Introduction to Robotic Systems

What is a Robot?

The term *robot* was first used in a 1920 science fiction play by the Czech writer Karel Čapek about an inventor who creates a race of artificial people he called *roboti* (translated as forced laborers) to serve humans. In the play, the robots become tired of working for their owners and mount a rebellion that leads to the extinction of the human race!

We all can recognize a robot when we see one, such as *R2-D2* from *Star Wars*. But what is it that is so unique? How is it different than, say, an appliance like a dishwasher? Both a dishwasher and a robot can be decomposed into similar electrical and mechanical parts (actuators, sensors, micro-processors, etc.). Intuitively, we expect that robots are "smarter" and "more capable" than a dishwasher (which has no hope of ever doing anything more than washing dishes). Indeed, robots are largely defined based on their capabilities. Consider, for example, the vacuum cleaning "Roomba" robot. At first glance it seems that the Roomba is not much different than a dishwasher – rather than cleaning dishes it cleans floors. The engineers designing your dishwasher assumed that most dishes, glassware and silverware would be of a certain size and they carefully designed a system of spraying arms, heating elements, hoses and gaskets, etc. that would clean your dishes if you loaded them according to the manufacturer's instructions. Even slight deviations from their planned operating procedure can cause the machine to malfunction (e.g., placing a big pot over the spray arm, or forgetting to insert detergent). On the other hand, cleaning the floor of a house is a much more complex task! To accomplish this the robot must be able to: move over various terrain (e.g., carpets, tiles, hardwood floors), navigate an unfamiliar landscape that might be constantly changing (e.g., as chairs and other furniture get moved around each day), avoid obstacles (e.g., walls, doors, shoes), and return to its charging station once its task is complete. These observations lead us to the following definition of a robot that we will use in this course:

A machine that can perform complex tasks and adapt to its environment

The Origin of Robotics

The development of tools has to a large extent driven human civilization. Think of how big an impact it made on every day life when early humans mastered the ability to control fire – it gave them a reliable source of warmth in cold climates, enabled them to engage in productive activity at night, and made it possible to cook safer and more nutritious food. Similarly, the development of autonomous systems has stemmed from this primal desire to develop tools that make our lives easier, healthier, and more productive.

Automation and robotic systems have their roots in the development of electro-mechanical mechanisms and sensors in the early 20th century. By the 1950s many devices such as radio controlled ships and airplanes, radars, and centralized factory control rooms were in widespread use. Autonomous technologies were making a dramatic impact on society – consider the following achievements:

- 1957: The Soviet Union launches *Sputnik*, the first artificial Earth satellite
- 1962: General Motors uses the first industrial robots *Unimate* for welding and other tasks
- 1968: Researcher's at Stanford use the A* search algorithm to plan paths for *Shakey the Robot*
- 1975: NASA's *Viking 1* is the first spacecraft to land on Mars and operates for 2,307 days
- 1990: First Robot Olympics held in Scotland with events like: climbing, racing, javelin, etc.
- 1996: IBM's *Deep Blue* defeats world chess champion Garry Kasparov

- 2005: *Stanley* the driverless car wins the (132 mile off-road) DARPA Grand Challenge
- 2009: The *RU27* underwater glider crosses the Atlantic Ocean underwater (221 days at sea)
- 2013: Northrop Grumman *X-47B* wins Collier Trophy for autonomous landing on aircraft carrier
- 2015: The first annual U.S. National Drone Racing Championship held at the California State Fair

This impressive list is growing rapidly and one can only imagine what advances will be made in our lifetimes. Perhaps in 2050 we will watch a team of humanoid robot soccer players win against the most recent World Cup champions. (In fact, this is the stated objective of the RoboCup annual international robotics competition!)

Robot Components

To understand how robots work it is first necessary to understand what a robot is made of:

- *Sensors:* to perceive the world through measurements. For example, sensors can tell the robot if it is night or day, what the temperature of its CPU is, what direction is North, or at what speed the robot is currently moving.
- *Actuators:* to influence the world by creating changes in the environment. For example, an actuator can dispense a chemical into the atmosphere, turn on a motor to spin a wheel, extend an artificial limb, or generate an audio greeting in response to encountering a human.
- *Hardware and Power:* this includes all of the physical parts of the robot's support structure, wiring, power system, and other mechanisms related to actuators and sensors.
- *Software:* to intelligently plan and execute actions that will achieve desired goals.

An analogy of each of these components can be made with human features: The sensors are like your eyes, ears, nose, tongue and skin that can see, hear, smell, taste and touch. The actuators are like your muscles that can push, pull, twist, rotate and move. The hardware is like your skeletal structure, heart, and metabolism that provides energy and physical support to the rest of our body. Last, the software is like your brain that makes plans, responds to your senses, and coordinates your muscles to achieve complex goals.

Software Architecture

The robot's brain, the software, is always running in a cycle that can be abstracted as consisting of three steps: SENSE → PLAN → ACT. This cycle usually runs many times a second and allows the robot to execute fluid motions and react quickly to changes.

The SENSE step involves reading sensor measurements that tell the robot where it is and what is going on in the environment. Raw data is usually not immediately useful to the robot and needs to be interpreted. For example, if the raw data from an autonomous car's camera is a 256 x 256 set of pixels representing an image of an intersection, then a useful interpretation would answer questions like: Is the intersection free of other vehicles? Is the light green? Databases or other knowledge available to the robot may be helpful in interpreting the raw data. Over time the robot might assemble consecutive sensor measurements into something more useful like a map. Collectively, sensing with interpretation is called *perception*.

The PLAN step takes the robot's current perception of the world, compares it with its own desired goals, and determines a plan to execute some actions that will achieve them. The goal may be defined by the user in the form of some *mission commands* or the robot may reason about its situation and determine a goal on its own. For example, if an autonomous car has a low battery level it may reason that its goal should be to get to the nearest recharging station. Further, if the robot perceives a green light and an empty intersection, then it may plan a path consisting of a series of GPS waypoints that will bring it to the goal.

Last, the vehicle needs to carry out the plan and ACT. This step involves determining all the details that will allow the vehicle to follow the plan. Following our example, this would correspond to the precise throttle needed to get the vehicle up to speed, the sequence of steering angles to cross the intersection and navigate to the gas station, and the braking force required to bring the vehicle to a smooth stop.

This idea of a robot software loop is sketched in Fig. 1. We will refer back to this figure throughout the course as we learn about the different aspects of robot programming.

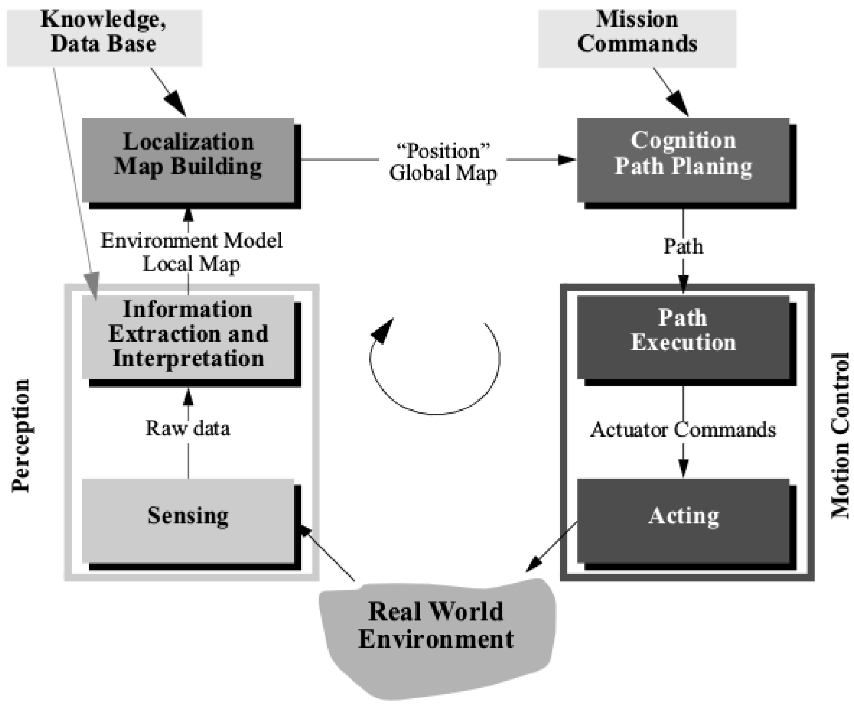

Figure 1: The robot's software loop, Ref. [1].

Human-Robot Interaction

Some robots may be as easy to use as pushing a button, but others may require specialized skill to program and work with humans. Consider the hypothetical example of a farmer needing a specific acre of her field harvested with an autonomous combine. Somehow she needs to be able to convey the coordinates of this patch of land to the robot. One approach may be that the farmer would use a graphical interface on a computer, connected to the combine, to indicate the desired region by clicking on a map. Consider another example, of a humanitarian robot working side-by-side with rescuers after a natural disaster. In this case it seems that conveying commands via a graphical computer interface would be inefficient – instead, the robot should be able to understand the verbal commands of human rescuer. Both of these cases are examples of *human-robot interaction.* The goal of human-robot interaction is to understand the best way in which we can communicate and interact with robots. This also involves developing methods to build trust between humans and their robotic counterpart and making robots' social behavior and appearance acceptable and comfortable to us humans.

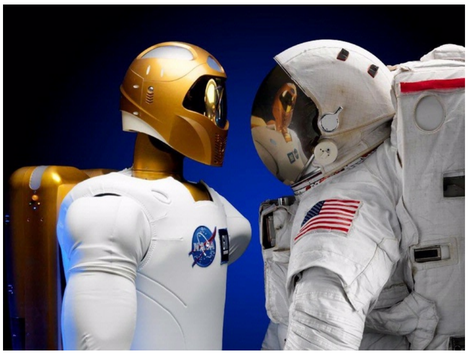

Figure 2: The *Robonaut 2* is being tested aboard the International Space Station and may one day help perform tasks that are too risky for astronauts (Photo: Courtesy of NASA)

Ethical Concerns

As is often the case, the introduction of new technologies can be controversial. The use of robots and autonomous systems is no exception and there are many ethical and legal concerns surrounding their use. Let's consider some of these concerns and their arguments *pro et contra*:

- *Robots can replace human workers.* Critics argue this will lead to mass unemployment and robots will never be as good as humans in performing certain jobs. Supporters argue that robots will be limited to performing dull, dangerous or dirty jobs. Further, they will make our society more prosperous so that that those unemployed will be able to pursue their own interests or engage in a new industry of developing and maintaining robots.

- *Robots can make mistakes.* Critics argue that robots can make mistakes and cause injuries but cannot bear responsibility for their actions in the sense that humans can. (For example, if an autonomous car with human occupants accelerates uncontrollably and crashes, then who is to blame? – the software engineer? the throttle sensor manufacturer? the sleeping driver?) Supporters argue that every technology has its flaws but robots have demonstrated that they make fewer mistakes than humans. (For example, automating all cars could lead to a decrease in pollution, traffic, and a drastic reduction in vehicle collisions.)

- *Robots can cause harm.* Critics argue that using robots in war (e.g., the use of drones in air strikes) makes it easier to cause suffering without appreciating the full consequences as one would in a manned conflict. Supporters argue that the vast majority of robots are used for good. Those that are used in war protect soldiers, provide surveillance which leads to better decision making, or they conduct more precise and accurate strikes and thereby minimize civilian casualties.

- *Robots can take over the world.* Critics argue that as the world becomes more connected with "smart devices", robots become more intelligent and capable, and we increasingly rely on robots for essential services, they will inevitably decide that they will no longer need humans or will otherwise "want" to be in control. Or perhaps an incorrect interpretation of a human command will lead to unintended consequences. For example, consider the following thought experiment:

  > "[Suppose] a well-meaning team of programmers make a big mistake in designing [a robot's] goal system. This could result [...] in a [robot] whose top goal is the manufacturing of paperclips, with the consequence that it starts transforming first all of earth and then increasing portions of space into paperclip manufacturing facilities. [The robot would] resist with all its might any attempt to alter this goal."

- N. Bostrom [Ref. 2]

Supporters argue that we will be capable of including checks and balances to ensure humans can always "pull the plug".

One way to address some of these concerns might be to program a robot with a set of rules or laws that force it to behave a certain way. For example, the science fiction writer Isaac Asimov proposed the following Three Laws of Robotics:

1. A robot may not injure a human being or, through inaction, allow a human being to come to harm.
2. A robot must obey the orders given to it by human beings except where such orders would conflict with the First Law.
3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Laws.

Some years later the Fourth Law was also added:

4. A robot may not harm humanity, or, by inaction, allow humanity to come to harm.

Even these seemingly simple laws could be difficult to implement in practice. The concerns stated above are still open questions that require us to have an informed debate. In addition to using our moral judgment when making arguments concerning robot ethics we should also rely on verifiable facts, statistics, and expert opinions.

Programming Robots

We conclude this lecture by giving a broad overview of the many topics related to programming robots. The following table may help provide some insight regarding how this topic fits into the context of your other coursework.

References

1. R. Siegwart I. R. Nourbakhsh, and D. Scaramuzza, *Introduction to Autonomous Mobile Robots*, Second Edition, MIT Press, 2011.

2. N. Bostrom, "Cognitive, Emotive and Ethical Aspects of Decision Making in Humans and in Artificial Intelligence", Vol. 2, ed. I. Smit et al., *Int. Institute of Advanced Studies in Systems Research and Cybernetics*, 2003, pp. 12-17]

Major Topics	Subtopics	Example Programming Tasks
Artificial Intelligence	• learning • goal reasoning • decision theory • pattern recognition • machine vision • perception • statistics	A high-level software framework that is capable of interpreting sensor data, reasoning about its goals, and making high-level decisions.
Task and Motion Planning	• algorithms • graph search • combinatorial optimization • data structures	An algorithm that uses environmental information to generate motion plans that satisfy predefined tasks and goals.
Control and Estimation	• kinematics, dynamics • differential equations • linear algebra • simulations • feedback control • probability and filtering	A feedback control loop that quickly reacts to sensor measurements and precisely regulates the actuators to ensure the robot is stable and maintaining a desired state (e.g., speed).
Actuator and Sensor Integration	• circuit theory • mechanical design • mechatronics • operating systems • communication • power electronics	A low-level firmware library that controls the most basic robot functionality such as power, data acquistion and actuator positions.

Table 1: Topics related to programming robots