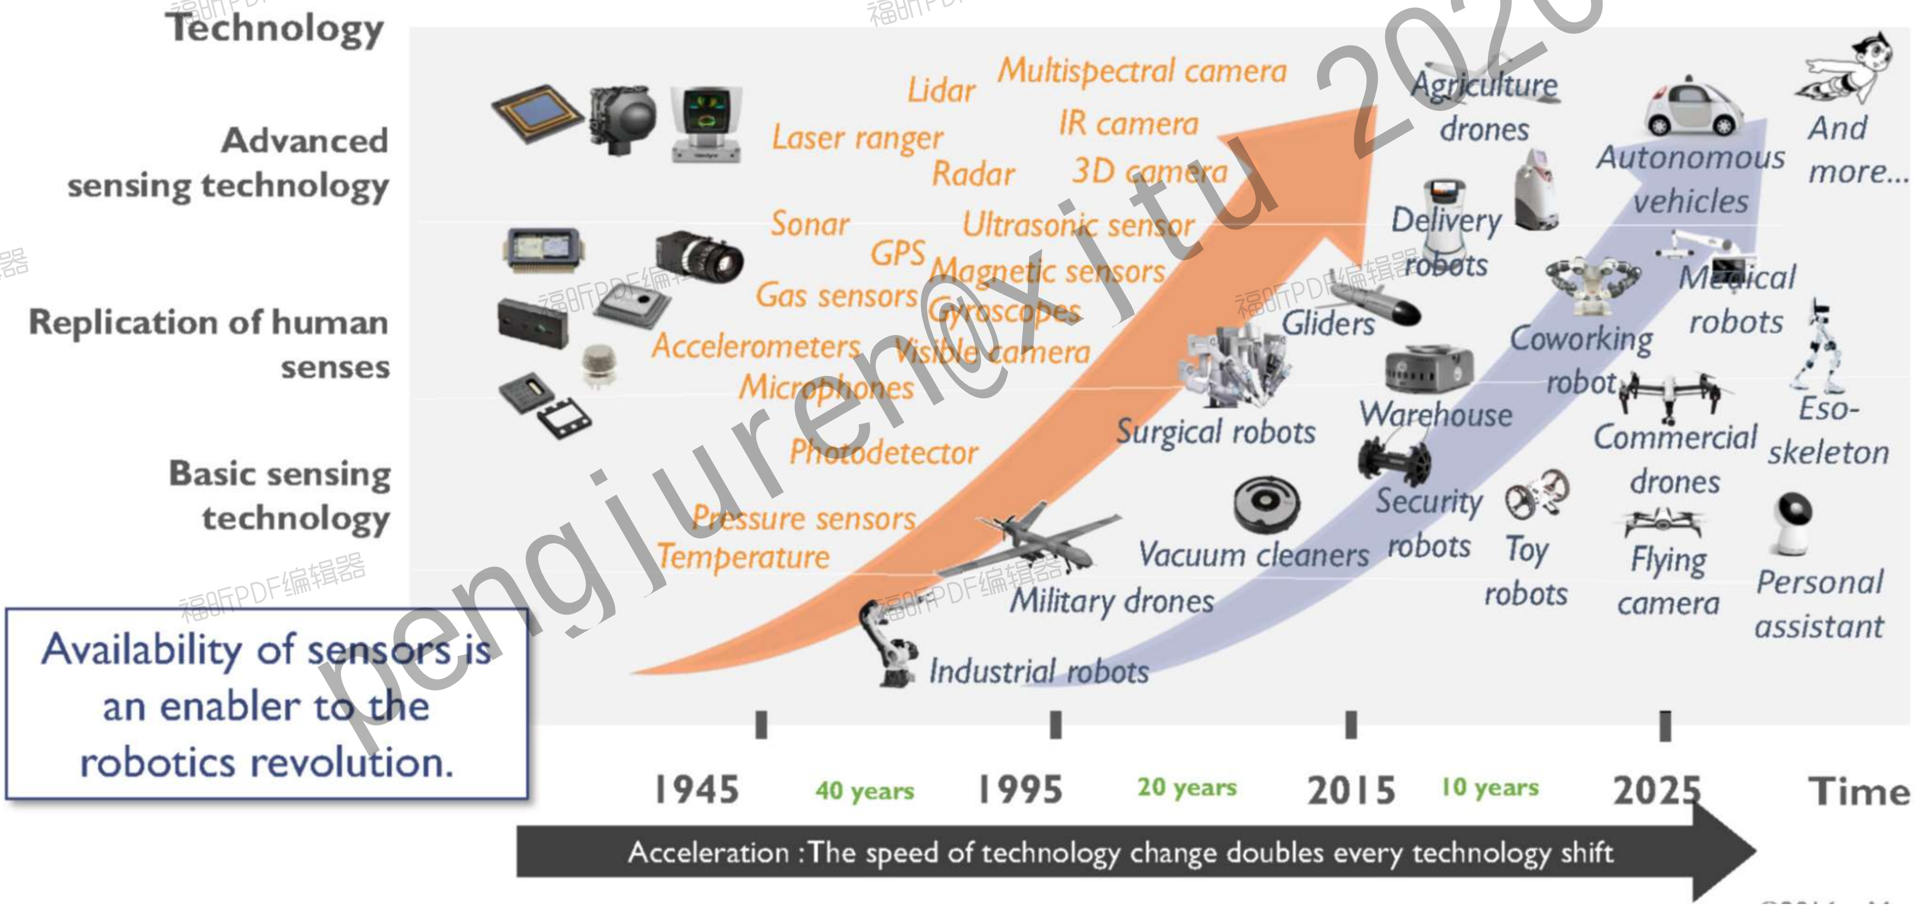

Acceleration : The speed of technology change doubles every technology shift

Availability of sensors is an enabler to the robotics revolution.

Why Self-driving Car has the potential to outperform Human Driver

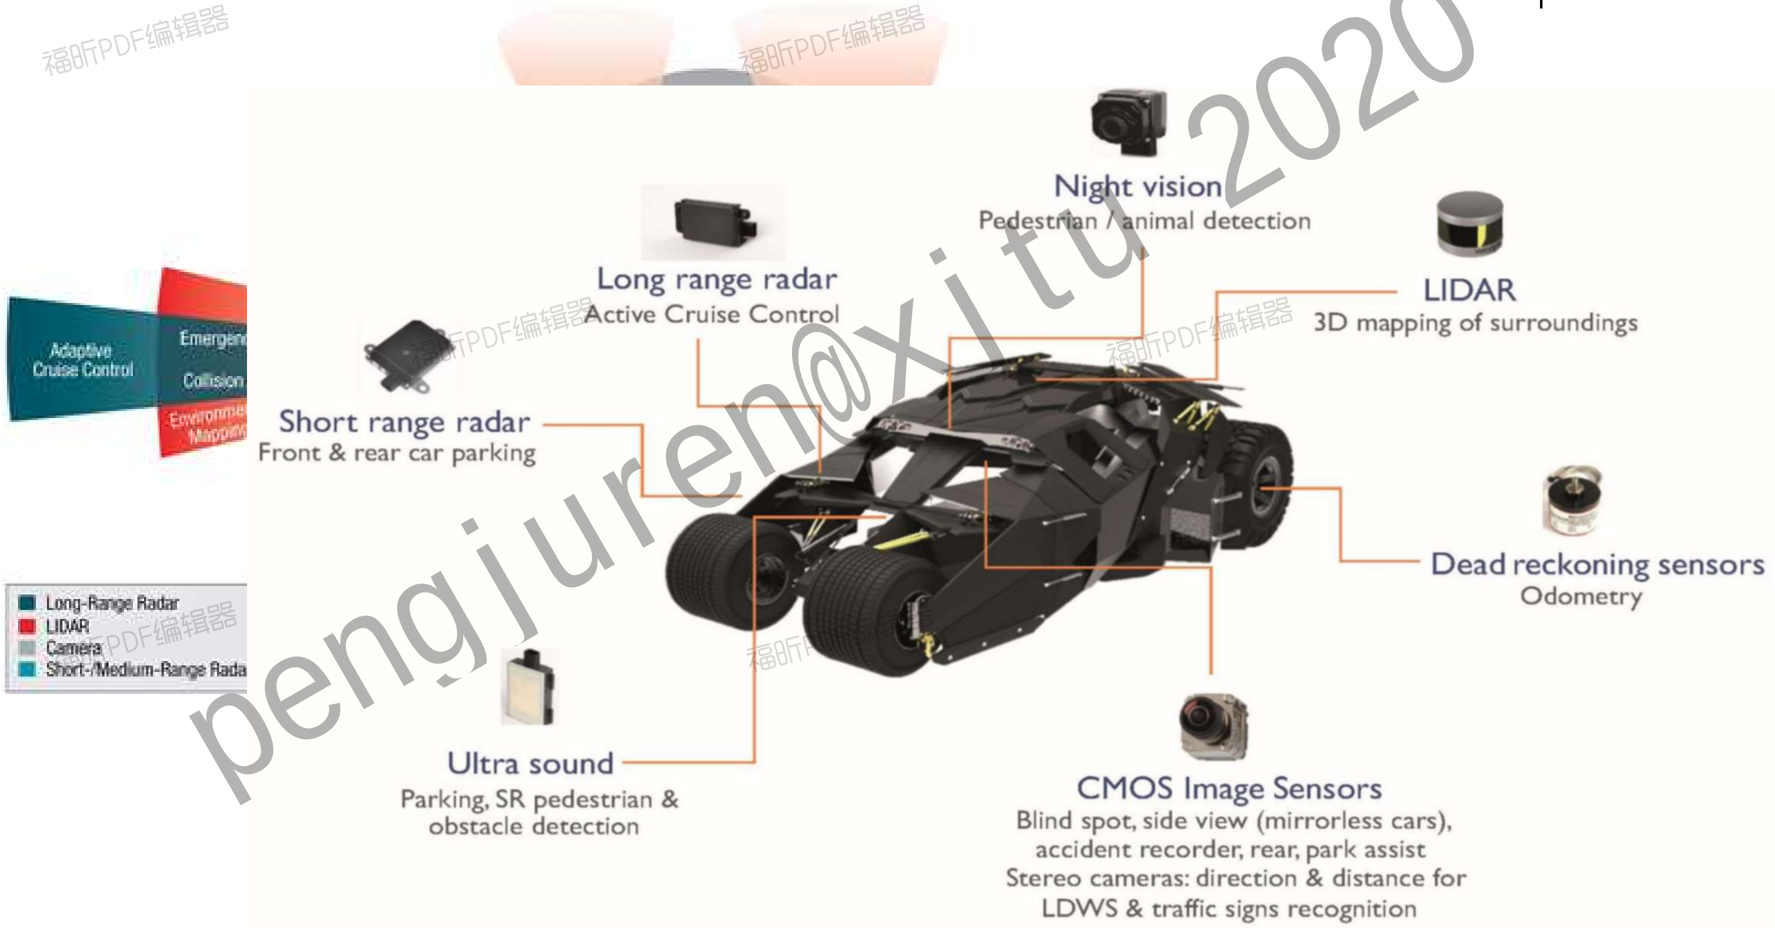

Still has a long way to go ……

e.g. Traffic Infrastructure, Legal issue, Computing capability

Why Self-driving Car has the potential to outperform Human Driver

Agents and Environment

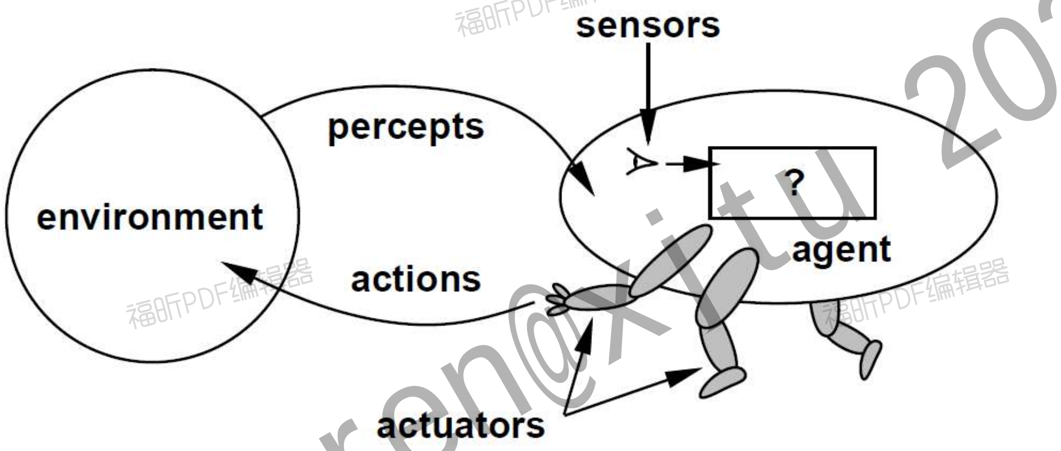

Agents include humans, robots, softbots, thermostats, etc.
The agent function maps from percept sequences to actions

[*f*: P* → A]

The agent program is the implementation of the produce *f*
Agent = Architecture + program

Agents functions and programs

- An agent is completely specified by the agent function mapping percept sequences to actions (Formulation)
- One agent function (or a small equivalence class) is rational
- Aim: find a way to implement the rational agent function concisely

For example: Table-lookup Agent

\input{algorithms/table-agent-algorithm}

Drawbacks:

- Huge table (Memory)
- Take a long time to build the table
- No autonomy
- Even with learning, need a long time to learn the table entries

Vacuum Cleaner

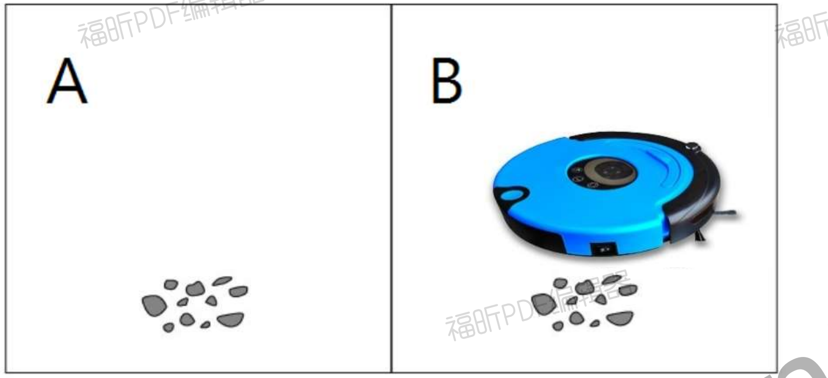

Percepts: Location and Content
e.g. [A, Dirty].

Actions: *Turn left, Turn Right, Suck, NoOp*

REFLEX-VACUUM-AGENT(*location*, *status*)

```
if status == Dirty
    return Suck
elseif location == A
    return Right
elseif location == B
    return Left
```

A Vacuum-cleaner Agent

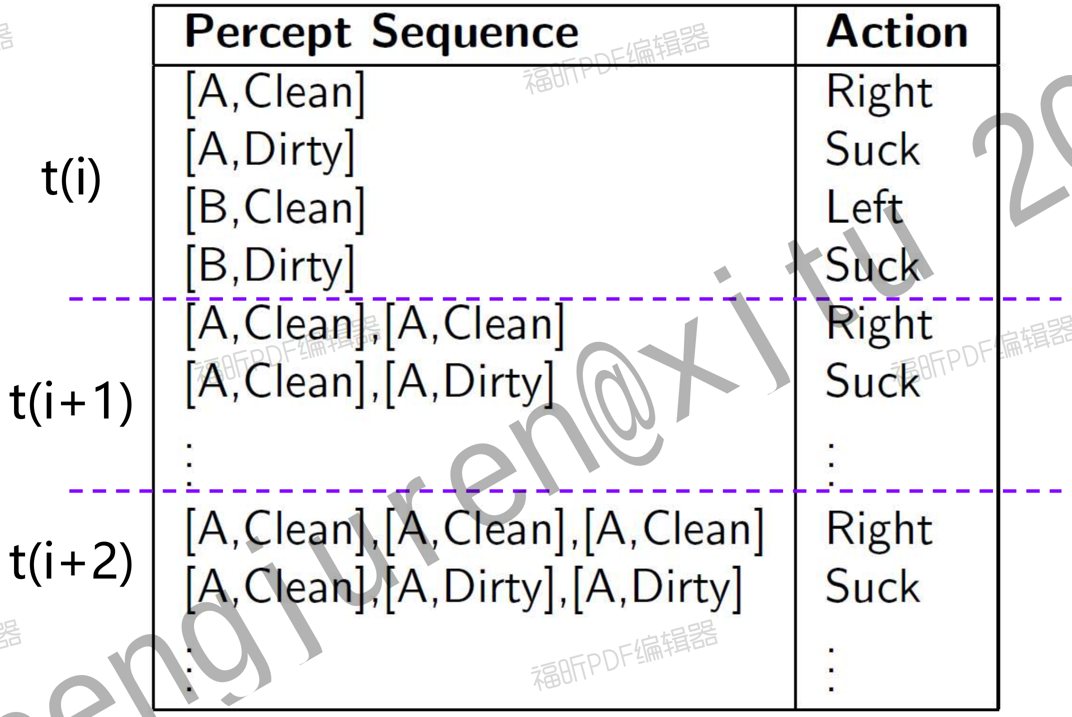

	Percept Sequence	Action
t(i)	[A,Clean]	Right
	[A,Dirty]	Suck
	[B,Clean]	Left
	[B,Dirty]	Suck
t(i+1)	[A,Clean],[A,Clean]	Right
	[A,Clean],[A,Dirty]	Suck
	⋮	⋮
t(i+2)	[A,Clean],[A,Clean],[A,Clean]	Right
	[A,Clean],[A,Dirty],[A,Dirty]	Suck
	⋮	⋮

- **What is the right way? (Look up Table or a small agent program ?)**
- **What makes an agent good or bad, intelligent or stupid ?**

Rationality

- ■ Fixed performance measure evaluates the environment sequence
  - ➢ one point per square cleaned up in time T?
  - ➢ one point per clean square per time step, minus one per move?
  - ➢ penalize for > k dirty squares?

- ■ A rational agent chooses whichever action maximizes the expected value of the performance measure given the percept sequence to date and whatever built-in knowledge the agent has.
- ■ Rational ≠ omniscient
  percepts may not supply all relevant information and all-knowing with infinite knowledge
- ■ Rational ≠ clairvoyant
  action outcomes may not be as expected
- ■ Hence, rational ≠ successful
- ■ Rational >> Information gathering,
  exploration, learning, autonomy

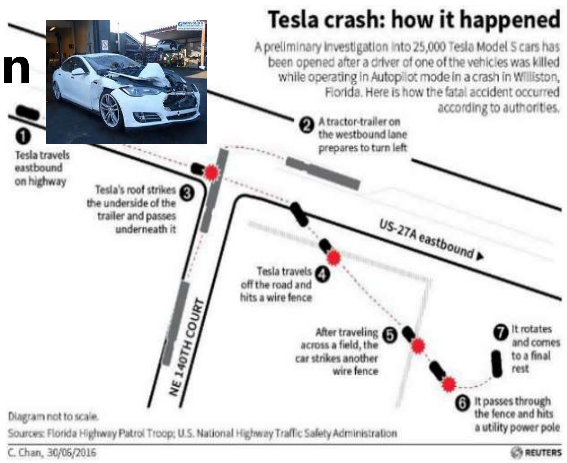

Tesla crash: how it happened

A preliminary investigation into 25,000 Tesla Model S cars has been opened after a driver of one of the vehicles was killed while operating in Autopilot mode in a crash in Williston, Florida. Here is how the fatal accident occurred according to authorities.

1 Tesla travels eastbound on highway
2 A tractor-trailer on the westbound lane prepares to turn left
3 Tesla's roof strikes the underside of the trailer and passes underneath it
4 Tesla travels off the road and hits a wire fence
5 After traveling across a field, the car strikes another wire fence
6 It passes through the fence and hits a utility power pole
7 It rotates and comes to a final rest
US-27A eastbound ▸
NE 140TH COURT

Diagram not to scale.
Sources: Florida Highway Patrol Troop; U.S. National Highway Traffic Safety Administration
C. Chan, 30/06/2016
REUTERS

Tasla‘s accident@2016

Some Cold Water: Tesla Autopilot Misclassifies Truck as Billboard

Problem: Why? How can you trust a blackbox?

A Smart Vacuum-cleaner Agent

PEAS

To design a rational agent, we must specify the task environment.

Agent Type	Performance Measure	Environment	Actuators	Sensors
Auto Vehicle	Safety Destination Profits Legality Comfort	Roads, Traffic lights, Pedestrians, Customers, Raining Snowing	Steering, Accelerator, Brake, Horn	Accelerometers, Camera, Engine Sensors, GPS, Laser
Medical diagnosis system	Healthy patient, minimize costs, lawsuits	Patient, hospital, staff	Screen display (questions, tests, diagnoses, treatments, referrals)	Keyboard (entry of symptoms, findings, patient's answers)

PEAS

To design a rational agent, we must specify the task environment.

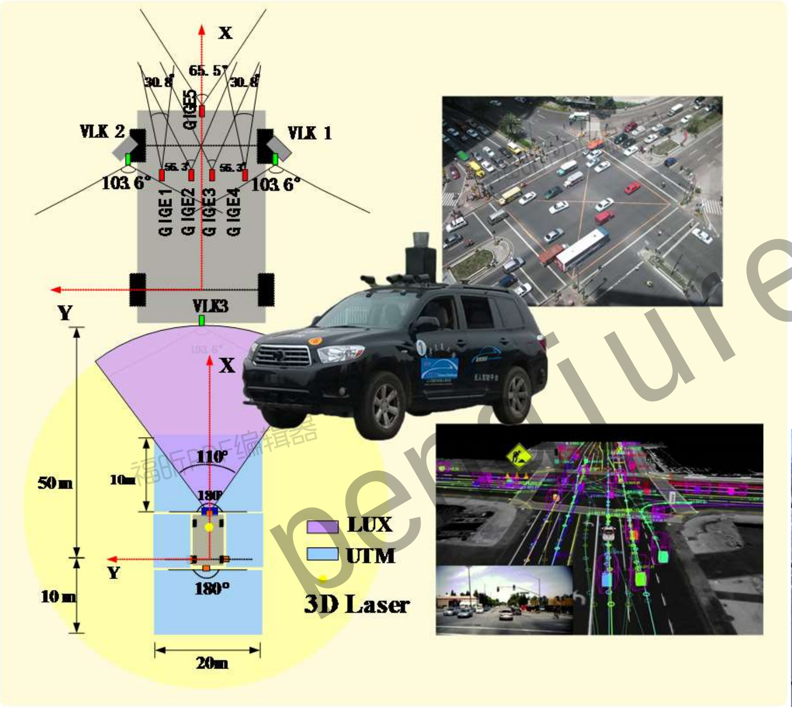

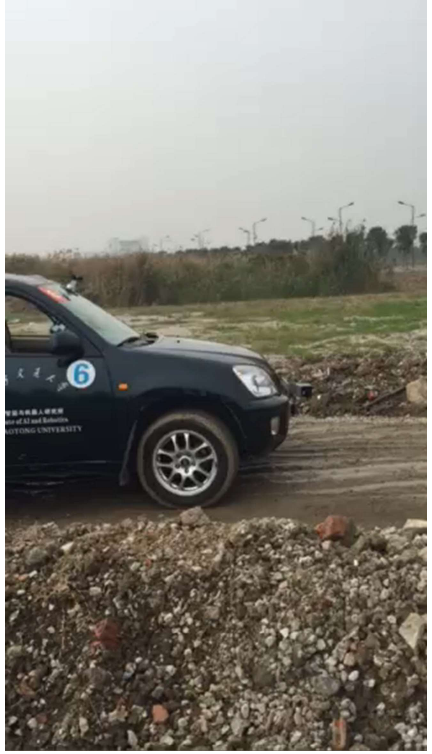

PEAS

To design a rational agent, we must specify the task environment.

Environment Types

- **Fully observable** (vs. **partially observable**): An agent's sensors give it access to the complete state of the environment at each point in time.
- **Deterministic** (vs. **stochastic**): The next state of the environment is completely determined by the current state and the action executed by the agent. (If the environment is deterministic except for the actions of other agents, then the environment is **strategic**)
- **Episodic** (vs. **sequential**): The agent's experience is divided into atomic "episodes" (each episode consists of the agent perceiving and then performing a single action), and the choice of action in each episode depends only on the episode itself.

Environment Types

Cooperative Navigation System of UAV and UGV

空地协同导航系统

Cognitive Architecture Group

Xian Jiaotong University

Environment Types

- Static (vs. dynamic): The environment is unchanged while an agent is deliberating. (The environment is semi-dynamic if the environment itself does not change with the passage of time but the agent's performance score does)
- Discrete (vs. continuous): A limited number of distinct, clearly defined percepts and actions.
- Single agent (vs. multiagent): An agent operating by itself in an environment.

Environment Types

	Chess	Go	Vehicle	Image analysis
Observability	Fully	Partially	Partially	Fully
Agents	Multi Competitive	Multi Competitive	Multi Competitive	Single
Deterministic/ Stochastic	Deterministic	Stochastic	Stochastic	Deterministic
Episodic /Sequential	Sequential	Sequential	Sequential	Episodic
Static/Dynamic	Static/Semi	Static/Semi	Dynamic	Static
Discrete /Continuous	Discrete	Discrete	Continuous	Discrete

- The environment type largely determines the agent design
- The real world is always partially observable, stochastic, sequential, dynamic, continuous, multi-agents.

Agent Types

Four basic types in order of increasing generality

- Simple reflex agents.
- Model-based agents.
- Goal-based agents.
- Utility-based agents.

All these can be turned into learning agents

Simple Reflex Agent

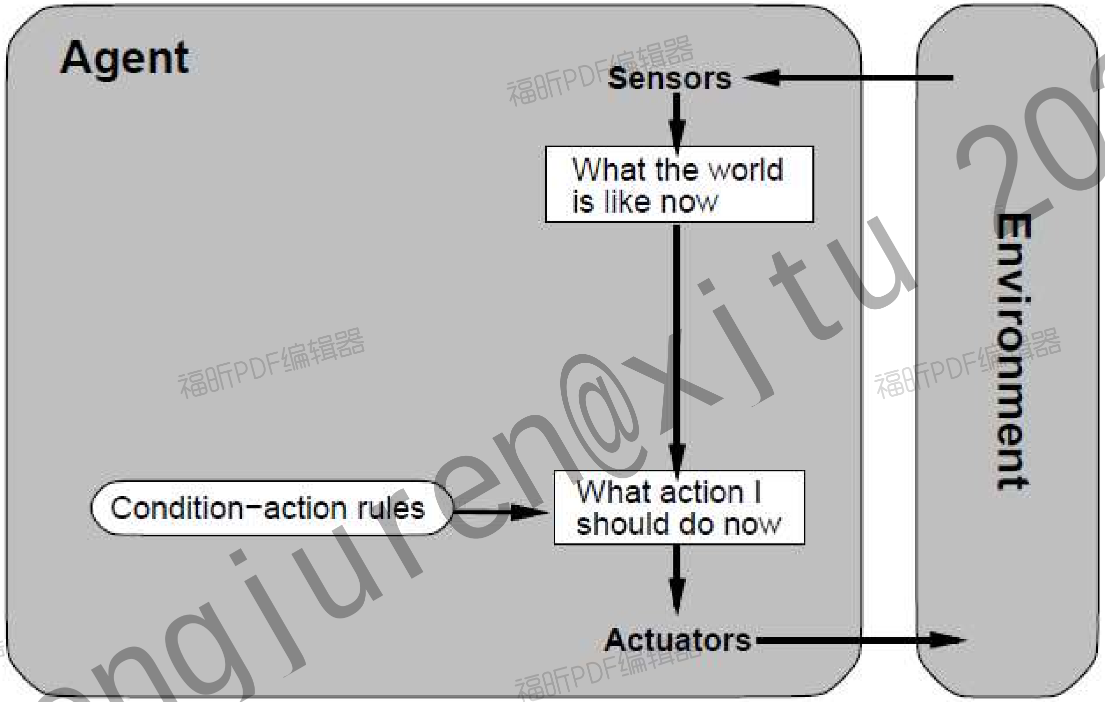

```
SIMPLE-REFLEX-AGENT(percept)
    state = INTERPRET-INPUT(percept)
    rule = RULE-MATCH(state, rules)
    action = rule.action
    return action
```

Simple Reflex Agent

- Reflex-Vacuum-Agent is a simple reflex agent.
- Actions rely purely on condition-action rules:

  if condition then action.
- Also called memory-less or state-less
- Works only if the correct decision can be made on the basis of only the current percept.
- Works only if the environment is fully observable.
- Often trapped in infinite loops if the environment is partial observable.

Model-Based Agent

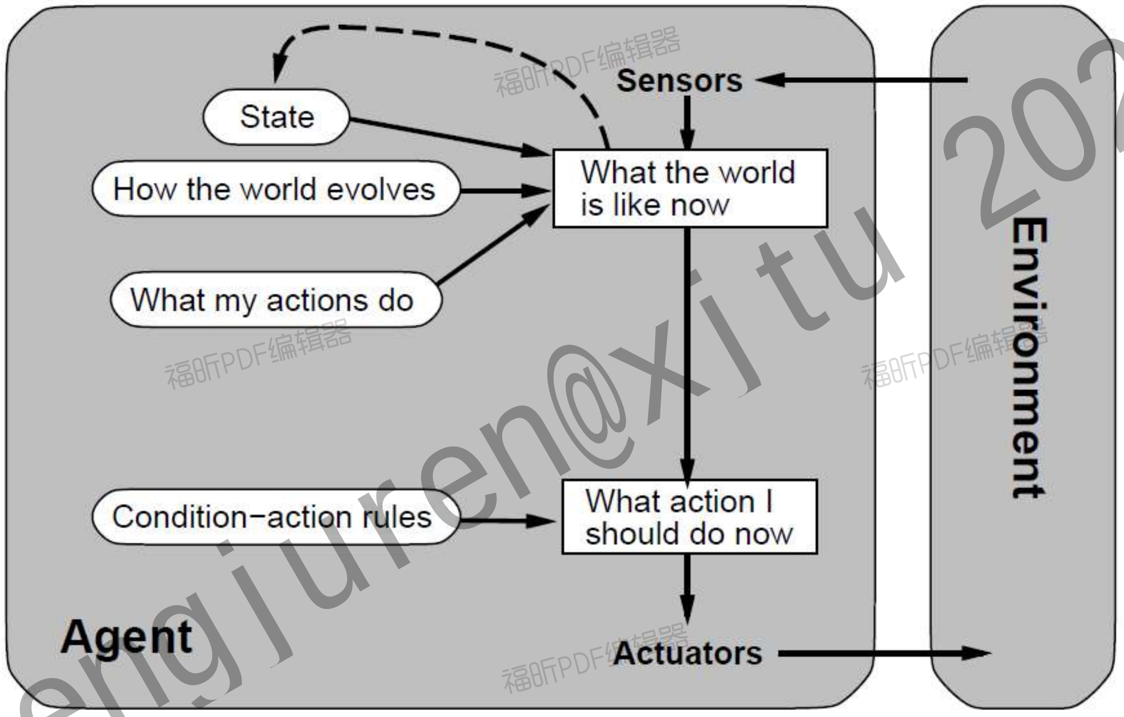

Sensors

State

How the world evolves

What my actions do

What the world is like now

Condition−action rules

What action I should do now

Agent

Actuators

Environment

```
MODEL-BASED-REFLEX-AGENT(percept)

state = UPDATE-STATE(state, action, percept, model)
rule = RULE-MATCH(state, rules)
action = rule.action
return action
```

Model-Based Agent

- Handle partial observability by keeping track of the part of the world it can't see now.
- Maintain internal states to model the world.
- The model of the world represents the agent's best guess(or prediction), can't be exact.
- Internal states can also be used to maintain the status of the agent instead of the world.

Goal-Based Agent

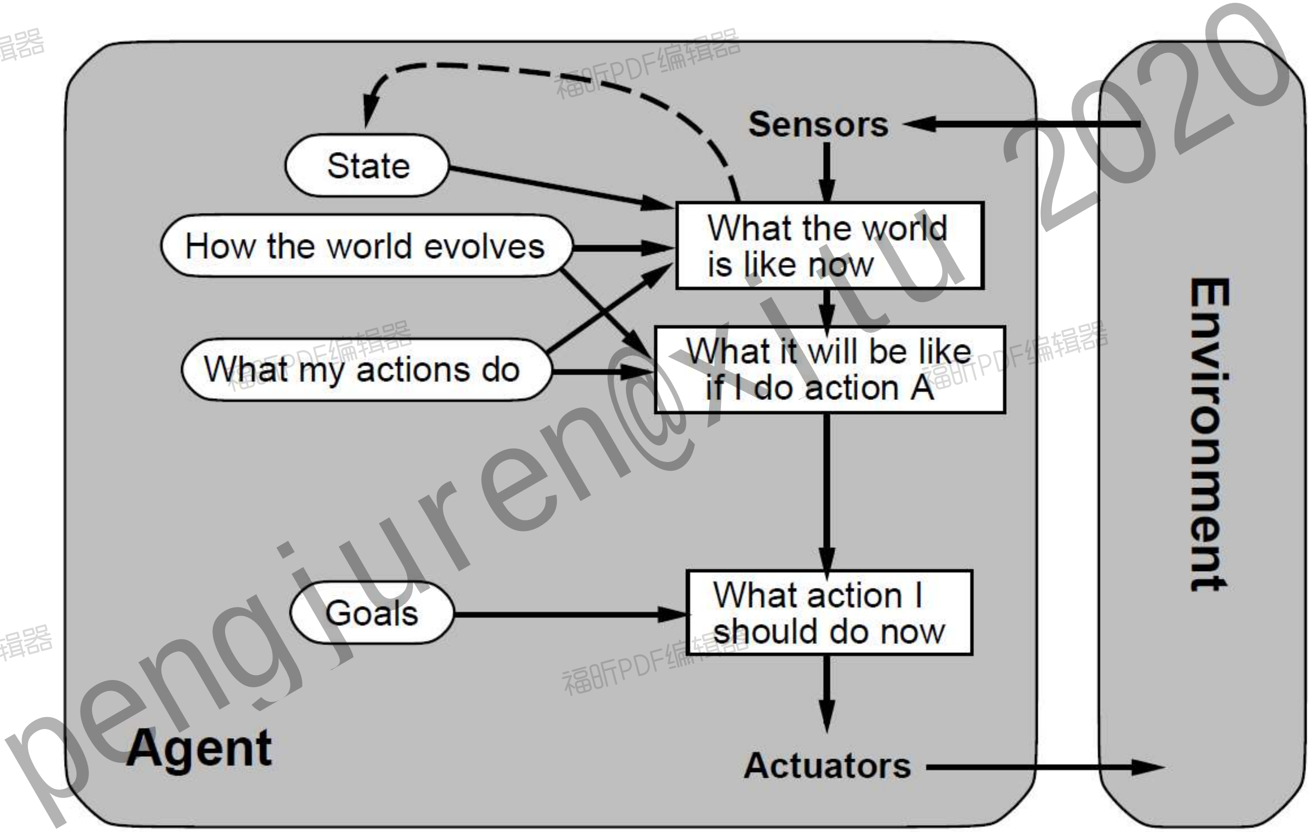
Sensors
State
How the world evolves
What my actions do
What the world is like now
What it will be like if I do action A
Goals
What action I should do now
Agent
Actuators
Environment

- Instead of using condition-action rules, the agent uses goals to decide what action it does.
- Search (Chapters 3, 4 and 5) and Planning (Chapters 10 and 11).

Utility-based Agent

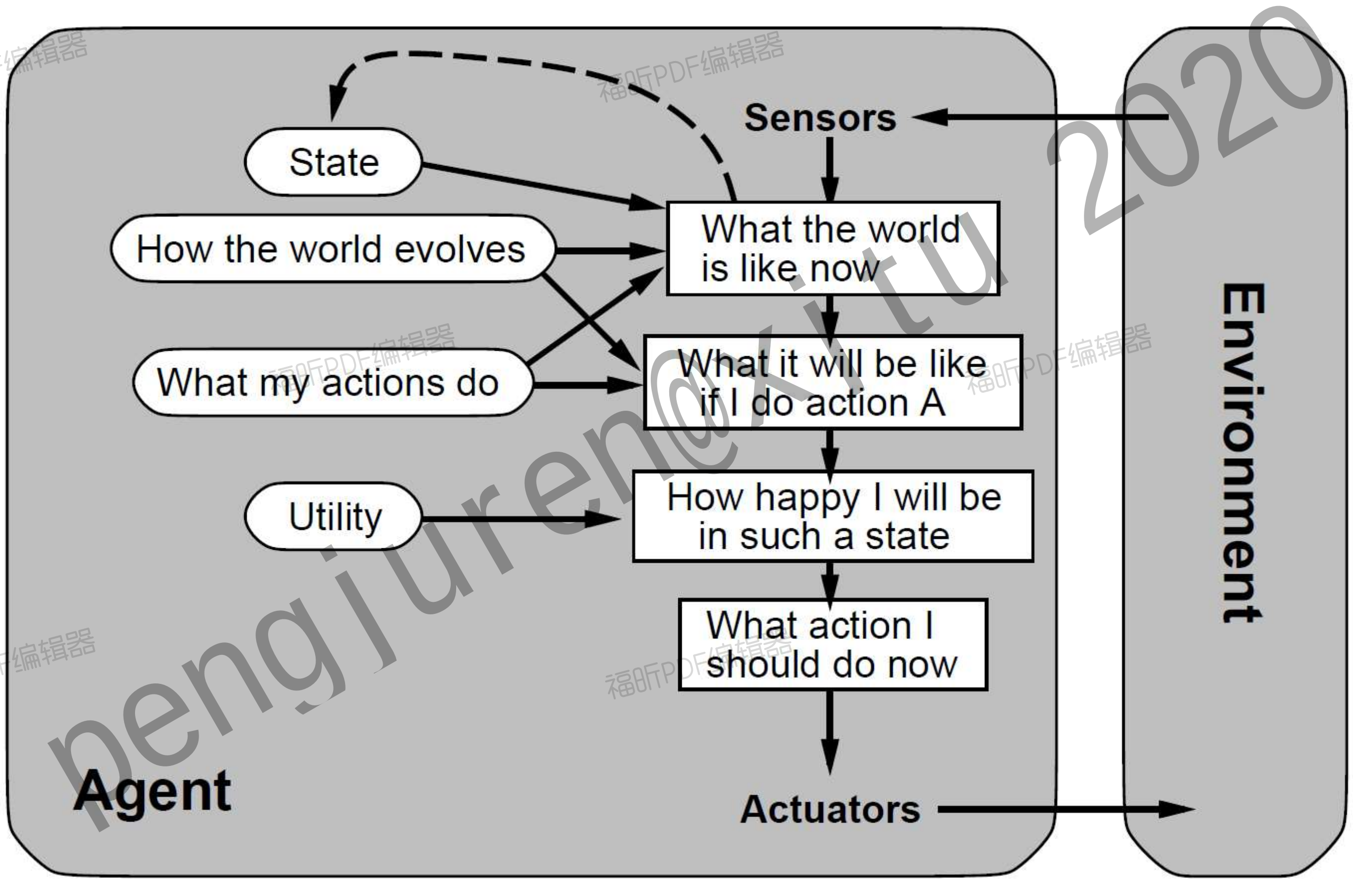
Sensors
State
How the world evolves
What my actions do
Utility
What the world is like now
What it will be like if I do action A
How happy I will be in such a state
What action I should do now
Agent
Actuators
Environment

- Utility function: Happiness of the agent.
- Maximizing the expected utility.

Learning Agents

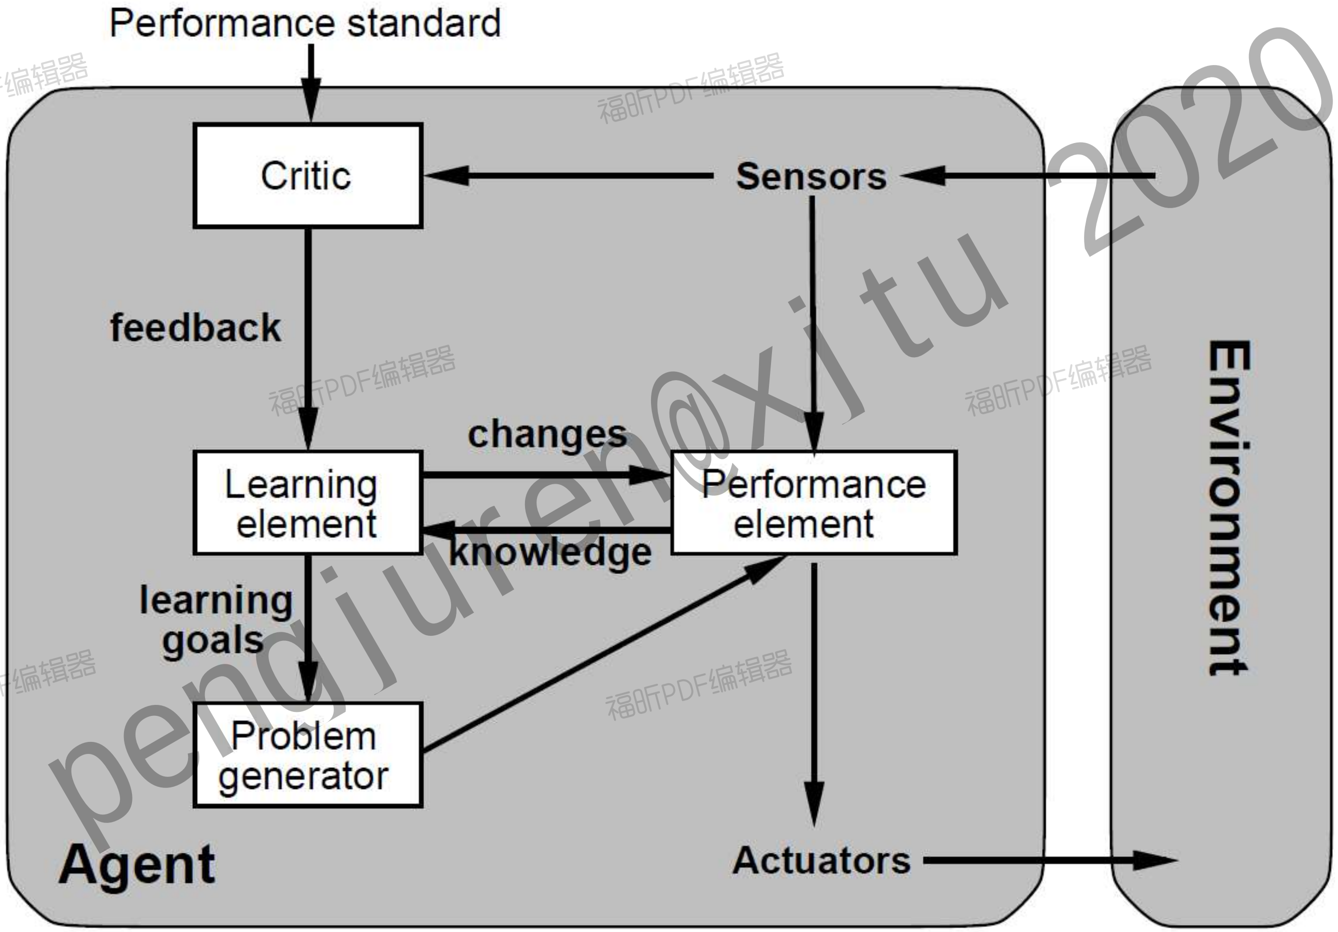

- Learning from rewards (or penalty).
- Learning techniques form another field called machine learning.

Three ways for representation

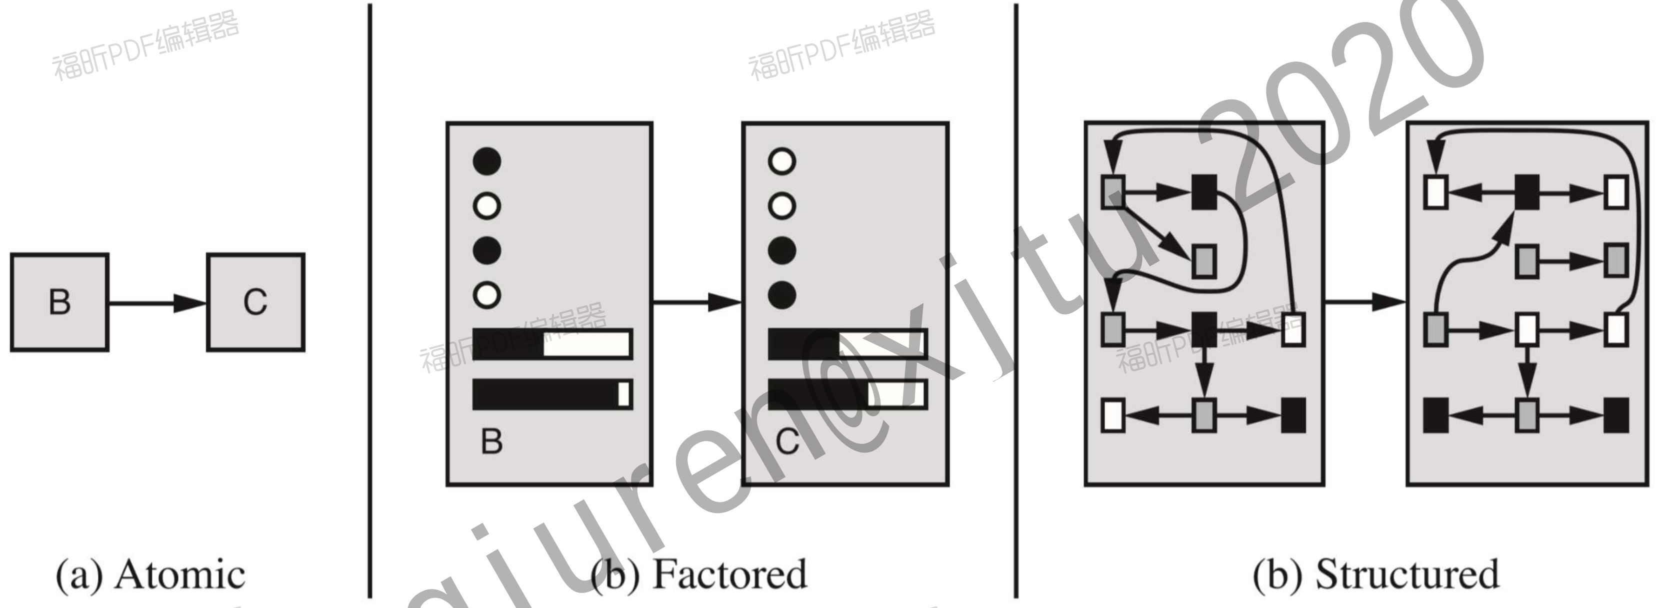

- **Atomic representation:** a state (such as B or C) is a black box with no internal structure;
- **Factored representation:** a state consists of a vector of attribute values; values can be Boolean, real-valued, or one of a fixed set of symbols.
- **Structured representation:** a state includes objects,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

Summary

- Agents interact with environments through actuators and sensors.
- The agent function describes what the agent does in all circumstances.
- The performance measure evaluates the environment sequence.
- A perfectly rational agent maximizes expected performance.
- Agent programs implement (some) agent functions.
- PEAS descriptions define task environments.
- Different environments & agent types.
  - observable? deterministic? episodic? static? discrete? single-agent?
  - reflex, reflex with state, goal-based, utility-based, learning agents
- All agents can improve their performances through learning.