Autonomous and Adaptive Systems

Autonomy

Mirco Musolesi

mircomusolesi@acm.org

Intelligent Robots

- A robot is a special type of agent: it is *physically situated* in the real world, while a software agent is situated in a virtual world like for example the World Wide Web, a simulation (think about videogames or confines of a software system).
- An *intelligent robot* is able to sense its environment and act on it in order to achieve its goal.

Typical Subsystems of Robots

- The five major subsystems of robots are:
  - **Effectors**: subsystems for movement, such as wheels or legs and manipulation/grasping.
  - **Perception**: subsystems for sensing that provide a robot with the equivalent of the five senses: sight, hearing, smell, taste and touch.
  - **Control**: subsystems for controlling inner state of the robots, planning and execution of actions, etc.
  - **Communications**: subsystems for communication with the other agents and operators.
  - **Power**: subsystem for providing power to all the other subsystems.

Unmanned Ground Systems

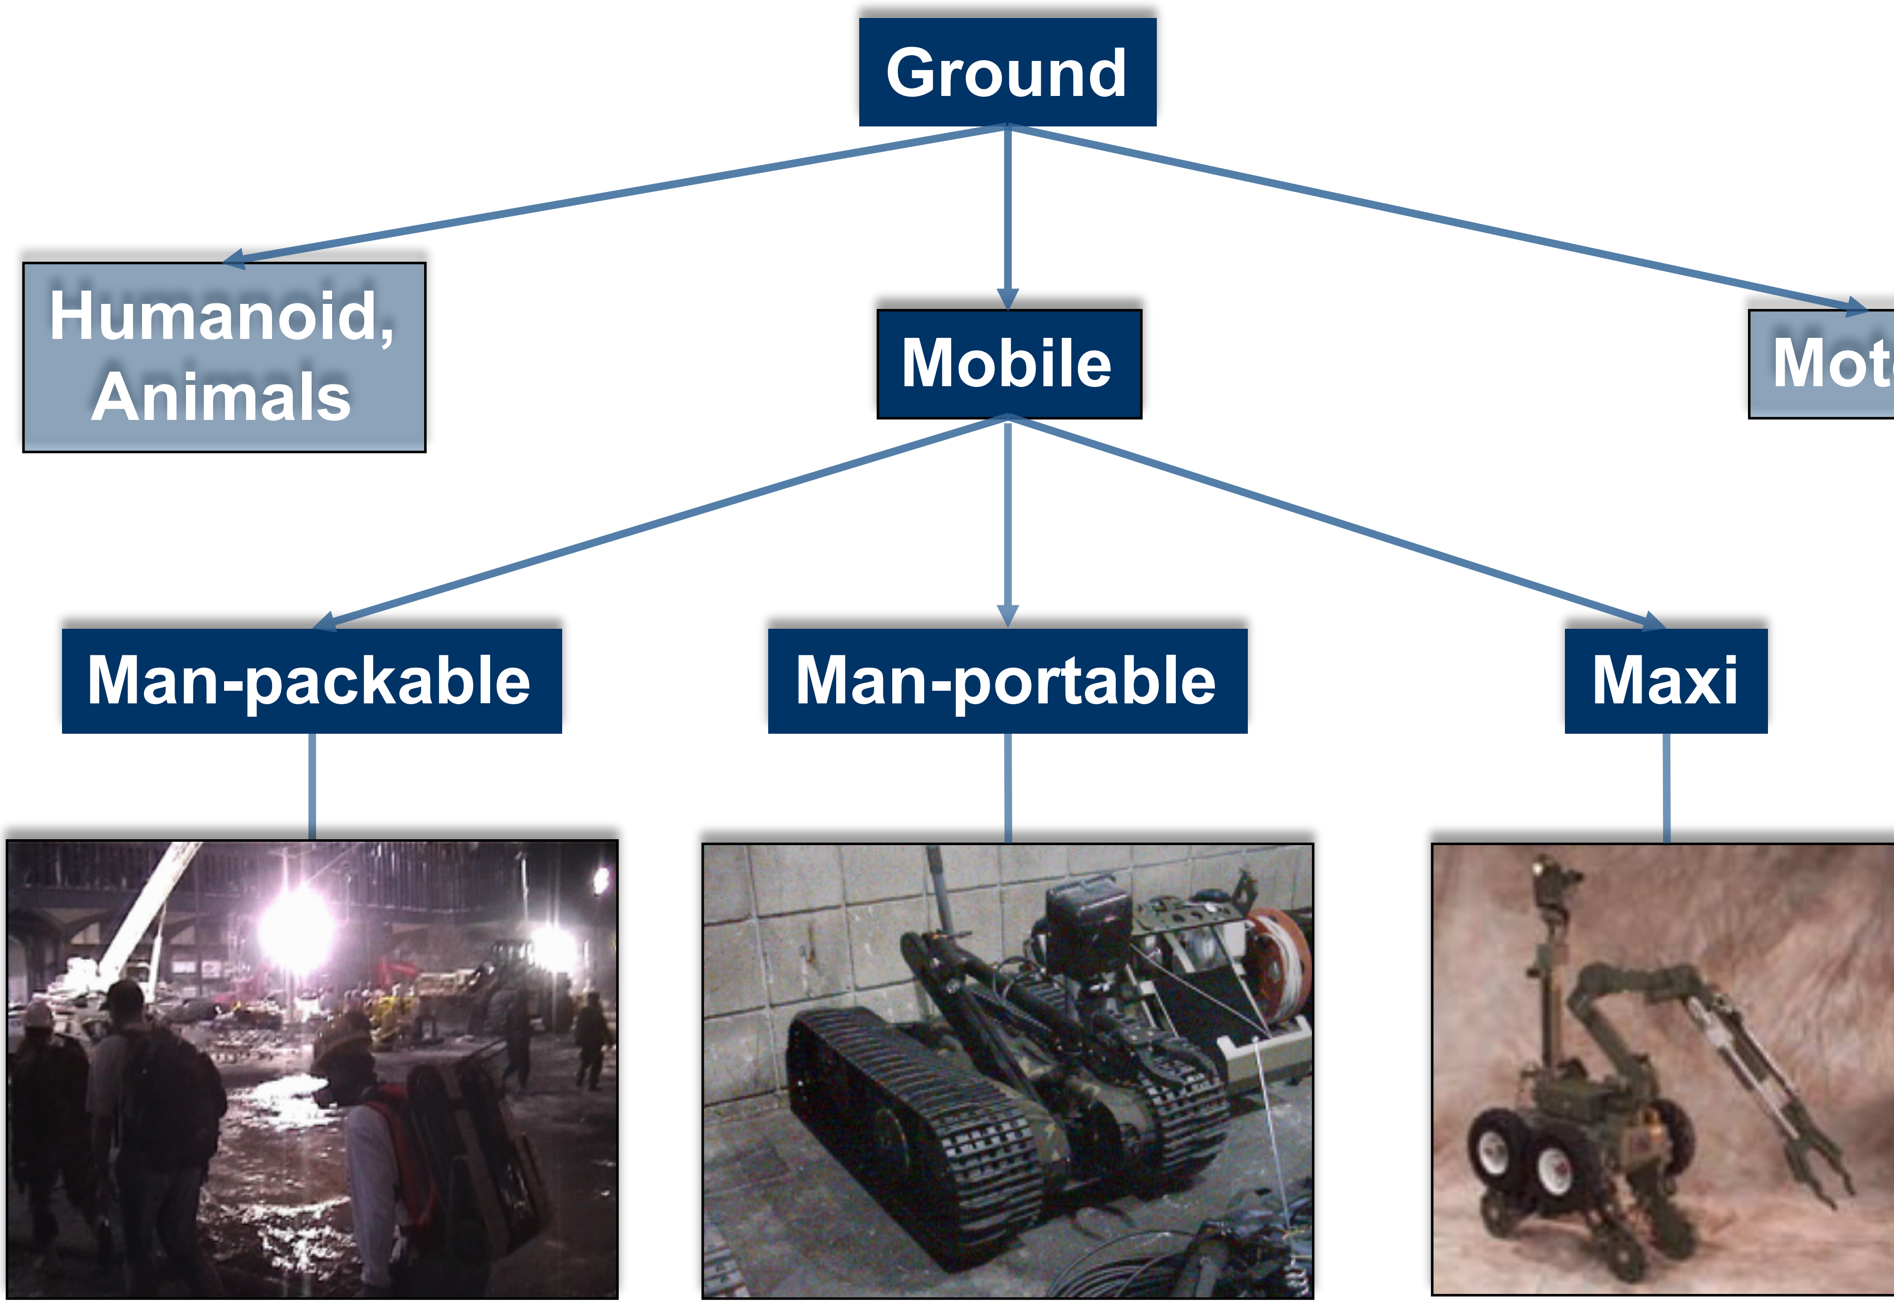

Source: Robin R. Murphy. Introduction to AI Robotics. MIT Press. 2019.

Modalities of Unmanned Ground Systems

- **Humanoid**: anthropomorphic robots (Honda P3, Sony Asimo, Alderaban Nao, etc.).
- **Mobile robots**: non anthropomorphic robots (Roomba, NASA Rovers, etc.).
- **Motes**: miniature robots (unattended ground sensors).

Unmanned Aerial Systems

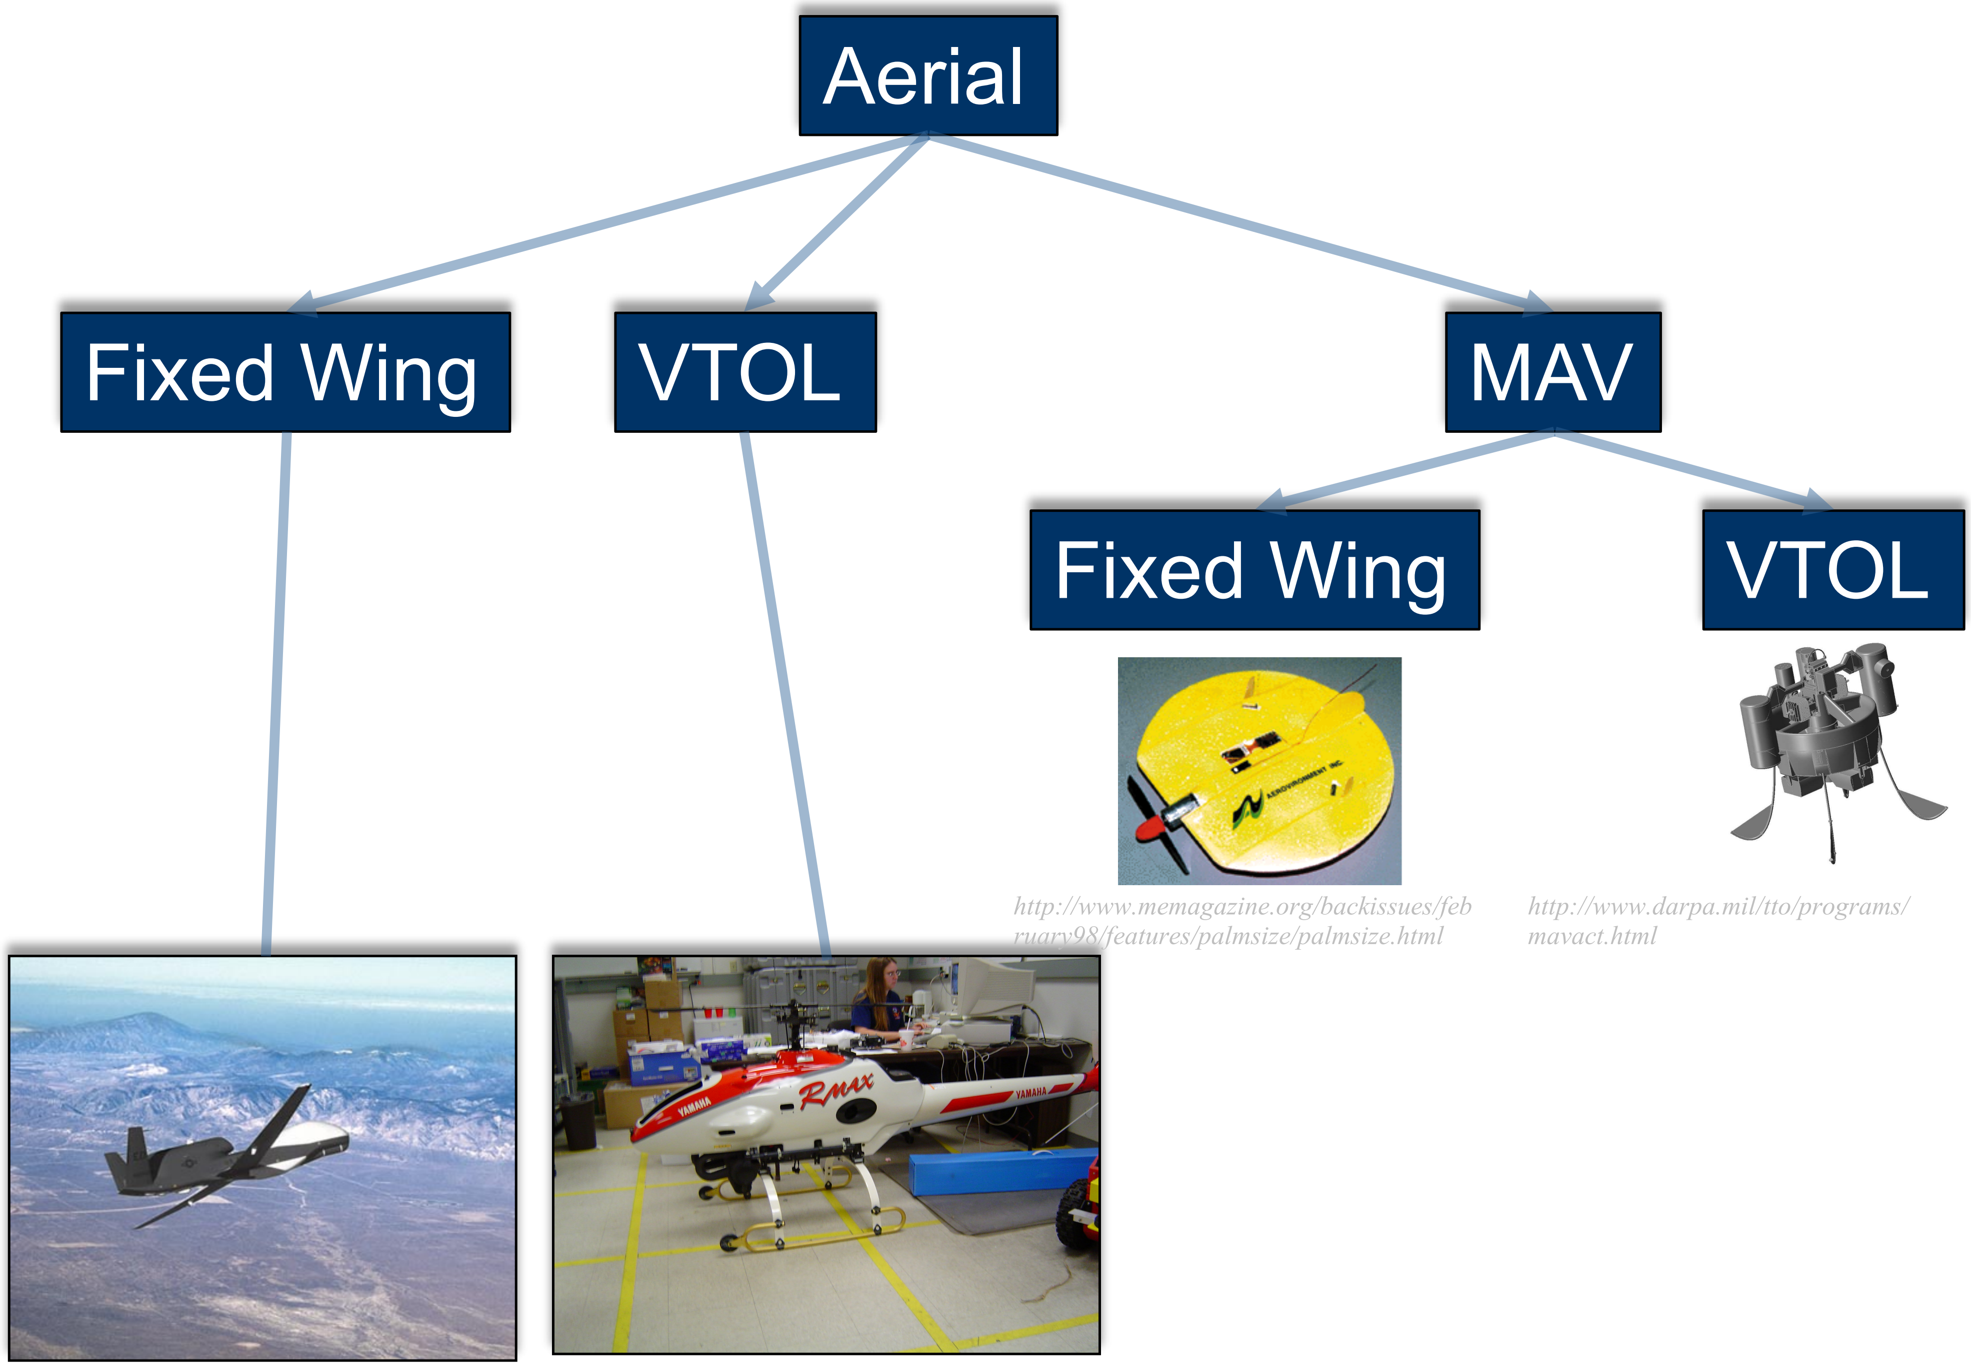

Source: Robin R. Murphy. Introduction to AI Robotics. MIT Press. 2019.

Modalities of Unmanned Aerial Systems

- **Fixed wing-aircrafts**: aircrafts that look and like as planes (US Defense Predators and Global Hawks).
- **Rotor-crafts**: vertical take-off and landing (VTOL) platforms (Yamaha RMAX)
- **Micro UAVs (MAVs)**: systems that less two meters in any characteristic dimension for example wing span, rotor blade length or fuselage. MAVs can be either fixed-wing or rotorcraft.

Unmanned Marine Systems

- MARINE
  - SURFACE
  - UNDERWATER
    - AUV Autonomous Underwater Vehicle
    - ROV Remotely Operated Vehicle
  - HYBRID

Source: Robin R. Murphy. Introduction to AI Robotics. MIT Press. 2019.

Modalities of Unmanned Marine Systems

- **Surface systems**: systems/vehicles that move over the surface of water.
- **Unmanned underwater systems**: systems/vehicles that submerge.
  - *Autonomous underwater vehicles (AUV)*: systems/vehicles that are not tethered and are not in constant communication with an operator.
  - *Remotely operated vehicle (ROV)*: systems/vehicles that are tethered and can be controlled in real time.
- **Hybrid systems**: systems/vehicles that are able to move on the surface and submerge.

Key Questions

- What is the difference between automation and autonomy?
- Why does it matter that there is a difference between autonomy and automation?

Examples

Automation

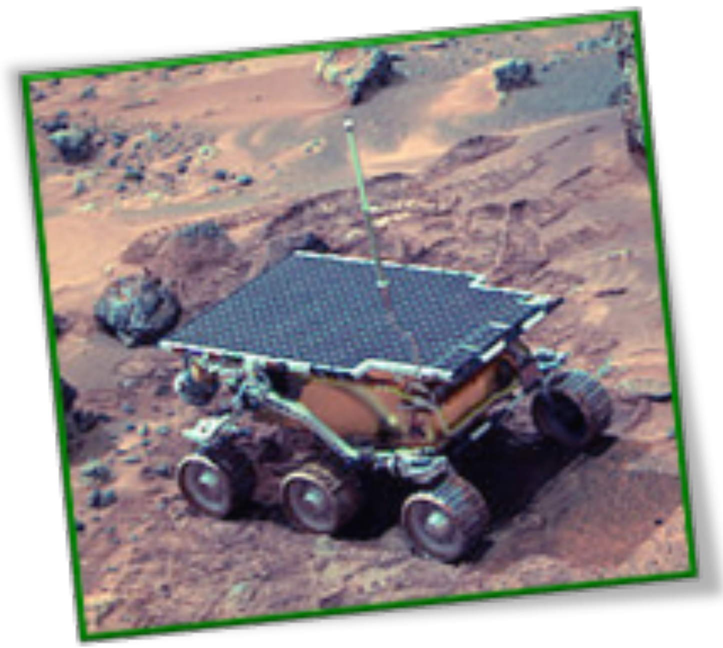

Autonomy

Source: Robin R. Murphy. Introduction to AI Robotics. MIT Press. 2019.

What is the Difference between Automation and Autonomy?

- One possible answer is that automation is about robots as tools and autonomy is about robots as agents.
- However, the difference is very blurred in practice. One possible way is to discuss the differences (in a continuum) in terms of four characteristics:
  - Does the system focus on execution of plans versus generation of plans?
  - Does the system use a close or an open-world model?
  - Does the system use deterministic or non-deterministic algorithms?
  - Does the system manipulates signals or symbols/concepts?

Choice between Automation and Autonomy

- This is a key design choice: elements of autonomy will lead to different types of failure modes, different types of programming, etc.
- Thinking about vulnerabilities and implications is fundamental.
- Autonomy also raises very important ethics questions about responsibility, liability, agency, etc.
- Most of the systems in reality will be a blend of both.

Choice between Automation and Autonomy

- We essentially lack design rules: it is very difficult to say in general when autonomy is better than automation.

- Autonomy is problematic in terms of testing (non-deterministic behaviour).

- Automation lacks adaptivity:

    - Think about a welding robot: if something is out of place, the robot will not adapt easily.

    - Adaptation (and the possibility of planning ahead also in presence of unexpected situations) is one of the most important characteristics of autonomous systems.

Models: Open vs Close World Models

- In order for a machine to understand the world, there has to be a computationally relevant representation called a world model.
- If the robots has only navigational capabilities, the world may consist of a map with objects.
- A representation of the world is not necessarily spatial.
  - If the robot interacts with other agents, it may need to represent what it thinks about the other agents (in terms of beliefs, intentions, etc.).
- The world can be a collection of world models.
- A world-model maybe pre-programmed into a robot, may be learned by a robot (or a combination of both).
- Examples of pre-programmed information are: rules, constraints, maps, etc.

Closed-World and Open-World Assumptions

- World models are classified as being *closed-world* or *open-world*.
- The close-world logic assumption says that everything possible is known.
    - In formal logic, this means that any object, condition or event that is not the database is false.
    - Closed-world assumption in *Jurassic Park*.
        - Also unpredictability of the system (see chaos theory)!
- An algorithm for a robot operating under the open-world assumption assumes that the list of possible states, objects or conditions cannot be completely specified.
    - See applications of machine learning for adding knowledge or function approximation in reinforcement learning for states that might not be present in a q-table.

The Frame Problem

- Related to the closed-world assumption is the frame problem, which refers to the problem of correctly identifying what is unchanging in the world and, thus, does not require constant updating, thereby reducing computation.

- Typically, control-theoretic approaches to robotics are based on closed-world models, whereas AI-based ones on open-world models.

Example: Industrial Robot

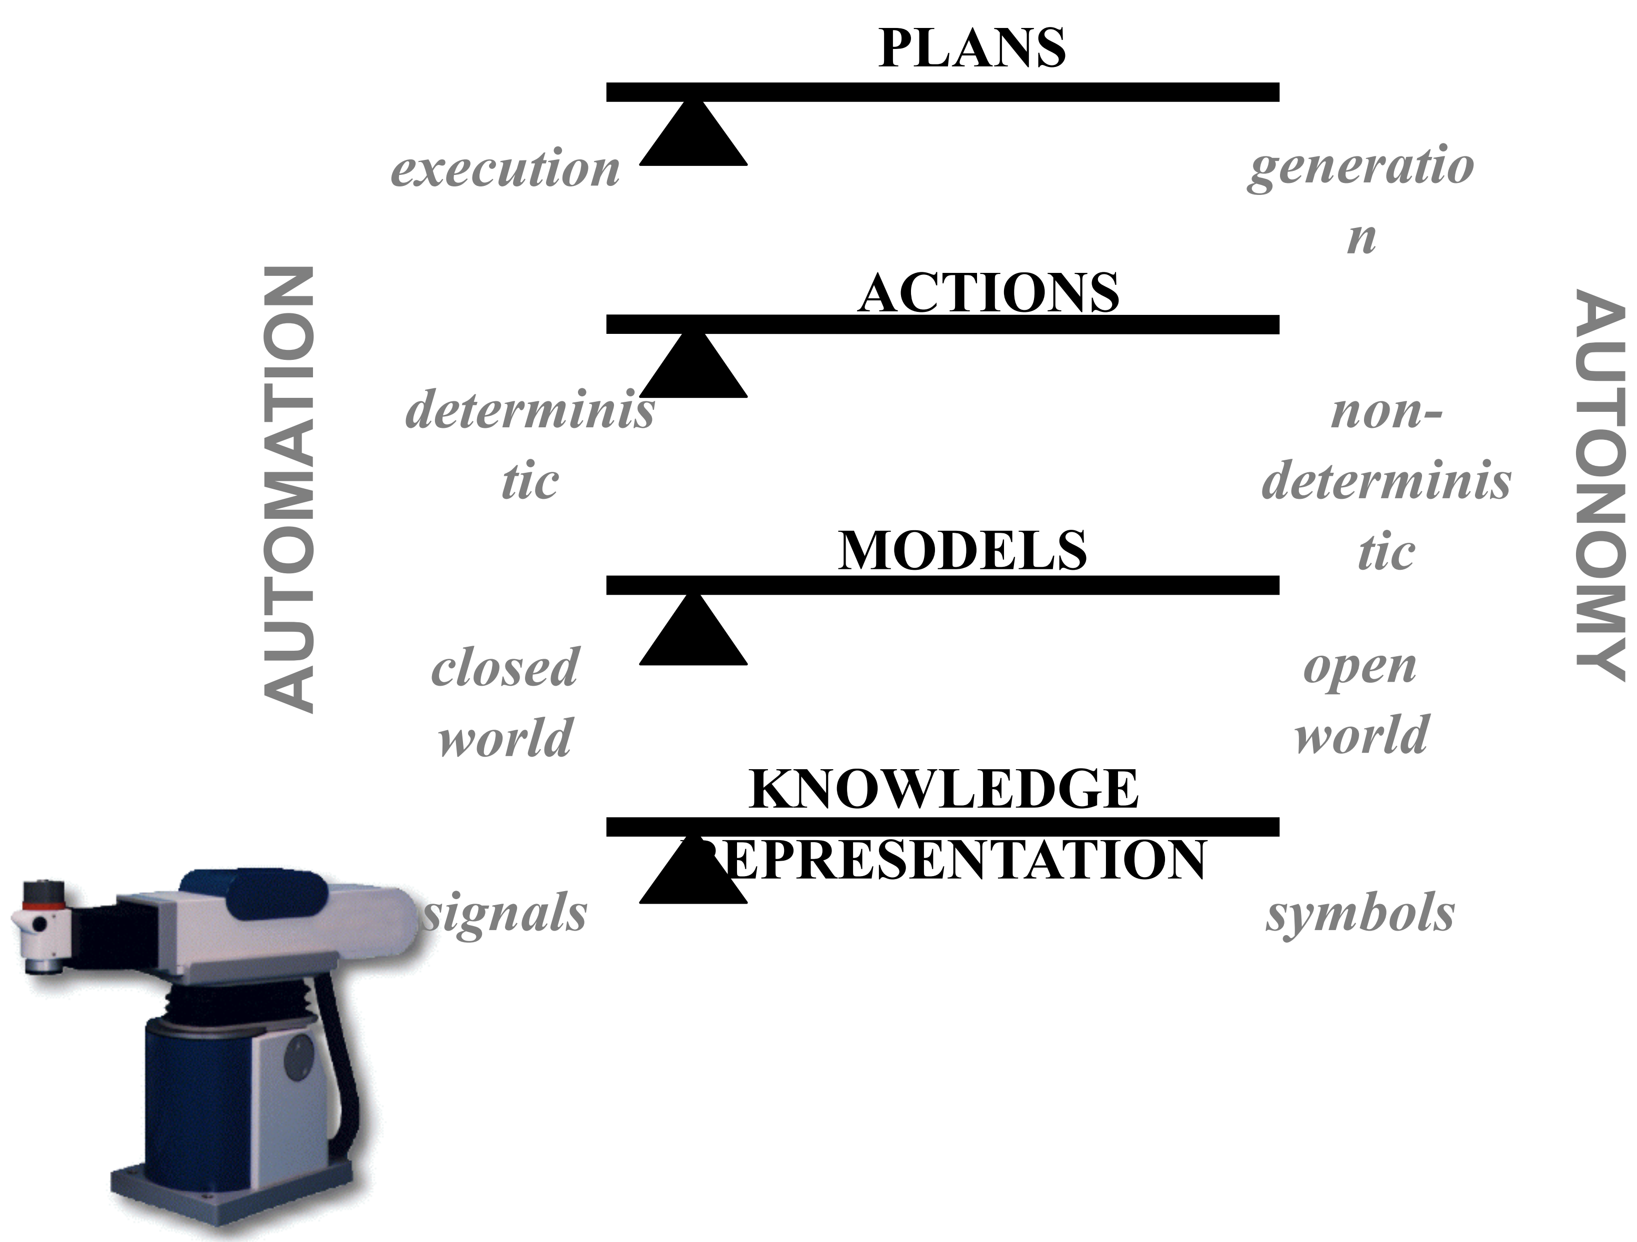

Source: Robin R. Murphy. Introduction to AI Robotics. MIT Press. 2019.

Example: Mars Rover

PLANS

execution — *generation*

ACTIONS

deterministic — *non-deterministic*

MODELS

closed world — *open world*

KNOWLEDGE REPRESENTATION

signals — *symbols*

AUTOMATION — AUTONOMY

Source: Robin R. Murphy. Introduction to AI Robotics. MIT Press. 2019.

Example: Roomba

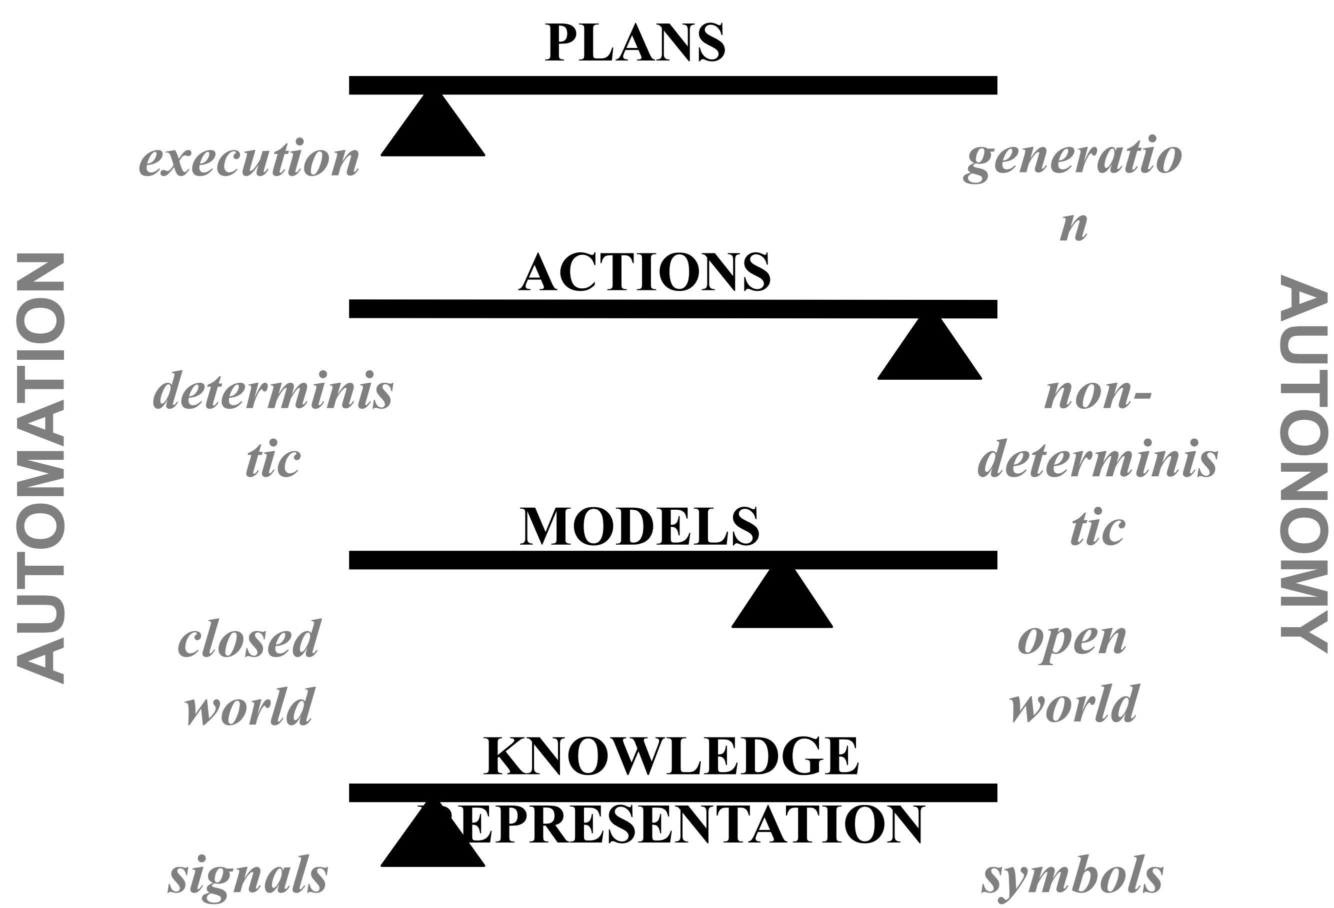

Source: Robin R. Murphy. Introduction to AI Robotics. MIT Press. 2019.

Bounded Rationality

- The concept of *bounded rationality* was introduced by Herbert Simon, one of the AI pioneers:
  - Decision making capability of all agents, human or artificial, is limited by how much information they have, their computational abilities and the amount of time they have available to make a decision.
- While a robot may dynamically adapt or replan to overcome the occurrence of unexpected events, it cannot go beyond to what it was programmed for (including machine learning algorithms).

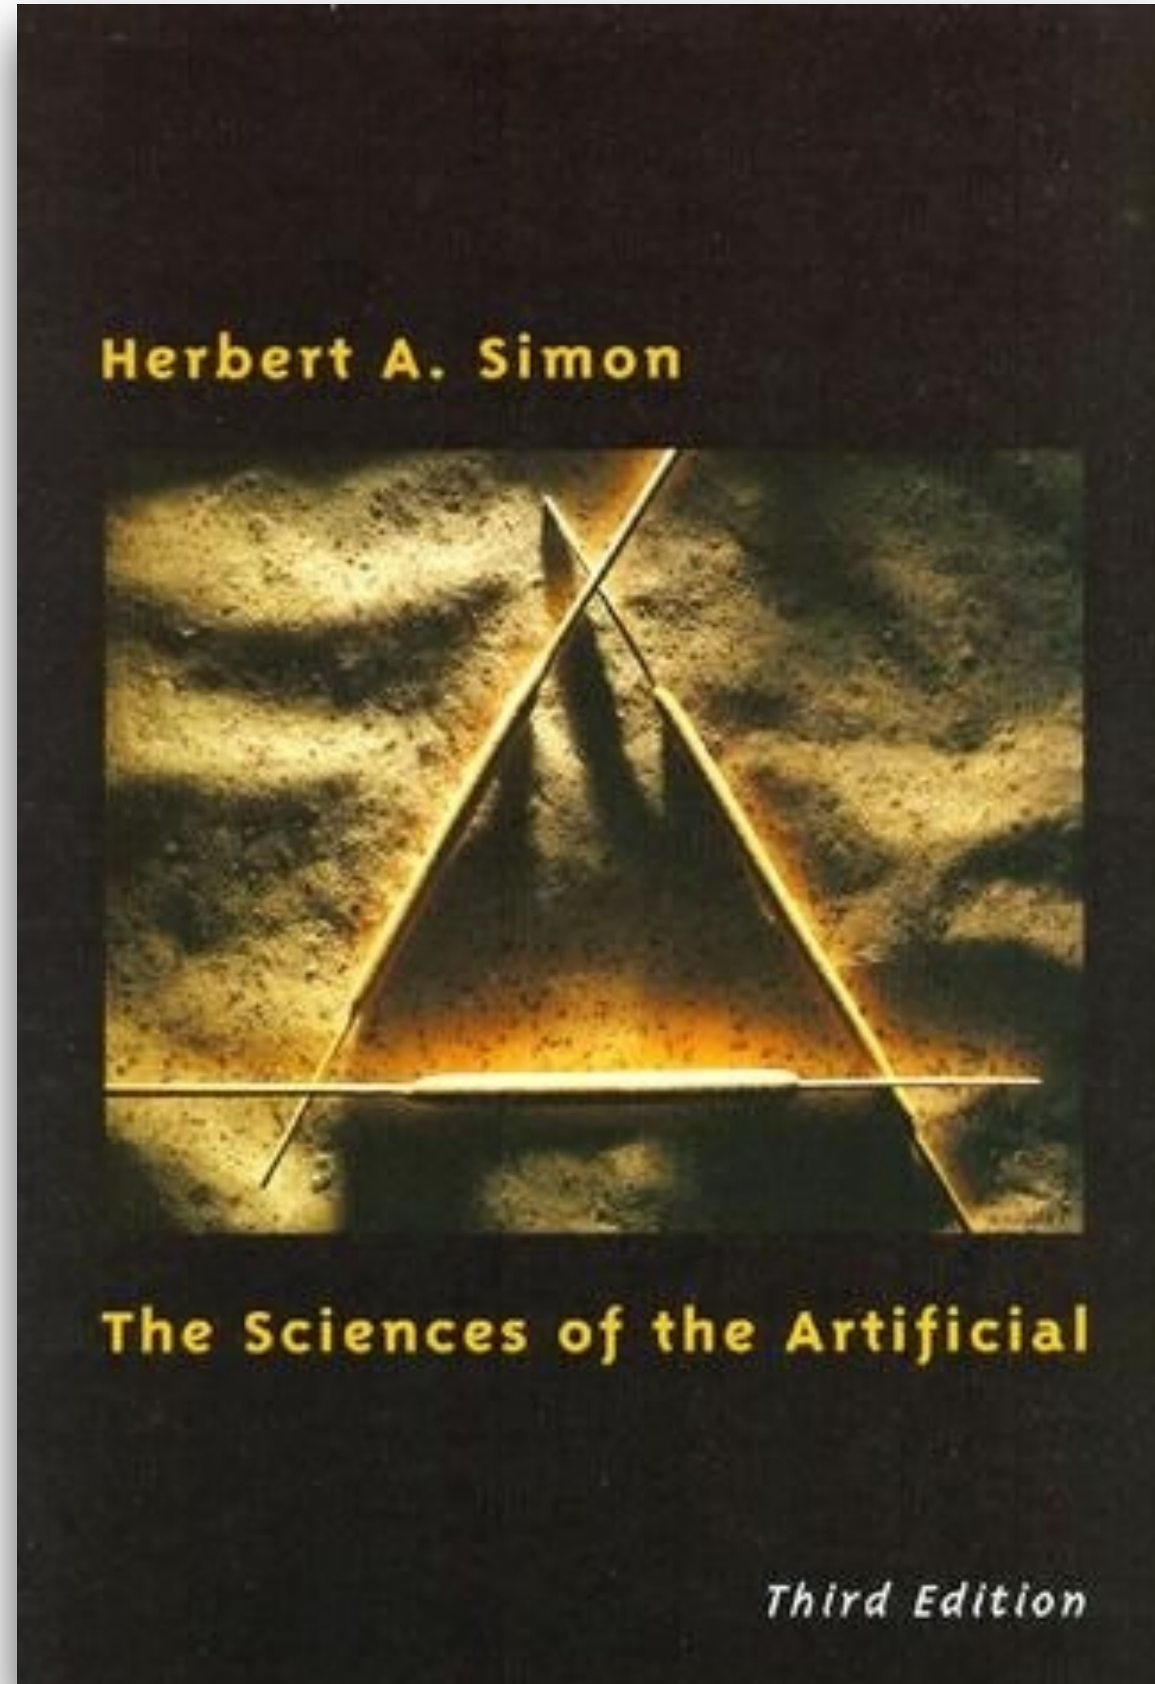
Herbert A. Simon
The Sciences of the Artificial
Third Edition

Self-driving Cars

- Self-driving cars are very special case of unmanned vehicles.
  - with a human (or humans) on them.
- One of the most interesting applications of autonomous systems technology.
- Big economic impact:
  - Urban planning (think about planning).
  - New applications also given the fact people will have more time during trips.

Source: America’s Electric Light and Power Companies (circa 1950s).

https://www.youtube.com/watch?v=TsaES--OTzM

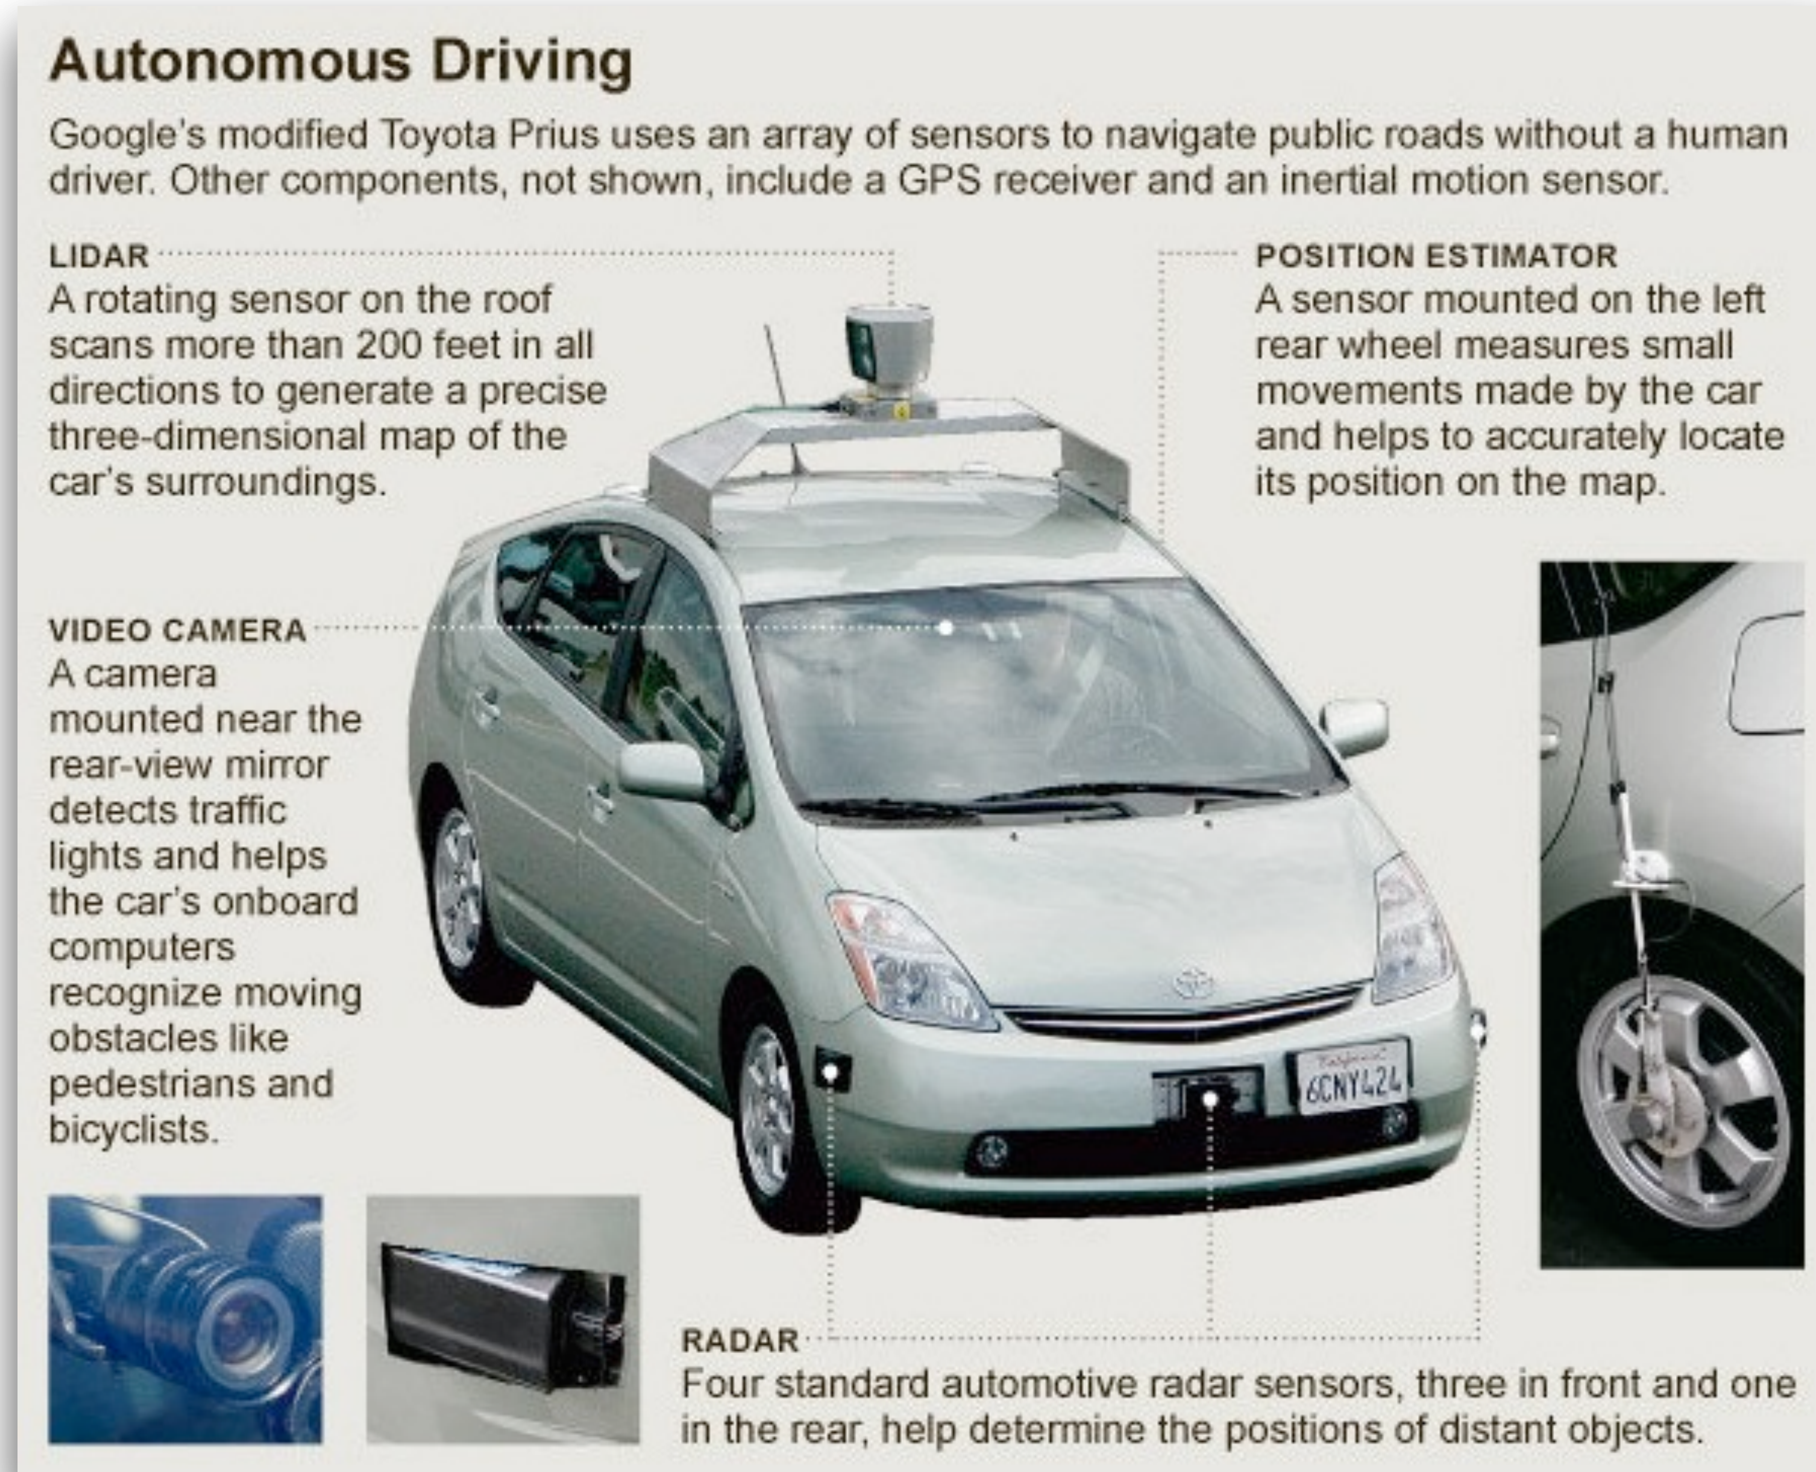
Autonomous Driving
Google's modified Toyota Prius uses an array of sensors to navigate public roads without a human driver. Other components, not shown, include a GPS receiver and an inertial motion sensor.
LIDAR
A rotating sensor on the roof scans more than 200 feet in all directions to generate a precise three-dimensional map of the car's surroundings.
POSITION ESTIMATOR
A sensor mounted on the left rear wheel measures small movements made by the car and helps to accurately locate its position on the map.
VIDEO CAMERA
A camera mounted near the rear-view mirror detects traffic lights and helps the car's onboard computers recognize moving obstacles like pedestrians and bicyclists.
RADAR
Four standard automotive radar sensors, three in front and one in the rear, help determine the positions of distant objects.

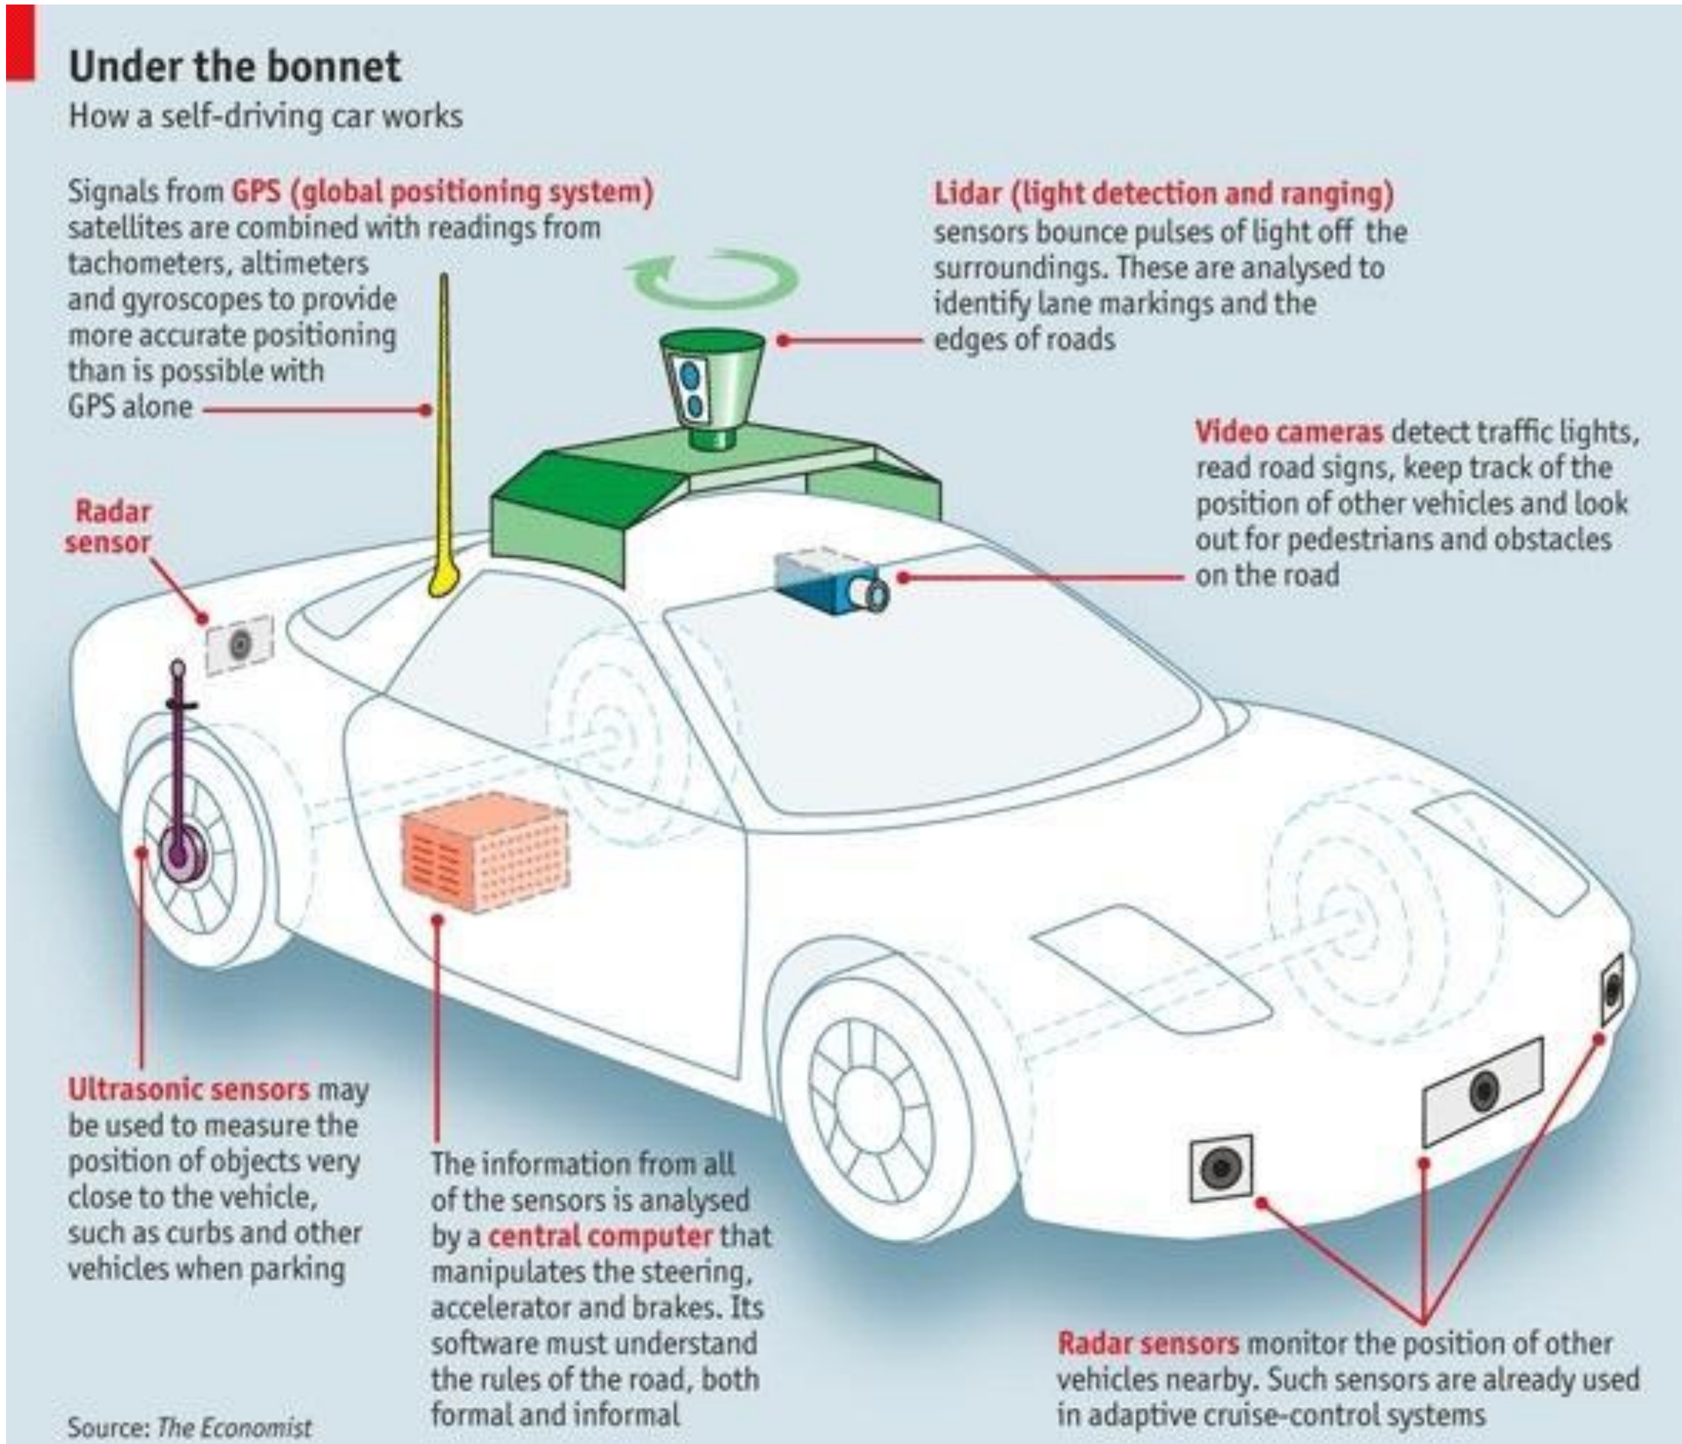

Under the bonnet
How a self-driving car works
Signals from GPS (global positioning system) satellites are combined with readings from tachometers, altimeters and gyroscopes to provide more accurate positioning than is possible with GPS alone
Lidar (light detection and ranging) sensors bounce pulses of light off the surroundings. These are analysed to identify lane markings and the edges of roads
Video cameras detect traffic lights, read road signs, keep track of the position of other vehicles and look out for pedestrians and obstacles on the road
Radar sensor
Ultrasonic sensors may be used to measure the position of objects very close to the vehicle, such as curbs and other vehicles when parking
The information from all of the sensors is analysed by a central computer that manipulates the steering, accelerator and brakes. Its software must understand the rules of the road, both formal and informal
Radar sensors monitor the position of other vehicles nearby. Such sensors are already used in adaptive cruise-control systems
Source: The Economist

References

- Robin R. Murphy. Introduction to AI Robotics. Second Edition. MIT Press. 2019.
- Hod Lipson and Melba Kurman. Driverless. Intelligent Cars and the Road Ahead. MIT Press. 2017.