University of Oxford

College of Computer Science

Elzein Salah Elzein

22/12/2021

PHD Report Draft

Written By Elzein Salah Elzein

Robotics and Autonomous Systems

1. Introduction

1 Locomotion

1.1 Introduction

A mobile robot needs locomotion mechanisms that enable it to move unbounded through- out its environment. But there are a large variety of possible ways to move, and so the selec- tion of a robot's approach to locomotion is an important aspect of mobile robot design. In the laboratory, there are research robots that can walk, jump, run, slide, skate, swim, fly, and, of course, roll. Most of these locomotion mechanisms have been inspired by their bio- logical counterparts (see figure 2.1).

There is, however, one exception: the actively powered wheel is a human invention that achieves extremely high efficiency on flat ground. This mechanism is not completely for- eign to biological systems. Our bipedal walking system can be approximated by a rolling polygon, with sides equal in length *d* to the span of the step (figure 2.2). As the step size decreases, the polygon approaches a circle or wheel. But nature did not develop a fully rotating, actively powered joint, which is the technology necessary for wheeled locomo- tion.

Biological systems succeed in moving through a wide variety of harsh environments. Therefore it can be desirable to copy their selection of locomotion mechanisms. However, replicating nature in this regard is extremely difficult for several reasons. To begin with,

mechanical complexity is easily achieved in biological systems through structural replica- tion. Cell division, in combination with specialization, can readily produce a millipede with several hundred legs and several tens of thousands of individually sensed cilia. In man- made structures, each part must be fabricated individually, and so no such economies of scale exist. Additionally, the cell is a microscopic building block that enables extreme min- iaturization. With very small size and weight, insects achieve a level of robustness that we have not been able to match with human fabrication techniques. Finally, the biological energy storage system and the muscular and hydraulic activation systems used by large ani- mals and insects achieve torque, response time, and conversion efficiencies that far exceed similarly scaled man-made systems.

Type of motion	Resistance to motion	Basic kinematics of motion
Flow in a Channel 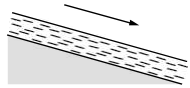	Hydrodynamic forces	Eddies 
Crawl 	Friction forces	Longitudinal vibration 
Sliding 	Friction forces	Transverse vibration 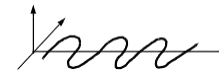
Running 	Loss of kinetic energy	Oscillatory movement of a multi-link pendulum 
Jumping 	Loss of kinetic energy	Oscillatory movement of a multi-link pendulum 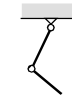
Walking 	Gravitational forces	Rolling of a polygon (see figure 2.2) 

Figure 2.1

Locomotion mechanisms used in biological systems.

Owing to these limitations, mobile robots generally locomote either using wheeled mechanisms, a well-known human technology for vehicles, or using a small number of articulated legs, the simplest of the biological approaches to locomotion (see figure 2.2).

In general, legged locomotion requires higher degrees of freedom and therefore greater mechanical complexity than wheeled locomotion. Wheels, in addition to being simple, are extremely well

suited to flat ground. As figure 2.3 depicts, on flat surfaces wheeled loco- motion is one to two orders of magnitude more efficient than legged locomotion. The rail- way is ideally engineered for wheeled locomotion because rolling friction is minimized on a hard and flat steel surface. But as the surface becomes soft, wheeled locomotion accumulates inefficiencies due to rolling friction whereas legged locomotion suffers much less because it consists only of point contacts with the ground. This is demonstrated in figure
2.3 by the dramatic loss of efficiency in the case of a tire on soft ground.

2. Literature Review

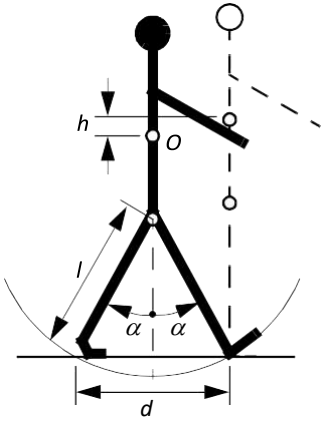

Figure 2.2

A biped walking system can be approximated by a rolling polygon, with sides equal in length *d* to the span of the step. As the step size decreases, the polygon approaches a circle or wheel with the radius *l*.

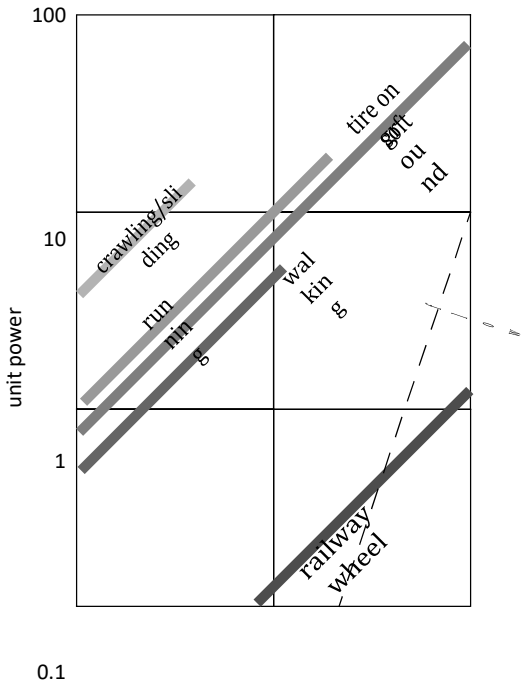

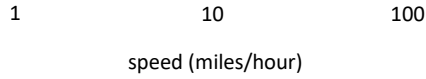

Figure 2.3

Specific power versus attainable speed of various locomotion mechanisms [33].

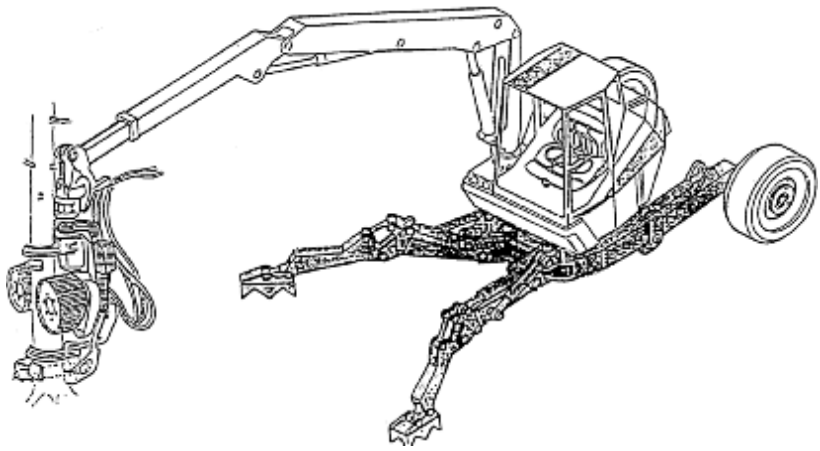

Figure 2.4

RoboTrac, a hybrid wheel-leg vehicle for rough terrain [130].

In effect, the efficiency of wheeled locomotion depends greatly on environmental qual- ities, particularly the flatness and hardness of the ground, while the efficiency of legged locomotion depends on the leg mass and body mass, both of which the robot must support at various points in a legged gait.

It is understandable therefore that nature favors legged locomotion, since locomotion systems in nature must operate on rough and unstructured terrain. For example, in the case of insects in a forest the vertical variation in ground height is often an order of magnitude greater than the total height of the insect. By the same token, the human environment fre- quently consists of engineered, smooth surfaces, both indoors and outdoors. Therefore, it is also understandable that virtually all industrial applications of mobile robotics utilize some form of wheeled locomotion. Recently, for more natural outdoor environments, there has been some progress toward hybrid and legged industrial robots such as the forestry robot shown in figure 2.4.

In the section 2.1.1, we present general considerations that concern

all forms of mobile robot locomotion. Following this, in sections 2.2 and 2.3, we present overviews of legged locomotion and wheeled locomotion techniques for mobile robots.

1.1.1 Key issues for locomotion

Locomotion is the complement of manipulation. In manipulation, the robot arm is fixed but moves objects in the workspace by imparting force to them. In locomotion, the environ- ment is fixed and the robot moves by imparting force to the environment. In both cases, the scientific basis is the study of actuators that generate interaction forces, and mechanisms

that implement desired kinematic and dynamic properties. Locomotion and manipulation thus share the same core issues of stability, contact characteristics, and environmental type:

- stability
  - number and geometry of contact points
  - center of gravity
  - static/dynamic stability
  - inclination of terrain
- characteristics of contact
  - contact point/path size and shape
  - angle of contact
  - friction
- type of environment
  - structure
  - medium, (e.g. water, air, soft or hard ground)

A theoretical analysis of locomotion begins with mechanics and physics. From this start- ing point, we can formally define and analyze all manner of mobile robot locomotion sys- tems. However, this book focuses on the mobile robot *navigation* problem, particularly stressing perception, localization, and cognition. Thus we will not delve deeply into the physical basis of locomotion. Nevertheless, the two remaining sections in this chapter present overviews of issues in legged locomotion [33] and wheeled locomotion. Then, chapter 3 presents a more detailed analysis of the kinematics and control of wheeled mobile robots.

1.2 Legged Mobile Robots

Legged locomotion is characterized by a series of point contacts between the robot and the ground. The key advantages include adaptability and maneuverability in rough terrain. Because only a set of point contacts is required, the quality of the ground between those points does not matter so long as the robot can maintain adequate ground clearance. In addi- tion, a walking robot is capable of crossing a hole or chasm so long as its reach exceeds the width of the hole. A final

advantage of legged locomotion is the potential to manipulate objects in the environment with great skill. An excellent insect example, the dung beetle, is capable of rolling a ball while locomoting by way of its dexterous front legs.

The main disadvantages of legged locomotion include power and mechanical complex- ity. The leg, which may include several degrees of freedom, must be capable of sustaining part of the robot's total weight, and in many robots must be capable of lifting and lowering the robot. Additionally, high maneuverability will only be achieved if the legs have a suf- ficient number of degrees of freedom to impart forces in a number of different directions.

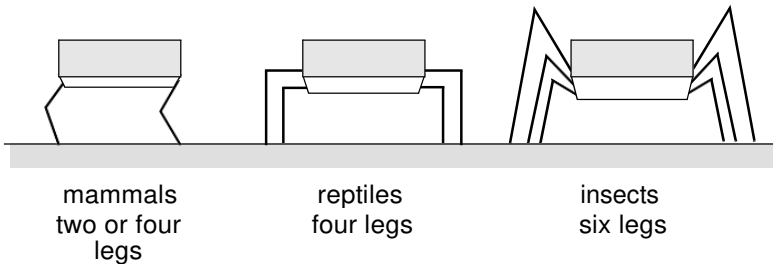

Figure 2.5

Arrangement of the legs of various animals.

1.2.1 Leg configurations and stability

Because legged robots are biologically inspired, it is instructive to examine biologically successful legged systems. A number of different leg configurations have been successful in a variety of organisms (figure 2.5). Large animals, such as mammals and reptiles, have four legs, whereas insects have six or more legs. In some mammals, the ability to walk on only two legs has been perfected. Especially in the case of humans, balance has progressed
to the point that we can even jump with one leg[1]. This exceptional maneuverability comes
at a price: much more complex active control to maintain balance.

In contrast, a creature with three legs can exhibit a static, stable pose provided that it can ensure that its center of gravity is within the tripod of ground contact. Static stability, dem- onstrated by a three-legged stool, means that balance is maintained with no need for motion. A small deviation from stability (e.g., gently pushing the stool) is passively cor- rected toward the stable pose when the upsetting force stops.

But a robot must be able to lift its legs in order to walk. In order to achieve static walk- ing, a robot must have at least six legs. In such a configuration, it is possible to design a gait in which a statically stable tripod of legs is in contact with the ground at all times (figure 2.8).

Insects and spiders are immediately able to walk when born. For them, the problem of balance during walking is relatively simple. Mammals, with four legs, cannot achieve static walking, but are able to stand easily on four legs. Fauns, for example, spend several minutes attempting to stand before they are able to do so, then spend several more minutes learning to walk without falling. Humans, with two legs, cannot even stand in one place with static stability. Infants require months to stand and walk, and even longer to learn to jump, run, and stand on one leg.

---

1. In child development, one of the tests used to determine if the child is acquiring advanced loco- motion skills is the ability to jump on one leg.

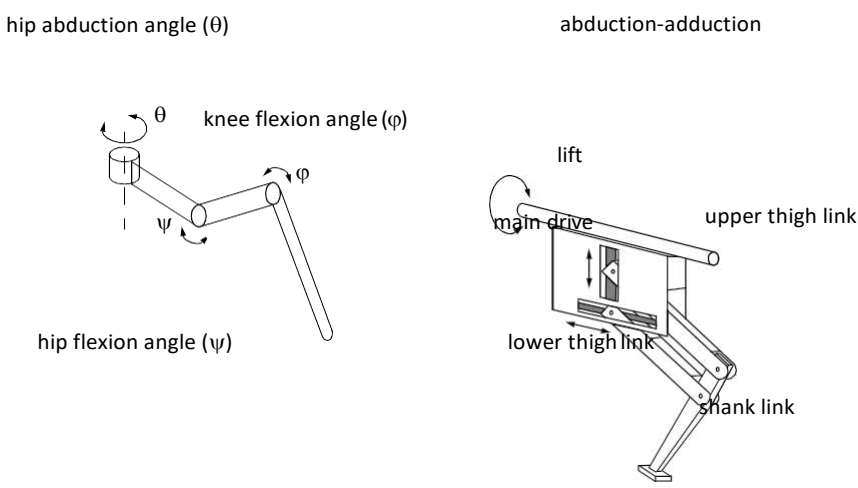

Figure 2.6

Two examples of legs with three degrees of freedom.

There is also the potential for great variety in the complexity of each individual leg. Once again, the biological world provides ample examples at both extremes. For instance, in the case of the caterpillar, each leg is extended using hydraulic pressure by constricting the body cavity and forcing an increase in pressure, and each leg is retracted longitudinally by relaxing the hydraulic pressure, then activating a single tensile muscle that pulls the leg in toward the body. Each leg has only a single degree of freedom, which is oriented longi- tudinally along the leg. Forward locomotion depends on the hydraulic pressure in the body, which extends the distance between pairs of legs. The caterpillar leg is therefore mechani- cally very simple, using a minimal number of extrinsic muscles to achieve complex overall locomotion.

At the other extreme, the human leg has more than seven major degrees of freedom, combined with further actuation at the toes. More than fifteen muscle groups actuate eight complex joints.

In the case of legged mobile robots, a minimum of two degrees of freedom is generally required to move a leg forward by lifting the leg and swinging it forward. More common is the addition of a third degree

of freedom for more complex maneuvers, resulting in legs such as those shown in figure 2.6. Recent successes in the creation of bipedal walking robots have added a fourth degree of freedom at the ankle joint. The ankle enables more consistent ground contact by actuating the pose of the sole of the foot.

In general, adding degrees of freedom to a robot leg increases the maneuverability of the robot, both augmenting the range of terrains on which it can travel and the ability of the robot to travel with a variety of gaits. The primary disadvantages of additional joints and actuators are, of course, energy, control, and mass. Additional actuators require energy and control, and they also add to leg mass, further increasing power and load requirements on existing actuators.

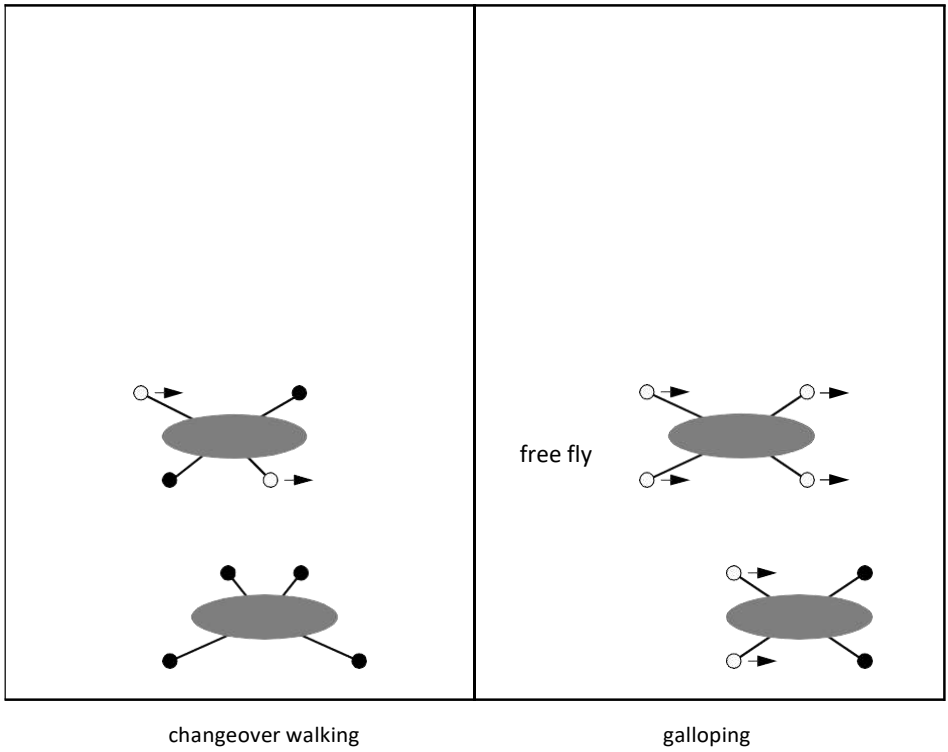

Figure 2.7

Two gaits with four legs. Because this robot has fewer than six legs, static walking is not generally possible.

In the case of a multilegged mobile robot, there is the issue of leg coordination for loco- motion, or gait control. The number of possible gaits depends on the number of legs [33]. The gait is a sequence of lift and release events for the individual legs. For a mobile robot with k legs, the total number of possible events N for a walking machine is

$$N = (2k - 1)! \tag{2.1}$$

For a biped walker $k = 2$ legs, the number of possible events N is

$$N = (2k - 1)! = 3! = 3 \cdot 2 \cdot 1 = 6 \tag{2.2}$$

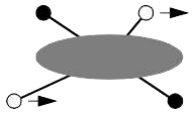

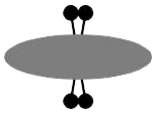

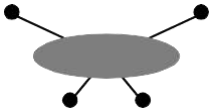

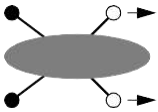

The six different events are

1. lift right leg;
2. lift left leg;
3. release right leg;
4. release left leg;
5. lift both legs together;
6. release both legs together.

Of course, this quickly grows quite large. For example, a robot with six legs has far more gaits theoretically:

$$N = 11! = 39916800 \quad (2.3)$$

Figures 2.7 and 2.8 depict several four-legged gaits and the static six-legged tripod gait.

1.2.2 Examples of legged robot locomotion

Although there are no high-volume industrial applications to date, legged locomotion is an important area of long-term research. Several interesting designs are presented below, beginning with the one-legged robot and finishing with six-legged robots. For a very good overview of climbing and walking robots, see *http://www.uwe.ac.uk/clawar/*.

1.2.2.1 One leg

The minimum number of legs a legged robot can have is, of course, one. Minimizing the number of legs is beneficial for several reasons. Body mass is particularly important to walking machines, and the single leg minimizes cumulative leg mass. Leg coordination is required when a robot has several legs, but with one leg no such coordination is needed. Perhaps most importantly, the one-legged robot maximizes the basic advantage of legged locomotion: legs have single points of contact with the ground in lieu of an entire track, as with wheels. A

single-legged robot requires only a sequence of single contacts, making it amenable to the roughest terrain. Furthermore, a hopping robot can dynamically cross a gap that is larger than its stride by taking a running start, whereas a multilegged walking robot that cannot run is limited to crossing gaps that are as large as its reach.

The major challenge in creating a single-legged robot is balance. For a robot with one leg, static walking is not only impossible but static stability when stationary is also impos- sible. The robot must actively balance itself by either changing its center of gravity or by imparting corrective forces. Thus, the successful single-legged robot must be dynamically stable.

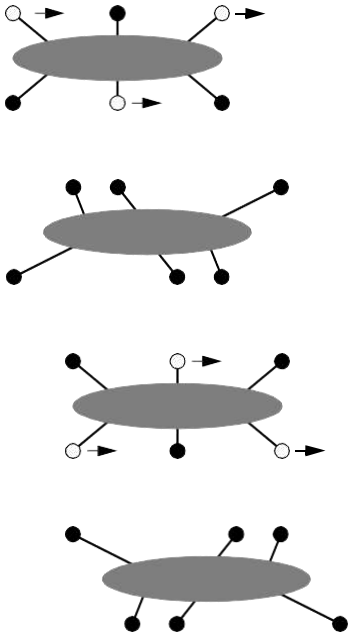

Figure 2.8

Static walking with six legs. A tripod formed by three legs always exists.

Figure 2.9 shows the Raibert hopper [28, 124], one of the most well-known single- legged hopping robots created. This robot makes continuous corrections to body attitude and to robot velocity by adjusting the leg angle with respect to the body. The actuation is hydraulic, including high-power longitudinal extension of the leg during stance to hop back into the air. Although powerful, these actuators require a large, off-board hydraulic pump to be connected to the robot at all times.

Figure 2.10 shows a more energy-efficient design developed more recently [46]. Instead of supplying power by means of an off-board hydraulic pump, the bow leg hopper is designed to capture the kinetic energy of the robot as it lands, using an efficient bow spring leg. This spring returns approximately 85% of the energy, meaning that stable

hopping requires only the addition of 15% of the required energy on each hop. This robot, which is constrained along one axis by a boom, has demonstrated continuous hopping for 20 minutes using a single set of batteries carried on board the robot. As with the Raibert hopper, the bow leg hopper controls velocity by changing the angle of the leg to the body at the hip joint.

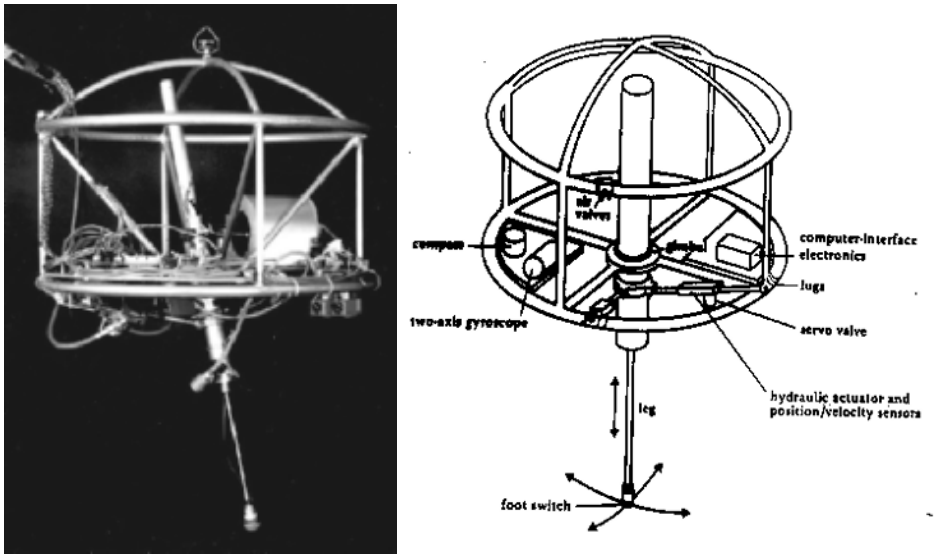

Figure 2.9

The Raibert hopper [28, 124]. Image courtesy of the LegLab and Marc Raibert. © 1983.

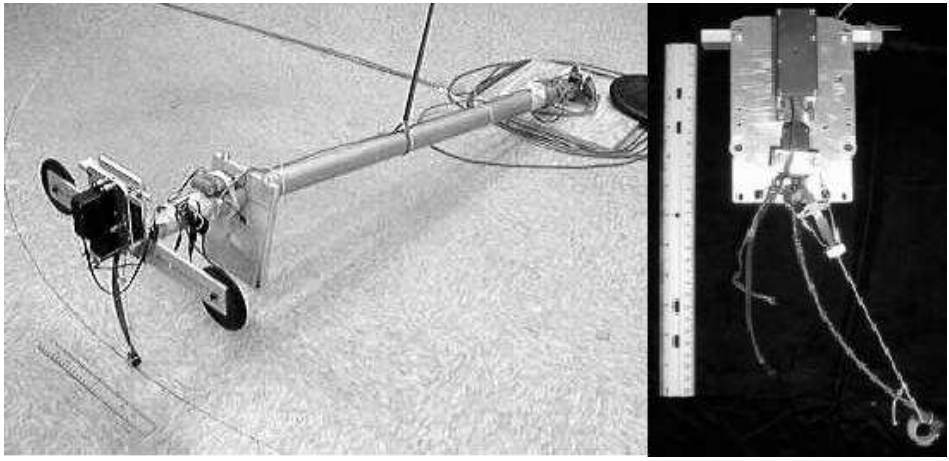

Figure 2.10

The 2D single bow leg hopper [46]. Image courtesy of H. Benjamin Brown and Garth Zeglin, CMU.

Specifications:

Weight:	7 kg
Height:	58 cm
Neck DOF:	4
Body DOF:	2
Arm DOF:	2 x 5
Legs DOF:	2 x 6 Five-finger Hands

Figure 2.11

The Sony SDR-4X II, © 2003 Sony Corporation.

The paper of Ringrose [125] demonstrates the very important duality of mechanics and controls as applied to a single-legged hopping machine. Often clever mechanical design can perform the same operations as complex active control circuitry. In this robot, the physical shape of the foot is exactly the right curve so that when the robot lands without being perfectly vertical, the proper corrective force is provided from the impact, making the robot vertical by the next landing. This robot is dynamically stable, and is furthermore passive. The correction is provided by physical interactions between the robot and its environment, with no computer or any active control in the loop.

1.2.2.2 Two legs (biped)

A variety of successful bipedal robots have been demonstrated over the past ten years. Two legged robots have been shown to run, jump, travel up and down stairways, and even do aerial tricks such as

somersaults. In the commercial sector, both Honda and Sony have made significant advances over the past decade that have enabled highly capable bipedal robots. Both companies designed small, powered joints that achieve power-to-weight per- formance unheard of in commercially available servomotors. These new “intelligent” servos provide not only strong actuation but also compliant actuation by means of torque sensing and closed-loop control.

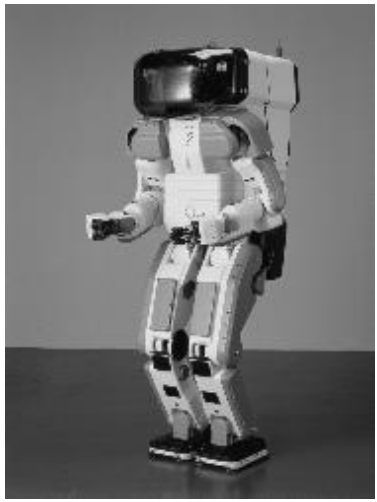

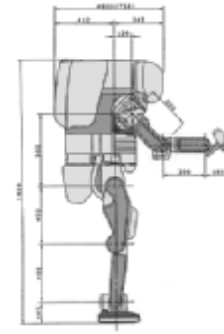

Specifications:

Maximum speed: 2 km/h Autonomy: 15 min
Weight: 210 kg
Height: 1.82 m
Leg DOF: 2 x 6
Arm DOF: 2 x 7

Figure 2.12

The humanoid robot P2 from Honda, Japan. © Honda Motor Corporation.

The Sony Dream Robot, model SDR-4X II, is shown in figure 2.11. This current model is the result of research begun in 1997 with the basic objective of motion entertainment and communication entertainment (i.e., dancing and singing). This robot with thirty-eight degrees of freedom has seven microphones for fine localization of sound, image-based person recognition, on-board miniature stereo depth-map reconstruction, and limited speech recognition. Given the goal of fluid and entertaining motion, Sony spent consider- able effort designing a motion prototyping application system to enable their engineers to script dances in a straightforward manner. Note that the SDR-4X II is relatively small, standing at 58 cm and weighing only 6.5 kg.

The Honda humanoid project has a significant history but, again, has tackled the very important engineering challenge of actuation. Figure 2.12 shows model *P2*, which is an immediate predecessor to the most recent Asimo model (advanced step in innovative mobility).

Note from this picture that the Honda humanoid is much larger than the SDR- 4X at 120 cm tall and 52 kg. This enables practical mobility in the human world of stairs and ledges while maintaining a nonthreatening size and posture. Perhaps the first robot to famously demonstrate biomimetic bipedal stair climbing and descending, these Honda humanoid series robots are being designed not for entertainment purposes but as human aids throughout society. Honda refers, for instance, to the height of Asimo as the minimum height which enables it to nonetheless manage operation of the human world, for instance, control of light switches.

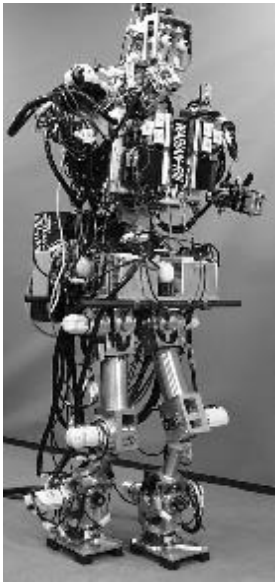

Specifications:

Weight: 131 [kg]
Height: 1.88 [m]

DOF in total: 43
Lower Limbs: 2 x 6
Trunk: 3
Arms: 2 x 10
Neck: 4
Eyes: 2 x 2

Figure 2.13

The humanoid robot WABIAN-RIII at Waseda University in Japan [75]. Image courtesy of Atsuo Takanishi, Waseda University.

An important feature of bipedal robots is their anthropomorphic shape. They can be built to have the same approximate dimensions as humans, and this makes them excellent vehi- cles for research in human-robot interaction. WABIAN is a robot built at Waseda Universities Japan (figure 2.13) for just such research [75]. WABIAN is designed to emulate human motion, and is even designed to dance like a human.

Bipedal robots can only be statically stable within some limits, and so robots such as P2 and WABIAN generally must perform continuous balance-correcting servoing even when standing still. Furthermore, each leg must have sufficient capacity to support the full weight of the robot. In the case of four-legged robots, the balance problem is facilitated along with the load requirements of each leg. An elegant design of a biped robot is the Spring Fla- mingo of MIT (figure 2.14).

This robot inserts springs in series with the leg actuators to achieve a more elastic gait. Combined with “kneecaps” that limit knee joint angles, the Fla- mingo achieves surprisingly biomimetic motion.

1.2.2.3 Four legs (quadruped)

Although standing still on four legs is passively stable, walking remains challenging because to remain stable the robot’s center of gravity must be actively shifted during the

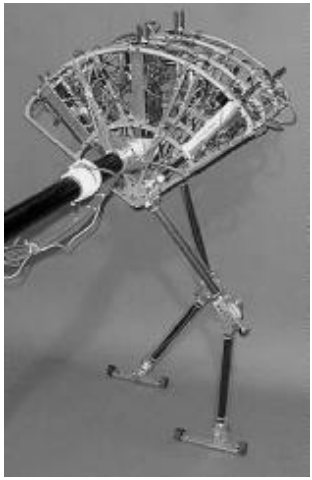

Figure 2.14

The Spring Flamingo developed at MIT [123]. Image courtesy of Jerry Pratt, MIT Leg Laboratory.

gait. Sony recently invested several million dollars to develop a four-legged robot called AIBO (figure 2.15). To create this robot, Sony produced both a new robot operating system that is near real-time and new geared servomotors that are of sufficiently high torque to sup- port the robot, yet back drivable for safety. In addition to developing custom motors and software, Sony incorporated a color vision system that enables AIBO to chase a brightly colored ball. The robot is able to function for at most one hour before requiring recharging. Early sales of the robot have been very strong, with more than 60,000 units sold in the first year. Nevertheless, the number of motors and the technology investment behind this robot dog resulted in a very high price of approximately $1500.

Four-legged robots have the potential to serve as effective artifacts for research in human-robot interaction (figure 2.16). Humans can treat the Sony robot, for example, as a pet and might develop an emotional relationship similar to that between man and dog. Fur- thermore, Sony has designed AIBO's walking style and general behavior to emulate learn- ing and maturation, resulting in dynamic behavior over time that

is more interesting for the owner who can track the changing behavior. As the challenges of high energy storage and motor technology are solved, it is likely that quadruped robots much more capable than AIBO will become common throughout the human environment.

1.2.2.4 Six legs (hexapod)

Six-legged configurations have been extremely popular in mobile robotics because of their static stability during walking, thus reducing the control complexity (figures 2.17 and 1.3).

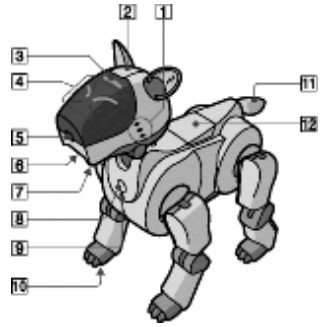

ERS-110 © 1999 Sony Corporation

ERS-210 © 2000 Sony Corporation

1 Stereo microphone: Allows AIBO to pick up surrounding sounds.
2 Head sensor: Senses when a person taps or pets AIBO on the head.
3 Mode indicator: Shows AIBO's operation mode.
4 Eye lights: These light up in blue-green or red to indicate AIBO's emotional state.
5 Color camera: Allows AIBO to search for objects and recognize them by color and movement.
6 Speaker: Emits various musical tones and sound effects.
7 Chin sensor: Senses when a person touches AIBO on the chin.
8 Pause button: Press to activate AIBO or to pause AIBO.
9 Chest light: Gives information about the status of the robot.
10 Paw sensors: Located on the bottom of each paw.
11 Tail light: Lights up blue or orange to show AIBO's emotional state.
12 Back sensor: Senses when a person touches AIBO on the back.

Figure 2.15

AIBO, the artificial dog from Sony, Japan.

In most cases, each leg has three degrees of freedom, including hip flexion, knee flexion, and hip abduction (see figure 2.6). Genghis is a commercially available hobby robot that has six legs, each of which has two degrees of freedom provided by hobby servos (figure 2.18). Such a robot, which consists only of hip flexion and hip abduction, has less maneu- verability in rough terrain but performs quite well on flat ground. Because it consists of a straightforward arrangement of servomotors and straight legs, such robots can be readily built by a robot hobbyist.

Insects, which are arguably the most successful locomoting

creatures on earth, excel at traversing all forms of terrain with six legs, even upside down. Currently, the gap between the capabilities of six-legged insects and artificial six-legged robots is still quite large. Interestingly, this is not due to a lack of sufficient numbers of degrees of freedom on the robots. Rather, insects combine a small number of active degrees of freedom with passive

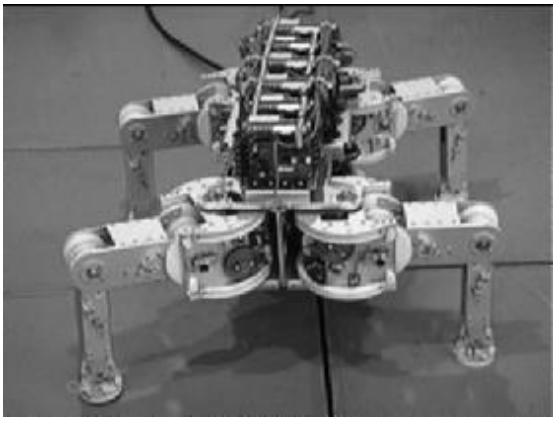

Specifications:

Weight:1 9 kg
Height: 0.25 m
DOF: 4 x 3

Figure 2.16

Titan VIII, a quadruped robot developed at Tokyo Institute of Technology. (http://mozu.mes.titech.ac.jp/research/walk/). © Tokyo Institute of Technology.

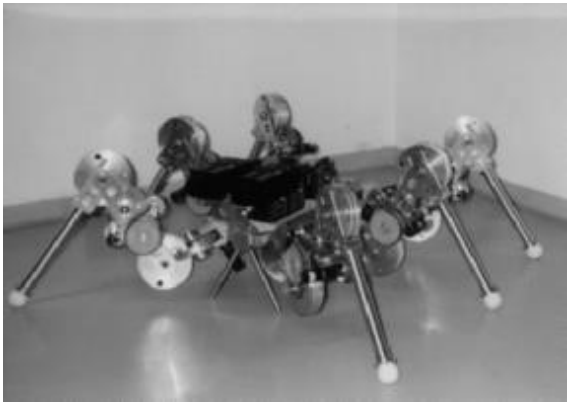

Specifications:

Maximum speed: 0.5 m/s Weight:1 6 kg
Height: 0.3 m
Length: 0.7 m
No. of legs: 6
DOF in total: 6 x 3 Power consumption:10 W

Figure 2.17

Lauron II, a hexapod platform developed at the University of Karlsruhe, Germany.

© University of Karlsruhe.

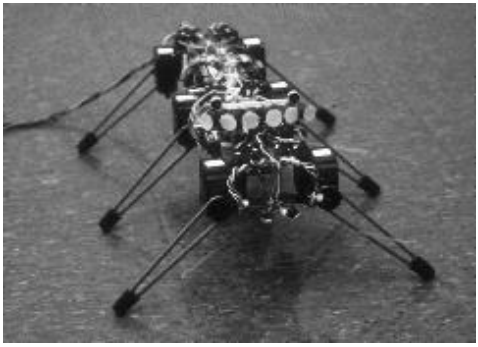

Figure 2.18

Genghis, one of the most famous walking robots from MIT, uses hobby servomotors as its actuators (http://www.ai.mit.edu/projects/genghis). © MIT AI Lab.

structures, such as microscopic barbs and textured pads, that increase the gripping strength of each leg significantly. Robotic research into such passive tip structures has only recently begun. For example, a research group is attempting to re-create the complete mechanical function of the cockroach leg [65].

It is clear from the above examples that legged robots have much progress to make before they are competitive with their biological equivalents. Nevertheless, significant gains have been realized recently, primarily due to advances in motor design. Creating actuation systems that approach the efficiency of animal muscles remains far from the reach of robotics, as does energy storage with the energy densities found in organic life forms.

2.3 Wheeled Mobile Robots

The wheel has been by far the most popular locomotion mechanism in mobile robotics and in man-made vehicles in general. It can achieve very good efficiencies, as demonstrated in figure 2.3, and does so with a relatively simple mechanical implementation.

In addition, balance is not usually a research problem in wheeled robot designs, because wheeled robots are almost always designed

so that all wheels are in ground contact at all times. Thus, three wheels are sufficient to guarantee stable balance, although, as we shall see below, two-wheeled robots can also be stable. When more than three wheels are used, a suspension system is required to allow all wheels to maintain ground contact when the robot encounters uneven terrain.

Instead of worrying about balance, wheeled robot research tends to focus on the prob- lems of traction and stability, maneuverability, and control: can the robot wheels provide

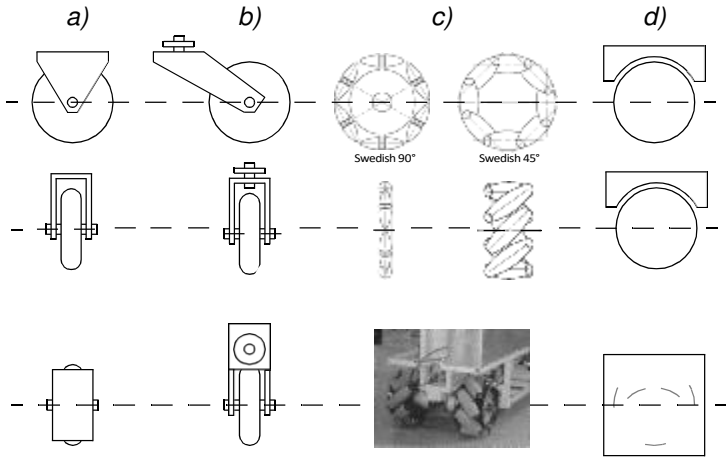

Figure 2.19

The four basic wheel types. (a) Standard wheel: two degrees of freedom; rotation around the (motorized) wheel axle and the contact point.(b) castor wheel: two degrees of freedom; rotation around an offset steering joint. (c) Swedish wheel: three degrees of freedom; rotation around the (motorized) wheel axle, around the rollers, and around the contact point. (d) Ball or spherical wheel: realization technically difficult.

sufficient traction and stability for the robot to cover all of the desired terrain, and does the robot's wheeled configuration enable sufficient control over the velocity of the robot?

2.3.1 Wheeled locomotion: the design space

As we shall see, there is a very large space of possible wheel configurations when one con- siders possible techniques for mobile robot locomotion. We begin by discussing the wheel in detail, as there are a number of different wheel types with specific strengths and weaknesses. Then, we examine complete wheel configurations that deliver particular forms of locomotion for a mobile robot.

2.3.1.1 Wheel design

There are four major wheel classes, as shown in figure 2.19. They differ widely in their kinematics, and therefore the choice of wheel type has a large effect on the overall kinemat- ics of the mobile robot. The

standard wheel and the castor wheel have a primary axis of rotation and are thus highly directional. To move in a different direction, the wheel must be steered first along a vertical axis. The key difference between these two wheels is that the standard wheel can accomplish this steering motion with no side effects, as the center of rotation passes through the contact patch with the ground, whereas the castor wheel rotates around an offset axis, causing a force to be imparted to the robot chassis during steering.

Figure 2.20

Navlab I, the first autonomous highway vehicle that steers and controls the throttle using vision and radar sensors [61]. Developed at CMU.

The Swedish wheel and the spherical wheel are both designs that are less constrained by directionality than the conventional standard wheel. The Swedish wheel functions as a normal wheel, but provides low resistance in another direction as well, sometimes perpen- dicular to the conventional direction, as in the Swedish 90, and sometimes at an intermedi- ate angle, as in the Swedish 45. The small rollers attached around the circumference of the wheel are passive and the wheel's primary axis serves as the only actively powered joint. The key advantage of this design is that, although the wheel rotation is powered only along the one principal axis (through the axle), the wheel can kinematically move with very little friction along many possible trajectories, not just forward and backward.

The spherical wheel is a truly omnidirectional wheel, often designed so that it may be actively powered to spin along any direction. One mechanism for implementing this spher- ical design imitates the computer mouse, providing actively powered rollers that rest against the top surface of the sphere and impart rotational force.

Regardless of what wheel is used, in robots designed for all-terrain

environments and in robots with more than three wheels, a suspension system is normally required to maintain wheel contact with the ground. One of the simplest approaches to suspension is to design flexibility into the wheel itself. For instance, in the case of some four-wheeled indoor robots that use castor wheels, manufacturers have applied a deformable tire of soft rubber to the wheel to create a primitive suspension. Of course, this limited solution cannot compete with a sophisticated suspension system in applications where the robot needs a more dynamic suspension for significantly non flat terrain.

2.3.1.2 Wheel geometry

The choice of wheel types for a mobile robot is strongly linked to the choice of wheel arrangement, or wheel geometry. The mobile robot designer must consider these two issues simultaneously when designing the locomoting mechanism of a wheeled robot. Why do wheel type and wheel geometry matter? Three fundamental characteristics of a robot are governed by these choices: maneuverability, controllability, and stability.

Unlike automobiles, which are largely designed for a highly standardized environment (the road network), mobile robots are designed for applications in a wide variety of situa- tions. Automobiles all share similar wheel configurations because there is one region in the design space that maximizes maneuverability, controllability, and stability for their stan- dard environment: the paved roadway. However, there is no single wheel configuration that maximizes these qualities for the variety of environments faced by different mobile robots. So you will see great variety in the wheel configurations of mobile robots. In fact, few robots use the Ackerman wheel configuration of the automobile because of its poor maneu- verability, with the exception of mobile robots designed for the road system (figure 2.20). Table 2.1 gives an overview of wheel configurations ordered by the number of wheels.

This table shows both the selection of particular wheel types and their geometric configu- ration on the robot chassis. Note that some of the configurations shown are of little use in mobile robot applications. For instance, the two-wheeled bicycle arrangement has moder- ate maneuverability and poor controllability. Like a single-legged hopping machine, it can never stand still. Nevertheless, this table provides an indication of the large variety of wheel configurations that are possible in mobile robot design.

The number of variations in table 2.1 is quite large. However, there are important trends and groupings that can aid in comprehending the advantages and disadvantages of each configuration. Below, we identify some of the key trade-offs in terms of the three issues we identified earlier: stability, maneuverability, and controllability.

2.3.1.3 Stability

Surprisingly, the minimum number of wheels required for static stability is two. As shown above, a two-wheel differential-drive robot can achieve static stability if the center of mass is below the wheel axle. Cye is a commercial mobile robot that uses this wheel configura- tion (figure 2.21).

However, under ordinary circumstances such a solution requires wheel diameters that are impractically large. Dynamics can also cause a two-wheeled robot to strike the floor with a third point of contact, for instance, with sufficiently high motor torques from stand- still. Conventionally, static stability requires a minimum of three wheels, with the addi- tional caveat that the center of gravity must be contained within the triangle formed by the ground contact points of the wheels. Stability can be further improved by adding more wheels, although once the number of contact points exceeds three, the hyperstatic nature of the geometry will require some form of flexible suspension on uneven terrain.

Table 2.1

Wheel configurations for rolling vehicles

# of wheels	Arrangement	Description	Typical examples
2		One steering wheel in the front, one traction wheel in the rear	Bicycle, motorcycle
	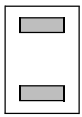	Two-wheel differential drive with the center of mass (COM) below the axle	Cye personal robot
3	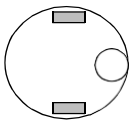	Two-wheel centered differen- tial drive with a third point of contact	Nomad Scout, smartRob EPFL
	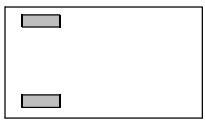	Two independently driven wheels in the rear/front, 1 unpowered omnidirectional wheel in the front/rear	Many indoor robots, including the EPFL robots Pygmalion and Alice
	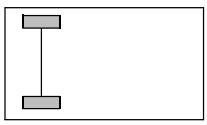	Two connected traction wheels (differential) in rear, 1 steered free wheel in front	Piaggio minitrucks
	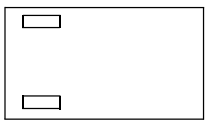	Two free wheels in rear, 1 steered traction wheel in front	Neptune (Carnegie Mellon University), Hero-1
	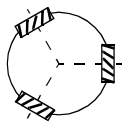	Three motorized Swedish or spherical wheels arranged in a triangle; omnidirectional move- ment is possible	Stanford wheel Tribolo EPFL, Palm Pilot Robot Kit (CMU)
	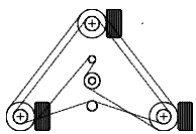	Three synchronously motorized and steered wheels; the orienta- tion is not controllable	“Synchro drive” Denning MRV-2, Geor- gia Institute of Technol- ogy, I-Robot B24, Nomad 200

Table 2.1

Wheel configurations for rolling vehicles

# of wheels	Arrangement	Description	Typical examples
4	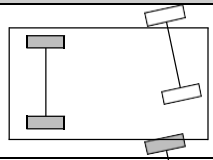	Two motorized wheels in the rear, 2 steered wheels in the front; steering has to be differ- ent for the 2 wheels to avoid slipping/skidding.	Car with rear-wheel drive
	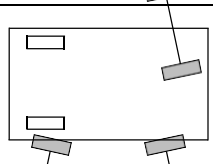	Two motorized and steered wheels in the front, 2 free wheels in the rear; steering has to be different for the 2 wheels to avoid slipping/skidding.	Car with front-wheel drive
	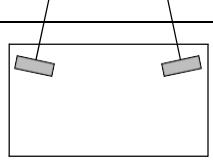	Four steered and motorized wheels	Four-wheel drive, four-wheel steering Hyperion (CMU)
	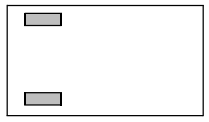	Two traction wheels (differen- tial) in rear/front, 2 omnidirec- tional wheels in the front/rear	Charlie (DMT-EPFL)
		Four omnidirectional wheels	Carnegie Mellon Uranus
	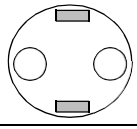	Two-wheel differential drive with 2 additional points of con- tact	EPFL Khepera, Hyperbot Chip
	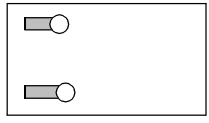	Four motorized and steered castor wheels	Nomad XR4000

Table 2.1

Wheel configurations for rolling vehicles

# of wheels	Arrangement	Description	Typical examples
6	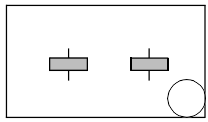	Two motorized and steered wheels aligned in center, 1 omnidirectional wheel at each corner	First
	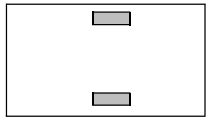	Two traction wheels (differen- tial) in center, 1 omnidirec- tional wheel at each corner	Terregator (Carnegie Mel- lon University)
Icons for the each wheel type are as follows:			
	unpowered omnidirectional wheel (spherical, castor, Swedish);		
	motorized Swedish wheel (Stanford wheel);		
	unpowered standard wheel;		
	motorized standard wheel;		
	motorized and steered castor wheel;		
	steered standard wheel;		
	connected wheels.		

2.3.1.4 Maneuverability

Some robots are omnidirectional, meaning that they can move at any time in any direction along the ground plane (x, y) regardless of the orientation of the robot around its vertical axis. This level of maneuverability requires wheels that can move in more than just one direction, and so omnidirectional robots usually employ Swedish or

spherical wheels that are powered. A good example is Uranus, shown in figure 2.24. This robot uses four Swed- ish wheels to rotate and translate independently and without constraints.

Figure 2.21

Cye, a commercially available domestic robot that can vacuum and make deliveries in the home, is built by Aethon Inc. (http://www.aethon.com). © Aethon Inc.

In general, the ground clearance of robots with Swedish and spherical wheels is some- what limited due to the mechanical constraints of constructing omnidirectional wheels. An interesting recent solution to the problem of omnidirectional navigation while solving this ground-clearance problem is the four-castor wheel configuration in which each castor wheel is actively steered and actively translated. In this configuration, the robot is truly omnidirectional because, even if the castor wheels are facing a direction perpendicular to the desired direction of travel, the robot can still move in the desired direction by steering these wheels. Because the vertical axis is offset from the ground-contact path, the result of this steering motion is robot motion.

In the research community, other classes of mobile robots are popular which achieve high maneuverability, only slightly inferior to that of the omnidirectional configurations. In such robots, motion in a particular direction may initially require a rotational motion. With a circular chassis and an axis of rotation at the center of the robot, such a robot can spin without changing its ground footprint. The most popular such robot is the two-wheel differential-drive robot where the two wheels rotate around the center point of the robot. One or two additional ground contact points may be used for stability, based on the

appli- cation specifics.

In contrast to the above configurations, consider the Ackerman steering configuration common in automobiles. Such a vehicle typically has a turning diameter that is larger than the car. Furthermore, for such a vehicle to move sideways requires a parking maneuver con- sisting of repeated changes in direction forward and backward. Nevertheless, Ackerman steering geometries have been especially popular in the hobby robotics market, where a robot can be built by starting with a remote control racecar kit and adding sensing and autonomy to the existing mechanism. In addition, the limited maneuverability of Ackerman

steering has an important advantage: its directionality and steering geometry provide it with very good lateral stability in high-speed turns.

2.3.1.5 Controllability

There is generally an inverse correlation between controllability and maneuverability. For example, the omnidirectional designs such as the four-castor wheel configuration require significant processing to convert desired rotational and translational velocities to individual wheel commands. Furthermore, such omnidirectional designs often have greater degrees of freedom at the wheel. For instance, the Swedish wheel has a set of free rollers along the wheel perimeter. These degrees of freedom cause an accumulation of slippage, tend to reduce dead-reckoning accuracy and increase the design complexity.

Controlling an omnidirectional robot for a specific direction of travel is also more diffi- cult and often less accurate when compared to less maneuverable designs. For example, an Ackerman steering vehicle can go straight simply by locking the steerable wheels and driv- ing the drive wheels. In a differential-drive vehicle, the two motors attached to the two wheels must be driven along exactly the same velocity profile, which can be challenging considering variations between wheels, motors, and environmental differences. With four- wheel omnidrive, such as the Uranus robot, which has four Swedish wheels, the problem is even harder because all four wheels must be driven at exactly the same speed for the robot to travel in a perfectly straight line.

In summary, there is no “ideal” drive configuration that simultaneously maximizes sta- bility, maneuverability, and controllability. Each mobile robot application places unique constraints on the robot design problem, and the designer’s task is to choose the most appropriate drive configuration possible from among this space of compromises.

2.3.2 Wheeled locomotion: case studies

Below we describe four specific wheel configurations, in order to demonstrate concrete applications of the concepts discussed above to mobile robots built for real-world activities.

2.3.2.1 Synchro drive

The synchro drive configuration (figure 2.22) is a popular arrangement of wheels in indoor mobile robot applications. It is an interesting configuration because, although there are three driven and steered wheels, only two motors are used in total. The one translation motor sets the speed of all three wheels together, and the one steering motor spins all the wheels together about each of their individual vertical steering axes. But note that the wheels are being steered with respect to the robot chassis, and therefore there is no direct way of reorienting the robot chassis. In fact, the chassis orientation does drift over time due to uneven tire slippage, causing rotational dead-reckoning error.

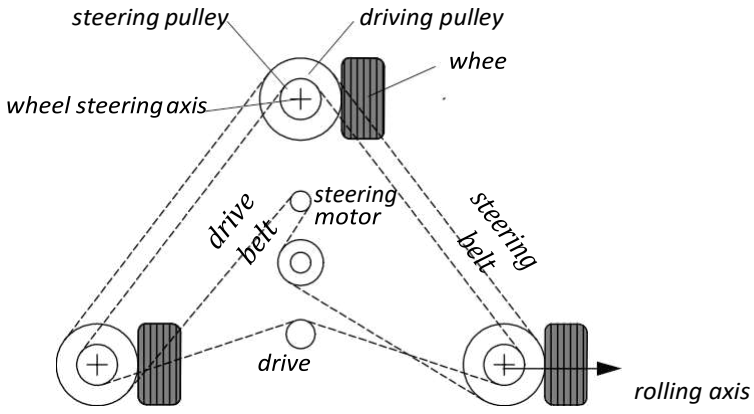

Figure 2.22

Synchro drive: The robot can move in any direction; however, the orientation of the chassis is not controllable.

Synchro drive is particularly advantageous in cases where omnidirectionality is sought. So long as each vertical steering axis is aligned with the contact path of each tire, the robot can always reorient its wheels and move along a new trajectory without changing its footprint. Of course, if the robot chassis has directionality and the designers intend to reorient the chassis purposefully, then synchro drive is only appropriate when combined with an independently rotating turret that attaches to the wheel chassis. Commercial research robots such as the Nomadics 150 or the RWI B21r have been sold with this configuration (figure 1.12).

In terms of dead reckoning, synchro drive systems are generally superior to true omni- directional configurations but inferior to differential-drive and Ackerman steering systems. There are two main reasons for this. First and foremost, the translation motor generally drives the three wheels using a single belt. Because of to slop and backlash in the drive train, whenever the drive motor engages, the closest wheel begins spinning before the fur- thest wheel, causing a

small change in the orientation of the chassis. With additional changes in motor speed, these small angular shifts accumulate to create a large error in ori- entation during dead reckoning. Second, the mobile robot has no direct control over the ori- entation of the chassis. Depending on the orientation of the chassis, the wheel thrust can be highly asymmetric, with two wheels on one side and the third wheel alone, or symmetric, with one wheel on each side and one wheel straight ahead or behind, as shown in figure
2.22. The asymmetric cases result in a variety of errors when tire-ground slippage can occur, again causing errors in dead reckoning of robot orientation.

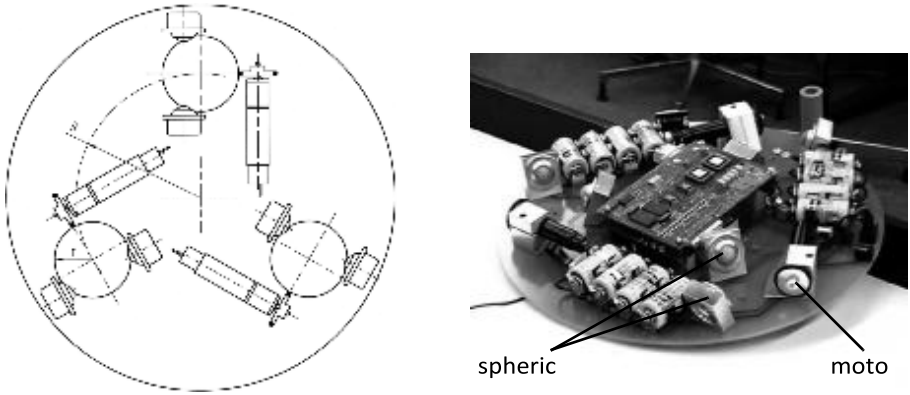

Figure 2.23

The Tribolo designed at EPFL (Swiss Federal Institute of Technology, Lausanne, Switzerland. Left: arrangement of spheric bearings and motors (bottom view). Right: Picture of the robot without the spherical wheels (bottom view).

2.3.2.2 Omnidirectional drive

As we will see later in section 3.4.2, omnidirectional movement is of great interest for com- plete maneuverability. Omnidirectional robots that are able to move in any direction (x, y, θ) at any time are also holonomic (see section 3.4.2). They can be realized by either using spherical, castor, or Swedish wheels. Three examples of such holonomic robots are presented below.

Omnidirectional locomotion with three spherical wheels. The omnidirectional robot depicted in figure 2.23 is based on three spherical wheels, each actuated by one motor. In this design, the spherical wheels are suspended by three contact points, two given by spher- ical bearings and one by a wheel connected to the motor axle. This concept provides excel- lent maneuverability and is simple in design. However, it is limited to flat surfaces and small loads, and it is quite difficult to find round wheels with high friction coefficients.

Omnidirectional locomotion with four Swedish wheels. The omnidirectional arrange- ment depicted in figure 2.24 has been used successfully on several research robots, includ- ing the Carnegie Mellon Uranus. This configuration consists of four Swedish 45-degree wheels, each driven by a separate motor. By varying the direction of rotation and relative speeds of the four wheels, the robot can be moved along any trajectory in the plane and, even more impressively, can simultaneously spin around its vertical axis.

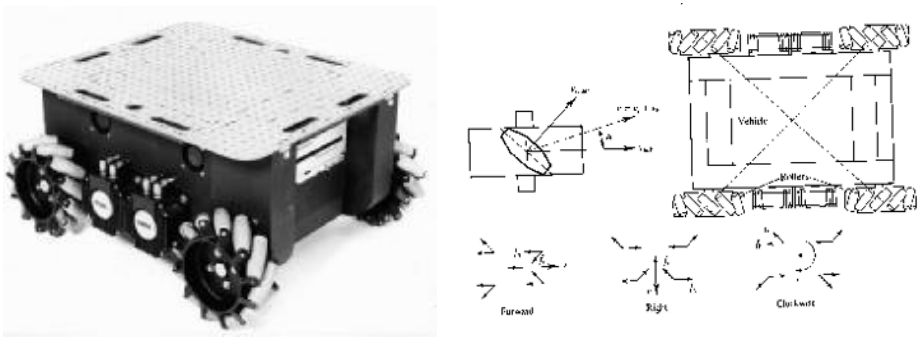

Figure 2.24

The Carnegie Mellon Uranus robot, an omnidirectional robot with four powered-swedish 45 wheels.

For example, when all four wheels spin “forward” or “backward” the robot as a whole moves in a straight line forward or backward, respectively. However, when one diagonal pair of wheels is spun in the same direction and the other diagonal pair is spun in the oppo- site direction, the robot moves laterally.

This four-wheel arrangement of Swedish wheels is not minimal in terms of control motors. Because there are only three degrees of freedom in the plane, one can build a three- wheel omnidirectional robot chassis using three Swedish 90-degree wheels as shown in table 2.1. However, existing examples such as Uranus have been designed with four wheels owing to capacity and stability considerations.

One application for which such omnidirectional designs are particularly amenable is mobile manipulation. In this case, it is desirable to reduce the degrees of freedom of the manipulator arm to save arm mass by using the mobile robot chassis motion for gross motion. As with humans, it would be ideal if the base could move omnidirectionally with- out greatly impacting the position of the manipulator tip, and a base such as Uranus can afford precisely such capabilities.

Omnidirectional locomotion with four castor wheels and eight motors. Another solution for omnidirectionality is to use castor wheels. This is done for the Nomad XR4000 from Nomadic Technologies (fig. 2.25), giving it excellent maneuverability. Unfortunately, Nomadic has ceased production of mobile robots.

The above three examples are drawn from table 2.1, but this is not an exhaustive list of all wheeled locomotion techniques. Hybrid approaches that combine legged and wheeled locomotion, or tracked and wheeled locomotion, can also offer particular advantages. Below are two unique designs created for specialized applications.

Figure 2.25

The Nomad XR4000 from Nomadic Technologies had an arrangement of four castor wheels for holonomic motion. All the castor wheels are driven and steered, thus requiring a precise synchronization and coordination to obtain a precise movement in x, y and θ .

2.3.2.3 Tracked slip/skid locomotion

In the wheel configurations discussed above, we have made the assumption that wheels are not allowed to skid against the surface. An alternative form of steering, termed slip/skid, may be used to reorient the robot by spinning wheels that are facing the same direction at different speeds or in opposite directions. The army tank operates this way, and the Nanokhod (figure 2.26) is an example of a mobile robot based on the same concept.

Robots that make use of tread have much larger ground contact patches, and this can sig- nificantly improve their maneuverability in loose terrain compared to conventional wheeled designs. However, due to this large ground contact patch, changing the orientation of the robot usually requires a skidding turn, wherein a large portion of the track must slide against the terrain.

The disadvantage of such configurations is coupled to the slip/skid steering. Because of the large amount of skidding during a turn, the exact center of rotation of the robot is hard to predict and the exact

change in position and orientation is also subject to variations depending on the ground friction. Therefore, dead reckoning on such robots is highly inac- curate. This is the trade-off that is made in return for extremely good maneuverability and traction over rough and loose terrain. Furthermore, a slip/skid approach on a high-friction surface can quickly overcome the torque capabilities of the motors being used. In terms of power efficiency, this approach is reasonably efficient on loose terrain but extremely inef- ficient otherwise.

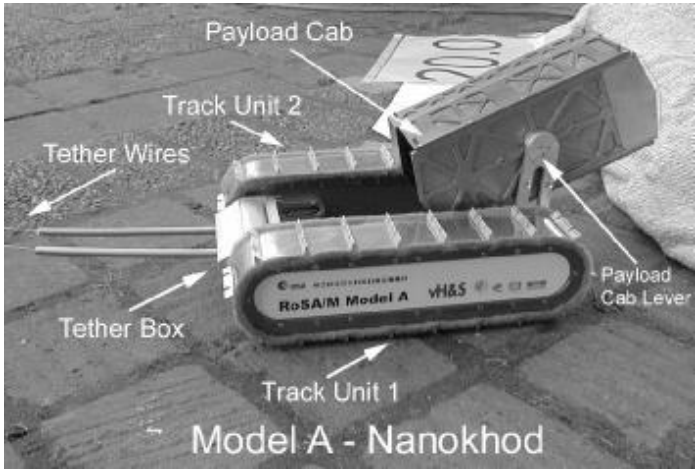

Figure 2.26

The microrover Nanokhod, developed by von Hoerner & Sulger GmbH and the Max Planck Institute, Mainz, for the European Space Agency (ESA), will probably go to Mars [138, 154].

2.3.2.4 Walking wheels

Walking robots might offer the best maneuverability in rough terrain. However, they are inefficient on flat ground and need sophisticated control. Hybrid solutions, combining the adaptability of legs with the efficiency of wheels, offer an interesting compromise. Solu- tions that passively adapt to the terrain are of particular interest for field and space robotics. The Sojourner robot of NASA/JPL (see figure 1.2) represents such a hybrid solution, able to overcome objects up to the size of the wheels. A more recent mobile robot design for similar applications has recently been produced by EPFL (figure 2.27). This robot, called Shrimp, has six motorized wheels and is capable of climbing objects up to two times its wheel diameter [97, 133]. This enables it to climb regular stairs though the robot is even smaller than the Sojourner. Using a rhombus configuration, the Shrimp has a steering wheel in the front and the rear, and two wheels arranged on a bogie on each side. The front wheel has a spring suspension to guarantee optimal ground

contact of all wheels at any time. The steering of the rover is realized by synchronizing the steering of the front and rear wheels and the speed difference of the bogie wheels. This allows for high-precision maneuvers and turning on the spot with minimum slip/skid of the four center wheels. The use of parallel articulations for the front wheel and the bogies creates a virtual center of rotation at the level of the wheel axis. This ensures maximum stability and climbing abilities even for very low friction coefficients between the wheel and the ground.

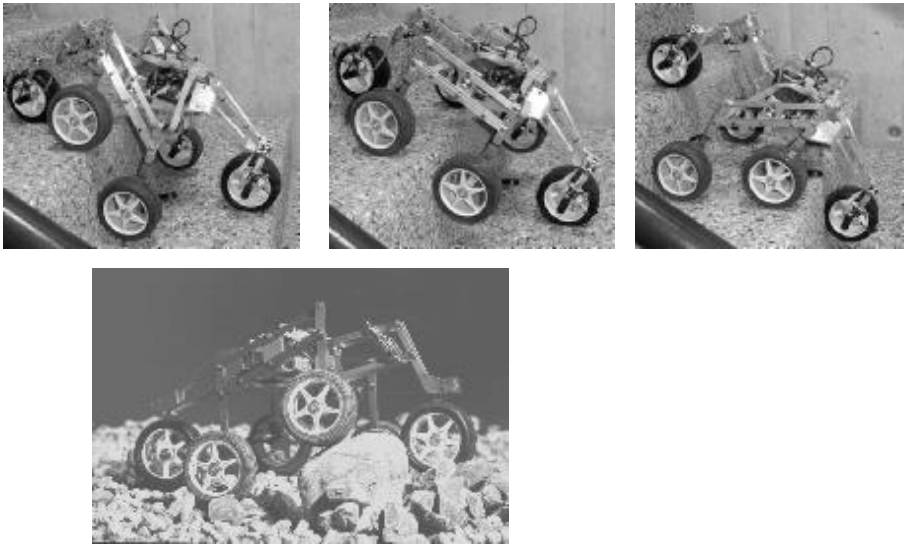

Figure 2.27

Shrimp, an all-terrain robot with outstanding passive climbing abilities (EPFL [97, 133]).

The climbing ability of the Shrimp is extraordinary in comparison to most robots of sim- ilar mechanical complexity, owing much to the specific geometry and thereby the manner in which the center of mass (COM) of the robot shifts with respect to the wheels over time. In contrast, the Personal Rover demonstrates active COM shifting to climb ledges that are also several times the diameter of its wheels, as demonstrated in figure 2.28. A majority of the weight of the Personal Rover is borne at the upper end of its swinging boom. A dedi- cated motor drives the boom to change the front/rear weight distribution in order to facili- tate step-climbing. Because this COM-shifting scheme is active, a control loop must explicitly decide how to move the boom during a climbing scenario. In this case the Per- sonal Rover accomplished this closed-loop control by inferring terrain based on measure- ments of current flowing to each independently driven wheel [66].

As mobile robotics research matures we find ourselves able to

design more intricate mechanical systems. At the same time, the control problems of inverse kinematics and dynamics are now so readily conquered that these complex mechanics can in general be controlled. So, in the near future, we can expect to see a great number of unique, hybrid mobile robots that draw together advantages from several of the underlying locomotion mechanisms that we have discussed in this chapter. They will each be technologically impressive, and each will be designed as the expert robot for its particular environmental niche.

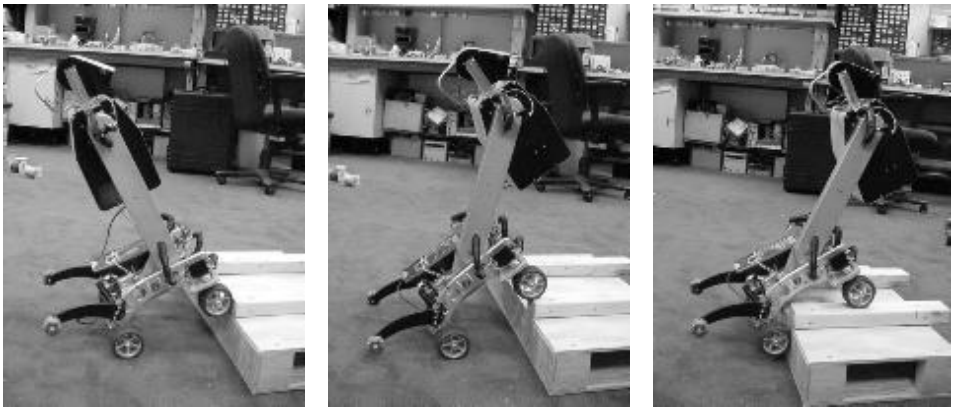

Figure 2.28

The Personal Rover, demonstrating ledge climbing using active center-of-mass shifting.

CHAPT

3. Research Overview

Chapter

Autonomous robots

Both animals and robots manipulate objects in their environment in order to achieve certain goals. Animals use their senses (e.g. vision, touch, smell) to probe the environment. The resulting information, in many cases also en- hanced by the information available from internal states (based on short-term or long-term memory), is processed in the brain, often resulting in an action carried out by the animal, with the use of its limbs.

Similary, robots gain information of the surroundings, using their sensors. The information is processed in the robot's brain[1], consisting of one or several processors, resulting in motor signals that are sent to the actuators (e.g. motors) of the robot.

In this course, the problem of providing robots with the ability of making rational, intelligent decisions will be central. Thus, the development of robotic brains is the main theme of the course. However, a robotic brain cannot op- erate in isolation: It needs sensory inputs, and it must produce motor output in order to influence objects in the environment. Thus, while it is the author's view that the main challenge in contemporary robotics lies with the devel- opment of robotic brains, consideration of the actual hardware, i.e. sensors, processors, motors etc., is certainly very important as well.

This chapter gives a brief overview of **robotic hardware**, i.e. the actual frame (body) of a robot, as well as its sensors, actuators, processors etc. The

[1]The term *control system* is commonly used (instead of the term **robotic brain**). However, this term is misleading, as it leads the reader to think of classical control theory. Concepts from classical control theory *are* relevant in robots; For example, the low-level control of the motors of robots is often taken care of by PI- or PID-regulators. However, **autonomous robots**, i.e. freely moving robots that operate without direct human supervision, are expected to function in complex, unstructured environments, and to make their own decisions concerning which action to take in any given situation. In such cases, systems based *only* on classical control theory are simply insufficient. Thus, hereafter, the term **robotic brain** (or, simply, *brain*) will be used when referring to the system that provides an autonomous robot, however simple, with the ability to process information and decide upon which actions to take.

1

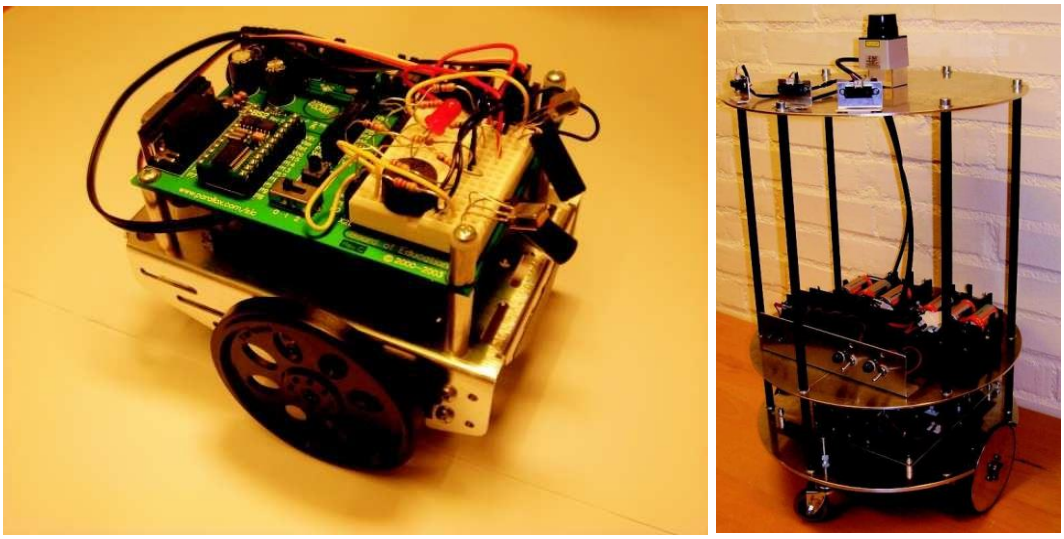

Figure 1.1: *Left panel: A Boe-bot. Right panel: A wheeled robot currently under construction in the Adaptive systems research group at Chalmers.*

various hardware-related issues will be studied in greater detail in the second half of the course, which will involve the construction of an actual robot of the kind shown in the left panel of Fig. 1.1.

1.1 Robot types

The are many different types of robots, and the taxonomy of such machines can be constructed in various ways. For example, one may classify different kinds of robots based on their complexity, their likeness to humans (or animals), their way of moving etc. In this course we shall limit ourselves to **mobile robots**, that is, robots that are able to move freely using, for example, wheels. The other main category of robots are stationary robotic arms, also referred to as **robotic manipulators**. Of course, as with any taxonomy, there are always ex- amples that do not fit neatly into any of the available categories. For example, a **smart home** equipped with a central computer and, perhaps, some form of manipulation capabilities, can also be considered a robot, albeit of a different kind.

Robotic manipulators constitute a very important class of robots and they are used extensively in many industries, for example in assembly lines in the vehicle industry. However, such robots normally follow a predefined move- ment sequence and are not equipped with behaviors (such as collision avoid- ance) designed to avoid harming people. While there is nothing preventing the use of, for instance, sonar proximity sensors on a robotic manipulator, such op- tions are rarely used. Instead, manipulators are confined to **robotic work cells**,

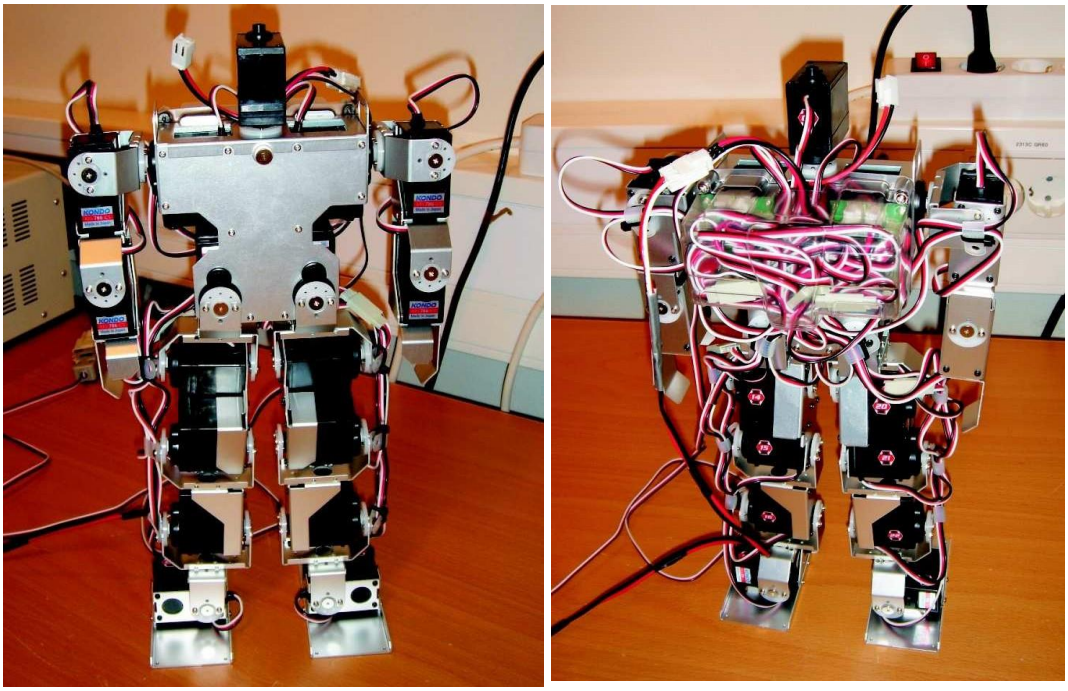

Figure 1.2: *A Kondo humanoid robot. Left panel: Front view. Right panel: Rear view.*

in which people are forbidden to enter while the manipulator is active.

By contrast, in this course, we shall consider **autonomous robots**, i.e. robots that are capable of making their own decisions (depending on the situation at hand) rather than merely executing a pre-defined sequence of motions. In fact, since most robots equipped with such decision-making capabilities are mo- bile, one may define an autonomous robot as a mobile robot with the ability to make decisions. Two examples of mobile robots can be seen in Fig. 1.1. The left panel shows a Boe-bot, which will be assembled and used in the second half of the course. Some of its main advantages are its small size (its length is around 0.14 m and its width 0.11 m) and its simplicity. Needless to say, the robot also has several limitations; for example, its onboard processor (micro- controller) is quite slow. However, on balance, the Boe-bot provides a good introduction to the field of mobile robots. The right panel of Fig. 1.1 shows a two-wheeled differentially steered robot which, although still under construction at the Adaptive systems research group at Chalmers, is already being used in several research projects. This robot has a diameter of 0.40 m and a height of around 1.00 m.

Robotic manipulators have long dominated the market for robots, but with the advent of low-cost mobile robots the situation is changing: In 2007, the number of mobile robots surpassed the number of manipulators for the first time, and the gap is expected to widen over the next decades.

The class of mobile robots can be further divided into subclasses, the most

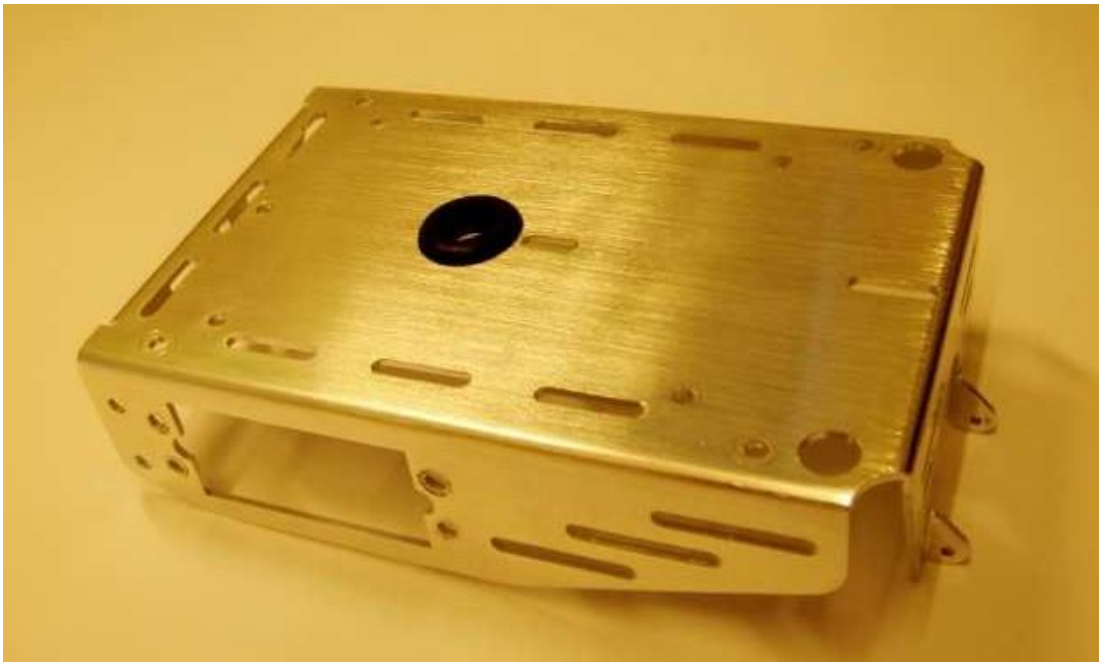

Figure 1.3: *The aluminium frame of a Boe-bot.*

important being **legged robots** and **wheeled robots**. Other kinds, such as fly- ing robots, exist as well, but will not be considered in this course. The class of legged robots can be subdivided based on the number of legs, the most common types being **bipedal robots** (with two legs) and **quadrupedal robots** (with four legs). Most bipedal robots resemble humans, at least to some extent; such robots are referred to as **humanoid robots**. An example of a humanoid robot is shown in Fig. 1.2. Humanoid robots that (unlike the robot shown in Fig. 1.2) not only have the approximate shape of a human, but have also been equipped with more detailed human-like features, e.g. artificial skin, artificial hair etc., are called **androids**. It should be noted that the term *humanoid* refers to the shape of the robot, not its size; in fact, many humanoid robots are quite small. For example, the Kondo robot shown in Fig. 1.2 is approximately 0.35 m tall.

Some robots are *partly* humanoid. For example, the wheeled robot shown in the right panel of Fig. 1.1 is currently being equipped with a humanoid up- per body. Unlike a fully humanoid robot, this robot need not be actively bal- anced, but will still exhibit many desirable features of humanoid robots, such as two arms for grasping and lifting objects, gesturing etc., as well as a head that will be equipped with two cameras for stereo vision and microphones providing capabilities for listening and speaking.

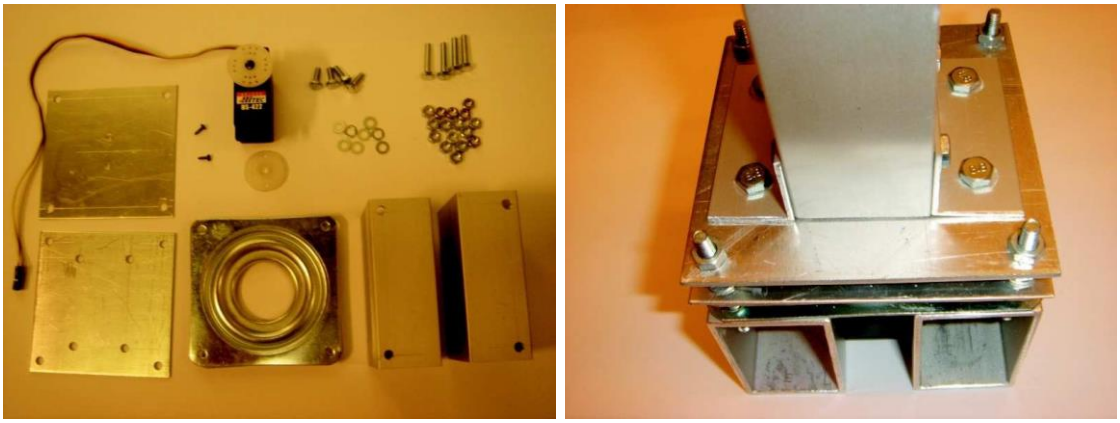

Figure 1.4: *Left panel: Aluminium parts used in the construction of a rotating base for a humanoid upper body. The servo motor used for rotating the base is also shown, as well as the screws, washers and nuts. Right panel: The assembled base.*

1.2 Robotic hardware

1.2.1 Construction material

Regarding the material used in the actual frame of the robot, several options are available, such as e.g. aluminium, steel, various forms of plastic etc. The frame of a robot should, of course, preferably be constructed using a material that is both sturdy and light and, for that reason, aluminium is often chosen. Albeit somewhat expensive, aluminium combines toughness with low weight in a near-optimal way, at least for small mobile robots. Steel is typically too heavy to be practical in a small robot, whereas many forms of plastic eas- ily break. The frame of the robot used in this course (the Boe-bot) is made in aluminium, and is shown in Fig. 1.3. The left panel of Fig 1.4 shows the aluminium parts used in a rotating base for a humanoid upper body. The as- sembled base, which can rotate around the vertical axis, is shown in the right panel.

1.2.2 Sensors

The purpose of robotic sensors is to measure either some physical characteris- tic of the robot (for example, its acceleration) or some aspect of its environment (for example, the detected intensity of a light source). The raw data thus ob- tained must then, in most cases, be processed further before being used in the brain of the robot. For example, an infrared (IR) proximity sensor may pro- vide a voltage (depending on the distance to the detected object) as its read- ing, which can then be converted to a distance, using the characteristics of the sensor available from its data sheet.

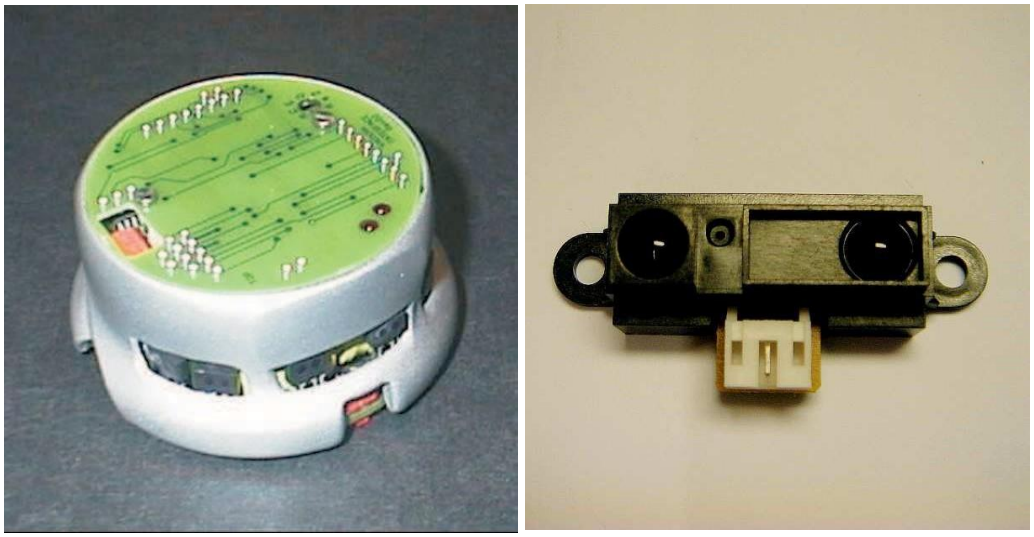

Figure 1.5: *Left panel: A Khepera II robot. Note the IR proximity sensors (small black rectangles around the periphery of the robot), consisting of an emitter and a detector. Right panel: A Sharp GP2D12 infrared sensor.*

Needless to say, there exists a great variety of sensors for mobile robots. Here, only a brief introduction will be given, focusing on a few fundamental sensor types.

Infrared proximity sensors

An **infrared proximity sensor** (or **IR sensor**, for short), consists of an emitter and a detector. The emitter, a light-emitting diode (LED), sends out infrared light, which bounces off nearby objects, and the reflected light is then mea- sured by the detector (e.g. a phototransistor). Some IR sensors can also be used for measuring the ambient light level, i.e. the light observed by the detector when the emitter is switched off. As an example, consider the Khepera robot (manufactured by K-Team, www.k-team.com), shown in the left panel Fig. 1.5. This robot is equipped with eight IR sensors, capable of measuring both am- bient and reflected light. The range of IR sensors is quite short, though. In the Khepera robot, reflected light measurements are only useful to a distance of around 0.050 m from the robot, i.e. approximately one robot diameter, even though other IR sensors have longer range. Another example is the Sharp GP2D12 IR sensor, shown in the right panel of Fig. 1.5. This sensor detects ob- jects in the range $[0.10, 0.80]$ m. It operates using a form of triangulation: Light is emitted from the sensor and, if an object is detected, the reflected light is re- ceived at an angle that depends on the distance to the detected object. The raw signal from the sensor consists of a voltage that can be mapped to a distance. The mapping is non-linear, and for very short distances, the sensor cannot give

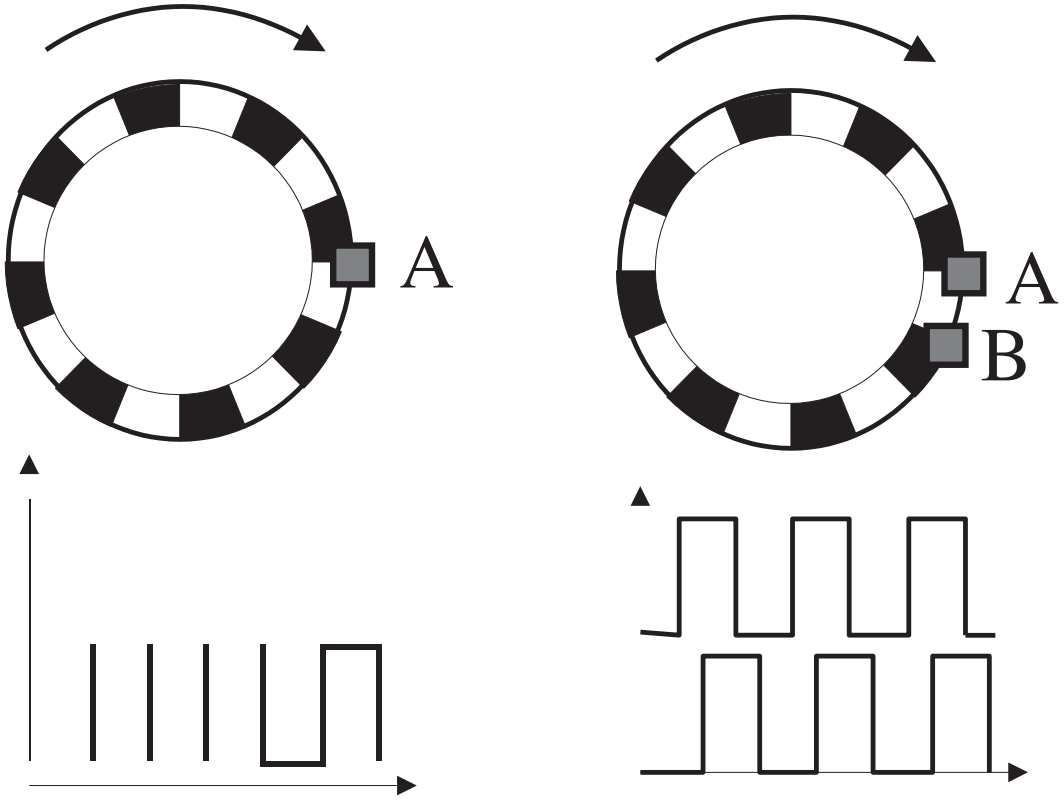

Figure 1.6: *The left panel shows a simple encoder, with a single detector (A), that measures the interruptions of a light beam, producing the curve shown below the encoder. In the right panel, two detectors are used, making it possible to determine also the* direction *of rotation.*

reliable readings (hence the lower limit of 0.10 m).

Digital optical encoders

In many applications, accurate position information is essential for a robot, and there are many different methods for positioning, e.g. inertial navigation, GPS navigation, landmark detection etc., some of which will be considered in a later chapter. One of the simplest forms of positioning, however, is **dead reck- oning**, in which the position of a robot is determined based on measurements of the distance travelled by each wheel of the robot. This information, when combined with knowledge of the robot's physical properties (i.e. its kinemat- ics, see Chapter 2) allows one to deduce the current position and heading. The process of measuring the rotation of the wheel of a robot is an example of **odometry**, and a sensor capable of such measurements is the **digital optical encoder** or, simply, **encoder**. Essentially, an encoder is a disc made of glass or plastic, with shaded regions that regularly interrupt a light beam. By count- ing the number of interruptions, the rotation of the wheel can be deduced, as shown in the left panel of Fig. 1.6. However, in order to determine also the *di- rection* of rotation, a second detector, placed at a quarter of a cycle out of phase

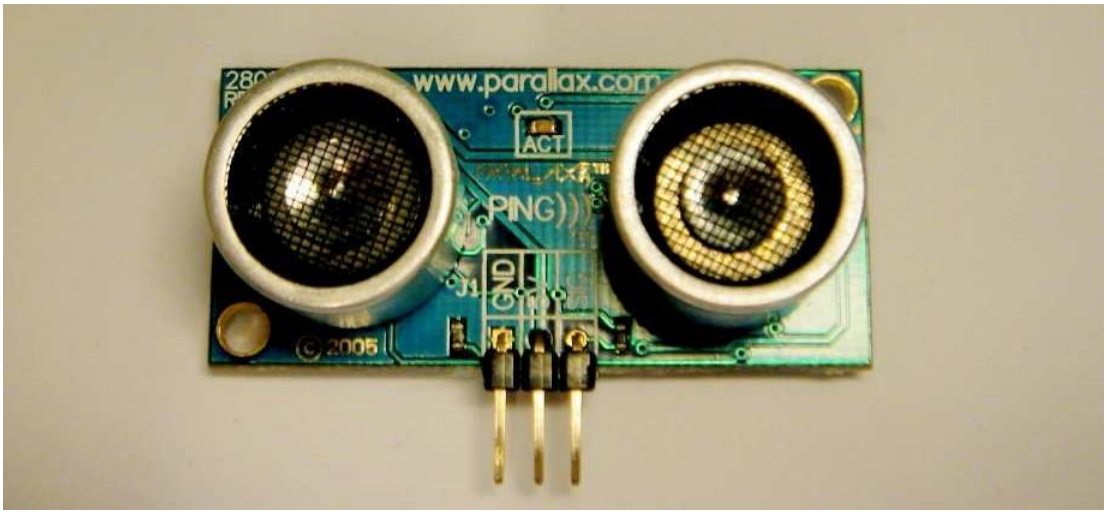

Figure 1.7: *A Ping ultrasonic distance sensor.*

with the first detector, is needed (such an arrangement is called **quadrature encoding**, and is shown in the right panel of Fig. 1.6).

Ultrasound (sonar) sensors

Ultrasound sensors, also known as **sonar sensors** or simply **sonars**, are based on time-of-flight measurement. Thus, in order to detect the distance to an ob- ject, a sonar emits a brief pulse of ultrasonic sound, typically in the frequency range 40-50 kHz[2]. The sensor then awaits the echo. Once the echo has been detected, the distance to the object can be obtained using the fact that sound travels at a speed of around 340 m/s. As in the case of IR sensors, there is both a lower and an upper limit for the detection range of a sonar sensor. If the distance to an object is too small, the sensor simply does not have enough time to switch from emission to listening, and the signal is lost. Similarly, if the distance is too large, the echo may be too weak to be detected.

Fig. 1.7 shows a Ping ultrasonic distance sensor, which is commonly used in connection with the Boe-bot. This sensor can detect distances to objects in the range $[0.02, 3.00]$ m.

Laser range finders

Laser range finders (LRFs) commonly rely, like sonar sensors, on time-of-flight measurements, but involve the speed of light rather than the speed of sound. Thus, a laser range finder emits pulses of laser light (in the form of thin beams),

[2]For comparison, a human ear can detect sounds in the range 20 hz to 20 kHz. Thus, the sound pulse emitted by a sonar sensor is not audible.

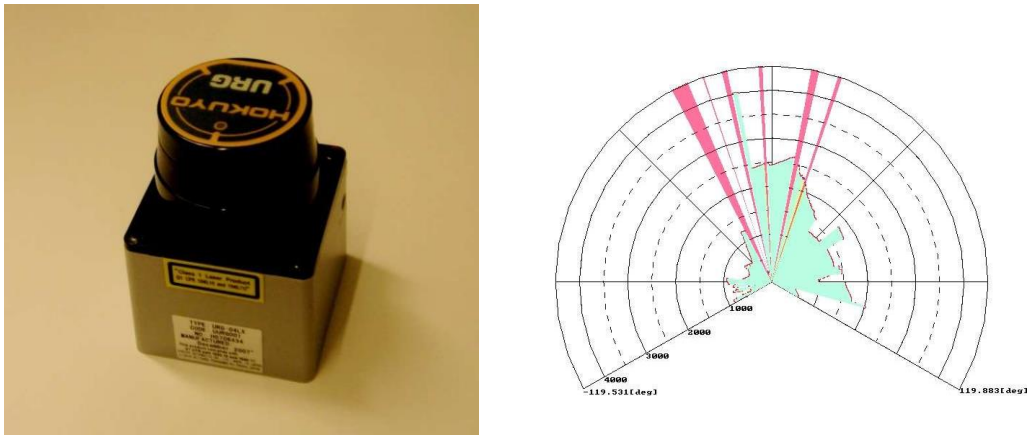

Figure 1.8: *Left panel: A Hokuyo URL-04LX laser range finder. Right panel: A typical reading, showing the distance to the nearest object in various directions. The pink rays indicate directions in which no detection is made. The maximum range of the sensor is 4 m.*

and measures the time it takes for the pulse to bounce off a target and return to the range finder. An LRF carries out a sweep over many directions[3] resulting in an accurate local map of distances to objects along the line-of-sight of each ray. LRFs are generally very accurate sensors, but they are also much more expensive than sonars sensors and IR sensors.

A Hokuyo URG-04LX LRF is shown in the left panel of Fig. 1.8. This sensor has a range of around four meters, with an accuracy of around 1 mm. It can generate readings in 683 different directions, with a frequency of around 10 Hz. As of the time of writing (Jan. 2010), a Hokuyo URG-04LX costs around 2,000 USD. The right panel of Fig. 1.8 shows a typical reading, obtained from the software delivered with the LRF.

Cameras

Cameras are used as the eyes of a robot. In many cases, two cameras are used, in order to provide the robot with binocular vision, allowing it to estimate the range to detected objects. There are many cameras available for robots, for example the CMUCam series which has been developed especially for use in mobile robots; The processor connected to the CMUCam is capable of basic image processing. At the time of writing (Jan. 2010), a CMUCam costs on the order of 150 USD. A low-cost alternative is to use ordinary webcams, for which prices start around 15 USD. Fig. 1.9 shows a simple robotic head consisting of two servo motors (see below) and a single webcam.

However, while the actual cameras may not be very costly, the use of cam- eras is *computationally* a very expensive procedure. Even at a low resolution,

[3] A typical angular interval for an LRF is around 180-240 degrees.

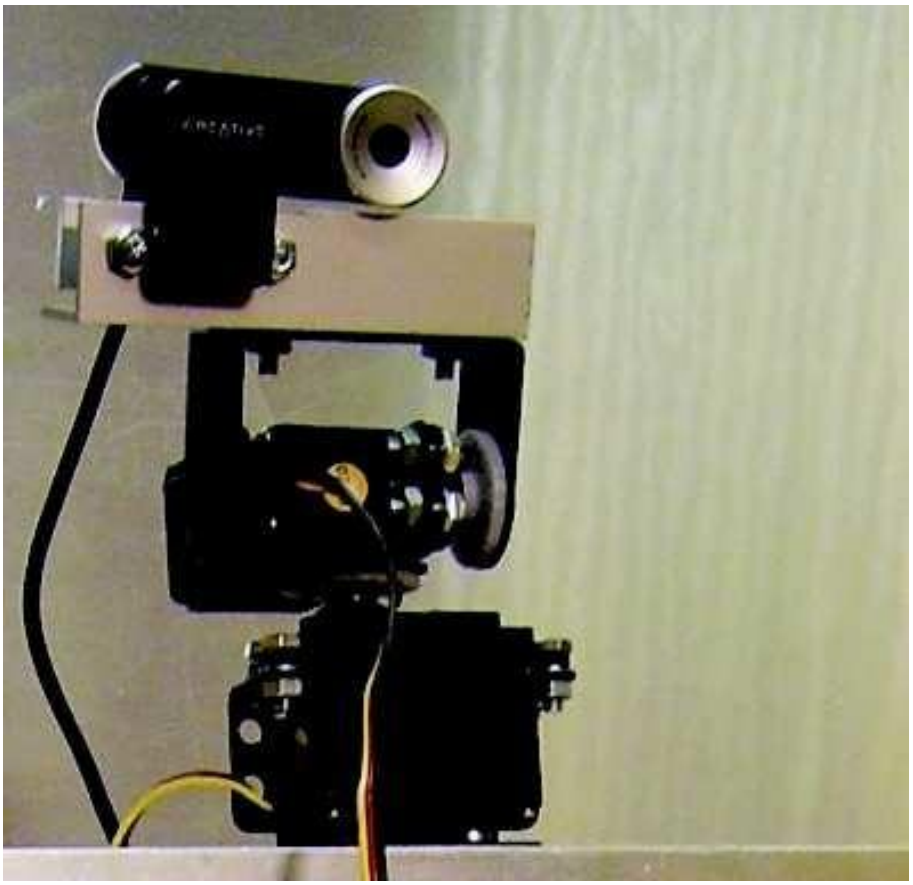

Figure 1.9: *A simple robotic head, consisting of two servo motors and a webcam.*

say 320×240 pixels, a webcam will deliver a flow of around 1.5 Mb/s, assum- ing a frame rate of 20 Hz and a single byte of data per pixel. The actual data transfer is easily handled by a Universal serial bus (USB), but the data must not only be transferred but also *analyzed*, something which is far from trivial. An introduction to image processing for robots will be given in a later chapter.

Other sensors

In addition to odometry based on digital optical encoders, robot positioning can be based on **inertial sensors**, i.e. sensors that measure the time derivatives of the position or heading angle of the robot. Examples of inertial sensors are **accelerometers**, measuring linear acceleration, and **gyroscopes**, measuring an- gular acceleration. Essentially, an accelerometer consists of a small object, with mass m, attached to a spring and damper, as shown in Fig. 1.10. As the system accelerates, the displacement z of the small object can be used to deduce the acceleration $\ddot{x}$ of the robot. Given continuous measurements of the accelera- tion, as a function of time, the position (relative to the starting position) can be

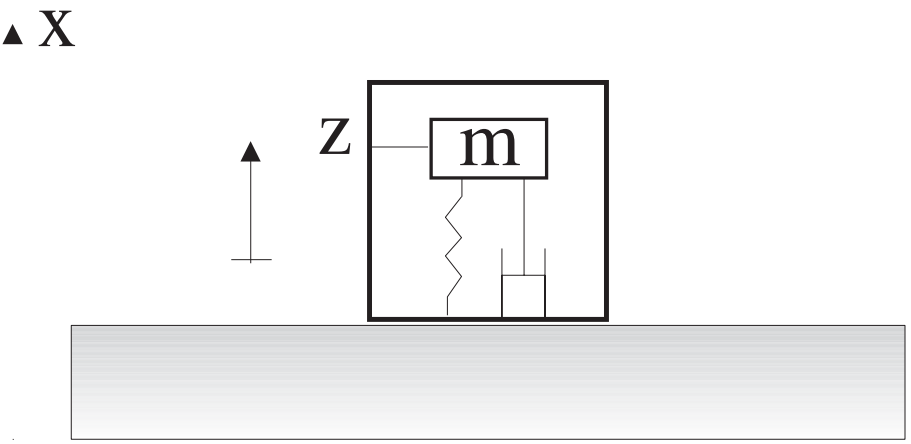

Figure 1.10: *An accelerometer. The motion of the small object (mass m) resulting from the acceleration of the larger object to which the accelerometer is attached can be used for deducing the acceleration.*

deduced. For robots operating in outdoor environments, positioning based on the **global positioning system (GPS)** is often a good alternative. The GPS re- lies on 24 satellites that transmit radio frequency signals which can be picked up by objects on Earth. Given the exact position of (at least) three satellites, rel- ative to the position of e.g. a robot, the absolute position (latitude, longitude, and altitude) of the robot can be deduced.

Other sensors include **strain gauge sensors** (measuring deformation), **tac- tile (touch) sensors** measuring physical contact between a robot and objects in its environment, and **compasses**, measuring the direction of movement.

1.2.3 Actuators

An **actuator** is a device that allows a robot to take action, i.e. to move or manip- ulate the surroundings in some other way. Motors, of course, are very common types of actuators. Other kinds of actuation include, for example, the use of microphones (for human-robot interaction).

Movements can be generated in various ways, using e.g. electrical motors, pneumatic or hydraulic systems etc. In this course, we shall only consider electrical, direct-current (DC) motors and, in particular, servo motors. Thus, when referring to actuation in this course, the use of such motors is implied.

Note that actuation normally requires the use of a **motor controller** in con- nection with the actual motor. This is so, since the microcontroller (see below) responsible for sending commands to the motor cannot, in general, provide sufficient current to drive the motor. The issue of motor control will be consid- ered briefly in connection with the discussion of servo motors below.

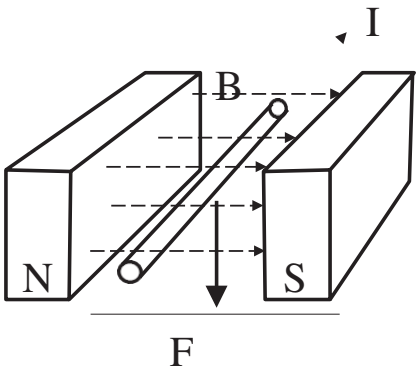

Figure 1.11: *A conducting wire in a magnetic field. B denotes the magnetic field strength and I the current through the wire. The Lorentz force* $\mathbf{F}$ *acting on the wire is given by* $\mathbf{F} = \mathbf{I} \times \mathbf{B}$.

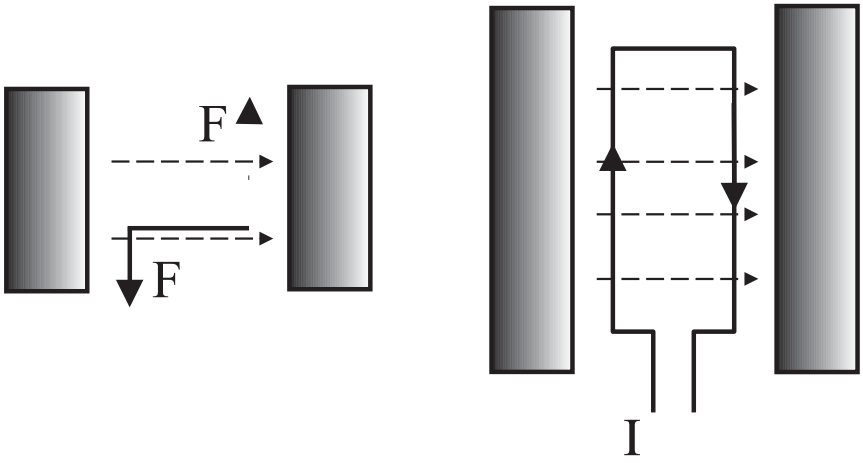

Figure 1.12: *A conducting loop of wire placed in a magnetic field. Due to the forces acting on the loop, it will begin to turn. The loop is shown from above in the right panel, and from the side in the left panel.*

DC motors

Electrical direct current (DC) motors are based on the principle that a force acts on a wire in a magnetic field if a current is passed through the wire, as illustrated in Fig. 1.11. If instead a current is passed through a closed loop of wire, as illustrated in Fig. 1.12, the forces acting on the two sides of the loop will point in opposite directions, making the loop turn. A standard DC motor consists of an outer stationary cylinder (the **stator**), providing the magnetic field, and an inner, rotating part (the **rotor**). From Fig. 1.12 it is clear that the loop will reverse its direction of rotation after a half-turn, unless the direction of the current is reversed. The role of the **commutator**, connected to the rotor of a DC motor, is to reverse the current through the motor every half-turn, thus allowing continuous rotation. Finally, carbon **brushes**, attached to the stator, complete the electric circuit of the DC motor. There are types of DC motors

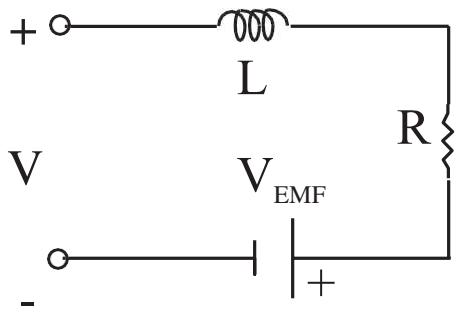

Figure 1.13: *The equivalent electrical circuit for a DC motor.*

that use electromagnets rather than a permanent magnet, and also types that are brushless. However, a detailed description of such motors are beyond the scope of this text.

DC motors are controlled by varying the applied voltage. The equations for DC motors can be divided into an electrical and a mechanical part. The motor can be modelled electrically by the equivalent circuit shown in Fig. 1.13. Letting V denote the applied voltage, and ω the angular speed of the motor shaft, the electrical equation takes the form

$$V = L\frac{\mathrm{d}i}{\mathrm{d}t} + Ri + V_{\text{EMF}}, \tag{1.1}$$

where i is the current flowing through the circuit, L is the inductance of the motor, R its resistance, and V_{EMF} the voltage (the **back EMF**) counteracting V. The back EMF depends on the angular velocity, and can be written as

$$V_{\text{EMF}} = c_e\omega, \tag{1.2}$$

where c_e is the **electrical constant** of the motor. For a DC motor, the generated torque τ_g is directly proportional to the current, i.e.

$$\tau_g = c_t i, \tag{1.3}$$

where c_t is the **torque constant** of the motor. Turning now to the mechanical equation, Newton's second law gives

$$I\frac{\mathrm{d}\omega}{\mathrm{d}t} = \sum \tau, \tag{1.4}$$

where I is the combined moment of inertia of the motor and its load, and $\sum\tau$ is the total torque acting on the motor. For the DC motor, the equation takes the form

$\mathrm{d}\omega$

$$I\frac{\mathrm{d}^2\theta}{\mathrm{d}t} = \tau_g - \tau_f - \tau, \tag{1.5}$$

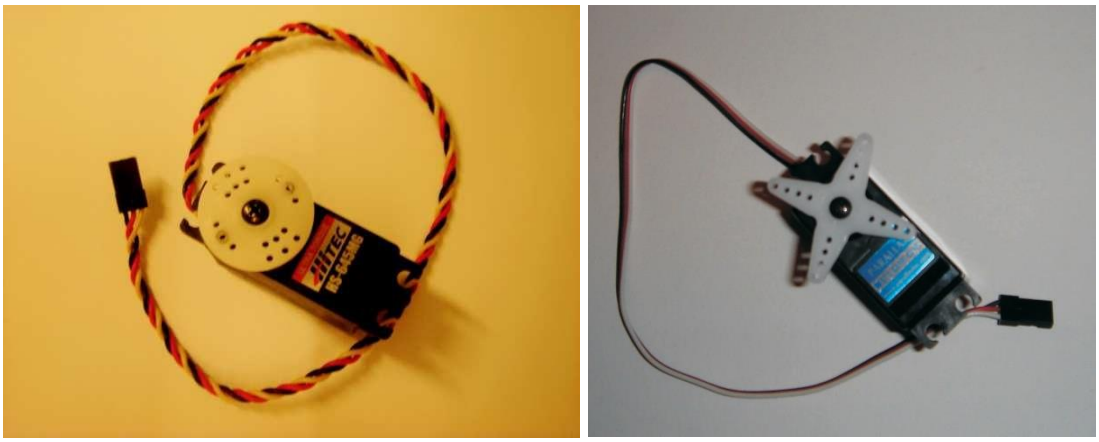

Figure 1.14: *Left panel: A HiTec 645MG servo. The suffix* MG *indicates that the servo is equipped with a metal gear train. Right panel: A Parallax servo, which has been modified for continuous rotation. Servos of this kind are used on the Boe-bot. The circular (left) and star-shaped (right) white plastic objects are the servo horns.*

where τ_f is the frictional torque opposing the motion and τ is the (output) torque acting on the load. The frictional torque can be divided into two parts, the **Coulomb friction** ($c_C \mathrm{sgn}(\omega)$) and the **viscous friction** ($c_v\omega$). Thus, the equations for the DC motor can now be summarized as

$$\tau_g = \frac{c_t}{R} V - \frac{c_t L}{R} \frac{\mathrm{d}i}{\mathrm{d}t} - \frac{c_e c_t}{R} \omega, \tag{1.6}$$

$$I \frac{\mathrm{d}\omega}{\mathrm{d}t} = \tau_g - c_C \mathrm{sgn}(\omega) - c_v \omega - \tau, \tag{1.7}$$

In many cases, the time constant of the electrical circuit is much shorter than that of the physical motion, so the inductance term can be neglected. Further- more, for simplicity, the dynamics of the mechanical part can *also* be neglected under certain circumstances (e.g. if the moment of inertia of the motor and load is small). Thus, setting $\mathrm{d}i/\mathrm{d}t$ and $\mathrm{d}\omega/\mathrm{d}t$ to zero, the steady-state DC mo- tor equations, determining the torque τ on the load for a given applied voltage V and a given angular velocity ω

$$\tau_g = \frac{c_t}{R} V - \frac{c_e c_t}{R} \omega, \tag{1.8}$$

$$\tau = \tau_g - c_C \mathrm{sgn}(\omega) - c_v \omega, \tag{1.9}$$

are obtained. In many cases, the axis of a DC motor rotates too fast and gener- ates a torque that is too weak for driving a robot. Thus, a **gear box** is commonly used, which reduces the rotation speed taken out from the motor (on the **sec- ondary drive shaft**) while, at the same time, increasing the torque. For an ideal (loss-free) gear box, the output torque and rotation speed are given by

$$\tau_{\mathrm{out}} = G\tau,$$

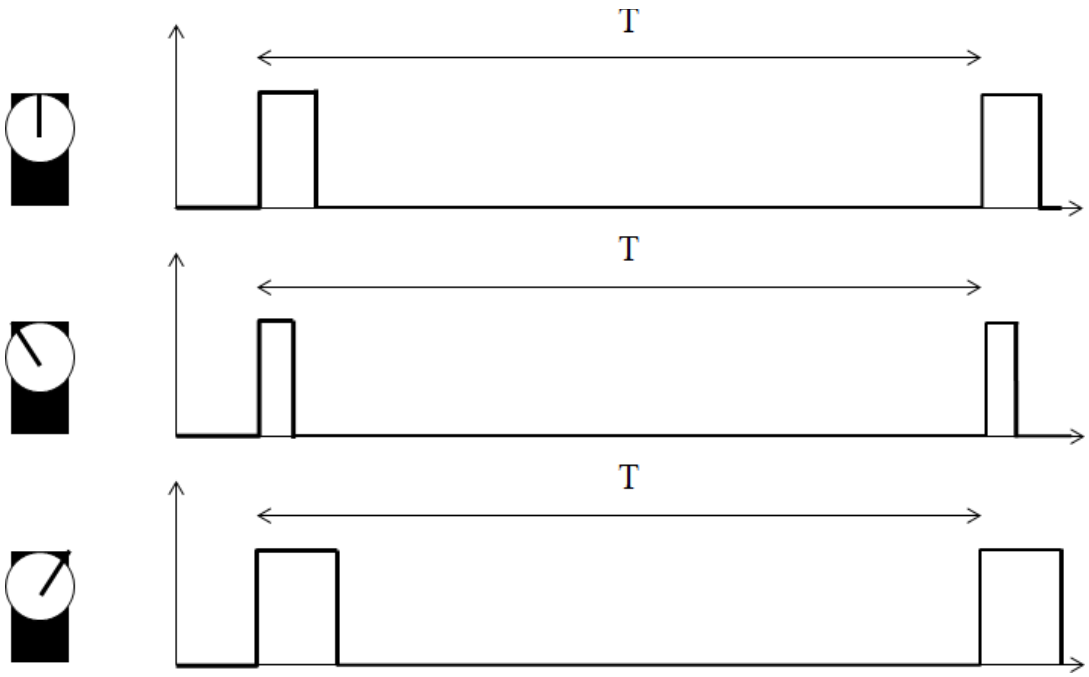

Figure 1.15: *Pulse width modulation control of a servo motor. The lengths of the pulses determine the requested position angle of the motor output shaft. The interval betwwn pulses (typically around 20 ms) is denoted* T *.*

$$\omega_{\text{out}} = \frac{1}{G}\omega, \tag{1.10}$$

where G is the **gear ratio**.

Servo motors

A **servo motor** is essentially a DC motor equipped with control electronics and a gear train (whose purpose is to increase the torque to the required level for moving the robot, as described above). The actual motor, the gear train, and the control electronics, are housed in a plastic container. A **servo horn** (either plastic or metal) makes it possible to connect the servo motor to a wheel or some other structure. Fig. 1.14 shows two examples of servo motors.

The angular position of a servo motor's output shaft is determined using a **potentiometer**. In a standard servo, the angle is constrained to a given range [α_{max}, α_{max}], and the role of the control electronics is to make sure that the servo rotates to a set position α (given by the user). A servo is fitted with a three-wire cable. One wire connects the servo to a power source (for exam- ple, a motor controller or, in some cases, a microcontroller board) and another wire connects it to *ground*. The third wire is responsible for sending signals to the servo motor. In servo motors, a technique called **pulse width modulation**

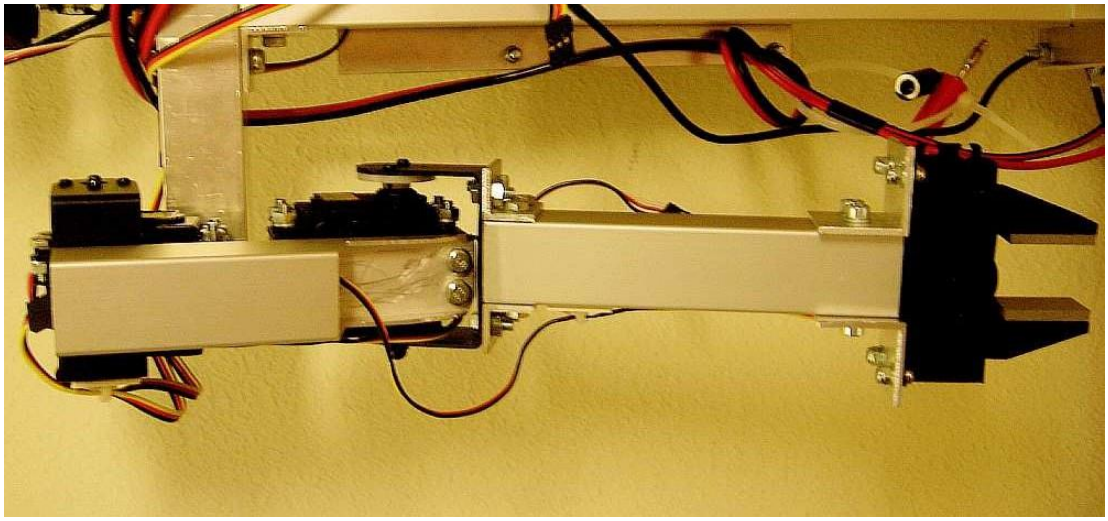

Figure 1.16: *An arm of a humanoid robot. The allowed rotation range of the elbow is around 100 degrees.*

(PWM) is used: Signals in the form of pulses are sent (e.g. from a microcon- troller) to the control electronics of the servo motor. The duration of the pulses determine the required position, to which the servo will (attempt to) rotate, as shown in Fig. 1.15. For a walking robot (or for a humanoid upper body), the limitation to a given angular range poses no problem: The allowed rotation range of a servo is normally sufficient for, say, an elbow or a shoulder joint. As an example, an arm of a humanoid robot is shown in Fig. 1.16. For this particu- lar robot, the rotation range for the elbow joint is around 100 degrees, i.e. easily within the range of a standard servo (around 180 degrees). The limitation is, of course, not very suitable for motors driving the wheels of a robot. Fortunately, servo motors can be modified to allow continuous rotation. The Boe-bot that will be built in the second half of the course uses **Parallax continuous rotation servos** (see the right panel of Fig. 1.14), rather than standard servos.

Other motors

There are many different types of motors, in addition to standard DC motors and servo motors. An example is the **stepper motor**, which is also a version of the DC motor, namely one that moves in fixed angular increments, as the name implies. However, in this course, only standard DC motors and servo motors will be considered.

1.2.4 Processors

Sensors and actuators are necessary for a robot to be able to perceive its environment and to move or manipulate the environment in various ways. How-

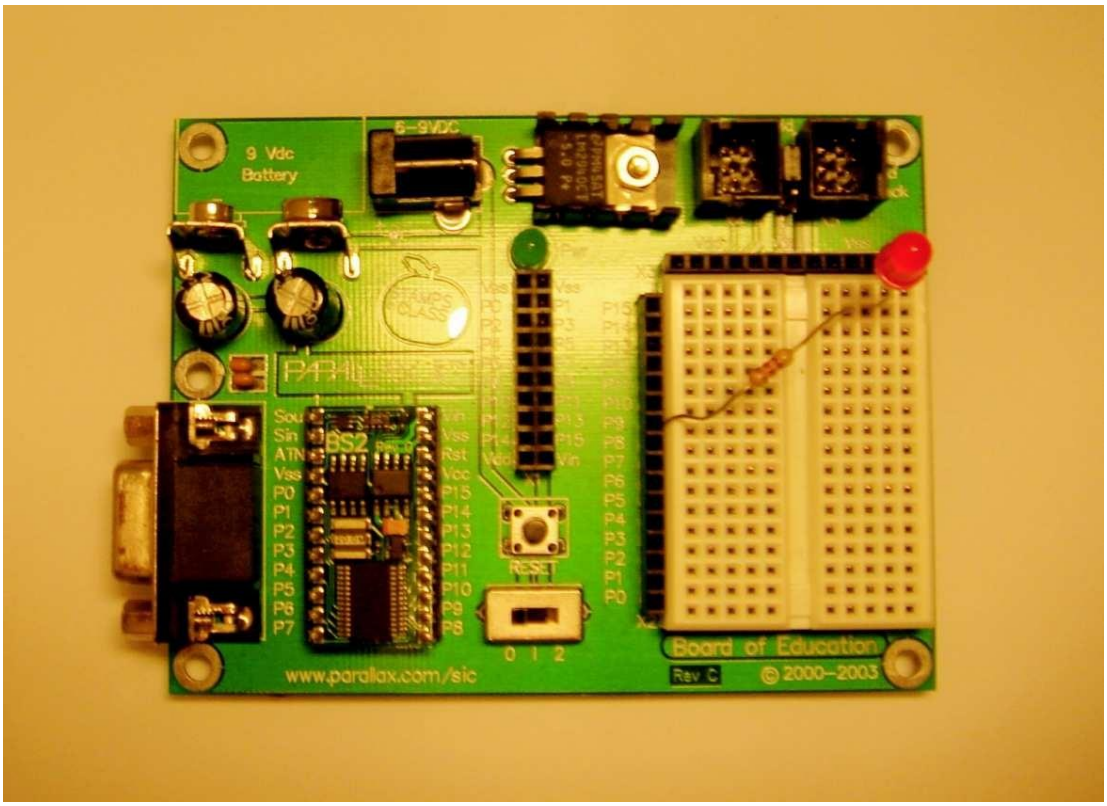

Figure 1.17: *A Board of Education (BOE) microcontroller board, with a Basic Stamp II (BS2) microcontroller attached. In addition to the microcontroller, the BOE has a serial port for communication with a PC (used, for example, when uploading a program onto the BS2), as well as sockets for attaching sensors and electronic circuits. In this case, a simple circuit involving a single LED, has been built on the BOE. The two black sockets in the upper right corner are used for connecting up to four servo motors.*

ever, in addition to sensors and actuators, there must also be a system for an- alyzing the sensory information, making decisions concerning what actions to take, and sending the necessary signals to the actuators.

In autonomous robots, it is common to use several processors to represent the brain of the robot. Typically, high-level tasks, such as decision-making, are carried out on a standard PC, for example a laptop computer mounted on the robot, whereas low-level tasks are carried out by microcontrollers, which will now be introduced briefly.

Microcontrollers

Essentially, a **microcontroller** is a single-chip computer, containing a **central processing unit** (CPU), read-only memory (ROM, for storing programs), random-access memory (RAM, for temporary storage, such as program variables), and

several input-output (I/O) ports. There exist many different microcontrollers, with varying degrees of complexity, and different price levels, down to a few USD for the simplest ones. An example is the **Basic Stamp II**[4] (BS2) microcon- troller, which costs around 50 USD.

While the BS2 is sufficient for the experimental work carried out in this course (in the next quarter), its speed is only around 4,000 operations per sec- ond (op/s) and it has a RAM memory (for program variables) of only 32 bytes and a ROM (for program storage) of 2 kilobytes (Kb).

However, many alternative microcontrollers are available for more advanced robots. Two examples, with roughly the same price as the BS2, are the BasicX and ZBasic microcontrollers, which are both compatible with the BOE micro- controller board used together with the BS2. The BasicX microcontroller has a RAM memory of 400 bytes and 32 Kb for program storage, whereas ZBasic has 4 Kb of RAM and 62 Kb for program storage. BasicX executes around 83,000 op/s, whereas (some versions of) ZBasic can reach up to 2.9 million op/s.

In many cases, microcontrollers are sold together with **microcontroller boards** (or **microcontroller modules**), containing sockets for wires connecting the mi- crocontroller to sensors and actuators as well as control electronics, power sup- ply etc. An example is the **Board of education** (BOE) microcontroller board.

The BOE, shown in Fig. 1.17, is equipped with a **solderless breadboard**, on which electronic circuits can be built without any soldering, which is very use- ful for prototyping.

Since microcontrollers do not have human-friendly interfaces such as a keyboard and a screen, the normal operating procedure is to write and compile programs on an ordinary computer (using, of course, a compiler adapted for the microcontroller in question), and then upload the programs onto the mi- crocontroller. In the case of the BS2 microcontroller, the language is a version of Basic called **PBasic**.

Robotic brain architectures

An autonomous robot must be capable of both high-level and low-level pro- cessing. The low-level processing consists, for example, of sending signals to motor controllers (see below) which, in turn, send (for example) PWM pulses to servo motors. Another low-level task is to retrieve raw data (e.g. a voltage value from an IR proximity sensor). The distinction between low-level and high-level tasks is a bit fuzzy. For example, the voltage value from an IR sen- sor (e.g. the Sharp GP2D12 mentioned above) can be mapped to a distance value, which of course normally is more relevant for decision-making than the raw voltage value. The actual conversion would normally be considered a low-level task but might as well also be carried out on the robot's onboard PC.

[4] Basic Stamp is a registered trademark of Parallax, inc., see www.parallax.com.

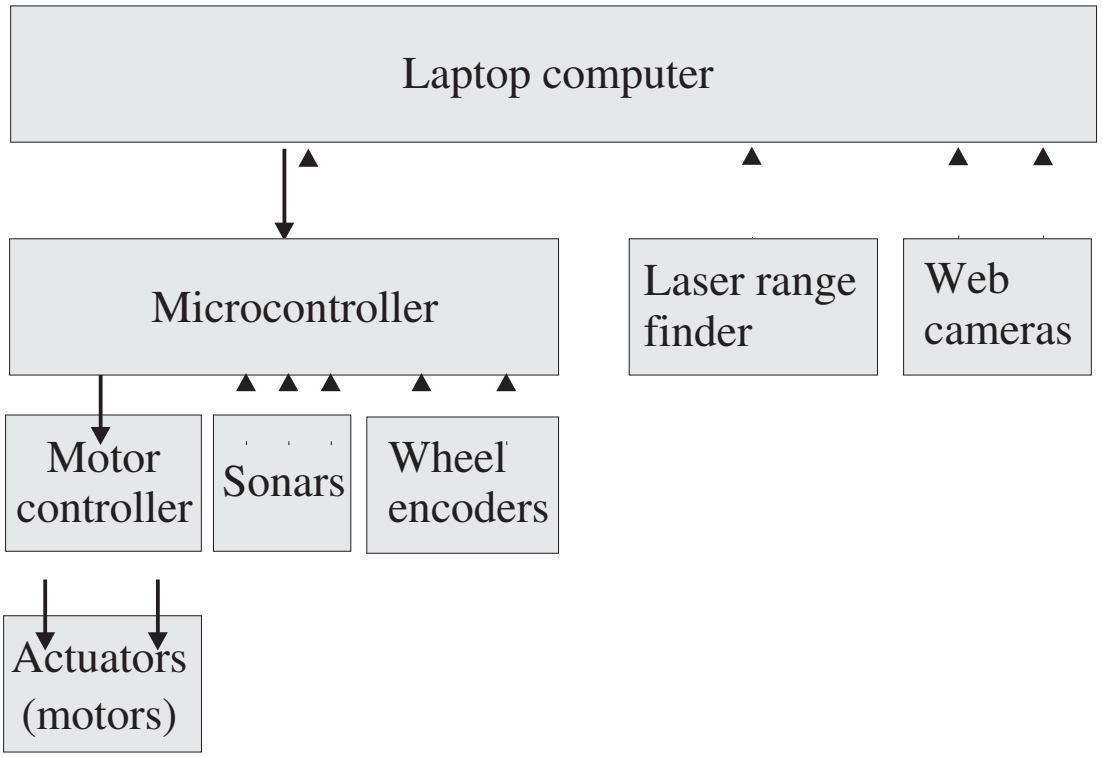

Figure 1.18: *An example of a typical robotic brain architecture, for a differentially steered two-wheeled robot equipped with wheel encoders, three sonar sensors, one LRF, and two web cameras.*

The hardware configuration providing a robot's processing capability is re- ferred to as the **robotic brain architecture**. An example of a typical robotic brain architecture is shown in Fig. 1.18. The robotic brain shown in the figure would be used in connection with a two-wheeled differentially steered robot. As can be seen in the figure, the microcontroller would handle low-level pro- cessing, such as measuring the pulse counts of the wheel encoders, collecting readings from the three sonars, and sending motor signals (e.g. desired set speeds) to the **motor controller**[5], which, in turn, would send signals to the motors. However, the LRF and the web cameras would be directly connected, via USB (or, possibly, serial) ports, to the main processor (on the laptop), since most microcontrollers would not be able to handle the massive data flow from such sensors.

The main program (i.e. the robotic brain), running on the laptop, would process the data from the various sensors. For example, the pulse counts from

[5]A separate motor controller (equipped with its own power source) is often used for robotics applications, since the power source for the microcontroller may not be able to de- liver sufficient current for driving the motors as well.

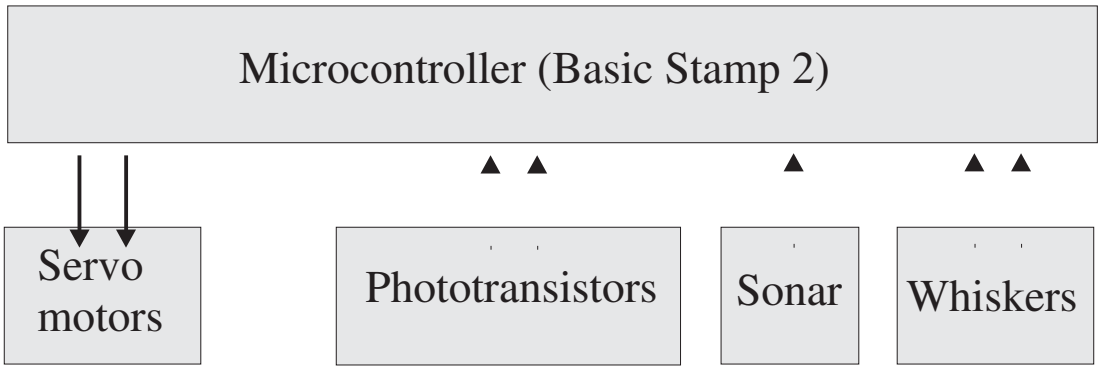

Figure 1.19: *An example of a robotic brain architecture for a Boe-bot.*

the wheel encoders would be translated to an estimate of position and head- ing, as described in Chapter 2. Given the processed sensory data, as well as in- formation stored in the (long-term or short-term) memory of the robotic brain (for example, a map of the arena in which the robot operates), the main pro- gram would determine the next action to be carried out by the robot, compute the appropiate motor commands and send them to the microcontroller.

Note that the figure only shows an *example*: Many other configurations could be used as well. For example, there are cameras developed specifically for robotics applications that, unlike standard web cameras, are able to carry out much of the relevant image processing (e.g. detecting and extracting faces), and then only sending *that* information (rather than the raw pixel values) to the laptop computer.

The robotic brain architecture shown in Fig. 1.18 would be appropriate for a rather complex (and costly!) robot. Such robots are beyond the scope of the experimental work carried out in the second half of this course. The ex- perimental work, which will be carried out using a Boe-bot (see the left panel of Fig. 1.1), involves a much simpler robotic brain architecture, illustrated in Fig. 1.19. As can be seen, in this case, the robot has a single processor, namely the BS2 microcontroller, which thus is responsible both for the low-level (sig- nal) processing and the high-level decision-making.

The microcontroller sends signals to the two servo motors and receives in- put from the sensors attached to the robot, for example, two photo-resistors, a sonar sensor, and whiskers. The whiskers are simple touch sensors that give a reading of either 0 (if no object is touched) or 1 (if the whisker touches an object). Of course, other sensors (such as IR sensors or simple wheel encoders) can be added as well, but one should keep in mind that the processing capabil- ity of the BS2 is very limited. Note that no motor controller is used: The BOE is capable of generating sufficient current for up to four Parallax servo motors.

4. Methodology

3

Simulation of autonomous robots

Simulations play an important role in research on (and development of) auto- nomous robots, for several reasons. First of all, testing a robot in a simulated environment can make it possible to detect whether or not the robot is prone to catastrophic failure in certain situations, so that the behavior of the robot can be altered before it is unleashed in the real world. Second, building a robot is often costly (for example, most laser range finders cost several thousand USD). Thus, through simulations, it is possible to test several designs before constructing an actual robot. Furthermore, it is common to use stochastic opti- mization methods, such as evolutionary algorithms, in connection with the de- velopment of autonomous robots. Such methods require that many different robotic brains be evaluated, which is *very* time-consuming if the work must be carried out in an actual robot. Thus, in such cases, simulations are often used, even though the resulting robotic brains must, of course, be thoroughly tested in real robots, a procedure which often requires several iterations involving simulated and actual robots. In this chapter, an introduction to some of the general issues pertaining to robotic simulations will be given, along with a brief description of (some of) the features of two particular simulators for mo- bile robots, namely GPRSim and ARSim. GPRSim is an advanced 3D simulator for automomous robots, which is used in certain research projects within the Adaptive systems group. ARSim is a simplified (2D) Matlab simulator used in this course.

3.1 Simulators

Over the years, several different simulators for mobile robots have appeared, with varying degrees of complexity. One of the most ambitious simulators to date is **Robotics studio** from Microsoft, which allows the user to simulate many of the commercially available mobile robots, or even to assemble a (vir-

tual) robot using generic parts.

Some simulators include not only general simulation of the kinematic and dynamics of robots, but also procedures for stochastic optimization. Some ex- amples of such simulators are `Webots`, which is manufactured by Cyberbotics (`www.cyberbotics.com`) and the open source package `Darwin2K`, which can be found at `darwin2k.sourceforge.net`.

The Adaptive systems research group at Chalmers has developed a simu- lator called the **General-purpose robotic simulator** (GPRSim), which is exten- sively used in our research projects. Unlike the other simulators mentioned above, GPRSim features, as an integral part of the simulator, an implementa- tion of the general-purpose robotic brain structure (GPRBS) (also developed in the Adaptive systems research group). The GPRBS, in turn, consists of a standardized representation of a robotic brain, consisting of a set of so called brain processes as well as a decision-making system. This structure allows re- searchers to build complex robotic brains involving many different behavioral aspects and also to export the resulting robotic brain for use in real (physical) robots. The existence of a standardized representation for robotic brains also makes it possible, for example, to reuse parts of a previously developed robotic brain in other applications than the original one.

However, GPRSim is primarily a research tool and, as such, it is not very user-friendly. Moreover, the underlying code is quite complex. Thus, in this course, a different simulator will be used, namely the **Autonomous robot sim- ulator** (ARSim), which is a 2D simulator written in Matlab. This simulator is generally too slow to be useful in research projects, but it is perfectly suited to most of the tasks considered in this course. Note also that, even though ARSim is greatly simplified, many parts of the code (for example the simulation of DC motors, IR sensors etc.) are essentially the same in GPRSim and ARSim

3.2 General simulation issues

In Fig. 3.1, the general flow of a single-robot simulation is shown. Basically, after initialization, the simulation proceeds in a stepwise fashion. In each step, the simulator reads the sensors of the robot, and the resulting signals are sent to the robotic brain, which computes appropriate motor signals that, finally, are sent to the motors. Given the motor signals, the acceleration of the robot can be updated, and new velocities and positions can be computed. Changes to the arena (if any) are then made, and the termination criteria are checked.

3.2.1 Timing of events

As mentioned earlier, simulation results in robotics must be validated in an actual robot. However, in order for this to be possible, some care must be

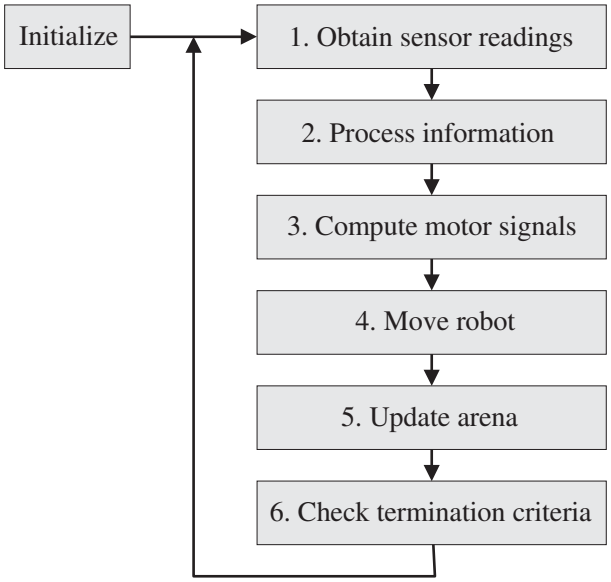

Figure 3.1: *The flow of a single-robot simulation. Steps 1 through 6 are carried out in each time step of the simulation.*

taken, particularly regarding steps 1-3. To be specific, one must make sure that these steps can be executed (on the real robot) in a time which does not exceed the time step length in the simulation. Here, it is important to distinguish between two different types of events, namely (1) those events that take a long time to complete in simulation, but would take a very short time in a real robot, and (2) those events that are carried out rapidly in simulation, but would take a long time to complete in a real robot.

An example of an event of type (1) is collision-checking. If performed in a straight-forward, brute-force way, the possibility of a collision between the (circular, say) body of the robot and an object must be checked by going through *all* lines in a 2D-projection of the arena. A better way (used, for ex- ample, in GPRSim) is to introduce an invisible grid, and only check for colli- sions between the robot and those objects that (partially) cover the grid cells that are also covered by the robot. However, even when such a procedure is used, collision-checking may nevertheless be very time-consuming in simula- tion whereas, in a real robot, it amounts simply to reading a bumper sensor (or, as on the Boe-bot, a whisker), and transferring the signal (which, in this case, is binary, i.e. a single bit of information) from the sensor to the brain of the robot. Events of this type cause no (timing) problems at the stage of transferring the results to a real robot, even though they may slow down the simulation con- siderably.

An example of an event of type (2) is the reading of sensors. For exam- ple, an IR sensor can be modelled using simple ray-tracing (see below) and, provided that the number of rays used is not too large, the update can be car-

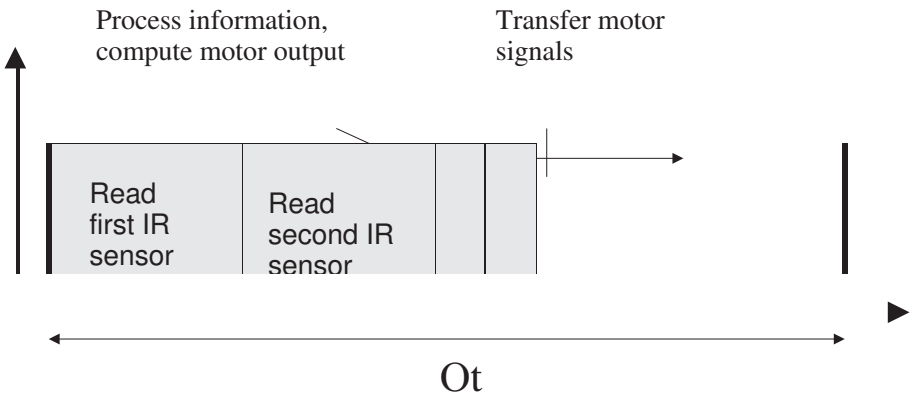

Figure 3.2: *A timing diagram. The boxes indicate the time required to complete the corresponding event in hardware, i.e. a* real *robot. In order for the simulation to be realistic, the time step* Δt *used in the simulation must be longer than the total duration (in hardware) of all events taking place within a time step.*

ried out in a matter of microseconds in a simulator. However, in a real robot it might take longer time. While the reading of an IR sensor involves a very limited signal flow compared to the reading of a camera with, say, 640×480 pixels, the *transfer* of the reading from the sensor to the robotic brain is a po- tential bottleneck. A common setup is to have a microcontroller (see Chapter
1) handling the low-level communication, i.e. obtaining sensor readings and sending signals to actuators, and a PC (for example, a laptop placed on the robot) handling high-level issues, such as decision-making, motion planning etc. Very often, the communication between the laptop and the microcontroller takes place through a serial port, operating with a speed of, say, 9600 or 38400 bits/s. If the onboard PC must read, for example, four proximity sensors (as- suming one byte per reading) and send signals to two motors (again assuming that each signal requires one byte), a total of $6 \times 8 = 48$ bits is needed, lim- iting the number of interactions between the PC and the microcontroller to $9600/48 = 200$ per second in the case of a serial port speed of 9600 bits/s. As another, more specific, example, consider the small mobile robot Khepera, shown in the left panel of Fig. 1.5. In its standard configuration, it is equipped with eight IR sensors, which are read in a sequential way every 2.5 ms, so that the processor of the robot receives an update of a given IR sensor's reading every 20 ms. The updating frequency of the sensors is therefore limited to 50 Hz. Thus, a simulation of a Khepera robot in which the simulated sensors are updated with a frequency of, say, 100 Hz would be unrealistic.

In practice, the problem of limited updating frequency in sensors can be solved by introducing a Boolean **readability state** for each (simulated) sensor. Thus, in the case of a Khepera simulation with a time step of 0.01s, the sensor values would be updated only every other time step. Step 2, i.e. the processing of information by the brain of the robot, must also, in a realistic simulation, be of limited complexity so that the three steps (1, 2, and 3) *together* can be carried

out within the duration Δt (the simulation time step) when transferred to the real robot. An example of a timing diagram for a generic robot (not Khepera) is shown in Fig. 3.2. In the case shown in the figure, two IR proximity sensors are read, the information is processed (for example, by being passed through an artificial neural network), and the motor signals (voltages, in the case of standard DC motors) are then transferred to the motors. The figure shows a case which could be realistically simulated, with the given time step length
Δt. However, if two additional IR sensors were to be added, the simulation would become unrealistic: The real robot would not be able to complete all steps during the time Δt.

For the simple robotic brains considered in this course, step 2 would gener- ally be carried out almost instantaneously (compared to step 1) in a real robot. Similarly, the transfer of motor signals to a DC motor is normally very rapid (note, however, that the dynamics of the *motors* may be such that it is pointless to send commands with a frequency exceeding a certain threshold).

To summarize, a sequence of events that takes, say, several seconds per time step to complete in simulation (e.g. the case of collision-checking in a very complex arena) may be perfectly simple to transfer to a real robot, whereas a sequence of events (such as the reading of a large set of IR sensors) that can be completed almost instantaneously in a simulated robot, may simply not be transferable to a real robot, unless a dedicated processor for signal processing and signal transfer is used.

3.2.2 Noise

Another aspect that should be considered in simulations is *noise*. Real sensors and actuators are invariably noisy, on several levels. Furthermore, even sen- sors that are supposed to be identical often show very different characteristics in practice. In addition, regardless of the noise level of a particular sensor, the frequency with which readings can be updated is limited, thus introducing another source of noise, in certain cases. For example, the limited sampling frequency of wheel encoders implies that, even in the (unrealistic) case where the kinematic model is perfect and there are no other sources of noise, the in- tegrals in the kinematic equations (Eqs. (2.7)-(2.9)) can only be approximately computed.

Thus, in any realistic robot simulation, noise must be added, at all relevant levels. Noise can be added in several different ways. A common method (used in GPRSim and ARSim) is to take the original reading S of a sensor and add noise to form the actual reading $\hat{S}$ as

$$\hat{S} = SN(1, \sigma), \tag{3.1}$$

where $N(1, \sigma)$ denotes the normal (Gaussian) distribution with mean 1 and

standard deviation σ. Of course, other distributions (e.g. a uniform distribu- tion) can be used as well.

An alternative method is to take some measurements of a *real* sensor and store the readings in a lookup table, which is then used by the simulated robot. For example, in the case of an IR sensor with a range of, say, 0.5 m, one may, for example, take 10 readings each at distances of $0.05, 0.10, \ldots, 0.50$ m, and store those readings in a matrix. In the simulator, when the IR sensor is used, the distance L to the nearest obstacle is determined, and the reading is then obtained by interpolating linearly between two samples from the lookup table. For example, if $L = 0.23$ m, a randomly chosen sample $\hat{s}_{20}$ is taken from the 10 readings stored for $L = 0.20$ m, and another randomly chosen sample $\hat{s}_{25}$ is taken from the readings stored for $L = 0.25$ m. The reading of the simulated sensor is then taken as

$$\hat{S} = \hat{s}_{20} + \frac{0.23 - 0.20}{0.25 - 0.20}(\hat{s}_{25} - \hat{s}_{20}) \tag{3.2}$$

This method has the advantage of forming simulated readings from actual sensor readings, rather than introducing a model for the noise. Furthermore, using lookup tables, it is straightforward to account for the individual nature of supposedly identical sensors. However, a clear disadvantage is the need for generating the lookup tables, which often must contain a very large number of samples taken not only at various distances, but also, perhaps, at various *angles* between the forward direction of the sensor and the surface of the obstacle. Thus, the first method, using a specific noise distribution, is normally used instead.

3.2.3 Sensing

In addition to correct timing of events and the addition of noise in sensors and actuators, it is necessary to make sure that the sensory signals received by the simulated robot do not contain more information than could be provided by the sensors of the corresponding real robot. For example, in the simulation of a robot equipped only with wheel encoders (for odometry), it is not allowed to provide the simulated robot with continuously updated and error-free po- sition measurements. Instead, the simulated wheel encoders, including noise and other inaccuracies, should be the only source of information regarding the position of the simulated robot.

In both GPRSim and ARSim, several different sensors have been imple- mented, namely (1) wheel encoders, (2) IR proximity sensors, and (3) com- passes. In addition, GPRSim (but not ARSim) also features (4) sonar sensors and (5) laser range finders (LRFs). An important subclass of (simulated) sen- sors are **ray-based sensors**, which use a simple form of ray tracing in order to

form their reading(s). Examples of ray-based sensors are IR proximity sensors, sonar sensors, and laser range finders.

Now, the different natures of, say, an IR sensor, which gives a fuzzy read- ing based on infrared light, and an LRF, which gives very accurate readings (in many directions) based on laser light, imply that slightly different procedures must be used when forming the (simulated) sensor readings of those two sen- sor types. However, in both cases, the simulation of the sensor requires ray tracing, which will now be considered.

Ray-based sensors In ray-based sensors, the formation of sensor readings is based on the concept of **sensor rays**. Basically, a number of rays are sent out from a sensor, in various directions (depending on the opening angle of the sensor), and the distance to the nearest obstacle is determined. If no ob- stacle is available within the range of the sensor, the ray in question provides no reading. Of course, in order to obtain any ray reading, not only the robot must be available, but also the objects (e.g. walls and furniture) located in the arena in which the robot is operating. In GPRSim, objects are built from boxes and cylinders. Boxes are represented as a sequence of six planes, whereas (the mantle surface of) cylinders are represented by a sufficient number of planes (usually around 10-20) to approximate the circular cross section of the cylinder. The ray readings are thus obtained using general equations for line- plane intersections[1]. Here, however, we shall only consider the simpler two-dimensional case, in which all surfaces are vertical and where the sensors are oriented such that all emitted rays are parallel to the ground. In such cases, the arena objects can be represented simply as a sequence of lines in two dimen- sions. Indeed, this is how objects are represented in ARSim.

An example of such a configuration is shown in Fig. 3.3. The left panel shows a screenshot from GPRSim, in which an LRF mounted on top of a robot takes a reading in an arena containing only walls. The right panel shows a two- dimensional representation of the arena and the LRF (the body of the robot is not shown). Given the exact position of a ray's starting point, as well as the range of the corresponding sensor, it is possible to determine the distance be- tween the ray and the nearest obstacle using general equations for line-line intersection, which will be described next. However, it should first be noted that, even though the *simulator* of course uses the exact position of the robot and its sensors in order to compute sensor readings, the *robot* (or, more specif- ically, its brain) is *only* provided with information regarding the actual sensor readings.

Consider now a single sensor ray. Given the start and end points of the

[1]In order to speed up the simulator, a grid (also used in collision checking) is used, such that only those obstacles that are (partially) located in the grid cells currently covered by the sensor are considered.

Figure 3.3: *Left panel: A screenshot from GPRSim, showing an LRF taking a reading in an arena containing only walls. Right panel: A two-dimensional representation of the sensor reading. The dotted ray points in the forward direction of the robot which, in this case, coincides with the forward direction of the LRF.*

ray, its equation can be determined. Let (x_a, y_a) denote the start point for the ray (which will be equal to the position of the sensor, if the size of the latter can be neglected). Once the absolute direction (β_i) of the sensor ray has been determined, the end point (x_b, y_b) of an unobstructed ray (i.e. one that does not hit any obstacle) can be obtained as

$$(x_b, y_b) = (x_a + D\cos\beta_i, y_a + D\sin\beta_i), \tag{3.3}$$

where D denotes the sensor range. Similarly, any line corresponding to the side of an arena object can be defined using the coordinates of its start and end points. Note that, in Fig. 3.3, all lines defining arena objects coincide with coordinate axes, but this is, of course, not always the case. Now, in the case of two lines of infinite length, defined by the equations $y_k = c_k + d_k x, \ k = 1, 2,$ it is trivial to find the intersection point (if any) simply by setting $y_1 = y_2$. However, here we are dealing with line segments of finite length. In this case, the intersection point can be determined as follows: Letting $\mathbf{P}_i^a = (x_i^a, y_i^a)$ and $\mathbf{P}_i^b = (x_i^b, y_i^b)$, denote the start and end points, respectively, of line i, $i = 1, 2,$

the equations for an arbitrary point $\mathbf{P}_i$ along the two line segments can be written

$$\mathbf{P}_1 = \mathbf{P}_1^a + t_1(\mathbf{P}_1^b - \mathbf{P}_1^a), \tag{3.4}$$

and

$$\mathbf{P}_2 = \mathbf{P}_2^a + u(\mathbf{P}_2^b - \mathbf{P}_2^a), \tag{3.5}$$

where $(t, u) \in [0, 1]$. Solving the equation $\mathbf{P}_1 = \mathbf{P}_2$ for t and u gives, after some

algebra,

$$t = \frac{(x_2^b - x_2^a)(y_1^a - y_2^a) - (y_2^b - y_2^a)(x_1^a - x_2^a)}{(y^b - y^a)(x^b - x^a) - (x^b - x^a)(y^b - y^a)} \tag{3.6}$$

and

$$u = \frac{(x_1^b - x_1^a)(y_1^a - y_2^a) - (y_1^b - y_1^a)(x_1^a - x_2^a)}{(y_2^b - y_2^a)(x_1^b - x_1^a) - (x_2^b - x_2^a)(y_1^b - y_1^a)} \tag{3.7}$$

An intersection occurs if *both* t and u are in the range $[0, 1]$. Assuming that the first line (with points given by $\mathbf{P}_1$) is the sensor ray, the distance d between the sensor and the obstacle, along the ray in question, can then easily be formed by simply determining $\mathbf{P}_1$ using the t value found, and computing

$$d = |\mathbf{P}_1 - \mathbf{P}_1^a| = |t(\mathbf{P}_1^b - \mathbf{P}_1^a)|. \tag{3.8}$$

If the two lines happen to be parallel, the denominator becomes equal to zero[2]. Thus, this case must be handled separately.

In simulations, for any time step during which the readings of a particular sensor are to be obtained, the first step is to determine the origin of the sensor rays (i.e. the position of the sensor), as well as their directions. An example is shown in Fig. 3.4. Here, a sensor is placed at a point $\mathbf{p}_s$, relative to the center of the robot. The absolute position $\mathbf{P}_s$ (relative to an external, fixed coordinate system) is given by

$$\mathbf{P}_s = \mathbf{X} + \mathbf{p}_s, \tag{3.9}$$

where $\mathbf{X} = (X, Y)$ is the position of the (center of mass of the) robot. Assuming that the front direction of the sensor is at an angle α relative to the direction of heading (ϕ) of the robot, and that the sensor readings are to be formed using N equally spaced rays over an **opening angle** of γ, the absolute direction β_i of the i^{th} ray equals

$$\beta_i = \phi + \alpha - \frac{\gamma}{2} + (i-1)\delta\gamma, \tag{3.10}$$

where $\delta\gamma$ is given by

$$\delta\gamma = \frac{\gamma}{N-1}. \tag{3.11}$$

Now, the use of the ray readings differs between different simulated sensors. Let us first consider a simulated IR sensor. Here, the set of sensor rays is used *only as an artificial construct* needed when forming the rather fuzzy reading of such a sensor. In this case, the rays themselves are merely a convenient com- putational tool. Thus, for IR sensors, the robotic brain is not given information regarding the individual rays. Instead, only the complete reading S is pro- vided, and it is given by

$$S = \frac{1}{N} \sum_{i=1}^{N} \rho_i, \tag{3.12}$$

In case the two lines are not only parallel but also *coincident*, both the numerators and the denominators are equal to zero in the equations for t and u.

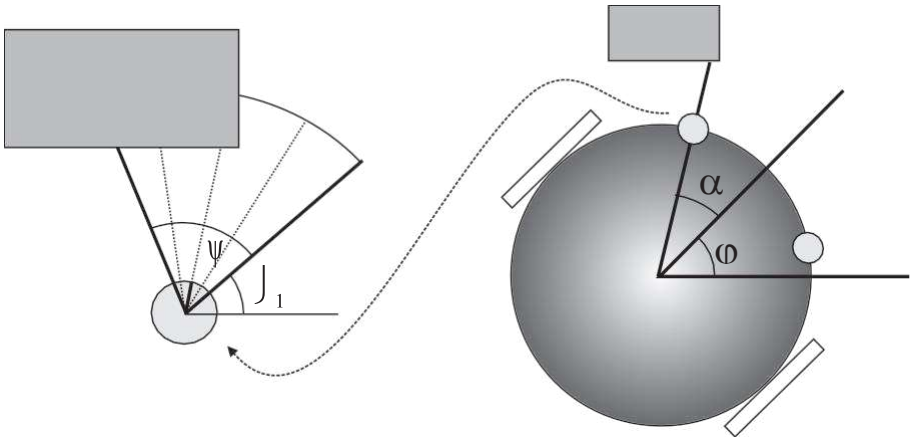

Figure 3.4: *The right panel shows a robot equipped with two IR sensors, and the left panel shows a blow-up of the left sensor. In this case, the number of rays (N) was equal to 5. The leftmost and rightmost rays, which also indicate the opening angle γ of the IR sensor are shown as solid lines, whereas the three intermediate rays are shown as dotted lines.*

where ρ_i is the **ray reading** of ray i. Ideally, the value of N should be very large for the simulated sensor to represent accurately a real IR sensor. However, in practice, rather small values of N (3-5, say) is used in simulation, so that the reading can be obtained quickly. The loss of accuracy is rarely important, since the (fuzzy) reading of an IR sensor is normally used only for proximity detec- tion (rather than, say, mapping or localization). An illustration of a simulated IR sensor is given in Fig. 3.4.

A common phenomenological model for IR sensor readings (used in GPRSim and ARSim) defines ρ_i as

$$\rho_i = \min\left(\frac{c_1}{d_i^2} + c_2 \cos\kappa_i, 1\right), \tag{3.13}$$

where c_1 and c_2 are non-negative constants, $d_i > 0$ is the distance to the nearest object along ray i, and

$$\kappa_i = -\frac{\gamma}{2} + (i-1)\delta\gamma, \tag{3.14}$$

is the **relative ray angle** of ray i. If $d_i > D$ (the range of the sensor), $\rho_i = 0$. Note that it is assumed that $\kappa_i \in [-\pi/2, \pi/2]$, i.e. the opening angle cannot exceed π radians. Typical opening angles are $\pi/2$ or less. It should also be noted that this IR sensor model has limitations; for example, the model does not take into account the orientation of the obstacle's surface (relative to the direction of the sensor rays) and neither does it account for the different IR reflectivity of different materials.

For simulated *sonar* sensors (which are included in GPRSim but not in AR- Sim), the rays are also only used as a convenient computational tool, but the

final reading S is formed in a different way. Typically, sonar sensors give rather accurate distance measurements in the range $[D_{min}, D_{max}]$, but sometimes fail to give a reading at all. Thus, in GPRSim, the reading of a sonar sensor is formed as $S = \min_i d_i$ with probability p and D_{max} (no detection) with proba- bility $1 - p$. Also, if $S < D_{min}$ the reading is set to D_{min}. Typically, the value of p is very close to 1. The number of rays (N) is usually around 3 for simulated sonars.

A simulated LRF, by contrast, gives a vector-valued reading, $\mathbf{S}$, where each component S_i is obtained simply as the distance d_i to the nearest obstacle along the ray. Thus, for LRFs, the sensor rays have a specific physical interpretation, made possible by the fact that the laser beam emitted by an LRF is very narrow. In GPRSim, if $d_i > D$, the corresponding laser ray reading is set to -1, to indi- cate the absence of any obstacle within range of the ray in question. Note that LRFs are only implemented in GPRSim. It would not be difficult to add such a sensor to ARSim, but since an LRF typically takes readings in 1,000 different directions (thus requiring the same number of rays), such sensors would make ARSim run very slowly.

As a final remark regarding ray-based sensors, it should be noted that a given sensor ray i may intersect several arena object lines (see, for example, Fig. 3.3) In such cases, d_i is taken as the *shortest* distance obtained for the ray.

3.2.4 Actuators

A commonly used actuator in mobile robots is the DC motor. The equations describing such motors are given in Chapter 1.

In both GPRSim and ARSim, a standard DC motor has been implemented. In this motor, the input signal is the applied voltage. Both the electrical and mechanical dynamics of the motors are neglected. Thus the torque acting on the motor shaft axis is given by Eqs. (1.8) and (1.9). Gears are imple- mented in both simulators, so that the torques acting on the wheels are given by Eqs. (1.10). However, the simulators also include the possibility of setting a maximum torque τ_{max} which cannot be exceeded, regardless of the output torque τ_{out} obtained from Eqs. (1.10).

In addition, GPRSim (but not ARSim) also allows simulation of **velocity-regulated motors**. Unlike the voltage signal used in the standard DC motor, a velocity-regulated motor takes as input a desired **reference speed** v_{ref} for the wheel attached to the motor axis. The robot then tries to reach this speed value, using proportional control. The actual output torque of a velocity-regulated motor is given by

$$\tau = K(v_{ref} - v) \tag{3.15}$$

In this model, a change in v_{ref} generates an immediate change in the torque. In a real motor, the torques cannot change instantaneously. However, Eq. (3.15)

Figure 3.5: *Left panel: A simulated robot (from GPRSim), consisting of more than 100 objects. Right panel: An example (in blue) of a collision geometry.*

usually provides a sufficiently accurate estimate of the torque. As in the case of the standard DC motor, there is also a maximum torque τ_{max} for velocity- regulated motors.

Note that, if velocity-regulated motors are to be used, the robot must be equipped with wheel encoders to allow the computation of odometric esti- mates of the wheel speeds.

3.2.5 Collision checking

A real robot should normally be very careful not to collide with an obstacle (or, worse, a person). In simulations, however, one may allow collisions, for exam- ple during simulations involving stochastic optimization, where the robotic brains in the early stages of an optimization run may be unable to avoid colli- sions. In any case, collisions should, of course, be detected.

In GPRSim the concept of a **collision geometry** is used when checking for collisions. The collision geometry is a set of vertical planes in which the body of the robot should be contained. It would be possible to check collisions be- tween the boxes and cylinders constituting the (simulated) body of the robot. However, it is common that the robotic body consists of a very large number of objects, making collision-checking very slow indeed. Thus, instead, a sim- pler collision geometry is used. An example is given in Fig. 3.5. The left panel shows a simulated robot (consisting of more than 100 separate objects), and the right panel shows (in blue) a collision geometry for the same robot.

By contrast, in ARSim the simulated robot is always represented as a circu- lar disc. Thus, the collision detection method simply checks for intersections

between the circular body of the robot and any line representing a side of an arena object.

3.2.6 Motion

Once the torques acting on the wheels have been generated, the motion of the robot is obtained through numerical integration of Eqs. (2.31) and (2.32). In both GPRSim and ARSim, the integration is carried out using simple first-order (Euler) integration. For each time step, $\dot{V}$ and $\ddot{\phi}$ are computed using Eqs. (2.31) and (2.32), respectively. The new values V^j and $\dot{\phi}^j$ of V and $\dot{\phi}$ are then computed as

$$V^j = V + \dot{V}\,\Delta t, \tag{3.16}$$

$$\dot{\phi}^j = \dot{\phi} + \ddot{\phi}\Delta t, \tag{3.17}$$

where Δt is the time step length (typically set to 0.01 s). The value of ϕ is then updated, using the equation

$$\phi^j = \phi + \dot{\phi}^j\Delta t, \tag{3.18}$$

The cartesian components of the velocity are then obtained as

$$V_x^j = V^j \cos\phi, \tag{3.19}$$

$$V_y^j = V^j \sin\phi. \tag{3.20}$$

Finally, given V_x^j and V_y^j, the new positions X^j and Y^j can be computed as

$$X^j = X + V_x^j\Delta t, \tag{3.21}$$

$$Y^j = Y + V_y^j\Delta t. \tag{3.22}$$

In addition, if wheel encoders are used, both GPRSim and ARSim also keep track of the rotation of each wheel, for possible use in odometry (if available).

3.2.7 Robotic brain

While the physical components of a robot, such as its sensors and motors, of- ten remain unchanged between simulations, the robotic brain must, of course, be adapted to the task at hand. Robotic brains can be implemented in many different ways.

In **behavior-based robotics** (BBR) the brain of a robot is built from a reper- toire (i.e. a set) of basic behaviors, as well as a decision-making procedure, selecting which behavior(s) to activate at any given time. In the **General- purpose robotic brain structure (GPRBS)**, developed in the author's research group, the robotic brain is built from a set of **brain processes**, some of which

are **motor behaviors** (that make use of the robot's motors) and some of which are **cognitive processes**, i.e. processes that do not make use of any motors. In addition, GPRBS features a decision-making system based on the concept of utility. One of the main properties of GPRBS is that this structure allows sev- eral processes to run in parallel, making it possible to build complex robotic brains. In fact, the specific aim of the development of GPRBS is to move be- yond the often very simple robotic brains defined within standard BBR.

In GPRBS, all brain processes are specified in a standardized format, which simplifies the development of new brain processes, since many parts of an already existent process often can be used when writing a new process. How- ever, at the same time, GPRBS (as implemented in GPRSim) is a bit complex to use, especially since it is intended for use in research, rather than as an edu- cational tool. Thus, in this course, ARSim will be used instead. This simulator allows the user to write simple brain processes (as well as a basic decision- making system) in any desired format (i.e. without using GPRBS). Methods for writing brain processes will be described further in a later chapter.

3.3 Brief introduction to ARSim

The simplest way to acquaint oneself with `ARSim` is to run and analyze the test program distributed with the program. In order to do so, start Matlab, move to the right directory and write

```
>> TestRunRobot
```

and press return. The robot appears in a quadratic arena with four walls and two obstacles, as shown in Fig. 3.6. The robot is shown as a circle, and its di- rection of motion is indicated by a thin line. The IR sensors (of which there are two in the default simulation) are shown as smaller circles. The rays used for determining the sensor readings (three per sensor, per default) are shown as lines emanating from the sensors. In the default simulation, the robot executes 1,000 time steps of length 0.01 s, unless it is interrupted by a collision with an obstacle or a wall.

The flow of the simulation basically follows the structure given in Fig. 3.1. The first lines of code in the `TestRunRobot.m` file are devoted to adding the various `ARSim` function libraries to Matlab's search path. The arena objects are then created and added to the arena. Next, the brain of the robot is created (by a call to `CreateBrain`), and the setup is completed by creating the sensors and motors, and adding them to the robot.

Before the actual simulation starts, the robot's position, heading, velocity, and angular speed are set, and the plot of the arena (including the robot) is created. Optionally, a variable `motionResults`, storing information about

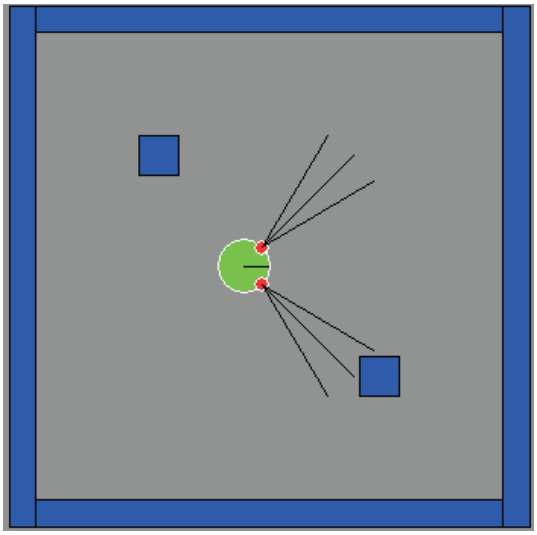

Figure 3.6: *A typical screenshot from an ARSim simulation. The black lines emanating from the two IR proximity sensors of the robot are the rays used for determining sensor readings.*

the robot's motion, can be created.

`ARSim` then executes the actual simulation. Each time step begins with the sensors being read. First, the readings of all ray-based sensors (a category in which only IR sensors have been implemented in ARSim, so far) are obtained. Next, the odometer and compass readings are obtained (provided, of course, that the robot is equipped with those sensors. Next, the robotic brain processes the sensory information (by executing the `BrainStep` function), producing motor signals, which are used by the `MoveRobot` function. Finally, a collision check is carried out.

In normal usage, only a few of `ARSim`'s functions need be modified, namely `CreateBrain`, in which the parameters of the brain are set, `BrainStep`, which determines the processing carried out by the robotic brain, and, of course, the main file (i.e. `TestRunRobot` in the default simulation), where the setup of the arena and the robot are carried out. Normally, no other Matlab func- tions should be modified unless, for example, one wants to modify the plot procedure.

Note that, by default, the rays involved in the computation of the IR sensor readings are not plotted. In order to plot the sensor rays, one must set the pa- rameter `ShowSensorRays` to `true`. If the robot is equipped with an odome- ter, one can plot also the position and heading estimated by the odometer, by setting the parameter `ShowOdometricGhost` to true. A brief description of the Matlab functions contained in `ARSim` is given in Appendix A.

6 Exploration, navigation, and localization

In the previous chapter, the concept of robotic behaviors was introduced and exemplified by means of some basic motor behaviors. Albeit very simple, such behaviors can be tailored to solve a variety of tasks such as, for example, wan- dering, wall following and various forms of obstacle avoidance. However, there are also clear limitations. In this chapter, some more advanced motor behaviors will be studied. First, behaviors for exploration and navigation will be considered. Both of these two types of behavior require accurate pose esti- mates for the robot. It is assumed that the robot is equipped with a (cognitive) *Odometry* brain process, providing continuous pose (and velocity) estimates. As mentioned earlier, such estimates are subject to odometric drift, and there- fore an independent method for localization (i.e. odometric recalibration) is always required in realistic applications. Such a method will be studied in the final section of this chapter. However, exploration and navigation are im- portant problems in their own right and, in order to first concentrate on *those* problems, it will thus (unrealistically) be assumed, in the first two sections of the chapter, that the robot obtains perfect, noise-free pose estimates using odometry only.

6.1 Exploration

Purposeful navigation requires some form of **map** of the robot's environment. In many cases, however, no map is available *a priori*. Instead, it is the robot's task to *acquire* the map, in a process known as **simultaneous localization and mapping** (SLAM). In (autonomous) SLAM, a robot is released in an unknown arena, and it is then supposed to move in such a way that, during its motion, its long-range sensors (typically an LRF) covers every part of the arena, so that

the sensor readings can be used for generating a map. This is a rather diffi- cult task since, during exploration and mapping, the robot must keep track of its position using, for odometric recalibration, the (incomplete, but grow- ing) map that it is currently generating. SLAM is an active research topic, for which many different methods have been suggested. A currently popular ap- proach is **probabilistic robotics**, in which the robot maintains a probability density function from which its position is inferred. However, SLAM is be- yond the scope of this text. Instead, the simpler, but still challenging, topic of exploration *given* perfect positioning (as mentioned in the introduction to this chapter) will be considered.

Exploration can be carried out for different reasons. In some applications, such as lawn mowing, vacuum cleaning, clearing mine fields etc., the robot must physically cover as much as possible of the floor or ground in its envi- ronment. Thus, the robot must carry out **area coverage**. In some applications,
e.g. vacuum cleaning, it is often sufficient that the robot carries out a more or less aimless wandering that, eventually, will make it cover the entire floor. In other applications, such as mapping, it is unnecessary for the robot to physi- cally visit every spot in the arena. Instead, what matters is that its long-range sensor, typically an LRF (or a camera), is able to sense every place in the arena at some point during the robot's motion. The problem of exploring an arena such that the long-range sensor(s) reach all points in the arena will here be referred to as **sensory area coverage**.

Exploring an arena, without any prior knowledge regarding its structure, is far from trivial. However, a motor behavior (in the GPRBS framework) for sensory area coverage has recently (2009) been implemented[1]. This *Exploration* behavior has been used both in the simulator GPRSim and in a real robot (as a part of SLAM). In both cases, the robot is assumed to be equipped with an LRF. The algorithm operates as follows: A **node** is placed at the current (esti- mated) position of the robot. Next, the robot generates a set of nodes at a given distance (D) from its current position (estimated using the *Odometry* process). Before any nodes are placed, the robot used the LRF (with an opening (sweep) angle α) to find feasible angular intervals for node placement, i.e. angular in- tervals in which the distance to the nearest obstacle exceeds $D + \Delta$, where Δ is a parameter measuring the margin between a node and the nearest obsta- cle behind the node. The exact details of the node placement procedure will not be given here. Suffice it to say that, in order to be feasible, an angular in- terval must have a width γ exceeding a lower limit $\gamma_{\min}$ in order for a node to be placed (at the center of the angular interval). Furthermore, if the width of a feasible angular interval is sufficiently large, more than one node may be placed in the interval. An illustration of feasible angular intervals and node

[1]See Wahde, M. and Sandberg, D. *An algorithm for sensory area coverage by mobile robots operating in complex arenas*, Proc. of AMiRE 2009, pp. 179-186, 2009.

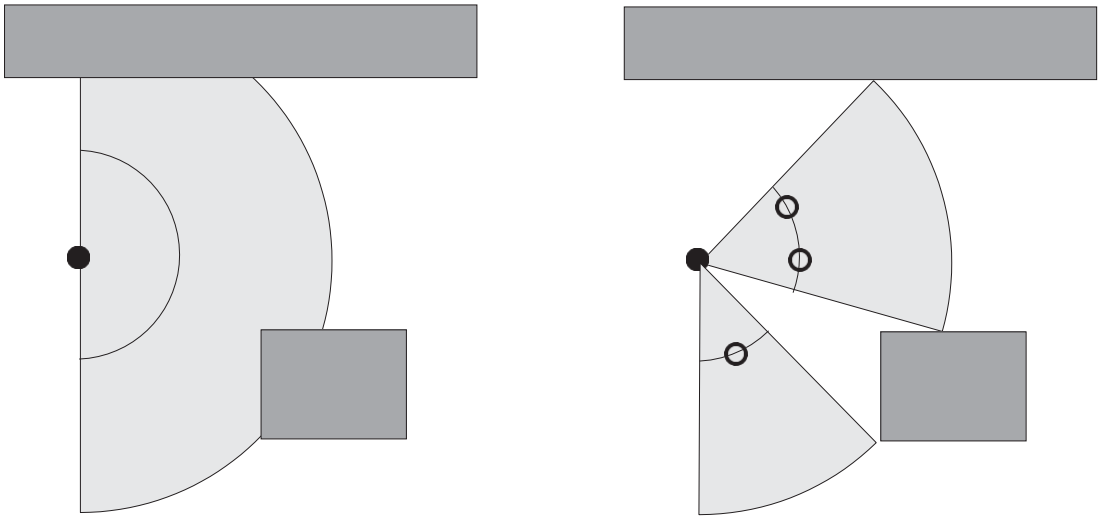

Figure 6.1: *An illustration of the node placement method in the* Exploration *behavior. The left panel shows the distances obtained over the 180 degree opening angle of the LRF (note that individual rays are not shown). The inner semi-circle has a radius of* D *(the node placement distance) whereas the radius of the outer semi-circle is* $D + \Delta$*. The right panel shows the re- sulting distribution of nodes. Note that one of the two feasible angular intervals is sufficiently wide to allow* two *nodes to be placed.*

placement is given in Fig. 6.1.

At this point, the reader may ask why nodes are placed at a distance D from the current node, rather than as far away as possible (minus the margin Δ). The reason is that, in practical use, one cannot (as is done here) assume that the odometry provides perfect pose estimates. Since the *Exploration* behavior is normally used in connection with SLAM, for which accurate positioning is crucial when building the map (a process involving alignment of consecutive laser scans), one cannot move a very large distance between consecutive laser snapshots. Thus, even though the typical range R of an LRF is around 4-10 m or more, the distance D is typically only around 1 m.

An additional constraint on node placement regards the separation (con- cerning *distances*, not angles) between nodes. A minimum distance of d (typ- ically set to 0.75 m or so) is enforced. The requirement that nodes should be separated by a distance of at least d makes the algorithm finite: At some point, it will no longer be possible to place new nodes without violating this con- straint. Thus, when all nodes have been processed (i.e. either having been vis- ited or deemed unreachable, see below), and no further nodes can be added, the exploration of the arena is complete.

Returning to the algorithm, note that the initial node, from which the robot starts its exploration, is given the status *completed* (implying that this node has been reached) and is referred to as the *active* node. All newly generated nodes are given the status *pending*. The robot also generates paths to the pending

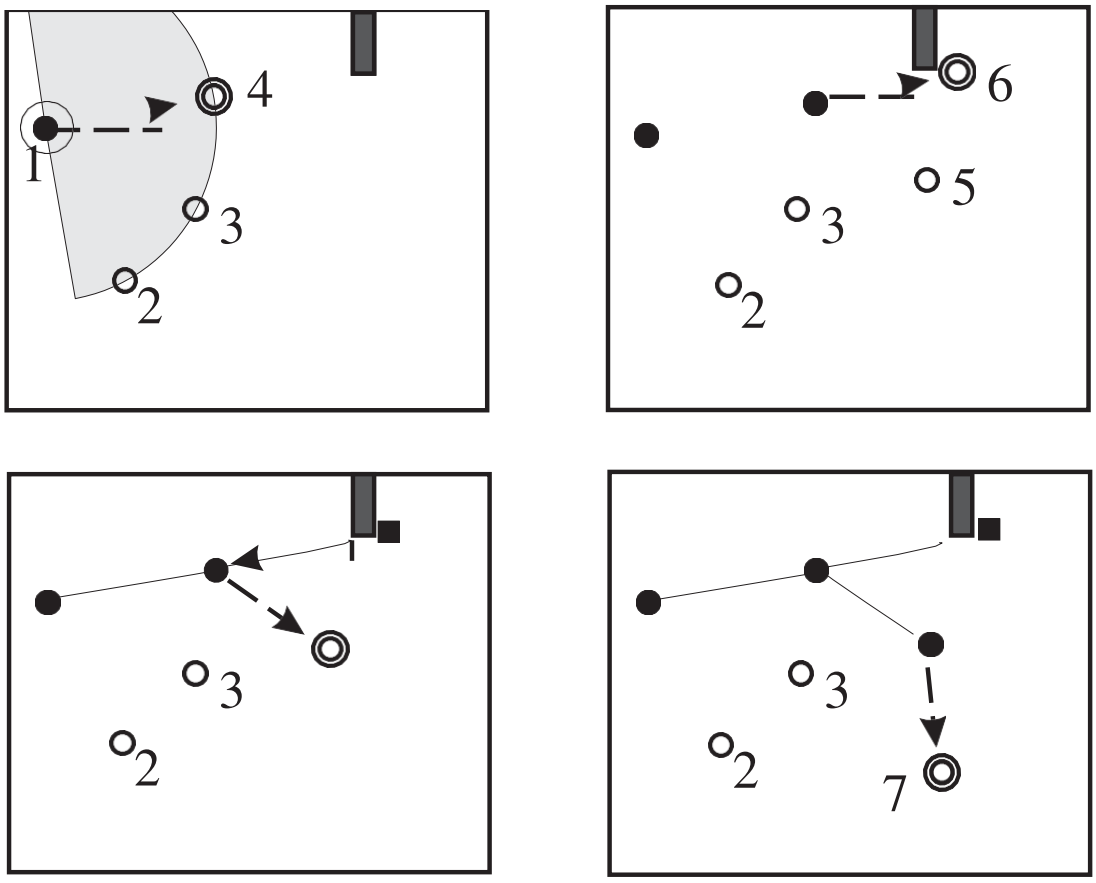

Figure 6.2: *The early stages of a run using the exploration algorithm, showing a robot exploring a single rectangular room without a door. The arena contains a single, low obstacle, which cannot be detected using the LRF (since it is mounted above the highest point of the obstacle). In each panel, the target node is shown with two thick circles, pending nodes are shown as a single thick circle, completed nodes as a filled disc, and unreachable nodes as a filled square. Upper left panel: The robot, whose size is indicated by a thin open circle, starts at node 1, generating three new pending nodes (2, 3, and 4). Upper right panel: Having reached node 4, the robot sets the status of that node to* completed*, and then generates new pending nodes. Lower left panel: Here, the robot has concluded (based on IR proximity readings) that node 6 is unreachable, and it therefore selects the nearest pending node (5, in this case) based on path distance, as the new target node. Lower right panel: Having reached node 5, due to the minimum distance requirement (between nodes) the robot can only generate one new node (7). It rotates to face that node, and then moves towards it etc.*

nodes. For example, if the robot is located at node 1 and generates three pend- ing nodes (2,3 and 4), the paths will be (1, 2), (1, 3) and (1, 4). The robot next selects the nearest node, based on the *path length* as the **target node**. In many cases (e.g. when more than one node can be placed), several nodes are at the same (estimated) distance from the robot. In such cases, the robot (arbitrarily)

selects one of those nodes as the target node. For the paths just described, the path length equals the cartesian distance between the nodes. If a path con- tains more than two elements, however, the path length will differ from the cartesian distance, unless all the nodes in the path lie along a straight line. The path length is more relevant since, when executing the exploration algorithm described here, the robot will generally follow the path, even though direct movement between the active node and a target node is also possible under certain circumstances; see below.

Next, the robot rotates to face the target node, and then proceeds towards it; see the upper left panel of Fig. 6.2. During the motion, one of two things can happen: Either (i) the robot reaches the target node or, (ii) using the output from a *Proximity detection* brain process (assumed available), it concludes that the target node cannot be reached along the current path. Note that, in order for the *Proximity detection* brain process to be useful, the sensors it uses should be mounted at a different (smaller) height compared to the LRF.

In case (i), illustrated in the upper right panel of Fig. 6.2, once the target node has been reached, it is given the status *completed* and is then set as the new active node. At this point, the paths to the remaining pending nodes are updated. Continuing with the example above, if the robot moves to node 4, the paths to the other pending nodes (2 and 3) will be updated to (4, 1, 2) and (4, 1, 3). Furthermore, having reached node 4, the robot generates new pend- ing nodes. Note that the robot need not be located exactly *on* node 4; instead, a node is considered to be reached when the robot passes within a distance a from it. The new nodes are added as described above. The minimum distance requirement between added nodes (see above) is also enforced. Proceeding with the example, the robot might, at this stage, add nodes 5 and 6, with the paths (4, 5) and (4 ,6). Again, the robot selects the nearest node based on the path length, rotates to face that node, and then starts moving towards it etc. Note that the robot can (and will) visit completed nodes more than once. How- ever, by the construction described above, only pending nodes can be target nodes.

In case (ii), i.e. when the target node cannot be reached, the node is assigned the status *unreachable*, and the robot instead selects another target node and proceeds towards it, along the appropriate path. This situation is illustrated in the lower left panel of Fig. 6.2: Here, using its *Proximity detection* brain process, the robot concludes that it cannot reach node 6. It therefore marks this node *unreachable*, sets it as the active node, and then sets the nearest pending node as the new target, in this case node 5. One may wonder why case (ii) can oc- cur, since the robot uses the LRF before assigning new nodes. The reason, of course, is that the LRF (which is assumed to be two-dimensional) only scans the arena at a given height, thus effectively only considering a horizontal slice of the arena. A low obstacle may therefore be missed, until the robot comes sufficiently close to it, so that the *Proximity detection* brain process can detect it.

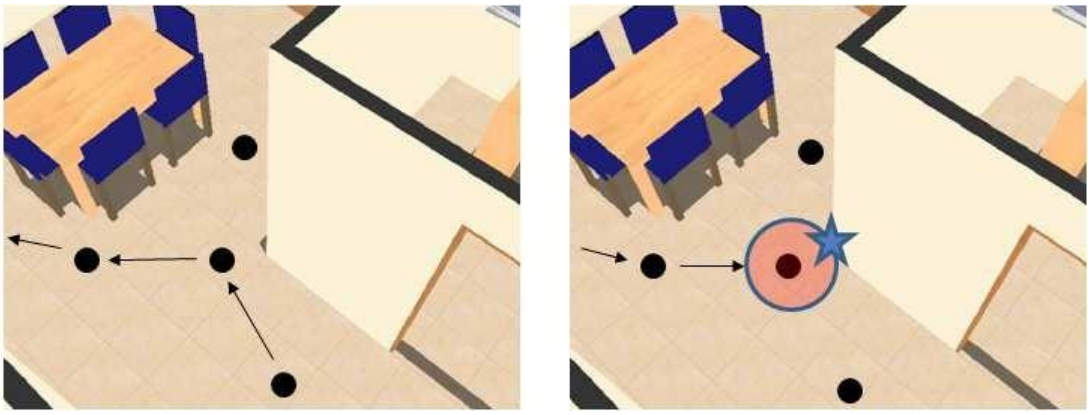

Figure 6.3: *An illustration of a problem that might occur during exploration. Moving in one particular direction (left panel) the robot is able to place and follow the nodes shown. However, upon returning (right panel), the robot may conclude that it will be unable to pass the node near the corner, due to the proximity detection triggered as the robot approaches the node, with the wall right in front of it.*

Note that unreachable nodes are exempt from the minimum distance require- ment. This is so, since a given node may be unreachable from *one* direction but perhaps reachable from some other direction. Thus, the exploration algorithm is allowed to place new pending nodes arbitrarily close to unreachable nodes.

One should note that robust exploration of any arena is more difficult than it might seem. An example of a problem that might occur is shown in Fig. 6.3. Here, the robot passes quite near a corner on its outbound journey (left panel), but no proximity detection is triggered. By contrast, upon returning (right panel) a proximity detection *is* triggered which, in turn, may force the robot to abandon its current path. In fact, the *Exploration* behavior contains a method (which will not be described here) for avoiding such deadlocks. In the (very rare) cases in which even the deadlock avoidance method fails, the robot sim- ply stops, and reports its failure.

Because of the path following strategy described above, the robot may sometimes take an unnecessarily long path from the active node to the tar- get node. However, this does not happen so often since, in most cases, the robot will proceed directly to a newly added node, for which the path length is the same as the cartesian distance. However, when the robot cannot place any more nodes (something that occurs, for example, when it reaches a cor- ner), a distant node may become the target node. Therefore, in cases where the path to the target node differs from the direct path, the robot rotates to face the target node (provided that it is located within a distance L, where L should be smaller than or equal to the range R of the LRF). Next, if the robot concludes (based on its LRF readings) that it can reach the target node, it then proceeds directly towards it, rather than following the path. However, also in this case,

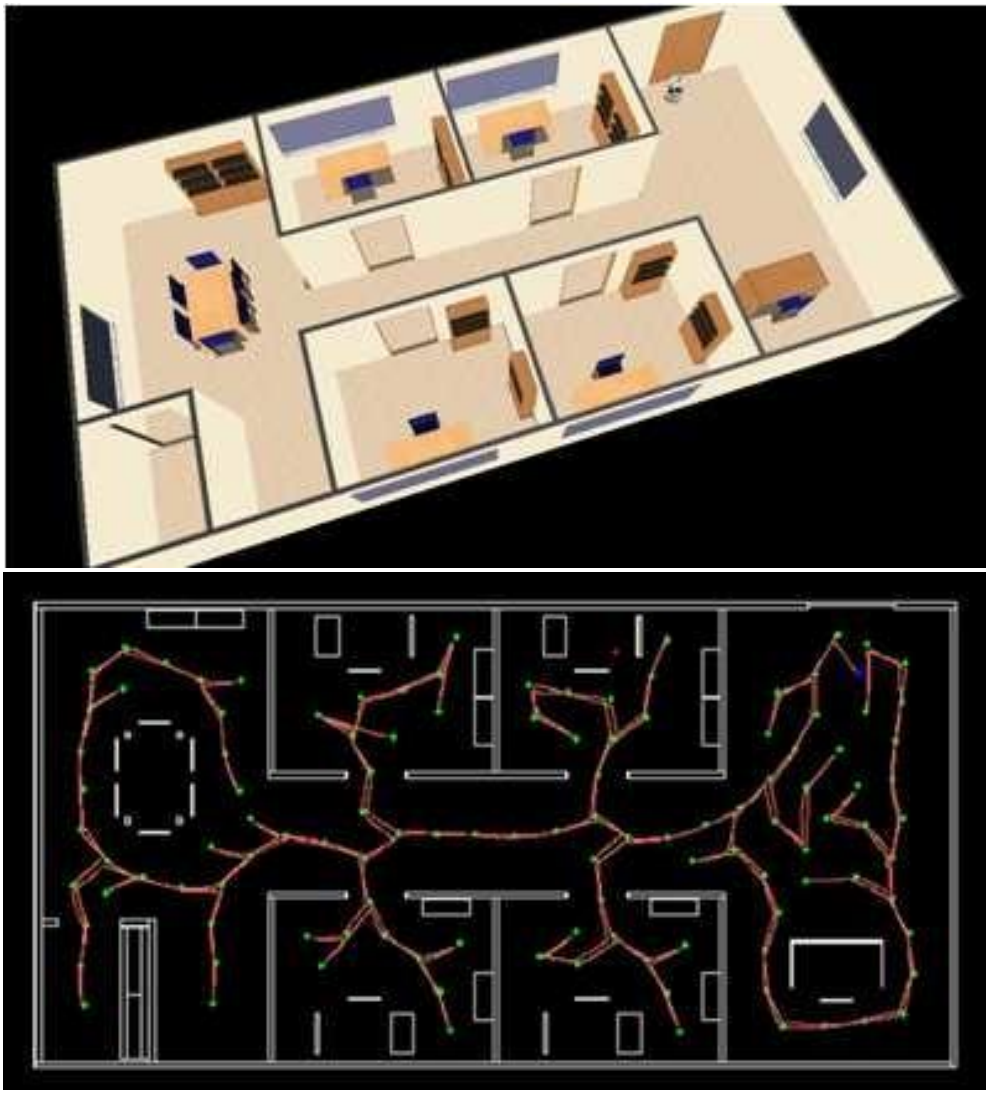

Figure 6.4: *Left panel: The robot in its initial position in an unexplored arena. Right panel: The final result obtained after executing the* Exploration *behavior. The completed (visited) nodes are shown as green dots, whereas the (single) unreachable node is shown as a red dot. The final target node (the last node visited) is shown as a blue dot. In this case, the robot achieved better than 99.5% sensory area coverage of the arena.*

it is possible that a (low) obstacle prevents the robot from reaching its target, in which case the robot instead switches to following the path as described above.

The robot continues this cycle of node placement and movement between nodes, until all nodes have been processed (i.e. either having been visited or deemed unreachable), at which point the exploration of the arena is complete. The *Exploration* behavior consists of an FSM with 17 states, which will not be

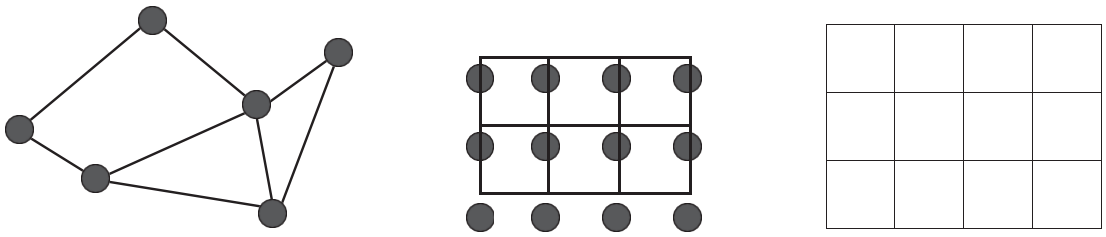

Figure 6.5: *Three examples of grids that can be used in connection with grid-based navigation methods. In the left panel, the nodes are not equidistant, unlike the middle panel which shows a regular lattice with equidistant nodes. The regular lattice can also be represented as grid cells, as shown in the right panel. Note that the right and middle panels show equivalent grids.*

described in detail here. A performance example is shown in Fig. 6.4. The left panel shows the robot at its starting point in a typical office arena. The right panel shows the final result, i.e. the path generated by the robot. The completed (visited) exploration nodes are shown as green dots, whereas the unreachable nodes (only one in this case) are shown as red dots. The final target node is shown as a blue dot. Note that the robot achieved a *sensory* area coverage (at the height of its LRF) of more than 99.5% during exploration.

6.2 Navigation

In this section, it will again be assumed that the robot has access to accurate estimates of its pose (from the *Odometry* brain process), and the question that will be considered is: Given that the robot knows its pose and velocity, how can it navigate between two arbitrary points in an arena? In the robotics liter- ature, many methods for navigation have been presented, three of which will be studied in detail in this section.

6.2.1 Grid-based navigation methods

In **grid-based navigation methods**, the robot's environment must be covered with an (artificial) grid, consisting of **nodes** (vertices) and **edges** connecting the nodes. The grid may have any shape, as illustrated in the left panel of Fig. 6.5, i.e. it need not be a rectangular lattice of the kind shown in the middle panel. However, if the grid happens to be a rectangular lattice, it is often rep- resented as shown in the right panel of the figure, where the nodes have been replaced by cells, and the edges are not shown[2]. Furthermore, the edges must be associated with a (non-negative) cost, which, in many cases is simply taken

[2]Note that in the cell representation in the right panel, the sides of each cell are *not* edges: The edges connect the centers of the grid cells to each other, as shown in the middle panel.

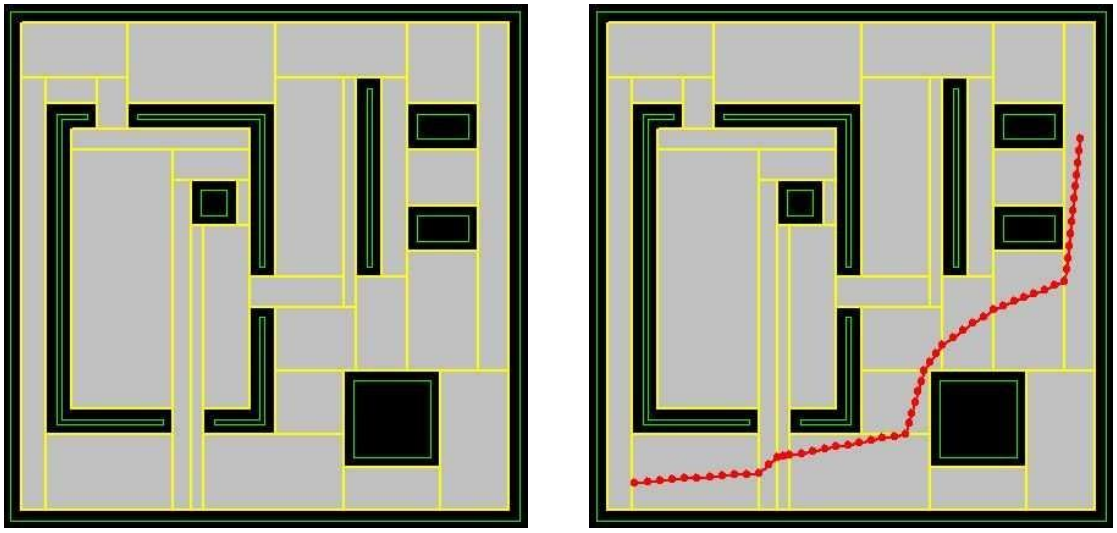

Figure 6.6: *Left panel: An example of automatic grid generation. The walls of the arena are shown as green thin lines. The black regions represent the forbidden parts of the arena, either unreachable locations or positions near walls and other obstacles. The grid cell boundaries are shown as thick yellow lines. Right panel: An example of a path between two points in the arena. The basic path (connecting grid cells) was generated using Dijkstra's algorithm (see below). The final path, shown in the figure, was adjusted to include changes of directions within grid cells, thus minimizing the length of the path. Note that all cells are convex, so that the path segments within a cell can safely be generated as straight lines between consecutive waypoints.*

as the euclidean distance between the nodes. Thus, for example, in the grids shown in the middle and right panels of Fig. 6.5, the cost of moving between adjacent nodes would be equal to 1 (length unit), whereas, in the grid shown in the left panel the cost would vary depending on which nodes are involved.

An interesting issue is the *generation* of a navigation grid, given a two-dimensional map of an arena. This problem is far from trivial, especially in complex arenas with many walls and other objects. Furthermore, the grid gen- eration should take the robot's size (with an additional margin) into account, in order to avoid situations where the robot must pass very close to a wall or some other object. The grid-based navigation methods described below gener- ate paths *between* grid cells. On a regular grid with small, quadratic cells (as in the examples below) it is sometimes sufficient to let the robot move on straight lines between the cell centers. However, the generated path may then become somewhat ragged. Furthermore, in more complex grids, where the cells are of different size, following a straight line between cell centers may result in an unnecessarily long path. Thus, in such cases, the robot must normally modify its heading *within* a cell, in order to find the shortest path.

When generating a grid, one normally requires the grid cells to be **convex**,

1. Place the robot at the start node, which then becomes the current node. Assign the status *unvisited* to all nodes.

2. Go through each of the cells a_i that are (i) unvisited and (ii) directly reach- able (via an edge) from the current node c. Such nodes are referred to as **neighbors** of the current node. Compute the cost of going from a_i to the target node t, using the heuristic $f(a_i)$.

3. Select the node $a_{\min}$ associated with the lowest cost, based on the cost values computed in Step 2.

4. Set the status of the current node c as *visited*, and move to $a_{\min}$ which then becomes the current node.

5. Return to Step 2.

Figure 6.7: *The best-first search algorithm.*

so that all points on a straight line between any two points in the cell also are part of the cell. One way of doing so is to generate a grid consisting of trian- gular cells, which will all be convex. However, such grids may not be optimal: The pointiness of the grid cells may force the robot to make many unnecessary (and sharp) turns. An algorithm for constructing, from a map, a general grid consisting of convex cells with four or more sides (i.e. non-triangular) exists as well[3]. Fig. 6.6 shows an example of a grid generated with this algorithm. Because of its complexity, the algorithm will not be considered in detail here. Instead, in the examples below, we shall consider grids consisting of small quadratic cells, and we will neglect changes of direction within grid cells.

Best-first search algorithm

In **best-first search** (BFS) algorithm the robot moves greedily towards the tar- get, as described in Fig. 6.7. As can be seen, the BFS method chooses the next node based on the (estimated) cost of going from that node n to the goal, which is estimated using a **heuristic** function $f(n)$. $f(n)$ can be chosen in different ways, the simplest being to use the euclidean distance between the node un- der consideration and the target. However, in that case, the BFS method may, in fact, get stuck. A more sophisticated heuristic function may, for example, add a penalty for each obstacle encountered on a straight-line path from the node under consideration to the target node.

The path can be generated by simply storing the list of visited nodes during

[3]See Wahde, M., Sandberg, D., and Wolff, K. *Reliable long-term navigation in indoor environments*, In: Topalov, A.V. (Ed.), Recent advances in Mobile Robots, InTech, 2011, pp. 261–286.

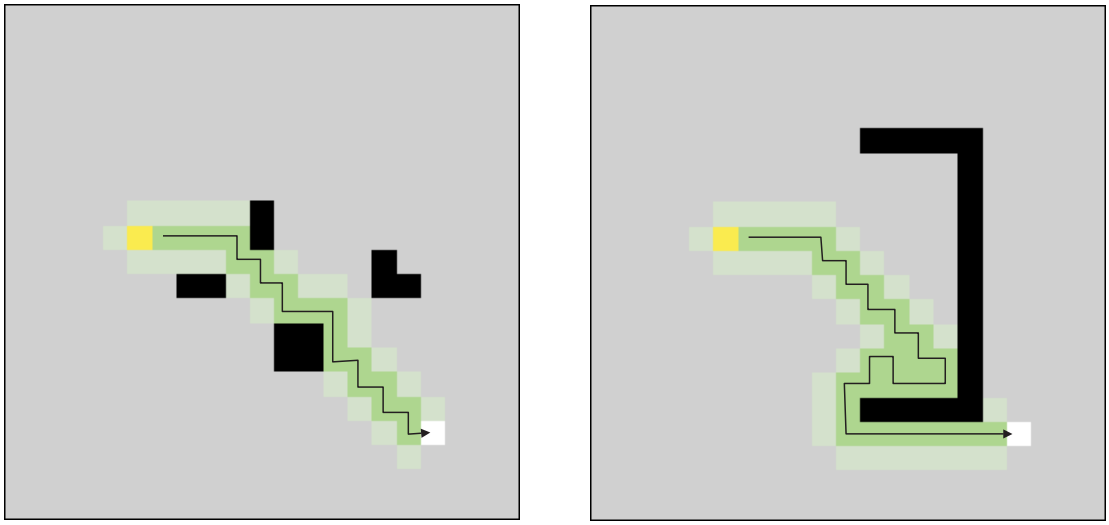

Figure 6.8: *Two examples of paths generated using the BFS algorithm. The cells (nodes) that were checked during path generation are shown in light green, whereas the actual path is shown in dark green and with a solid line. The yellow cell is the start node and the white cell is the target node.*

path generation. The BFS method is very efficient in the absence of obstacles or when the obstacles are few, small, and far apart. An example of such a path generated with BFS is shown in the left panel of Fig. 6.8. As can be seen, the robot quickly moves from the start node to the target node. However, if there are extended obstacles between the robot's current position and the target node, the BFS algorithm will not find the shortest path, as shown in the right panel of Fig. 6.8. Because of its greedy approach to the target, the robot will find itself in front of the obstacle, and must then make a rather long detour to arrive at the target node.

Dijkstra's algorithm

Like BFS, **Dijkstra's algorithm** also relies on a grid in which the edges are associated with non-negative costs. Here, the cost will simply be taken as the euclidean distance between nodes. Instead of focusing on the (estimated) cost of going from a given node to the target note, Dijkstra's algorithm considers the distance between the *start* node and the node under consideration, as de- scribed in Fig. 6.9. In Step 2, the distance from the start node s to any node a_i is computed using the (known) distance from the initial node to the current node c and simply adding the distance between c and a_i. This algorithm will check a large number of nodes, in an expanding pattern from the start node, as shown in Fig. 6.10. In order to determine the actual path to follow, whenever a new node a is checked, a note is made regarding the predecessor node p, i.e. the

1. Place the robot at the start node s, which then becomes the current node. Assign the distance value 0 to the start node, and infinity to all other nodes (in practice, use a very large, finite value). Set the status of all nodes to *unvisited*.

2. Go through all the unvisited, accessible (i.e. empty) neighbors a_i of the current node c, and compute their distance d from the *start* node s. If d is smaller than the previously stored distance d_i (initially infinite, see Step 1), then (i) update the stored distance, i.e. set $d_i = d$ and (ii) assign the current node as the predecessor node of a_i.

3. After checking all the neighbors of the current node, set its status to vis- ited.

4. Select the node (among *all* the unvisited, accessible nodes in the grid) with the smallest distance from the start node, and set it as the new cur- rent node.

5. Return to Step 2, unless the target has been reached.

6. When the target has been reached, use the predecessor nodes to trace a path from the target node to the start node. Finally, reverse the order of the nodes to find the path from the start node to the target node.

Figure 6.9: *Dijkstra's algorithm.*

node that was the current node when checking node a. When the target has been found, the path connecting it to the initial node can be obtained by going through the predecessor nodes backwards, from the target node to the initial node.

Unlike the BFS algorithm, Dijkstra's algorithm is guaranteed to find the shortest path[4] from the start node to the target node. However, a drawback with Dijkstra's algorithm is that it typically searches many nodes that, in the end, turn out to be quite irrelevant. Looking at the search patterns in Figs. 6.8 and 6.10, one may hypothesize that a *combination* of the two algorithms would be useful. Indeed, there is an algorithm, known as **A*** that combines the BFS and Dijkstra algorithms. Like Dijkstra's algorithm, A* is guaranteed to find the shortest path. Moreover, it does so more efficiently than Dijkstra's algorithm. However, A* is beyond the scope of this text.

[4]There may be more than one such path: Dijkstra's algorithm will select one of them.

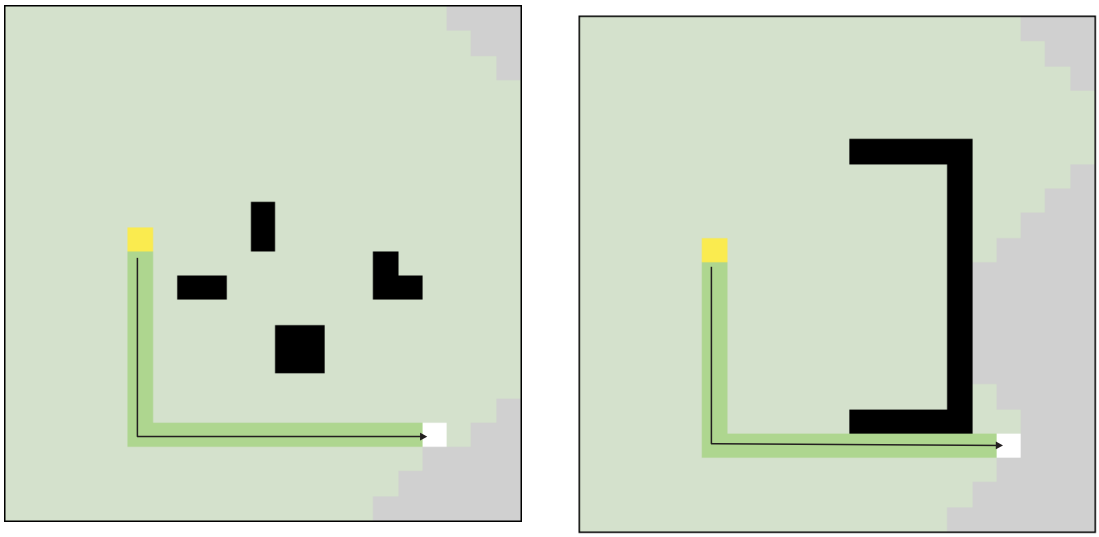

Figure 6.10: *Two examples of paths generated using Dijkstra's algorithm. The cells (nodes) that were checked during path generation are shown in light green, whereas the actual path is shown in dark green and with a solid line. The yellow cell is the start node and the white cell is the target node.*

6.2.2 Potential field navigation

Unlike the algorithms described above, the **potential field method** does not require a grid. In the potential field method, a robot obtains its desired direc- tion of motion as the negative gradient of an artificial potential field, generated by potentials assigned to the navigation target and to objects in the arena.

Potential fields

As shown in Fig. 6.11, a potential field can be interpreted as a landscape with hills and valleys, and the motion of a robot can be compared to that of a ball rolling through this landscape. The navigation target is assigned a potential corresponding to a gentle downhill slope, whereas obstacles should generate potentials corresponding to steep hills.

In principle, a variety of different equations could be used for defining dif- ferent kinds of potentials. An example, namely a potential with ellipsoidal equipotential surfaces, and exponential variation with (ellipsoidal) distance from the center of the potential, takes the mathematical form

$$\varphi(x, y; x_p, y_p, \alpha, \beta, \gamma) = \alpha e^{-\left(\frac{x-x_p}{\beta}\right)^2-\left(\frac{y-y_p}{\gamma}\right)^2}, \quad (6.1)$$

where (x, y) is the current (estimated) position at which the potential is calcu- lated, (x_p, y_p) is the position of the object generating the potential, and α, β and γ are constants (not to be confused with the constants defined in connection

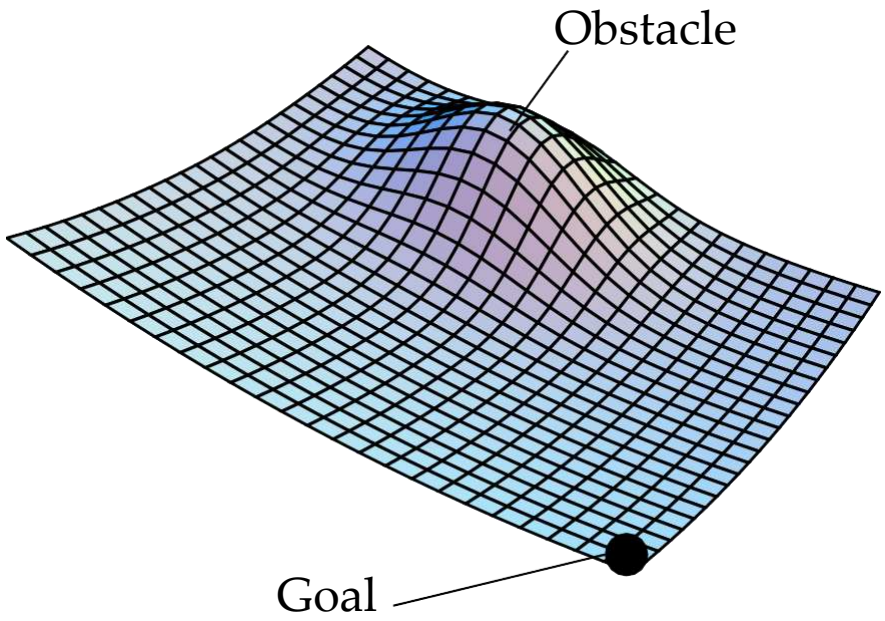

Figure 6.11: *A potential field containing a single obstacle and a navigation goal.*

with the equations of motion in Chapter 2 and the sensor equations in Chapter 3). Now, looking at the mathematical form of the potentials, one can see that an attractive potential (a valley) is formed if α is negative, whereas a positive value of α will generate a repulsive potential (a hill).

Normally, the complete potential field contains many potentials of the form given in Eq. (6.1), so that the total potential becomes

$$\Phi(x, y) = \sum_{i=1}^{k} \varphi_i(x, y; x_{p_i}, y_{p_i}, \alpha_i, \beta_i, \gamma_i), \tag{6.2}$$

where k is the number of potentials. An example of a potential field, for a simple arena with four central pillars, is shown in Fig. 6.12.

Navigating in a potential field

Once the potential field has been defined, the desired direction of motion $\hat{\mathbf{r}}$ of the robot can be computed as the negative of the normalized gradient of the field

$$\hat{\mathbf{r}} = -\frac{\nabla\Phi}{|\nabla\Phi|} \equiv -\frac{\left(\frac{\partial\Phi}{\partial x}, \frac{\partial\Phi}{\partial y}\right)}{\sqrt{\left(\frac{\partial\Phi}{\partial x}\right)^2 + \left(\frac{\partial\Phi}{\partial y}\right)^2}} \tag{6.3}$$

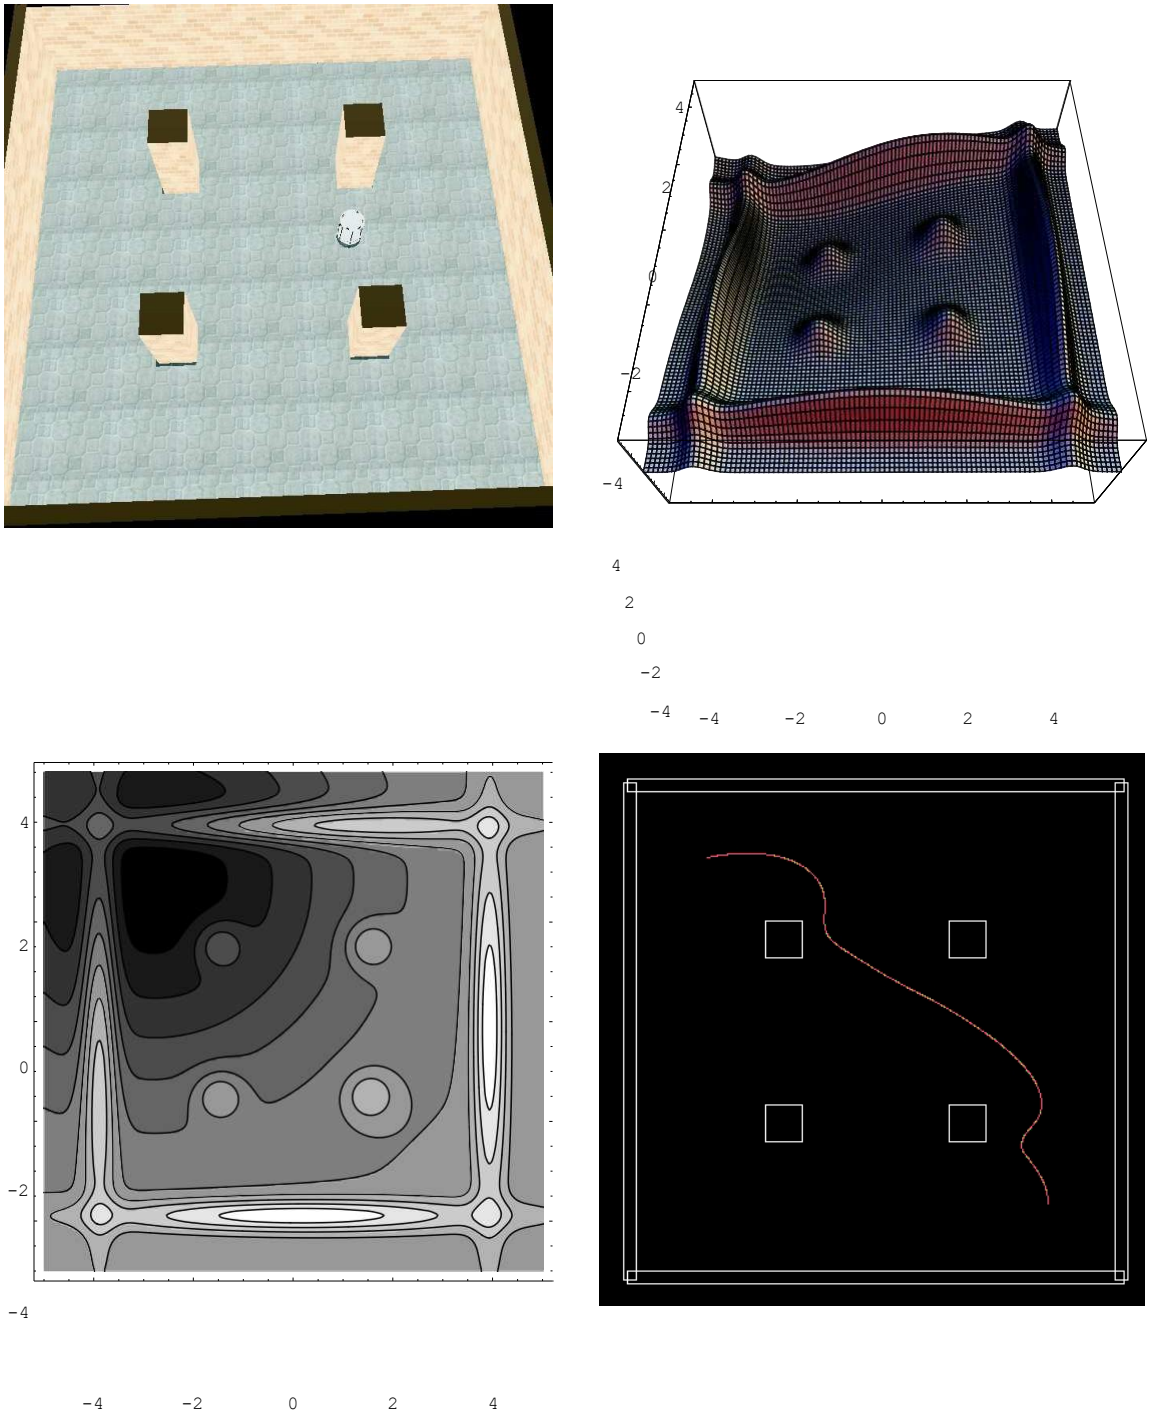

Figure 6.12: *An illustration of potential field navigation in GPRSim. Upper left panel: A simple arena, with a robot following a potential field toward a target position in the upper left corner of the arena. Upper right panel: The corresponding potential field, generated by a total of*

nine potentials (one for the target, one for each of the walls, and one for each pillar). Lower left panel: A contour plot of the potential field, in which the target position can be seen in the upper left corner. Lower right panel: The trajectory followed by the robot. Note that, in this simulation, the odometric readings were (unrealistically) noise-free.

In order to integrate the equations of motion of the robot, it is not sufficient only to know the desired direction: The magnitude of the force acting on the robot must also be known. In principle, the negative gradient of the potential field could be taken (without normalization) as the *force* acting on the robot, providing both magnitude and direction. However, in that case, the magni- tude of the force would vary quite strongly with the position of the robot, making the robot a dangerous moving object (if it is large). Thus, the poten- tial field is only used for providing the direction, as in Eq. (6.3). The robot's speed v (i.e. the magnitude of its velocity vector $\mathbf{v}$) can be assigned in various ways. For example, one may use proportional control to try to keep the speed constant[5].

An example of a trajectory generated during potential field navigation is shown in the lower right panel of Fig. 6.12. In the experiment in which this figure was generated, the noise in the odometric readings was (unrealistically) set to zero, since the aim here is simply to illustrate potential field navigation. However, in a realistic application, one would have to take into account the fact that the robot's estimate of its pose will never be error-free. Thus, when setting up a potential field, it is prudent to make the potentials slightly larger[6] than the physical objects that they represent. At the same time, in narrow corridors, one must be careful not to make the potentials (for walls on opposite sides of the corridor, say) so wide that the robot will be unable to pass.

In fact, the definition of a potential field for a given arena is something of an art. In addition to the problem of determining the effective extension of the potentials, one also has to decide whether a given object should be repre- sented by one or several potentials (for instance of the form given in Eq. (6.1)). For example, an extended object (for example, a long wall) *can* be represented as a single potential (typically with very different values of the parameters β and γ), but it can also be represented as a sequence of potentials. In complex environments, one may resort to stochastic optimization of the potential field, as well as the details of the robot's motion in the field[7].

Aspects of potential field navigation

A gradient-following method, such as the potential field method, always suf- fers the risk of encountering local minima in the field. Of course, in potential

[5]The procedure for assigning the robot's speed in potential field navigation will be de- scribed below.

[6]Of course, since the exponential potentials defined in Eq. (6.1) have infinite extension, the corresponding force never drops exactly to zero, but beyond a distance of a few d, where $d = \max(\beta, \gamma)$, the force is negligible.

[7]For an example of such an approach, see Savage *et al.*, *Optimization of waypoint-guided potential field navigation using evolutionary algorithms*, Proceedings of the 2004 IEEE/RSJ Inter- national Conference on Intelligent Robots and Systems (IROS 2004), 3463-3468, 2004.

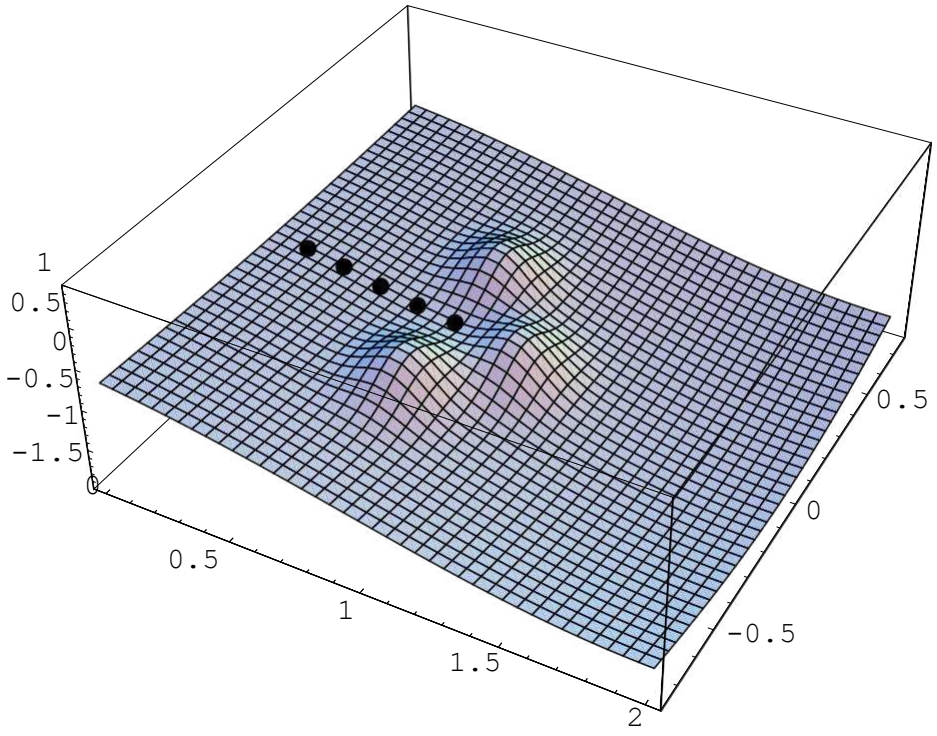

Figure 6.13: *The locking phenomenon. Following the gradient of the potential field the robot, whose trajectory is shown as a sequence of black dots, moves along the x-axis toward the goal, located at* (2, 0)*. However, because of the local minimum in the potential field, the robot eventually gets stuck.*

field navigation, the goal is to reach the local minimum represented by the navigation target. However, depending on the shape of the arena (and there- fore the potential field), there may also appear one or several unwanted local minima in the field, at which the robot may become trapped.

This is called the **locking phenomenon** and it is illustrated in Fig. 6.13. Here, a robot encounters a wedge-shaped obstacle represented by three poten- tials. At a large distance from the obstacle, the robot will be directed toward the goal potential, which is located behind the obstacle as seen from the start- ing position of the robot. However, as the robot approaches the obstacles their repulsive potentials will begin to be noticeable. Attracted by the goal, the robot will thus eventually find itself stuck inside the wedge, at a local minimum of the potential.

In order to avoid locking phenomena, the path between the robot and the goal can be supplied with **waypoints**, represented by attractive potentials (for example, of the form given in Eq. (6.1)) with rather small extension. Of course, the introduction of waypoints leads to the problem of determining where to

put them. An analysis of such methods will not be given here[8]. Suffice it to say that the problem of waypoint placement can be solved in various ways to aid the robot in its navigation. A waypoint should be removed once the robot has passed within a distance L from it, to avoid situations in which the robot finds itself stuck at a *waypoint*.

The potential field method also has several advantages, one of them being that the direction of motion is obtained simply by computing the gradient of the potential field at the current position, without the need to generate an en- tire path from the current position to the navigation target. Furthermore, the potential field is defined for all points in the arena. Thus, if the robot tem- porarily must suspend its navigation (for example, in order to avoid a moving obstacle), it can easily resume the navigation from wherever it happens to be located when the *Obstacle avoidance* behavior is deactivated.

In the discussion above, only stationary obstacles were considered. Of course, moving obstacles can be included as well. In fact, the potential field method is commonly used in conjunction with, say, a grid-based navigation method, such that the latter generates the nominal path of the robot, whereas the potential field method is used for adjusting the path to avoid moving ob- stacles. However, methods for reliably *detecting* moving obstacles are beyond the scope of this text.

Using the potential field method

As mentioned above, the potential field only provides the current desired *di- rection* of motion. In order to specify a potential field navigation behavior com- pletely, one must also provide a method for setting the speed of the robot. This can be done as follows: Given the robot's estimated (from odometry) angle of heading ϕ_{est} and the desired (reference) direction ϕ_{ref} (obtained from the po- tential field), one can form the quantity $\Delta\phi$ as

$$\Delta\phi = \phi_{\text{ref}} - \phi_{\text{est}}. \tag{6.4}$$

The desired speed differential ΔV (the difference between the right and left wheel speeds) can then be set according to

$$\Delta V = K_{\text{p}} V_{\text{nav}} \Delta\phi, \tag{6.5}$$

where K_{p} is a regulatory constant (P-regulation is used) and V_{nav} is the (de- sired) speed of the robot during normal navigation. Once ΔV has been com- puted, reference speeds are sent to the (velocity-regulated) motors according to

$$v_{\text{L}} = V_{\text{nav}} - \frac{\Delta V}{2}, \tag{6.6}$$

[8]See, however, the paper by Savage *et al.* mentioned in Footnote 6.

$$v_R = V_{nav} + \frac{\Delta V}{2}, \tag{6.7}$$

where v_R and v_L are the reference speeds of the left and right wheels, respec- tively. Note that one can of course only set the *desired* (reference) speed values; the actual speed values obtained depend on the detailed dynamics of the robot and its motors.

If the reference angle differs strongly from the estimated heading (which can happen, for example, in situations where the robot comes sufficiently close to an obstacle whose potential generates a steep hill), the robot may have to suspend its normal navigation and instead carry out a pure rotation, setting $v_L = V_{rot}$, $v_R = V_{rot}$ for a left (counterclockwise) rotation, where V_{rot} is the ro- tation velocity, defined along with V_{nav} (and the other constants) during setup. In case the robot should carry out a clockwise rotation, the signs are reversed. The direction of rotation is, of course, determined by the sign of the differ- ence between the reference angle and the (estimated) heading. In this case, the robot should turn until the difference between the reference angle and the esti- mated heading drops below a user-specified threshold, at which point normal navigation can resume.

6.3 Localization

In Sects. 6.1 and 6.2, it was (unrealistically) assumed that the robot's odome- try would provide perfect estimates of the pose. In reality, this will never be the case, and therefore the problem of recalibrating the odometric readings, from time to time, is a fundamental problem in robotics. Doing so requires a method for localization independent from odometry, and such methods usu- ally involve LRFs (even though cameras are also sometimes used), sensors that are difficult to simulate in ARSim (because of the large number of rays which would slow down the simulation considerably). Therefore, in this section, lo- calization will be described as it is implemented in the simulator GPRSim and in GPRBS, where LRFs are used.

Robot localization requires two brain processes: The cognitive *Odometry* process and an independent process for odometric recalibration, which both in GPRSim and in GPRBS goes under the name *Laser localization*, since the behavior for odometric recalibration uses the readings of an LRF, together with a map, to infer its current location using scan matching, as described below.

In fact, the problem of localization can be approached in many different ways. For outdoor applications, a robot may be equipped with GPS, which in many cases will give sufficiently accurate position estimates. However, in indoor applications (standard) GPS cannot be used, since the signal is too weak to penetrate the walls of a building. Of course, it is possible to set up a local GPS system, for example by projecting IR beacons on the ceiling, using which

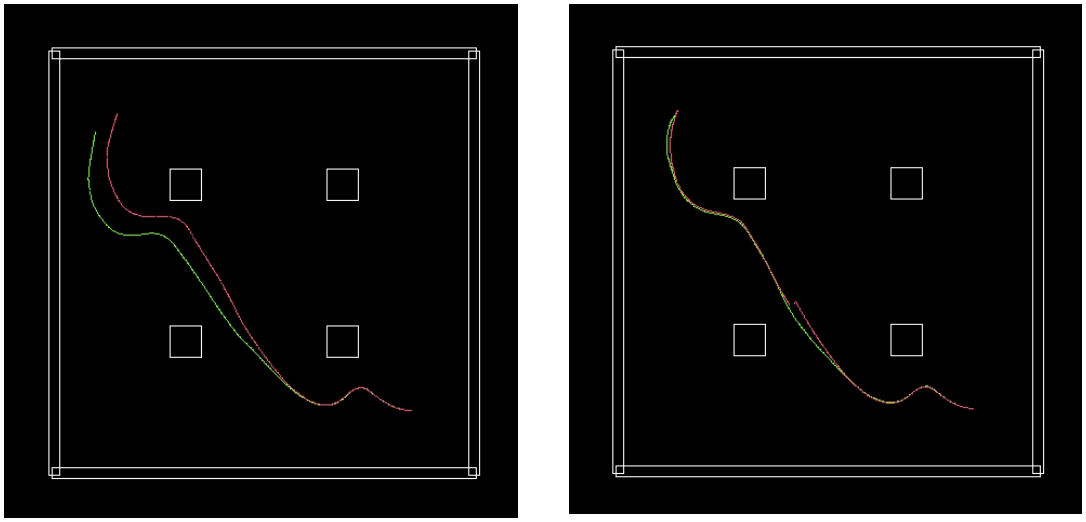

Figure 6.14: *An illustration of the need for localization in mobile robot navigation. In the left panel, the robot navigates using odometry only. As a result, the odometric trajectory (red) deviates quite significantly from the actual (green) trajectory. In the right panel,* Laser localization *was activated periodically, leading to much improved odometric estimates.*

the robot can deduce its position by means of triangulation[9]. However, such a system requires that the arena should be adapted to the robot, something that might not always be desirable or even possible.

The localization method (combining odometry and laser scan matching) that will be described here is normally used together with some navigation behavior. Thus, the robotic brain will consist of at least two motor behav- iors, in which case decision-making also becomes important. This topic will be studied in a later chapter: For now, the *Laser localization* behavior will be considered separately.

6.3.1 Laser localization

The behavior is intended for localization in arenas for which a map has been provided to the robot (in the form of a sequence of lines). The map can ei- ther be obtained using a robot (executing a *Mapping* behavior) or, for example, from the floor plan of a building. The behavior relies on scans of the arena using a two-dimensional LRF and, like many methods for localization in auto- nomous robots, it assumes that all scans are carried out in a horizontal plane, thus limiting the behavior to planar (i.e. mostly indoor) environments. In fact, the name *Laser localization* is something of a misnomer: The behavior does not actually carry out (continuous) localization. Instead, when activated, the behavior takes as input the current pose estimate and tries to improve it. If

[9]This is the method used in the **Northstar®** system, developed by Evolution Robotics, inc.

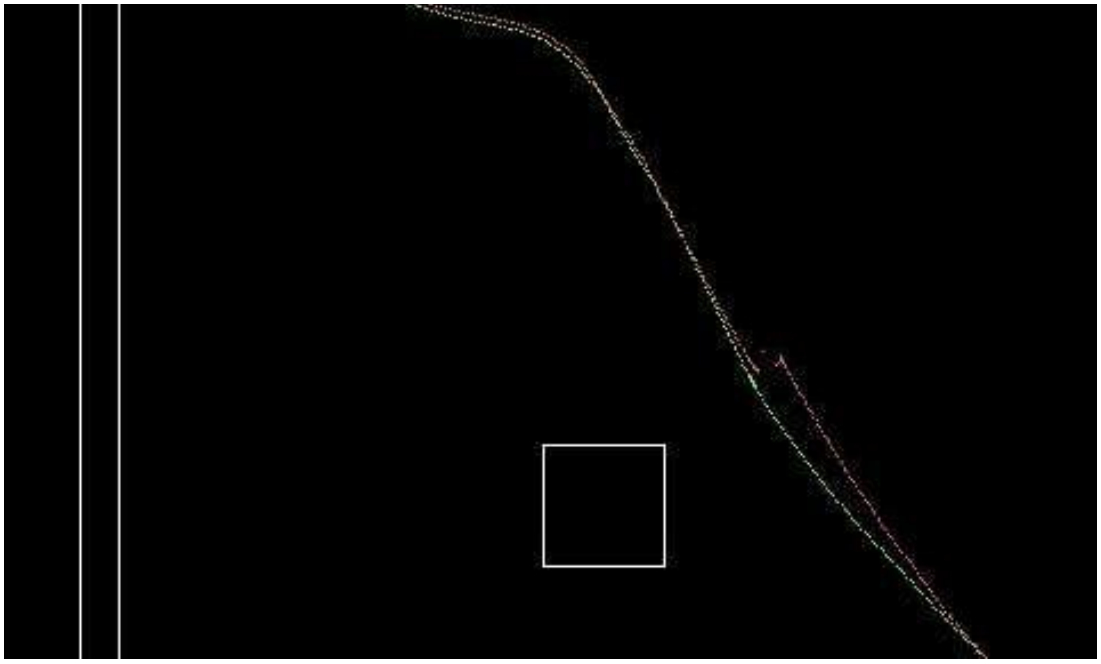

Figure 6.15: *An enlargement of the most significant correction in odometric readings (in the right panel of Fig. 6.14) resulting from the* Laser localization *behavior.*

successful, the odometric pose is reset to the position suggested by the *Laser localization* behavior.

The left panel of Fig. 6.14 illustrates the need for localization: The nav- igation task shown in Fig. 6.12 was considered again (with the same start- ing point, but a different initial direction of motion), this time with realistic (i.e. non-zero) levels of noise in the wheel encoders and, therefore, also in the odometry. As can be seen in the figure, the odometric drift causes a rather large discrepancy between the actual trajectory (green) and the odometric es- timate (red). In the right panel, the robotic brain contained two behaviors (in addition to the cognitive *Odometry* process), namely *Potential field navigation* and *Laser localization*. The *Laser localization* behavior was activated periodically (thus deactivating the *Potential field navigation* behavior), each time recalibrat- ing (if necessary) the odometric readings. As can be seen in the right panel of Fig. 6.14, with laser localization in place, the discrepancy between the odo- metric and actual trajectories is reduced significantly. At one point, the *Laser localization* behavior was required to make a rather large correction of the odo- metric readings. That particular event is shown enlarged in Fig. 6.15. As can be seen, the odometric readings undergo a discrete step at the moment of lo- calization.

When activated, the localization behavior[10] considered here first stops the

[10]See Sandberg, D., Wolff, K., and Wahde, M. *A robot localization method based on laser scan*

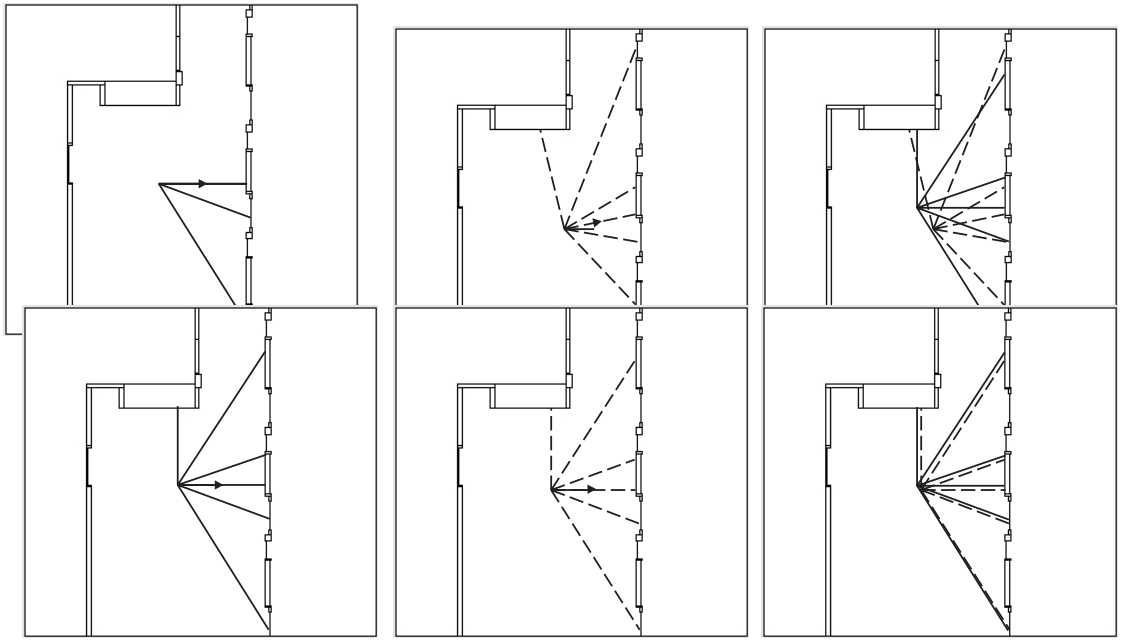

Figure 6.16: *Two examples of scan matching. The leftmost panel in each row shows a few rays (solid lines) from an actual LRF reading (plotted in the map used by the virtual LRF), and the middle panels show the virtual LRF readings (dotted lines) in a case in which the estimated pose differs quite strongly from the correct one (upper row), and one in which the difference is small (bottom row). The direction of heading is illustrated with arrows. The right panel in each row shows both the actual LRF rays and the virtual ones. The figure also illustrates the map, which consists of a sequence of lines.*

robot, and then takes a reading of the LRF. Next, it tries to match this reading to a *virtual* reading taken by placing a virtual LRF (hereafter: vLRF) at various positions *in the map*. Two examples of scan matching are shown in Fig. 6.16. The three panels in the upper row show a situation in which the odometry has drifted significantly. The upper left panel shows the readings (i.e. laser ray distances) from an actual LRF mounted on top of the robot (not shown). Note that, for clarity, the figure only shows a few of the many (typically hundreds) laser ray directions. The upper middle panel shows the readings of the vLRF, placed at the initial position and heading obtained from odometry. As can be seen in the upper right panel, the two scans match rather badly. By contrast, the three panels of the bottom row show a situation in which the pose error is small. The purpose of the search algorithm described below is to be able to correct the odometry, i.e. to reach a situation similar to the one shown in the bottom row of Fig. 6.16. Fig. 6.17 shows another example of a good (left panel) and a bad (right panel) scan match. In the case shown in the left panel, the

matching, Proc. of AMiRE 2009, pp. 171-178, 2009.

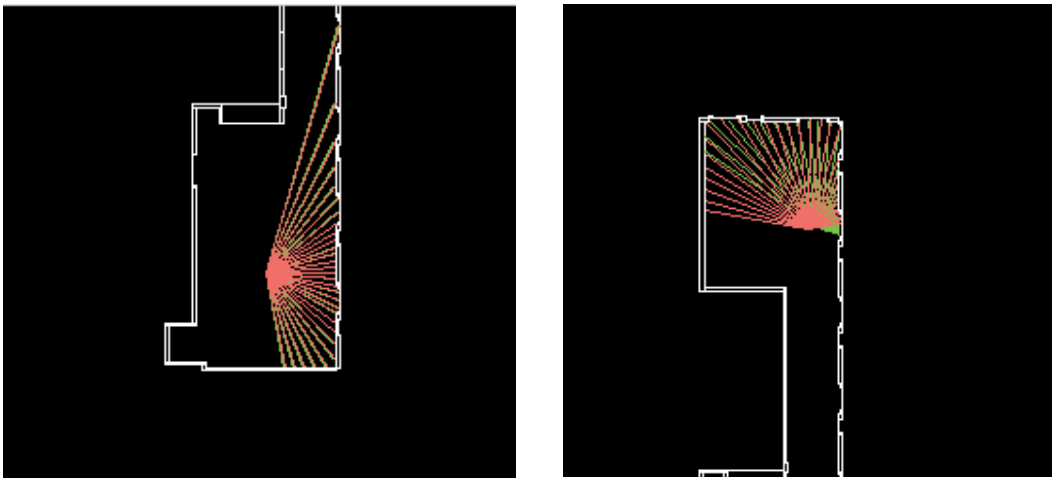

Figure 6.17: *Matching of LRF rays (in a different arena than the one used in the examples above). The readings of the actual LRF are shown in green, and those of the virtual LRF are shown in red. Left panel: An almost exact match. Right panel: In this case, the odometry has drifted enough to cause a large discrepancy between the actual and virtual LRF rays.*

odometric pose estimate is quite good, so that the rays from the actual LRF (green) match those of the vLRF quite well, at the current pose estimate. By constrast, in the situation shown in the right panel, the odometry has drifted significantly.

Scan matching algorithm

Let $\mathbf{p} = (x, y, \phi)$ denote a pose (in the map) of the vLRF. The distances between the vLRF and an obstacle, along ray i, are obtained using the map[11] and are de- noted δ_i. Similarly, the distances obtained for the *real* LRF (at its current pose, which normally differs from $\mathbf{p}$ when the localization behavior is activated) are denoted d_i.

The matching error s between two scans can be defined in various ways. For rays that do not intersect an obstacle, the corresponding reading (d_i or δ_i) is (arbitrarily) set to -1. Such rays should be excluded when computing the error. Thus, the matching error is taken as

$$s = \frac{1}{\nu} \sum_{i=1}^{n} \chi_i \left(1 - \frac{\delta_i}{d_i}\right)^2, \tag{6.8}$$

where n is the number of LRF rays used[12]. The parameter χ_i is equal to one

[11] In practice, the ray reading δ_i of the vLRF is obtained by checking for intersection between the lines in the map and a line of length R (the range of the LRF) pointing in the direction of the ray, and then choosing the shortest distance thus obtained (corresponding to the nearest obstacle along the ray). If no intersection is found, the corresponding reading is set to -1.

[12] For example, in the case of a Hokuyo URG-04LX LRF, a maximum of 682 rays are available.

for those indices i for which *both* the real LRF and the vLRF detect an obstacle (i.e. obtain a reading different from -1) whereas χ_i is equal to zero for indices i such that either the real LRF or the vLRF (or both) do not detect any obstacle (out to the range R of the LRF). v denotes the number of rays actually used in forming the error measure, i.e. the number of rays for which χ_i is equal to one. As can be seen, s is a measure of the (normalized) average relative deviation in detected distances between the real LRF and the vLRF.

Since d_i are given and δ_i depend on the pose of the vLRF, one may write $s = s(\mathbf{p})$. Now, if the odometric pose estimate happens to be exact, the virtual and actual LRF scans will be (almost) identical (depending on the accuracy of the map and the noise level in the real LRF), resulting in a very small matching error, in which case the localization behavior can be deactivated and the robot may continue its navigation. However, if the error exceeds a user-defined threshold T, the robot can conclude that its odometric estimates are not suf- ficiently accurate, and it must therefore try to minimize the matching error by trying various poses in the map, i.e. by carrying out a number of virtual scans, in order to determine the *actual* pose of the robot. The scan matching problem can thus be formulated as the optimization problem of finding the pose $\mathbf{p} = \mathbf{p}_v$ that minimizes $s = s(\mathbf{p})$. Once this pose has been found, the new odometric pose $\mathbf{p}^{\text{new}}$ is set equal to $\mathbf{p}_v$.

Note that it is assumed that the robot is standing still during localization.
This restriction (which, in principle, can be removed) is introduced in order to (i) avoid having to correct for the motion occuring during the laser sweep, which typically lasts between 0.01 and 0.1 s and (ii) avoid having to correct for the motion that would otherwise take place during scan matching procedure, which normally takes a (non-negligible) fraction of a second. Thus, only *one* scan needs to be carried out using the real LRF mounted on the robot. The re- maining work consists of generating virtual scans in the map, at a sequence of poses, and to match these to the actual LRF readings. Unlike some other scan matching methods, the method used here does not attempt to fit lines to the LRF readings. Instead, the LRF rays (actual and virtual, as described above) are used directly during scan matching. The sequence of poses for the vLRF is generated as follows: First the actual LRF scan is carried out, generating the distances d_i. Next, a virtual scan is carried out (in the map) at the current estimated position $\mathbf{p}_0$. If the error $s_0 = s(\mathbf{p}_0)$ is below the threshold T, local- ization is complete. If not, the algorithm picks a random pose $\mathbf{p}_j$ (where $j = 1$ in the first iteration) in a rectangular box of size $L_x \quad L_y \quad L_\phi$, centered on $\mathbf{p}_0$ in pose space, and computes the matching error $s_j = s(\mathbf{p}_j)$. The constants L_x and L_y are typically set to around 0.1 m and the constant L_ϕ is set to around
0.1 radians.

The process is repeated until, for some $j = j_1$, an error $s_{j_1} < s_0$ is found. At this point, the rectangular box is re-centered to $\mathbf{p}_{j_1}$, and the search continues, now picking a random pose in the rectangular box centered on $\mathbf{p}_{j_1}$. Once a po-

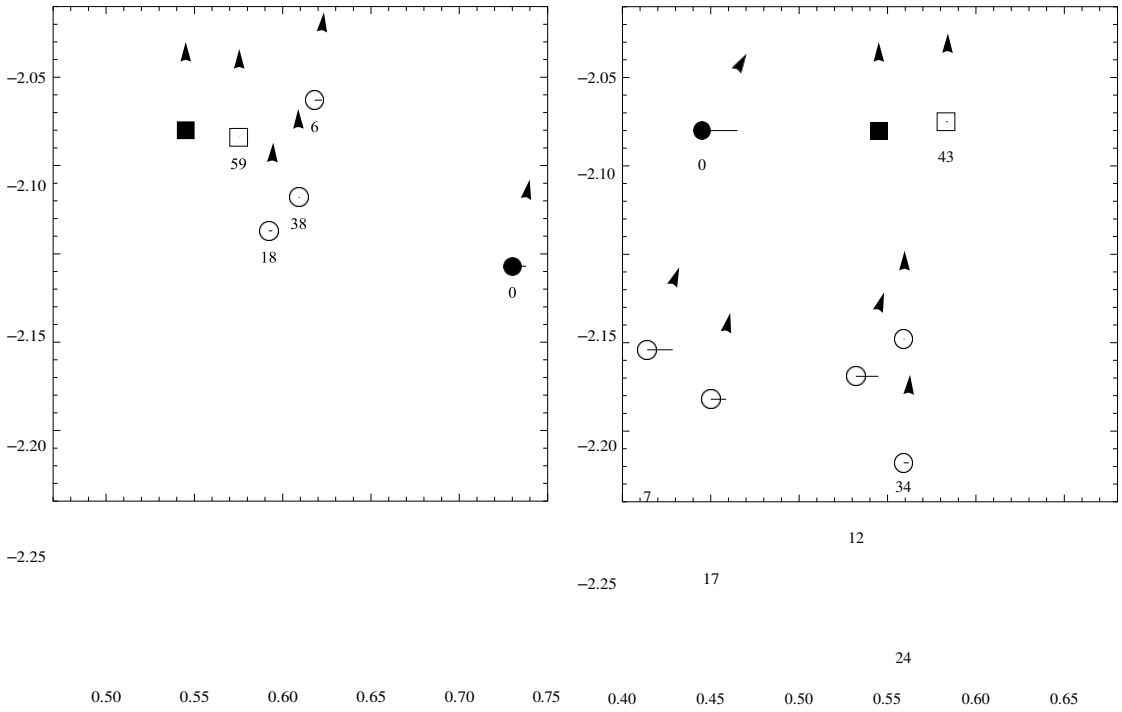
−2.05
−2.10
−2.15
−2.20
−2.25
6
59
38
18
0
0.50
0.55
0.60
0.65
0.70
0.75
−2.05
−2.10
−2.15
−2.20
−2.25
0
43
34
7
12
17
24
0.40
0.45
0.50
0.55
0.60
0.65

Figure 6.18: *An illustration of the sequence of poses generated during two executions of the search algorithm (with arrows indicating the direction of heading). In each panel, the actual position (measured in meters) of the robot is indicated with a filled square. The initial estimated position (i.e. from odometry, before correction) is shown as a filled disc, and the final estimated position is visualized as an open square. The intermediate points generated during the search are represented as open discs and are shown together with the corresponding iteration number. Note that, for clarity, only some of the intermediate points are shown.*

sition $\mathbf{p}_{j_2}$ is found for which $s_{j_2} < s_{j_1}$, the rectangular box is again re-centered etc. The procedure is repeated for a given number (N) of iterations[13].

Even though the algorithm is designed to improve both the position and the heading simultaneously, in practice, the result of running the algorithm is usually to correct the heading first (which is easiest, since an error in heading typically has a larger effect on the scan match than a position error), as can be seen clearly in the right panel of Fig. 6.18. At this stage, the estimated pose can make a rather large excursion in position space. However, once a fairly correct heading has been found, the estimated position normally converges quite rapidly to the correct position.

5. Related Work

Chapter

Autonomous robots

Both animals and robots manipulate objects in their environment in order to achieve certain goals. Animals use their senses (e.g. vision, touch, smell) to probe the environment. The resulting information, in many cases also en- hanced by the information available from internal states (based on short-term or long-term memory), is processed in the brain, often resulting in an action carried out by the animal, with the use of its limbs.

Similary, robots gain information of the surroundings, using their sensors. The information is processed in the robot's brain[1], consisting of one or several processors, resulting in motor signals that are sent to the actuators (e.g. motors) of the robot.

In this course, the problem of providing robots with the ability of making rational,

intelligent decisions will be central. Thus, the development of robotic brains is the main theme of the course. However, a robotic brain cannot op- erate in isolation: It needs sensory inputs, and it must produce motor output in order to influence objects in the environment. Thus, while it is the author's view that the main challenge in contemporary robotics lies with the devel- opment of robotic brains, consideration of the actual hardware, i.e. sensors, processors, motors etc., is certainly very important as well.

This chapter gives a brief overview of **robotic hardware**, i.e. the actual frame (body) of a robot, as well as its sensors, actuators, processors etc. The

[1]The term *control system* is commonly used (instead of the term **robotic brain**). However, this term is misleading, as it leads the reader to think of classical control theory. Concepts from classical control theory *are* relevant in robots; For example, the low-level control of the motors of robots is often taken care of by PI- or PID-regulators. However, **autonomous robots**, i.e. freely moving robots that operate without direct human supervision, are expected to function in complex, unstructured environments, and to make their own decisions concerning which action to take in any given situation. In such cases, systems based *only* on classical control theory are simply insufficient. Thus, hereafter, the term **robotic brain** (or, simply, *brain*) will be used when referring to the system that provides an autonomous robot, however simple, with the ability to process information and decide upon which actions to take.

1

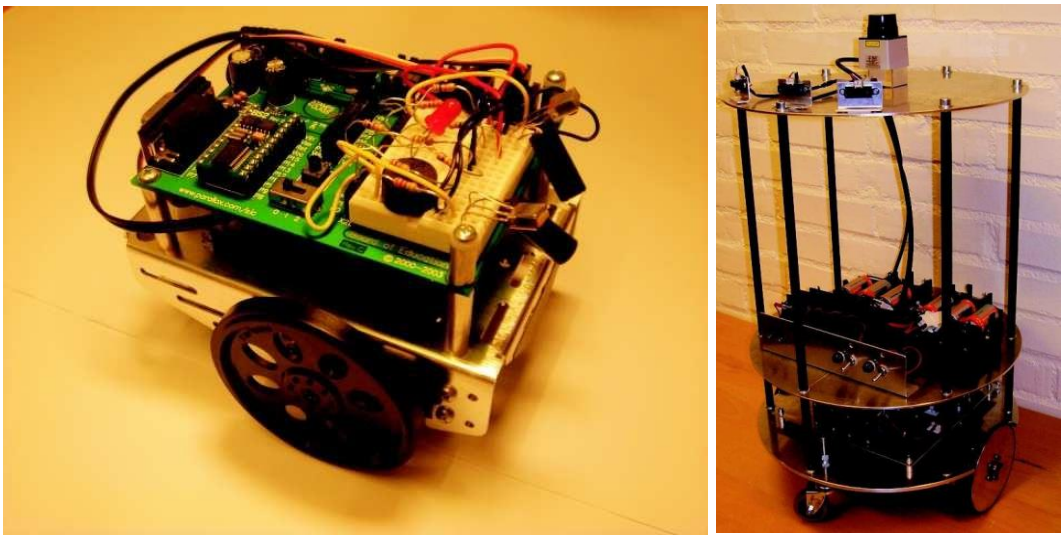

Figure 1.1: *Left panel: A Boe-bot. Right panel: A wheeled robot currently under construction in the Adaptive systems research group at Chalmers.*

various hardware-related issues will be studied in greater detail in the second half of the course, which will involve the construction of an actual robot of the kind shown in the left panel of Fig. 1.1.

1.3 Robot types

The are many different types of robots, and the taxonomy of such machines can be constructed in various ways. For example, one may classify different kinds of robots based on their complexity, their likeness to humans (or animals), their way of moving etc. In this course we shall limit ourselves to **mobile robots**, that is, robots that are able to move freely using, for example, wheels. The other main category of robots are stationary robotic arms, also referred to as **robotic manipulators**. Of course, as with any taxonomy, there are always ex- amples that do not fit neatly into any of the available categories. For example, a **smart home** equipped with a central computer and, perhaps, some form of manipulation capabilities, can also be considered a robot, albeit of a different kind.

Robotic manipulators constitute a very important class of robots and they are used extensively in many industries, for example in assembly lines in the vehicle industry. However, such robots normally follow a predefined move- ment sequence and are not equipped with behaviors (such as collision avoid- ance) designed to avoid harming people. While there is nothing preventing the use of, for instance, sonar proximity sensors on a robotic manipulator, such op- tions are rarely used. Instead, manipulators are confined to **robotic work cells**,

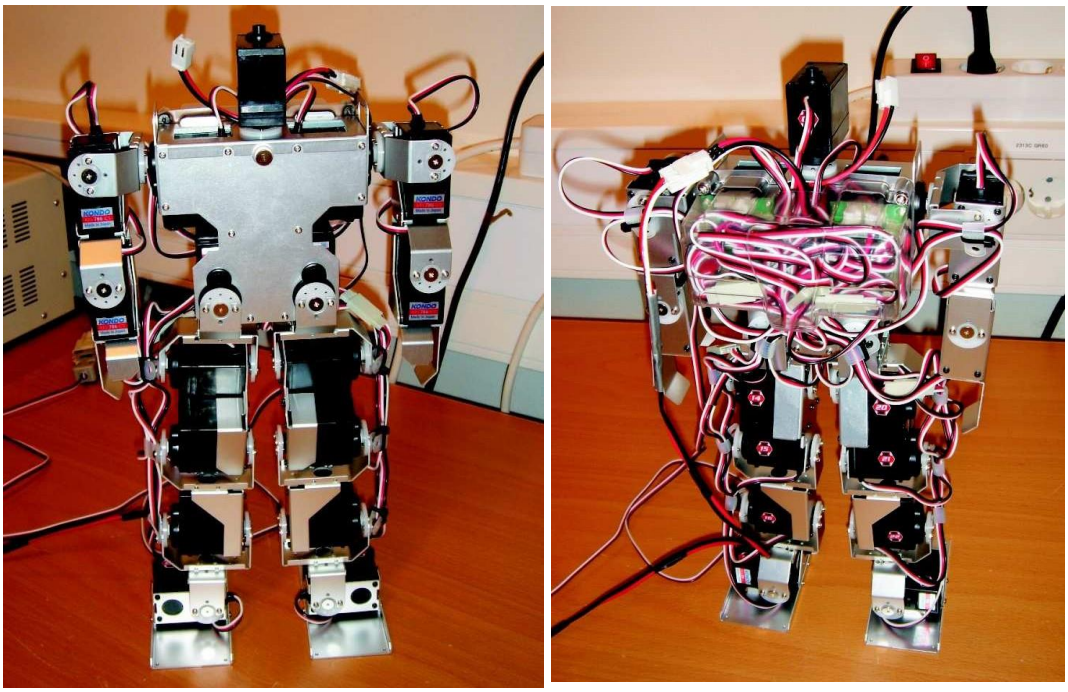

Figure 1.2: *A Kondo humanoid robot. Left panel: Front view. Right panel: Rear view.*

in which people are forbidden to enter while the manipulator is active.

By contrast, in this course, we shall consider **autonomous robots**, i.e. robots that are capable of making their own decisions (depending on the situation at hand) rather than merely executing a pre-defined sequence of motions. In fact, since most robots equipped with such decision-making capabilities are mo- bile, one may define an autonomous robot as a mobile robot with the ability to make decisions. Two examples of mobile robots can be seen in Fig. 1.1. The left panel shows a Boe-bot, which will be assembled and used in the second half of the course. Some of its main advantages are its small size (its length is around 0.14 m and its width 0.11 m) and its simplicity. Needless to say, the robot also has several limitations; for example, its onboard processor (micro- controller) is quite slow. However, on balance, the Boe-bot provides a good introduction to the field of mobile robots. The right panel of Fig. 1.1 shows a two-wheeled differentially steered robot which, although still under construction at the Adaptive systems research group at Chalmers, is already being used in several research projects. This robot has a diameter of 0.40 m and a height of around 1.00 m.

Robotic manipulators have long dominated the market for robots, but with the advent of low-cost mobile robots the situation is changing: In 2007, the number of mobile robots surpassed the number of manipulators for the first time, and the gap is expected to widen over the next decades.

The class of mobile robots can be further divided into subclasses, the most

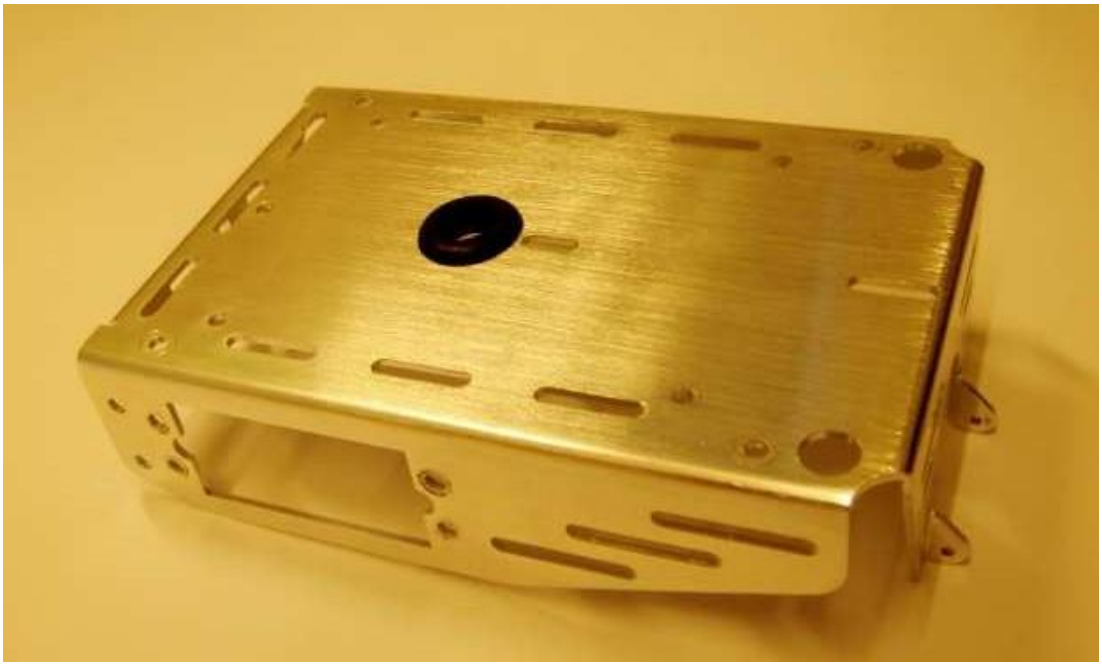

Figure 1.3: *The aluminium frame of a Boe-bot.*

important being **legged robots** and **wheeled robots**. Other kinds, such as fly- ing robots, exist as well, but will not be considered in this course. The class of legged robots can be subdivided based on the number of legs, the most common types being **bipedal robots** (with two legs) and **quadrupedal robots** (with four legs). Most bipedal robots resemble humans, at least to some extent; such robots are referred to as **humanoid robots**. An example of a humanoid robot is shown in Fig. 1.2. Humanoid robots that (unlike the robot shown in Fig. 1.2) not only have the approximate shape of a human, but have also been equipped with more detailed human-like features, e.g. artificial skin, artificial hair etc., are called **androids**. It should be noted that the term *humanoid* refers to the shape of the robot, not its size; in fact, many humanoid robots are quite small. For example, the Kondo robot shown in Fig. 1.2 is approximately 0.35 m tall.

Some robots are *partly* humanoid. For example, the wheeled robot shown in the right panel of Fig. 1.1 is currently being equipped with a humanoid up- per body. Unlike a fully humanoid robot, this robot need not be actively bal- anced, but will still exhibit many desirable features of humanoid robots, such as two arms for grasping and lifting objects, gesturing etc., as well as a head that will be equipped with two cameras for stereo vision and microphones providing capabilities for listening and speaking.

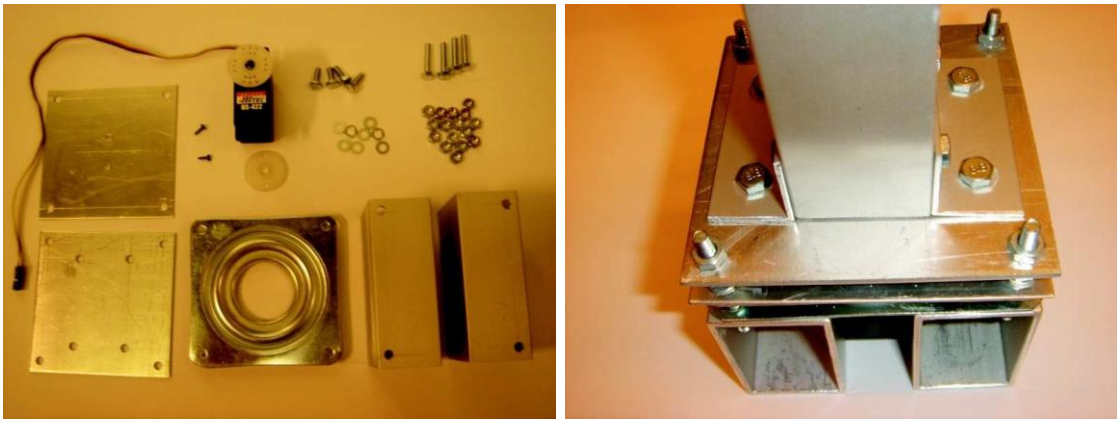

Figure 1.4: *Left panel: Aluminium parts used in the construction of a rotating base for a humanoid upper body. The servo motor used for rotating the base is also shown, as well as the screws, washers and nuts. Right panel: The assembled base.*

1.4 Robotic hardware

1.4.1 Construction material

Regarding the material used in the actual frame of the robot, several options are available, such as e.g. aluminium, steel, various forms of plastic etc. The frame of a robot should, of course, preferably be constructed using a material that is both sturdy and light and, for that reason, aluminium is often chosen. Albeit somewhat expensive, aluminium combines toughness with low weight in a near-optimal way, at least for small mobile robots. Steel is typically too heavy to be practical in a small robot, whereas many forms of plastic eas- ily break. The frame of the robot used in this course (the Boe-bot) is made in aluminium, and is shown in Fig. 1.3. The left panel of Fig 1.4 shows the aluminium parts used in a rotating base for a humanoid upper body. The as- sembled base, which can rotate around the vertical axis, is shown in the right panel.

1.4.2 Sensors

The purpose of robotic sensors is to measure either some physical characteris- tic of the robot (for example, its acceleration) or some aspect of its environment (for example, the detected intensity of a light source). The raw data thus ob- tained must then, in most cases, be processed further before being used in the brain of the robot. For example, an infrared (IR) proximity sensor may pro- vide a voltage (depending on the distance to the detected object) as its read- ing, which can then be converted to a distance, using the characteristics of the sensor available from its data sheet.

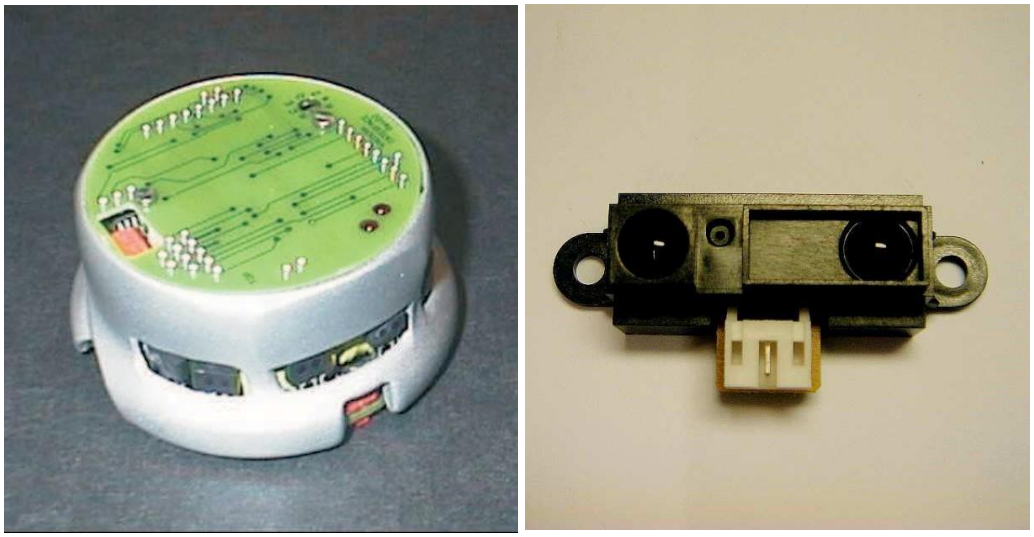

Figure 1.5: *Left panel: A Khepera II robot. Note the IR proximity sensors (small black rectangles around the periphery of the robot), consisting of an emitter and a detector. Right panel: A Sharp GP2D12 infrared sensor.*

Needless to say, there exists a great variety of sensors for mobile robots. Here, only a brief introduction will be given, focusing on a few fundamental sensor types.

Infrared proximity sensors

An **infrared proximity sensor** (or **IR sensor**, for short), consists of an emitter and a detector. The emitter, a light-emitting diode (LED), sends out infrared light, which bounces off nearby objects, and the reflected light is then mea- sured by the detector (e.g. a phototransistor). Some IR sensors can also be used for measuring the ambient light level, i.e. the light observed by the detector when the emitter is switched off. As an example, consider the Khepera robot (manufactured by K-Team, www.k-team.com), shown in the left panel Fig. 1.5. This robot is equipped with eight IR sensors, capable of measuring both am- bient and reflected light. The range of IR sensors is quite short, though. In the Khepera robot, reflected light measurements are only useful to a distance of around 0.050 m from the robot, i.e. approximately one robot diameter, even though other IR sensors have longer range. Another example is the Sharp GP2D12 IR sensor, shown in the right panel of Fig. 1.5. This sensor detects ob- jects in the range $[0.10, 0.80]$ m. It operates using a form of triangulation: Light is emitted from the sensor and, if an object is detected, the reflected light is re- ceived at an angle that depends on the distance to the detected object. The raw signal from the sensor consists of a voltage that can be mapped to a distance. The mapping is non-linear, and for very short distances, the sensor cannot give

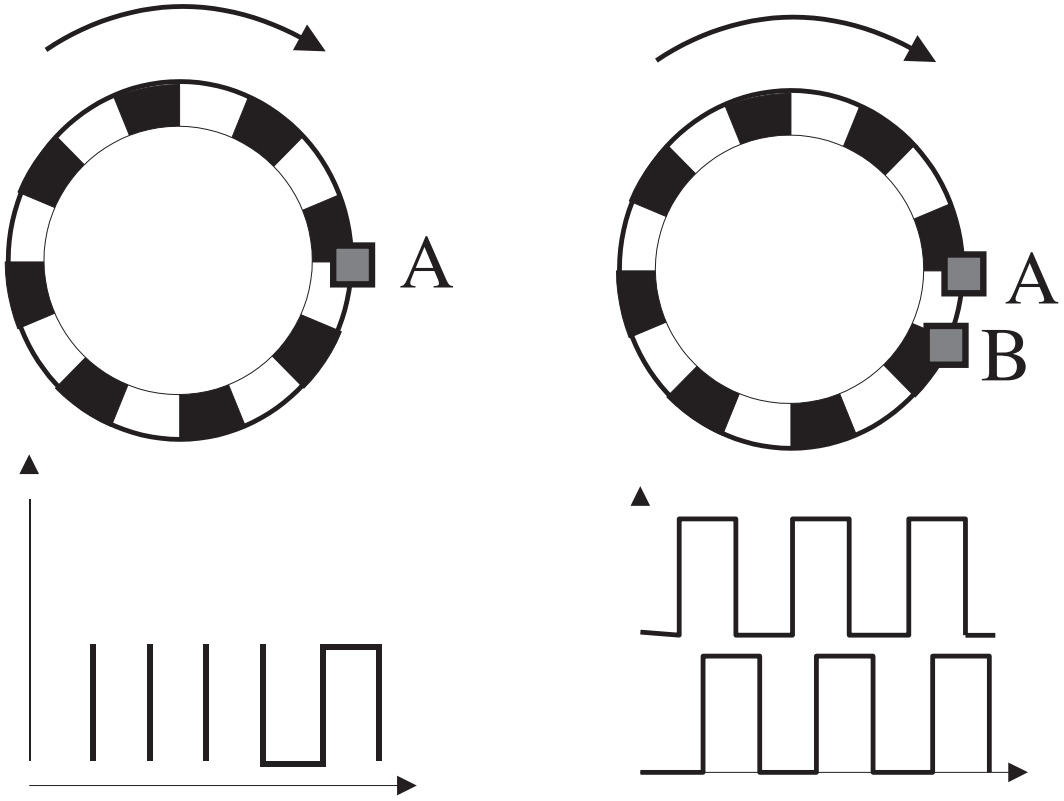

Figure 1.6: *The left panel shows a simple encoder, with a single detector (A), that measures the interruptions of a light beam, producing the curve shown below the encoder. In the right panel, two detectors are used, making it possible to determine also the* direction *of rotation.*

reliable readings (hence the lower limit of 0.10 m).

Digital optical encoders

In many applications, accurate position information is essential for a robot, and there are many different methods for positioning, e.g. inertial navigation, GPS navigation, landmark detection etc., some of which will be considered in a later chapter. One of the simplest forms of positioning, however, is **dead reck- oning**, in which the position of a robot is determined based on measurements of the distance travelled by each wheel of the robot. This information, when combined with knowledge of the robot's physical properties (i.e. its kinemat- ics, see Chapter 2) allows one to deduce the current position and heading. The process of measuring the rotation of the wheel of a robot is an example of **odometry**, and a sensor capable of such measurements is the **digital optical encoder** or, simply, **encoder**. Essentially, an encoder is a disc made of glass or plastic, with shaded regions that regularly interrupt a light beam. By count- ing the number of interruptions, the rotation of the wheel can be deduced, as shown in the left panel of Fig. 1.6. However, in order to determine also the *di- rection* of rotation, a second detector, placed at a quarter of a cycle out of phase

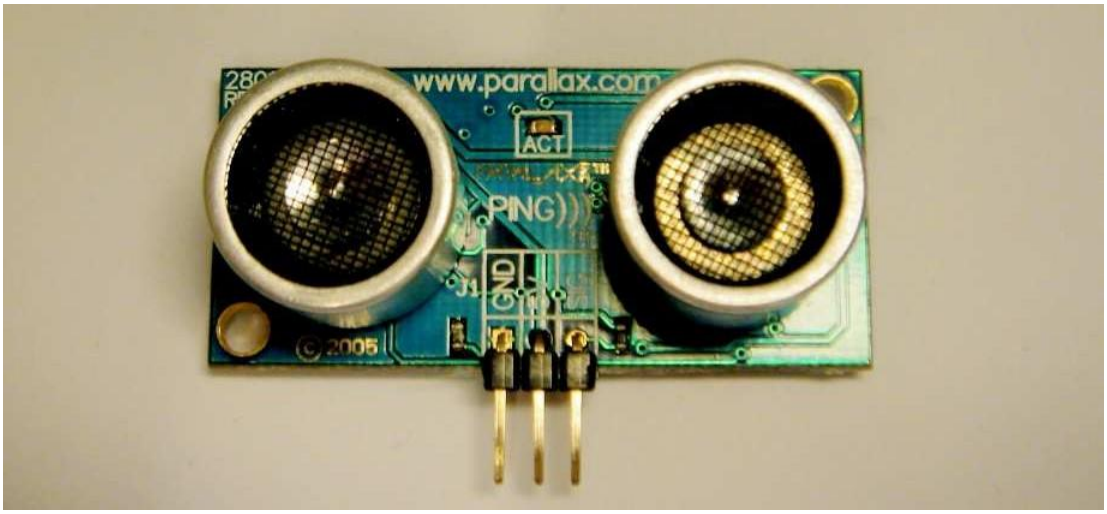

Figure 1.7: *A Ping ultrasonic distance sensor.*

with the first detector, is needed (such an arrangement is called **quadrature encoding**, and is shown in the right panel of Fig. 1.6).

Ultrasound (sonar) sensors

Ultrasound sensors, also known as **sonar sensors** or simply **sonars**, are based on time-of-flight measurement. Thus, in order to detect the distance to an ob- ject, a sonar emits a brief pulse of ultrasonic sound, typically in the frequency range 40-50 kHz[2]. The sensor then awaits the echo. Once the echo has been detected, the distance to the object can be obtained using the fact that sound travels at a speed of around 340 m/s. As in the case of IR sensors, there is both a lower and an upper limit for the detection range of a sonar sensor. If the distance to an object is too small, the sensor simply does not have enough time to switch from emission to listening, and the signal is lost. Similarly, if the distance is too large, the echo may be too weak to be detected.

Fig. 1.7 shows a Ping ultrasonic distance sensor, which is commonly used in connection with the Boe-bot. This sensor can detect distances to objects in the range $[0.02, 3.00]$ m.

Laser range finders

Laser range finders (LRFs) commonly rely, like sonar sensors, on time-of-flight measurements, but involve the speed of light rather than the speed of sound. Thus, a laser range finder emits pulses of laser light (in the form of thin beams),

[2]For comparison, a human ear can detect sounds in the range 20 hz to 20 kHz. Thus, the sound pulse emitted by a sonar sensor is not audible.

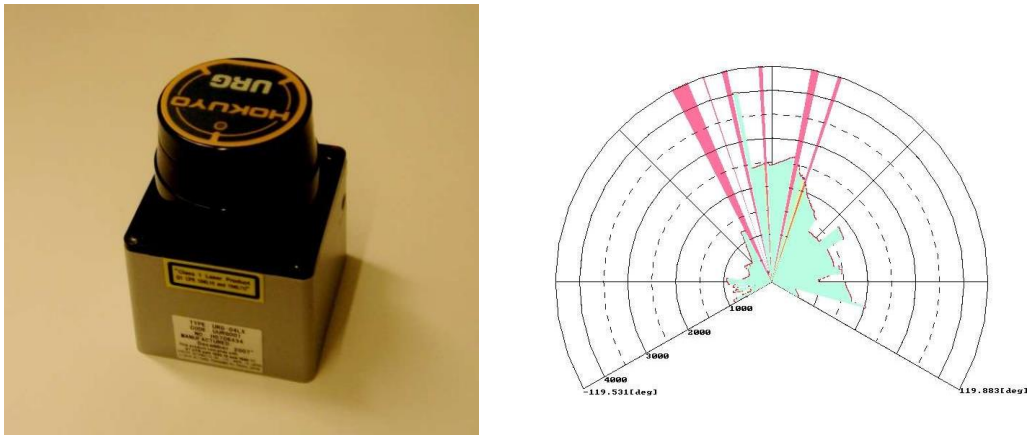

Figure 1.8: *Left panel: A Hokuyo URL-04LX laser range finder. Right panel: A typical reading, showing the distance to the nearest object in various directions. The pink rays indicate directions in which no detection is made. The maximum range of the sensor is 4 m.*

and measures the time it takes for the pulse to bounce off a target and return to the range finder. An LRF carries out a sweep over many directions[3] resulting in an accurate local map of distances to objects along the line-of-sight of each ray. LRFs are generally very accurate sensors, but they are also much more expensive than sonars sensors and IR sensors.

A Hokuyo URG-04LX LRF is shown in the left panel of Fig. 1.8. This sensor has a range of around four meters, with an accuracy of around 1 mm. It can generate readings in 683 different directions, with a frequency of around 10 Hz. As of the time of writing (Jan. 2010), a Hokuyo URG-04LX costs around 2,000 USD. The right panel of Fig. 1.8 shows a typical reading, obtained from the software delivered with the LRF.

Cameras

Cameras are used as the eyes of a robot. In many cases, two cameras are used, in order to provide the robot with binocular vision, allowing it to estimate the range to detected objects. There are many cameras available for robots, for example the CMUCam series which has been developed especially for use in mobile robots; The processor connected to the CMUCam is capable of basic image processing. At the time of writing (Jan. 2010), a CMUCam costs on the order of 150 USD. A low-cost alternative is to use ordinary webcams, for which prices start around 15 USD. Fig. 1.9 shows a simple robotic head consisting of two servo motors (see below) and a single webcam.

However, while the actual cameras may not be very costly, the use of cam- eras is *computationally* a very expensive procedure. Even at a low resolution,

[3] A typical angular interval for an LRF is around 180-240 degrees.

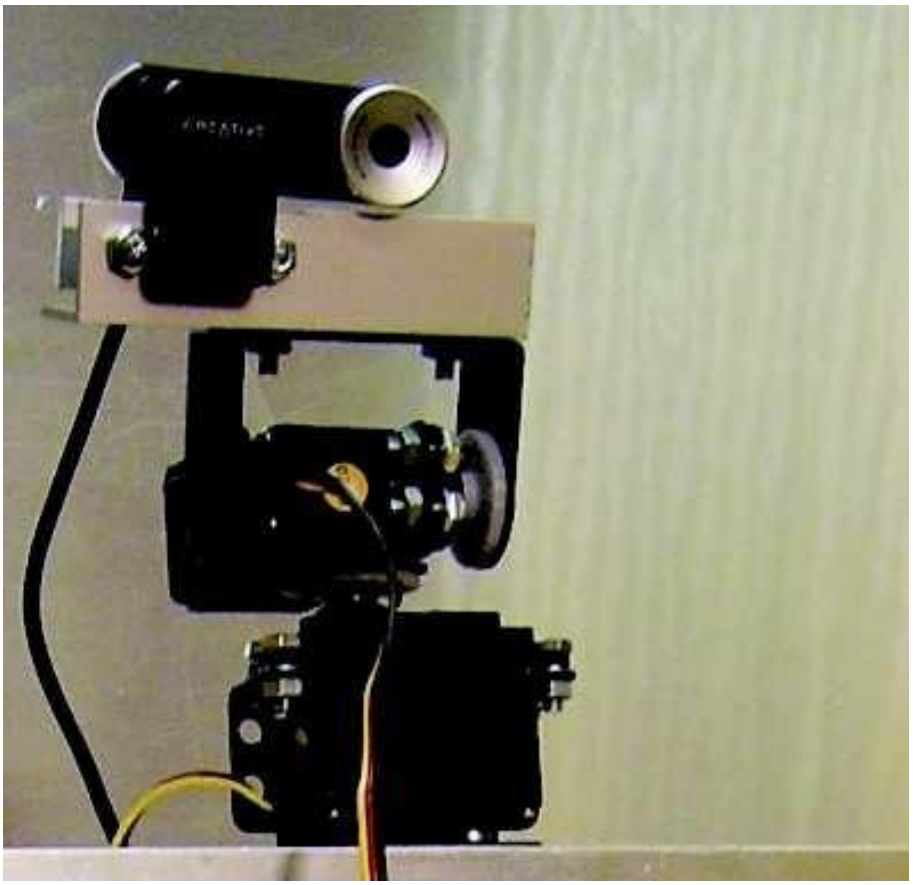

Figure 1.9: *A simple robotic head, consisting of two servo motors and a webcam.*

say 320×240 pixels, a webcam will deliver a flow of around 1.5 Mb/s, assum- ing a frame rate of 20 Hz and a single byte of data per pixel. The actual data transfer is easily handled by a Universal serial bus (USB), but the data must not only be transferred but also *analyzed*, something which is far from trivial. An introduction to image processing for robots will be given in a later chapter.

Other sensors

In addition to odometry based on digital optical encoders, robot positioning can be based on **inertial sensors**, i.e. sensors that measure the time derivatives of the position or heading angle of the robot. Examples of inertial sensors are **accelerometers**, measuring linear acceleration, and **gyroscopes**, measuring an- gular acceleration. Essentially, an accelerometer consists of a small object, with mass m, attached to a spring and damper, as shown in Fig. 1.10. As the system accelerates, the displacement z of the small object can be used to deduce the acceleration $\ddot{x}$ of the robot. Given continuous measurements of the accelera- tion, as a function of time, the position (relative to the starting position) can be

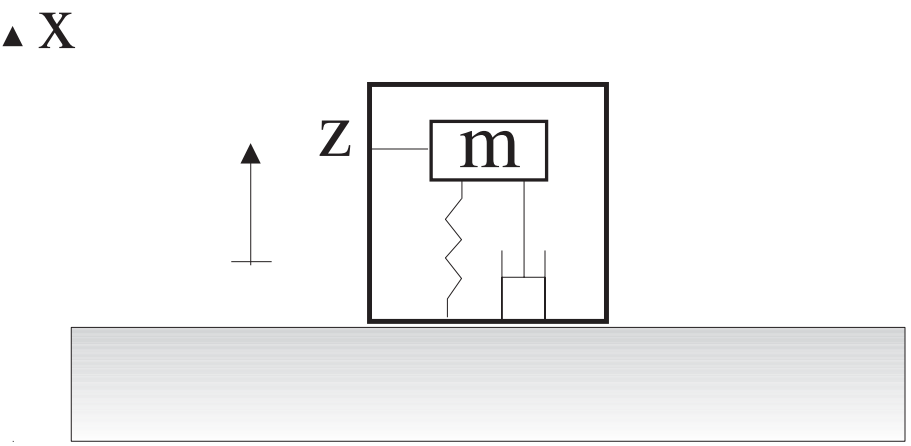

Figure 1.10: *An accelerometer. The motion of the small object (mass m) resulting from the acceleration of the larger object to which the accelerometer is attached can be used for deducing the acceleration.*

deduced. For robots operating in outdoor environments, positioning based on the **global positioning system (GPS)** is often a good alternative. The GPS re- lies on 24 satellites that transmit radio frequency signals which can be picked up by objects on Earth. Given the exact position of (at least) three satellites, rel- ative to the position of e.g. a robot, the absolute position (latitude, longitude, and altitude) of the robot can be deduced.

Other sensors include **strain gauge sensors** (measuring deformation), **tac- tile (touch) sensors** measuring physical contact between a robot and objects in its environment, and **compasses**, measuring the direction of movement.

1.4.3 Actuators

An **actuator** is a device that allows a robot to take action, i.e. to move or manip- ulate the surroundings in some other way. Motors, of course, are very common types of actuators. Other kinds of actuation include, for example, the use of microphones (for human-robot interaction).

Movements can be generated in various ways, using e.g. electrical motors, pneumatic or hydraulic systems etc. In this course, we shall only consider electrical, direct-current (DC) motors and, in particular, servo motors. Thus, when referring to actuation in this course, the use of such motors is implied.

Note that actuation normally requires the use of a **motor controller** in con- nection with the actual motor. This is so, since the microcontroller (see below) responsible for sending commands to the motor cannot, in general, provide sufficient current to drive the motor. The issue of motor control will be consid- ered briefly in connection with the discussion of servo motors below.

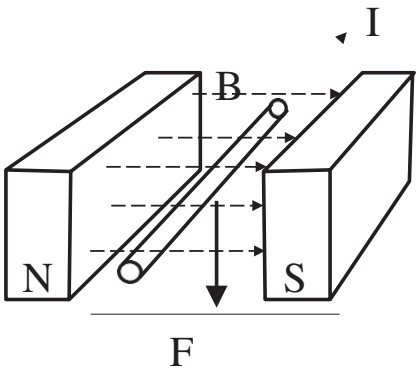

Figure 1.11: *A conducting wire in a magnetic field. B denotes the magnetic field strength and I the current through the wire. The Lorentz force $\mathbf{F}$ acting on the wire is given by $\mathbf{F} = \mathbf{I} \times \mathbf{B}$.*

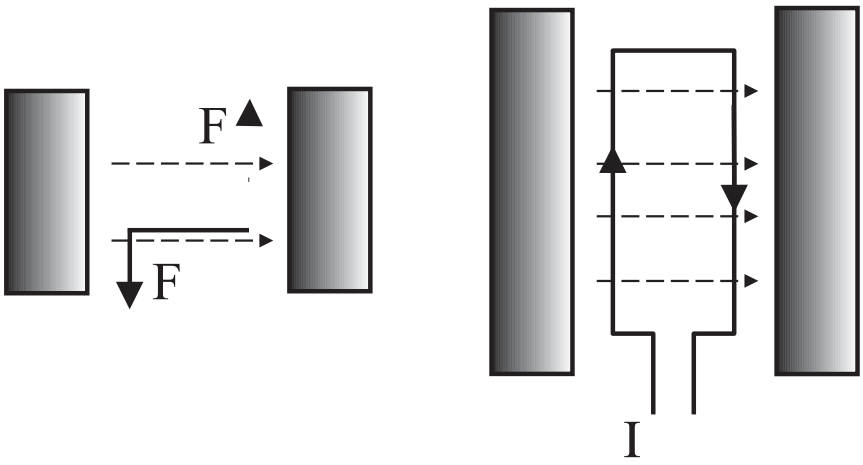

Figure 1.12: *A conducting loop of wire placed in a magnetic field. Due to the forces acting on the loop, it will begin to turn. The loop is shown from above in the right panel, and from the side in the left panel.*

DC motors

Electrical direct current (DC) motors are based on the principle that a force acts on a wire in a magnetic field if a current is passed through the wire, as illustrated in Fig. 1.11. If instead a current is passed through a closed loop of wire, as illustrated in Fig. 1.12, the forces acting on the two sides of the loop will point in opposite directions, making the loop turn. A standard DC motor consists of an outer stationary cylinder (the **stator**), providing the magnetic field, and an inner, rotating part (the **rotor**). From Fig. 1.12 it is clear that the loop will reverse its direction of rotation after a half-turn, unless the direction of the current is reversed. The role of the **commutator**, connected to the rotor of a DC motor, is to reverse the current through the motor every half-turn, thus allowing continuous rotation. Finally, carbon **brushes**, attached to the stator, complete the electric circuit of the DC motor. There are types of DC motors

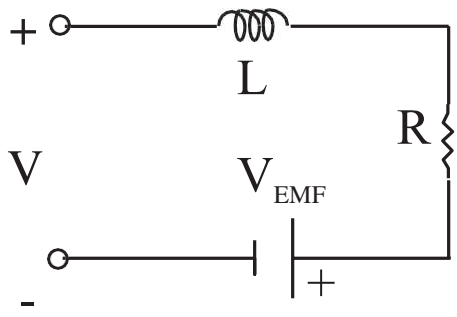

Figure 1.13: *The equivalent electrical circuit for a DC motor.*

that use electromagnets rather than a permanent magnet, and also types that are brushless. However, a detailed description of such motors are beyond the scope of this text.

DC motors are controlled by varying the applied voltage. The equations for DC motors can be divided into an electrical and a mechanical part. The motor can be modelled electrically by the equivalent circuit shown in Fig. 1.13. Letting V denote the applied voltage, and ω the angular speed of the motor shaft, the electrical equation takes the form

$$V = L\frac{\mathrm{d}i}{\mathrm{d}t} + Ri + V_{\text{EMF}}, \tag{1.1}$$

where i is the current flowing through the circuit, L is the inductance of the motor, R its resistance, and V_{EMF} the voltage (the **back EMF**) counteracting V. The back EMF depends on the angular velocity, and can be written as

$$V_{\text{EMF}} = c_e\omega, \tag{1.2}$$

where c_e is the **electrical constant** of the motor. For a DC motor, the generated torque τ_g is directly proportional to the current, i.e.

$$\tau_g = c_t i, \tag{1.3}$$

where c_t is the **torque constant** of the motor. Turning now to the mechanical equation, Newton's second law gives

$$I\frac{\mathrm{d}\omega}{\mathrm{d}t} = \sum \tau, \tag{1.4}$$

where I is the combined moment of inertia of the motor and its load, and $\sum\tau$ is the total torque acting on the motor. For the DC motor, the equation takes the form

$$\frac{\mathrm{d}\omega}{}$$

$$I\frac{\mathrm{d}^2\theta}{\mathrm{d}t} = \tau_g - \tau_f - \tau, \tag{1.5}$$

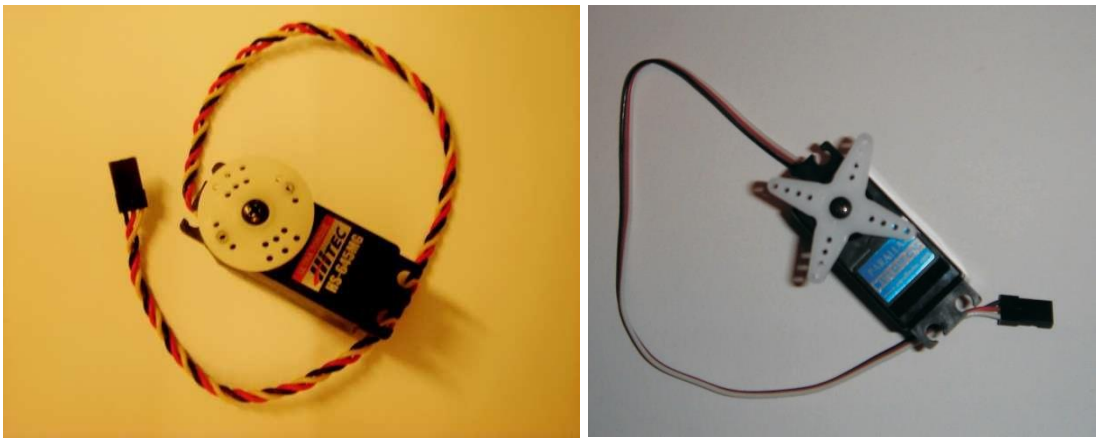

Figure 1.14: *Left panel: A HiTec 645MG servo. The suffix* MG *indicates that the servo is equipped with a metal gear train. Right panel: A Parallax servo, which has been modified for continuous rotation. Servos of this kind are used on the Boe-bot. The circular (left) and star-shaped (right) white plastic objects are the servo horns.*

where τ_f is the frictional torque opposing the motion and τ is the (output) torque acting on the load. The frictional torque can be divided into two parts, the **Coulomb friction** ($c_C \mathrm{sgn}(\omega)$) and the **viscous friction** ($c_v\omega$). Thus, the equations for the DC motor can now be summarized as

$$\tau_g = \frac{c_t}{R} V - \frac{c_t L}{R} \frac{\mathrm{d}i}{\mathrm{d}t} - \frac{c_e c_t}{R} \omega, \tag{1.6}$$

$$I \frac{\mathrm{d}\omega}{\mathrm{d}t} = \tau_g - c_C \mathrm{sgn}(\omega) - c_v \omega - \tau, \tag{1.7}$$

In many cases, the time constant of the electrical circuit is much shorter than that of the physical motion, so the inductance term can be neglected. Further- more, for simplicity, the dynamics of the mechanical part can *also* be neglected under certain circumstances (e.g. if the moment of inertia of the motor and load is small). Thus, setting $\mathrm{d}i/\mathrm{d}t$ and $\mathrm{d}\omega/\mathrm{d}t$ to zero, the steady-state DC mo- tor equations, determining the torque τ on the load for a given applied voltage V and a given angular velocity ω

$$\tau_g = \frac{c_t}{R} V - \frac{c_e c_t}{R} \omega, \tag{1.8}$$

$$\tau = \tau_g - c_C \mathrm{sgn}(\omega) - c_v \omega, \tag{1.9}$$

are obtained. In many cases, the axis of a DC motor rotates too fast and gener- ates a torque that is too weak for driving a robot. Thus, a **gear box** is commonly used, which reduces the rotation speed taken out from the motor (on the **sec- ondary drive shaft**) while, at the same time, increasing the torque. For an ideal (loss-free) gear box, the output torque and rotation speed are given by

$$\tau_{\mathrm{out}} = G\tau,$$

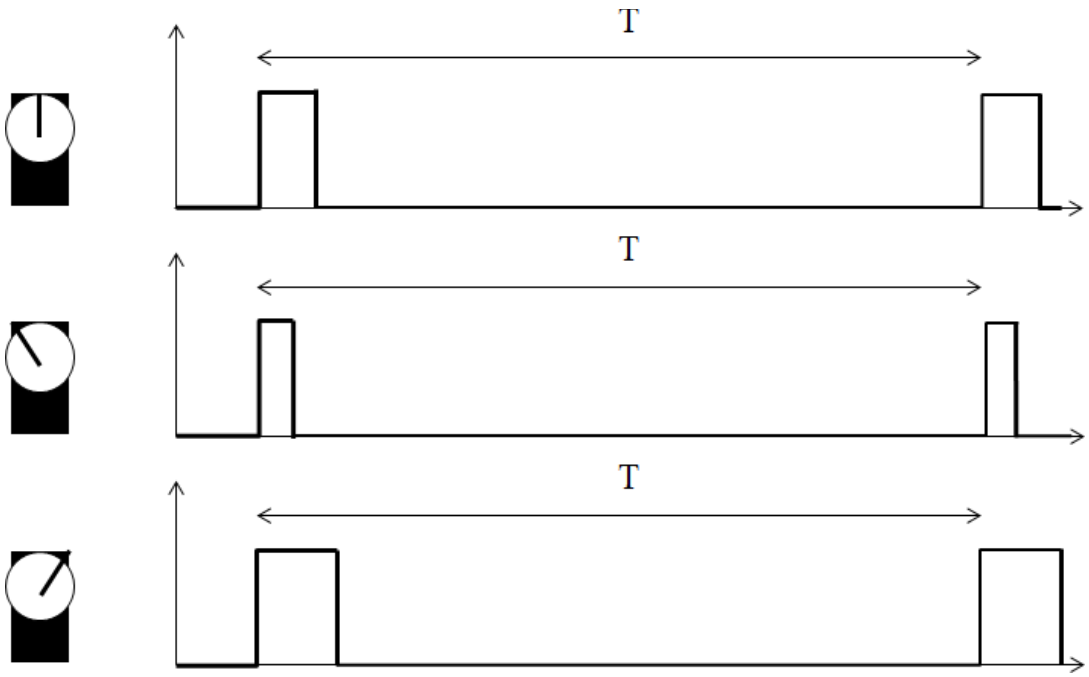

Figure 1.15: *Pulse width modulation control of a servo motor. The lengths of the pulses determine the requested position angle of the motor output shaft. The interval betwwn pulses (typically around 20 ms) is denoted* T *.*

$$\omega_{\text{out}} = \frac{1}{G}\omega, \tag{1.10}$$

where G is the **gear ratio**.

Servo motors

A **servo motor** is essentially a DC motor equipped with control electronics and a gear train (whose purpose is to increase the torque to the required level for moving the robot, as described above). The actual motor, the gear train, and the control electronics, are housed in a plastic container. A **servo horn** (either plastic or metal) makes it possible to connect the servo motor to a wheel or some other structure. Fig. 1.14 shows two examples of servo motors.

The angular position of a servo motor's output shaft is determined using a **potentiometer**. In a standard servo, the angle is constrained to a given range [α_{max}, α_{max}], and the role of the control electronics is to make sure that the servo rotates to a set position α (given by the user). A servo is fitted with a three-wire cable. One wire connects the servo to a power source (for exam- ple, a motor controller or, in some cases, a microcontroller board) and another wire connects it to *ground*. The third wire is responsible for sending signals to the servo motor. In servo motors, a technique called **pulse width modulation**

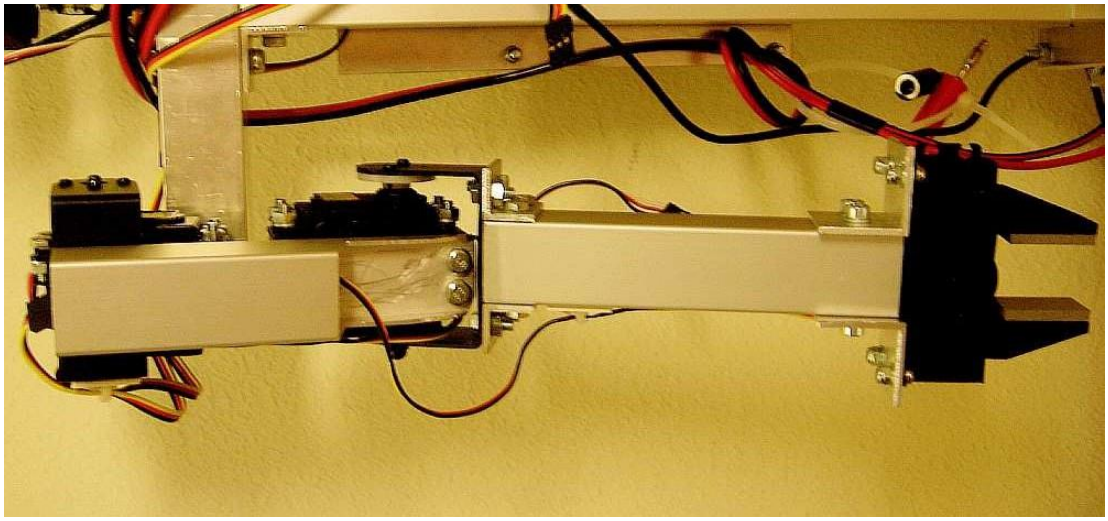

Figure 1.16: *An arm of a humanoid robot. The allowed rotation range of the elbow is around 100 degrees.*

(PWM) is used: Signals in the form of pulses are sent (e.g. from a microcon- troller) to the control electronics of the servo motor. The duration of the pulses determine the required position, to which the servo will (attempt to) rotate, as shown in Fig. 1.15. For a walking robot (or for a humanoid upper body), the limitation to a given angular range poses no problem: The allowed rotation range of a servo is normally sufficient for, say, an elbow or a shoulder joint. As an example, an arm of a humanoid robot is shown in Fig. 1.16. For this particu- lar robot, the rotation range for the elbow joint is around 100 degrees, i.e. easily within the range of a standard servo (around 180 degrees). The limitation is, of course, not very suitable for motors driving the wheels of a robot. Fortunately, servo motors can be modified to allow continuous rotation. The Boe-bot that will be built in the second half of the course uses **Parallax continuous rotation servos** (see the right panel of Fig. 1.14), rather than standard servos.

Other motors

There are many different types of motors, in addition to standard DC motors and servo motors. An example is the **stepper motor**, which is also a version of the DC motor, namely one that moves in fixed angular increments, as the name implies. However, in this course, only standard DC motors and servo motors will be considered.

1.4.4 Processors

Sensors and actuators are necessary for a robot to be able to perceive its envi- ronment and to move or manipulate the environment in various ways. How-

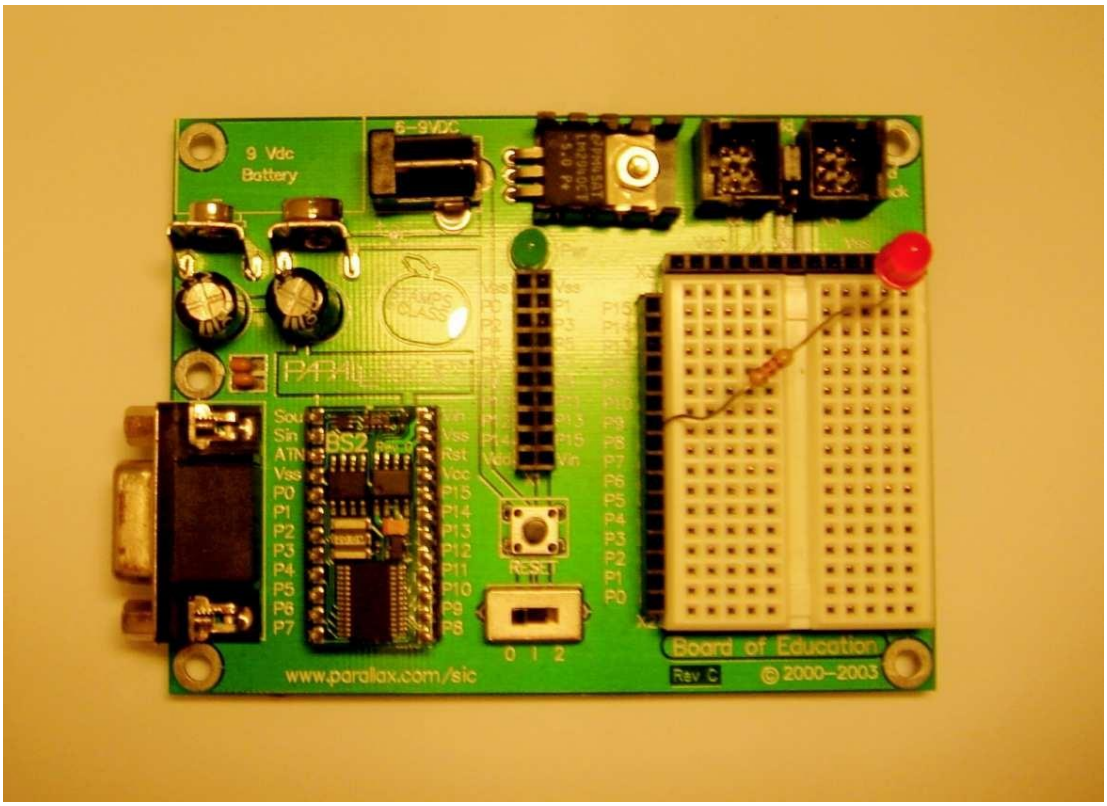

Figure 1.17: *A Board of Education (BOE) microcontroller board, with a Basic Stamp II (BS2) microcontroller attached. In addition to the microcontroller, the BOE has a serial port for communication with a PC (used, for example, when uploading a program onto the BS2), as well as sockets for attaching sensors and electronic circuits. In this case, a simple circuit involving a single LED, has been built on the BOE. The two black sockets in the upper right corner are used for connecting up to four servo motors.*

ever, in addition to sensors and actuators, there must also be a system for an- alyzing the sensory information, making decisions concerning what actions to take, and sending the necessary signals to the actuators.

In autonomous robots, it is common to use several processors to represent the brain of the robot. Typically, high-level tasks, such as decision-making, are carried out on a standard PC, for example a laptop computer mounted on the robot, whereas low-level tasks are carried out by microcontrollers, which will now be introduced briefly.

Microcontrollers

Essentially, a **microcontroller** is a single-chip computer, containing a **central processing unit** (CPU), read-only memory (ROM, for storing programs), random-access memory (RAM, for temporary storage, such as program variables), and

several input-output (I/O) ports. There exist many different microcontrollers, with varying degrees of complexity, and different price levels, down to a few USD for the simplest ones. An example is the **Basic Stamp II**[4] (BS2) microcon- troller, which costs around 50 USD.

While the BS2 is sufficient for the experimental work carried out in this course (in the next quarter), its speed is only around 4,000 operations per sec- ond (op/s) and it has a RAM memory (for program variables) of only 32 bytes and a ROM (for program storage) of 2 kilobytes (Kb).

However, many alternative microcontrollers are available for more advanced robots. Two examples, with roughly the same price as the BS2, are the BasicX and ZBasic microcontrollers, which are both compatible with the BOE micro- controller board used together with the BS2. The BasicX microcontroller has a RAM memory of 400 bytes and 32 Kb for program storage, whereas ZBasic has 4 Kb of RAM and 62 Kb for program storage. BasicX executes around 83,000 op/s, whereas (some versions of) ZBasic can reach up to 2.9 million op/s.

In many cases, microcontrollers are sold together with **microcontroller boards** (or **microcontroller modules**), containing sockets for wires connecting the mi- crocontroller to sensors and actuators as well as control electronics, power sup- ply etc. An example is the **Board of education** (BOE) microcontroller board.

The BOE, shown in Fig. 1.17, is equipped with a **solderless breadboard**, on which electronic circuits can be built without any soldering, which is very use- ful for prototyping.

Since microcontrollers do not have human-friendly interfaces such as a keyboard and a screen, the normal operating procedure is to write and compile programs on an ordinary computer (using, of course, a compiler adapted for the microcontroller in question), and then upload the programs onto the mi- crocontroller. In the case of the BS2 microcontroller, the language is a version of Basic called **PBasic**.

Robotic brain architectures

An autonomous robot must be capable of both high-level and low-level pro- cessing. The low-level processing consists, for example, of sending signals to motor controllers (see below) which, in turn, send (for example) PWM pulses to servo motors. Another low-level task is to retrieve raw data (e.g. a voltage value from an IR proximity sensor). The distinction between low-level and high-level tasks is a bit fuzzy. For example, the voltage value from an IR sen- sor (e.g. the Sharp GP2D12 mentioned above) can be mapped to a distance value, which of course normally is more relevant for decision-making than the raw voltage value. The actual conversion would normally be considered a low-level task but might as well also be carried out on the robot's onboard PC.

[4] Basic Stamp is a registered trademark of Parallax, inc., see www.parallax.com.

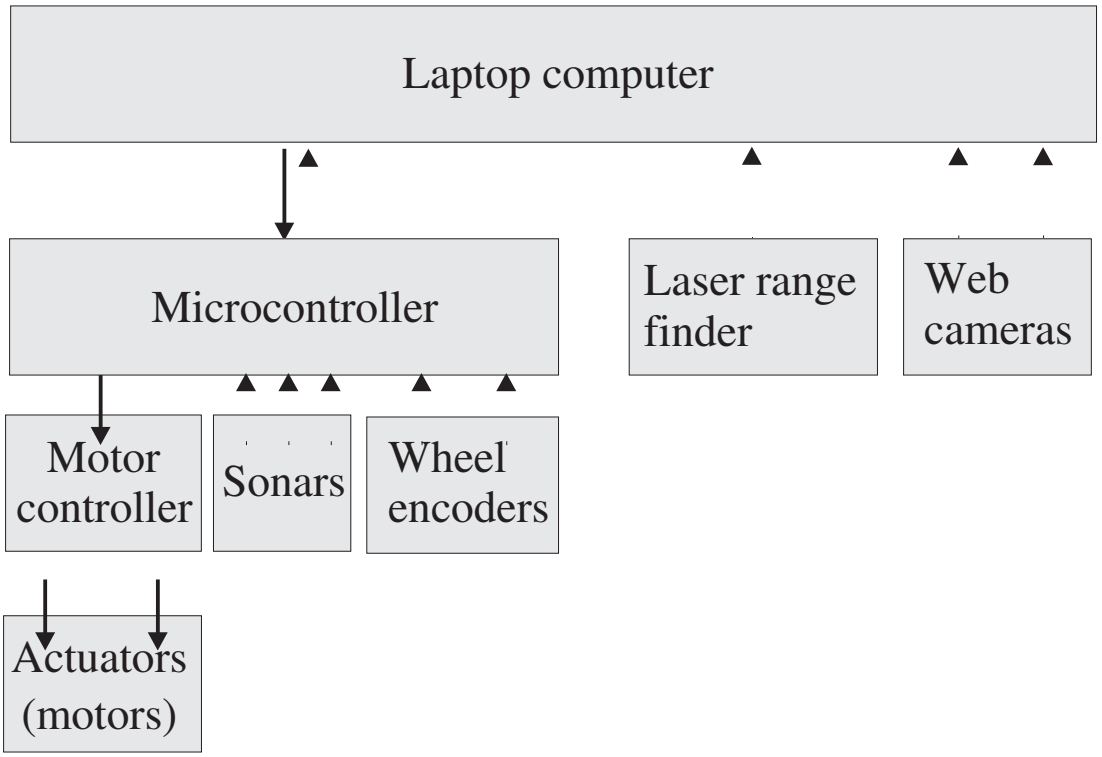

Figure 1.18: *An example of a typical robotic brain architecture, for a differentially steered two-wheeled robot equipped with wheel encoders, three sonar sensors, one LRF, and two web cameras.*

The hardware configuration providing a robot's processing capability is re- ferred to as the **robotic brain architecture**. An example of a typical robotic brain architecture is shown in Fig. 1.18. The robotic brain shown in the figure would be used in connection with a two-wheeled differentially steered robot. As can be seen in the figure, the microcontroller would handle low-level pro- cessing, such as measuring the pulse counts of the wheel encoders, collecting readings from the three sonars, and sending motor signals (e.g. desired set speeds) to the **motor controller**[5], which, in turn, would send signals to the motors. However, the LRF and the web cameras would be directly connected, via USB (or, possibly, serial) ports, to the main processor (on the laptop), since most microcontrollers would not be able to handle the massive data flow from such sensors.

The main program (i.e. the robotic brain), running on the laptop, would process the data from the various sensors. For example, the pulse counts from

[5]A separate motor controller (equipped with its own power source) is often used for robotics applications, since the power source for the microcontroller may not be able to de- liver sufficient current for driving the motors as well.

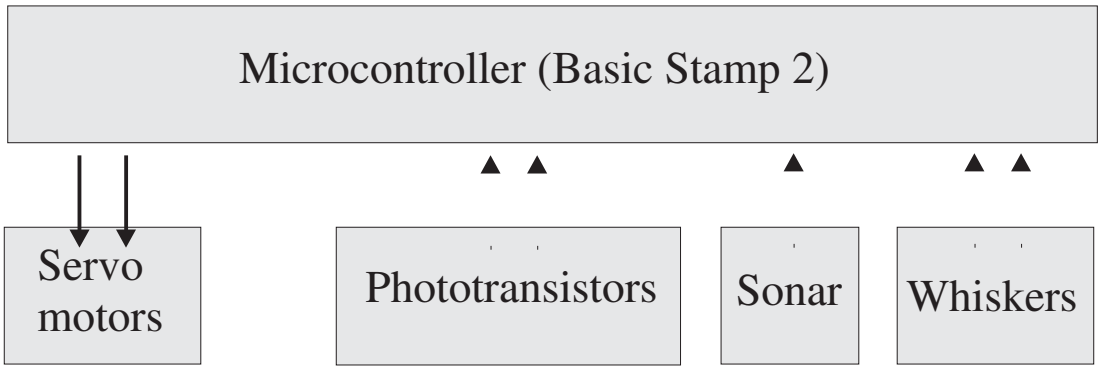

Figure 1.19: *An example of a robotic brain architecture for a Boe-bot.*

the wheel encoders would be translated to an estimate of position and head- ing, as described in Chapter 2. Given the processed sensory data, as well as in- formation stored in the (long-term or short-term) memory of the robotic brain (for example, a map of the arena in which the robot operates), the main pro- gram would determine the next action to be carried out by the robot, compute the appropiate motor commands and send them to the microcontroller.

Note that the figure only shows an *example*: Many other configurations could be used as well. For example, there are cameras developed specifically for robotics applications that, unlike standard web cameras, are able to carry out much of the relevant image processing (e.g. detecting and extracting faces), and then only sending *that* information (rather than the raw pixel values) to the laptop computer.

The robotic brain architecture shown in Fig. 1.18 would be appropriate for a rather complex (and costly!) robot. Such robots are beyond the scope of the experimental work carried out in the second half of this course. The ex- perimental work, which will be carried out using a Boe-bot (see the left panel of Fig. 1.1), involves a much simpler robotic brain architecture, illustrated in Fig. 1.19. As can be seen, in this case, the robot has a single processor, namely the BS2 microcontroller, which thus is responsible both for the low-level (sig- nal) processing and the high-level decision-making.

The microcontroller sends signals to the two servo motors and receives in- put from the sensors attached to the robot, for example, two photo-resistors, a sonar sensor, and whiskers. The whiskers are simple touch sensors that give a reading of either 0 (if no object is touched) or 1 (if the whisker touches an object). Of course, other sensors (such as IR sensors or simple wheel encoders) can be added as well, but one should keep in mind that the processing capabil- ity of the BS2 is very limited. Note that no motor controller is used: The BOE is capable of generating sufficient current for up to four Parallax servo motors.

2 Kinematics and dynamics

2.1 Kinematics

Kinematics is the process of determining the range of possible movements for a robot, without consideration of the forces acting on the robot, but tak- ing into account the various constraints on the motion. The kinematic equa- tions for a robot depend on the robot's structure, i.e. the number of wheels, the type of wheels used etc. Here, only the case of differentially steered two- wheeled robots will be considered. For balance, a two-wheeled robot must also have one or several supporting wheels (or some other form of ground contact, such as a ball in a material with low friction). The influence of the supporting wheels on the kinematics and dynamics will not be considered.

2.1.1 The differential drive

A schematic view of a differentially steered robot is shown in Fig. 2.1. The Boe- bot that will be considered in the second half of the course (see the left panel

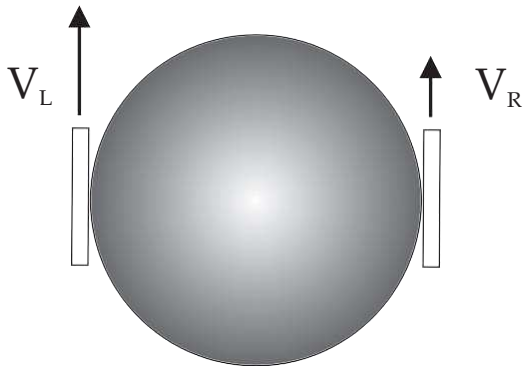

Figure 2.1: *A schematic representation of a two-wheeled, differentially steered robot.*

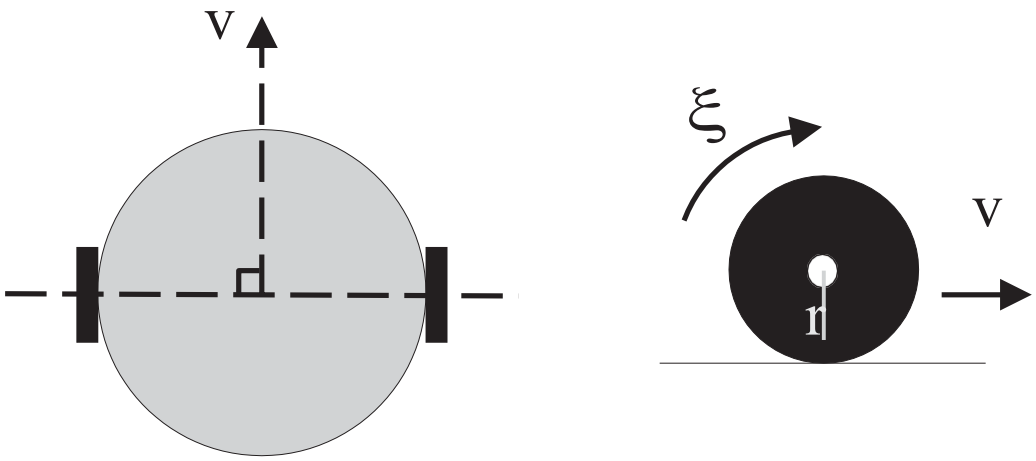

Figure 2.2: *Left panel: Kinematical constraints force a differentially steered robot to move in a direction perpendicular to a line through the wheel axes. Right panel: For a wheel that rolls without slipping, the equation $v = \omega r$ holds.*

of Fig. 1.1) is an example of such a robot.

A differentially steered robot is equipped with two independently steered wheels. The position of the robot is given by the two coordinates x and y, and its direction of motion is denoted ϕ.

It will be assumed that the wheels are only capable of moving in the direc- tion perpendicular to the wheel axis (see the left panel of Fig. 2.2). Furthermore, it will be assumed that the wheels roll without slipping, as illustrated in the right panel of Fig. 2.2. For such motion, the forward speed v of the wheel is related to the angular velocity ω through the equation

$$v = \omega r, \tag{2.1}$$

where r is the radius of the wheel.

The **forward kinematics** of the robot, i.e. the variation of x, y and ϕ, given the speeds v_L and v_R of the left and right wheel, respectively, can be obtained by using the constraints on the motion imposed by the fact that the frame of the robot is a rigid body. For any values of the wheel speeds, the motion of the robot can be seen as a pure rotation, with angular velocity $\omega = \dot{\phi}$ around the **instantaneous center of rotation** (ICR). Letting L denote the distance from the ICR to the center of the robot, the speeds of the left and right wheels can be written

$$v_L = \omega(L - R), \tag{2.2}$$

and

$$v_R = \omega(L + R), \tag{2.3}$$

where R is the radius of the robot's body (which is assumed to be circular and with a circularly symmetric mass distribution). The speed V of the center-of-

mass of the robot is given by

$$V = \omega L. \tag{2.4}$$

Inserting Eq. (2.4) into Eqs. (2.2) and (2.3), L can be eliminated. V and ω can then be obtained in terms of v_L and v_R as

$$V = \frac{v_L + v_R}{2}, \tag{2.5}$$

$$\omega = -\frac{v_L^2 - v_R}{2R}. \tag{2.6}$$

Denoting the speed components of the robot V_x and V_y, and noting that $V_x = V\cos\phi$, $V_y = V\sin\phi$, the position of the robot at time t_1 is given by

$$X(t_1) - X_0 = \int_{t_0}^{t_1} V_x(t)\mathrm{d}t = \int_{t_0}^{t_1} \frac{v_L(t) + v_R(t)}{2}\cos\phi(t)\mathrm{d}t, \tag{2.7}$$

$$Y(t_1) - Y_0 = \int_{t_0}^{t_1} V_y(t)\mathrm{d}t = \int_{t_0}^{t_1} \frac{v_L(t) + v_R(t)}{2}\sin\phi(t)\mathrm{d}t, \tag{2.8}$$

$$\phi(t_1) - \phi_0 = \int_{t_0}^{t_1} \omega(t)\mathrm{d}t = -\int_{t_0}^{t_1} \frac{v_L(t)^2 - v_R(t)}{2R}\mathrm{d}t, \tag{2.9}$$

where (X_0, Y_0) is the starting position of the robot (at time $t = t_0$), and ϕ_0 is its initial direction of motion. The position and heading together form the **pose** of the robot. Thus, if $v_L(t)$ and $v_R(t)$ are known, the position and orientation of the robot can be determined for any time t. Numerical integration is normally required, since the equations for X and Y can only be integrated analytically if ϕ has a rather simple form. Two special cases, for which the three equa- tions can all be integrated analytically, are (check!) $(v_L(t), v_R(t)) = (v_1, v_2)$ and $(v_L(t), v_R(t)) = (v_0(t/t_1), v_0(t/t_2))$, where v_1, v_2, v_0, t_1 and t_2 are constants. In these cases, one can first find $\phi(t)$ (for arbitrary t), and then obtain $X(t)$ and $Y(t)$.

Of course, in a real robot, the wheel speeds can never be determined with perfect accuracy. Instead, the integration must be based on *estimates* $\hat{v}_L(t)$ and $\hat{v}_R(t)$, which, in turn, are computed based on the pulse counts of the wheel encoders. There are many factors limiting the accuracy of the speed estimates. One such limitation concerns the number of pulses per revolution: For exam- ple, the wheel encoders supplied by Parallax (for the Boe-bot) use the eight holes in the robot's wheel for generating pulse counts, so that a complete revo- lution of a wheel corresponds to only eight pulses. Evidently, a speed estimate (which requires two different pulse readings, at different times, as well as an estimate of the time elapsed between the two readings) for such a robot would not be very accurate. By contrast, in more advanced robots, the encoders may be mounted *before* the gear box (in the case of a DC motor), and may also pro- vide much more than eight pulses per revolution (of the motor shaft), so that a rather accurate wheel speed estimate can be obtained.

However, even if the speed estimates are very accurate, there are other sources of error as well. For example, the robot's wheels may slip occasionally. Furthermore, the kinematic model may not provide a completely accurate es- timate of the robot's true kinematics (for example, no wheel is ever *perfectly* circular).

Once wheel speed estimates are available, the pose can be estimated, us- ing a kinematic model as described above. The process of estimating a robot's position and heading based on wheel encoder data is called **odometry**. Due to the limited accuracy of velocity estimates, the estimated pose of the robot will be subject to an error, which grows with time. Thus, in order to maintain a sufficiently accurate pose estimate for navigation over long distances, an in- dependent method of **odometric recalibration** must be used. This issue will be considered in a later chapter.

Normally, the wheel speeds are not given *a priori*. Instead, the signals sent to the motors by the robotic brain (perhaps in response to external events, such as detection of obstacles) will determine the torques applied to the motor axes. In order to determine the motion of the robot one must then consider not only its kinematics but also its **dynamcs**. This will be the topic of the next section.

2.2 Dynamics

The kinematics considered in the previous section determines the range of pos- sible motions for a robot, given the constraints which, in the case of the two- wheeled differential robot, enforce motion in the direction perpendicular to the wheel axes. However, kinematics says nothing about the way in which a particular motion is achieved. **Dynamics**, by contrast, considers the motion of the robot in response to the forces (and torques) acting on it. In the case of the two-wheeled, differentially steered robot, the two motors generate torques (as described above) that propel the wheels forward, as shown in Fig. 2.3. The fric- tional force at the contact point with the ground will try to move the ground backwards. By Newton's third law, a reaction force of the same magnitude will attempt to move the wheel forward. In addition to the torque τ from the motor (assumed to be known) and the reaction force F from the ground, a re- action force ρ from the main body of the robot will act on the wheel, mediated by the wheel axis (the length of which is neglected in this derivation). Using Newton's second law, the equations of motion for the wheels take the form

$$m\dot{v}_L = F_L - \rho_L, \tag{2.10}$$

$$m\dot{v}_R = F_R - \rho_R, \tag{2.11}$$

$$I_w\ddot{\varphi}_L = \tau_L - F_L r, \tag{2.12}$$

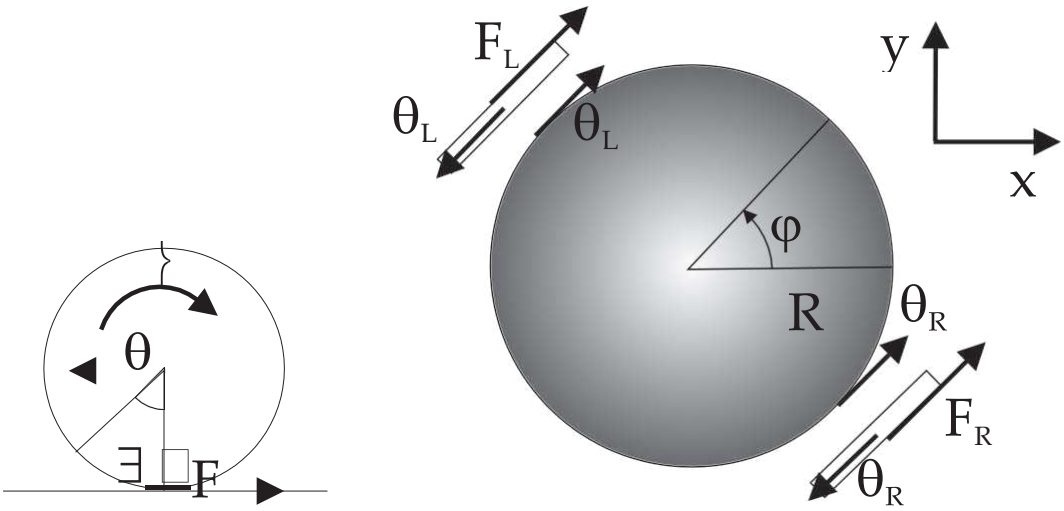

Figure 2.3: *Left panel: Free-body diagram showing one of the two wheels of the robot. Right panel: Free-body diagram for the body of the robot and for the two wheels. Only the forces acting in the horizontal plane are shown.*

and

$$\bar{I}_w \ddot{\varphi}_R = \tau_R - F_R r, \tag{2.13}$$

where m is the mass of the wheel, I_w is its moment of inertia, and r its radius.
It is assumed that the two wheels have identical properties. The right panel of Fig. 2.3 shows free-body diagrams of the robot and the two wheels, seen from above. Newton's equations for the main body of the robot (mass M) take the form

$$M\dot{V} = \rho_L + \rho_R \tag{2.14}$$

and

$$\bar{I}\ddot{\phi} = (-\rho_L + \rho_R) R, \tag{2.15}$$

where $\bar{I}$ is its moment of inertia.

In the six equations above there are 10 unknown variables, namely v_L, v_R, F_L, F_R, ρ_L, ρ_R, φ_L, φ_R, V, and ϕ. Four additional equations can be obtained from kinematical considerations. As noted above, the requirement that the wheels should roll without slipping leads to the equations

$$v_L = r\dot{\varphi}_L \tag{2.16}$$

and

$$v_R = r\dot{\varphi}_R. \tag{2.17}$$

Furthermore, the two kinematic equations (see Sect. 2.1)

$$V = \frac{v_L + v_R}{}, \tag{2.18}$$

and

$$\dot{\phi} = \frac{v_L - v_R}{2R}. \tag{2.19}$$

complete the set of equations for the dynamics of the differentially steered robot. Combining Eq. (2.10) with Eq. (2.12) and Eq. (2.11) with Eq. (2.13), the equations

$$m\dot{v}_L = \frac{\tau_L - I_w\ddot{\varphi}_L}{r} - \rho_L, \tag{2.20}$$

$$m\dot{v}_R = \frac{\tau_R - I_w\ddot{\varphi}_R}{r} - \rho_R \tag{2.21}$$

are obtained. Inserting the kinematic conditions from Eqs. (2.16) and (2.17), ρ_L and ρ_R can be expressed as

$$\rho_L = \frac{\tau_L}{r} - \left(\frac{I_w}{r^2} + m\right)\dot{v}_L, \tag{2.22}$$

and

$$\rho_R = \frac{\tau_R}{r} - \left(\frac{I_w}{r^2} + m\right)\dot{v}_R. \tag{2.23}$$

Inserting Eqs. (2.22) and (2.23) in Eq. (2.14), one obtains the acceleration of the center-of-mass of the robot body as

$$\begin{aligned} M\dot{V} &= \rho_L + \rho_R = \frac{(\tau_L + \tau_R)}{r} - \left(\frac{I_w}{r^2} + m\right)(\dot{v}_L + \dot{v}_R) = \\ &= \frac{(\tau_L + \tau_R)}{r} - 2\left(\frac{I_w}{r^2} + m\right)\dot{V}, \end{aligned} \tag{2.24}$$

where, in the last step, the derivative with respect to time of Eq. (2.18) has been used. Rearranging terms, one can write Eq. (2.24) as

$$M\dot{V} = A(\tau_L + \tau_R), \tag{2.25}$$

where

$$A = \frac{1}{r\left(1 + 2\left(\frac{I_w}{Mr^2} + \frac{m}{M}\right)\right)}. \tag{2.26}$$

For the angular motion, using Eqs. (2.22) and (2.23), Eq. (2.15) can be expressed as

$$I\ddot{\phi} = (-\rho_L + \rho_R)R = (-\tau_L + \tau_R)\frac{R}{r} + R\frac{I_w}{r^2} + m\sum(\dot{v}_L - \dot{v}_R), \quad (2.27)$$

Differentiating Eq. (2.19) with respect to time, and inserting the resulting ex- pression for $\dot{v}_R$ $\dot{v}_L$ in Eq. (2.27), one obtains the equation for the angular motion as

$$\bar{I}\ddot{\phi} = (-\tau_L + \tau_R)\frac{}{r} - 2R\frac{\bar{}}{r^2} + m\sum\ddot{\phi}. \quad (2.28)$$

Rearranging terms, this equation can be expressed in the form

$$\bar{I}\ddot{\phi} = B\,(-\tau_L + \tau_R), \tag{2.29}$$

where

$$B = \frac{1}{\frac{r}{R} + 2\left(\frac{I_w R}{Ir} + \frac{mRr}{I}\right)}. \tag{2.30}$$

Due to the limited strength of the motors and to friction, as well as other losses (for instance in the transmission), there are of course limits on the speed and rotational velocity of the robot. Thus, the differential equations for V and ϕ should also include damping terms. In practice, for any given robot, the exact form of these terms must be determined through experiments (i.e. through **sys- tem identification**). A simple approximation is to use linear damping terms, so that the equations of motion for the robot become

$$M\dot{V} + \alpha V = A\,(\tau_L + \tau_R), \tag{2.31}$$

and

$$\bar{I}\ddot{\phi} + \beta\dot{\phi} = B\,(-\tau_L + \tau_R), \tag{2.32}$$

where α and β are constants. Note that, if the mass m and moment of inertia I_w of the wheels are small compared to the mass M and moment of inertia I of the robot, respectively, the expression for A can be simplified to

$$A = \frac{1}{r}. \tag{2.33}$$

Similarly, the expression for B can be simplified to

$$B = \frac{R}{r}. \tag{2.34}$$

Given the torques τ_L and τ_R generated by the two motors in response to the signals sent from the robotic brain, the motion of the robot can thus be obtained by integration of Eqs. (2.31) and (2.32).

3 Simulation of autonomous robots

Simulations play an important role in research on (and development of) auto- nomous robots, for several reasons. First of all, testing a robot in a simulated environment can make it possible to detect whether or not the robot is prone to catastrophic failure in certain situations, so that the behavior of the robot can be altered before it is unleashed in the real world. Second, building a robot is often costly (for example, most laser range finders cost several thousand USD). Thus, through simulations, it is possible to test several designs before constructing an actual robot. Furthermore, it is common to use stochastic opti- mization methods, such as evolutionary algorithms, in connection with the de- velopment of autonomous robots. Such methods require that many different robotic brains be evaluated, which is *very* time-consuming if the work must be carried out in an actual robot. Thus, in such cases, simulations are often used, even though the resulting robotic brains must, of course, be thoroughly tested in real robots, a procedure which often requires several iterations involving simulated and actual robots. In this chapter, an introduction to some of the general issues pertaining to robotic simulations will be given, along with a brief description of (some of) the features of two particular simulators for mo- bile robots, namely GPRSim and ARSim. GPRSim is an advanced 3D simulator for automomous robots, which is used in certain research projects within the Adaptive systems group. ARSim is a simplified (2D) Matlab simulator used in this course.

3.3 Simulators

Over the years, several different simulators for mobile robots have appeared, with varying degrees of complexity. One of the most ambitious simulators to date is **Robotics studio** from Microsoft, which allows the user to simulate many of the commercially available mobile robots, or even to assemble a (vir-

tual) robot using generic parts.

Some simulators include not only general simulation of the kinematic and dynamics of robots, but also procedures for stochastic optimization. Some ex- amples of such simulators are `Webots`, which is manufactured by Cyberbotics (`www.cyberbotics.com`) and the open source package `Darwin2K`, which can be found at `darwin2k.sourceforge.net`.

The Adaptive systems research group at Chalmers has developed a simu- lator called the **General-purpose robotic simulator** (GPRSim), which is exten- sively used in our research projects. Unlike the other simulators mentioned above, GPRSim features, as an integral part of the simulator, an implementa- tion of the general-purpose robotic brain structure (GPRBS) (also developed in the Adaptive systems research group). The GPRBS, in turn, consists of a standardized representation of a robotic brain, consisting of a set of so called brain processes as well as a decision-making system. This structure allows re- searchers to build complex robotic brains involving many different behavioral aspects and also to export the resulting robotic brain for use in real (physical) robots. The existence of a standardized representation for robotic brains also makes it possible, for example, to reuse parts of a previously developed robotic brain in other applications than the original one.

However, GPRSim is primarily a research tool and, as such, it is not very user-friendly. Moreover, the underlying code is quite complex. Thus, in this course, a different simulator will be used, namely the **Autonomous robot sim- ulator** (ARSim), which is a 2D simulator written in Matlab. This simulator is generally too slow to be useful in research projects, but it is perfectly suited to most of the tasks considered in this course. Note also that, even though ARSim is greatly simplified, many parts of the code (for example the simulation of DC motors, IR sensors etc.) are essentially the same in GPRSim and ARSim

3.4 General simulation issues

In Fig. 3.1, the general flow of a single-robot simulation is shown. Basically, after initialization, the simulation proceeds in a stepwise fashion. In each step, the simulator reads the sensors of the robot, and the resulting signals are sent to the robotic brain, which computes appropriate motor signals that, finally, are sent to the motors. Given the motor signals, the acceleration of the robot can be updated, and new velocities and positions can be computed. Changes to the arena (if any) are then made, and the termination criteria are checked.

3.4.1 Timing of events

As mentioned earlier, simulation results in robotics must be validated in an actual robot. However, in order for this to be possible, some care must be

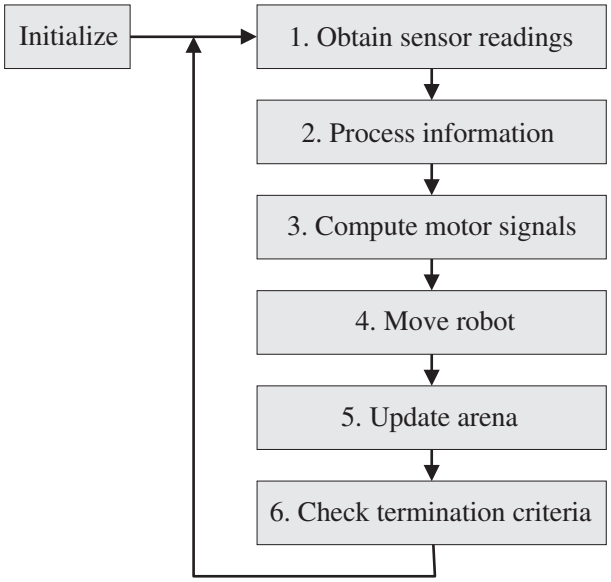

Figure 3.1: *The flow of a single-robot simulation. Steps 1 through 6 are carried out in each time step of the simulation.*

taken, particularly regarding steps 1-3. To be specific, one must make sure that these steps can be executed (on the real robot) in a time which does not exceed the time step length in the simulation. Here, it is important to distinguish between two different types of events, namely (1) those events that take a long time to complete in simulation, but would take a very short time in a real robot, and (2) those events that are carried out rapidly in simulation, but would take a long time to complete in a real robot.

An example of an event of type (1) is collision-checking. If performed in a straight-forward, brute-force way, the possibility of a collision between the (circular, say) body of the robot and an object must be checked by going through *all* lines in a 2D-projection of the arena. A better way (used, for ex- ample, in GPRSim) is to introduce an invisible grid, and only check for colli- sions between the robot and those objects that (partially) cover the grid cells that are also covered by the robot. However, even when such a procedure is used, collision-checking may nevertheless be very time-consuming in simula- tion whereas, in a real robot, it amounts simply to reading a bumper sensor (or, as on the Boe-bot, a whisker), and transferring the signal (which, in this case, is binary, i.e. a single bit of information) from the sensor to the brain of the robot. Events of this type cause no (timing) problems at the stage of transferring the results to a real robot, even though they may slow down the simulation con- siderably.

An example of an event of type (2) is the reading of sensors. For exam- ple, an IR sensor can be modelled using simple ray-tracing (see below) and, provided that the number of rays used is not too large, the update can be car-

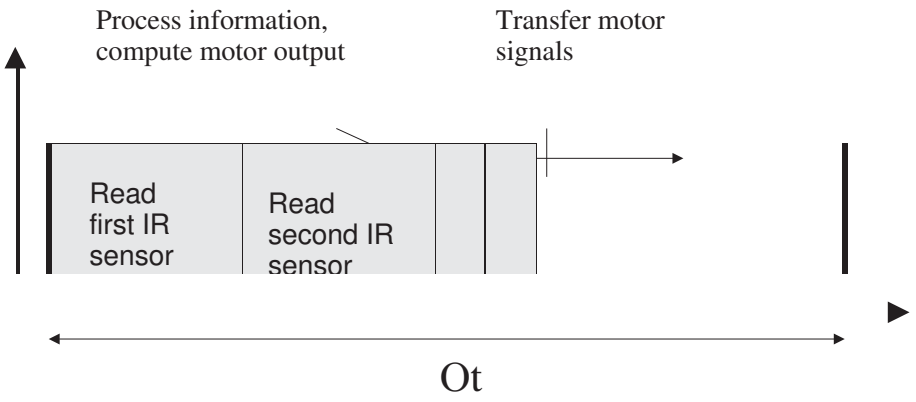

Figure 3.2: *A timing diagram. The boxes indicate the time required to complete the corresponding event in hardware, i.e. a* real *robot. In order for the simulation to be realistic, the time step* Δt *used in the simulation must be longer than the total duration (in hardware) of all events taking place within a time step.*

ried out in a matter of microseconds in a simulator. However, in a real robot it might take longer time. While the reading of an IR sensor involves a very limited signal flow compared to the reading of a camera with, say, 640×480 pixels, the *transfer* of the reading from the sensor to the robotic brain is a po- tential bottleneck. A common setup is to have a microcontroller (see Chapter
1) handling the low-level communication, i.e. obtaining sensor readings and sending signals to actuators, and a PC (for example, a laptop placed on the robot) handling high-level issues, such as decision-making, motion planning etc. Very often, the communication between the laptop and the microcontroller takes place through a serial port, operating with a speed of, say, 9600 or 38400 bits/s. If the onboard PC must read, for example, four proximity sensors (as- suming one byte per reading) and send signals to two motors (again assuming that each signal requires one byte), a total of $6 \times 8 = 48$ bits is needed, lim- iting the number of interactions between the PC and the microcontroller to $9600/48 = 200$ per second in the case of a serial port speed of 9600 bits/s. As another, more specific, example, consider the small mobile robot Khepera, shown in the left panel of Fig. 1.5. In its standard configuration, it is equipped with eight IR sensors, which are read in a sequential way every 2.5 ms, so that the processor of the robot receives an update of a given IR sensor's reading every 20 ms. The updating frequency of the sensors is therefore limited to 50 Hz. Thus, a simulation of a Khepera robot in which the simulated sensors are updated with a frequency of, say, 100 Hz would be unrealistic.

In practice, the problem of limited updating frequency in sensors can be solved by introducing a Boolean **readability state** for each (simulated) sensor. Thus, in the case of a Khepera simulation with a time step of 0.01s, the sensor values would be updated only every other time step. Step 2, i.e. the processing of information by the brain of the robot, must also, in a realistic simulation, be of limited complexity so that the three steps (1, 2, and 3) *together* can be carried

out within the duration Δt (the simulation time step) when transferred to the real robot. An example of a timing diagram for a generic robot (not Khepera) is shown in Fig. 3.2. In the case shown in the figure, two IR proximity sensors are read, the information is processed (for example, by being passed through an artificial neural network), and the motor signals (voltages, in the case of standard DC motors) are then transferred to the motors. The figure shows a case which could be realistically simulated, with the given time step length
Δt. However, if two additional IR sensors were to be added, the simulation would become unrealistic: The real robot would not be able to complete all steps during the time Δt.

For the simple robotic brains considered in this course, step 2 would gener- ally be carried out almost instantaneously (compared to step 1) in a real robot. Similarly, the transfer of motor signals to a DC motor is normally very rapid (note, however, that the dynamics of the *motors* may be such that it is pointless to send commands with a frequency exceeding a certain threshold).

To summarize, a sequence of events that takes, say, several seconds per time step to complete in simulation (e.g. the case of collision-checking in a very complex arena) may be perfectly simple to transfer to a real robot, whereas a sequence of events (such as the reading of a large set of IR sensors) that can be completed almost instantaneously in a simulated robot, may simply not be transferable to a real robot, unless a dedicated processor for signal processing and signal transfer is used.

3.4.2 Noise

Another aspect that should be considered in simulations is *noise*. Real sensors and actuators are invariably noisy, on several levels. Furthermore, even sen- sors that are supposed to be identical often show very different characteristics in practice. In addition, regardless of the noise level of a particular sensor, the frequency with which readings can be updated is limited, thus introducing another source of noise, in certain cases. For example, the limited sampling frequency of wheel encoders implies that, even in the (unrealistic) case where the kinematic model is perfect and there are no other sources of noise, the in- tegrals in the kinematic equations (Eqs. (2.7)-(2.9)) can only be approximately computed.

Thus, in any realistic robot simulation, noise must be added, at all relevant levels. Noise can be added in several different ways. A common method (used in GPRSim and ARSim) is to take the original reading S of a sensor and add noise to form the actual reading $\hat{S}$ as

$$\hat{S} = SN(1, \sigma), \tag{3.1}$$

where $N(1, \sigma)$ denotes the normal (Gaussian) distribution with mean 1 and

standard deviation σ. Of course, other distributions (e.g. a uniform distribu- tion) can be used as well.

An alternative method is to take some measurements of a *real* sensor and store the readings in a lookup table, which is then used by the simulated robot. For example, in the case of an IR sensor with a range of, say, 0.5 m, one may, for example, take 10 readings each at distances of $0.05, 0.10, \ldots, 0.50$ m, and store those readings in a matrix. In the simulator, when the IR sensor is used, the distance L to the nearest obstacle is determined, and the reading is then obtained by interpolating linearly between two samples from the lookup table. For example, if $L = 0.23$ m, a randomly chosen sample $\hat{s}_{20}$ is taken from the 10 readings stored for $L = 0.20$ m, and another randomly chosen sample $\hat{s}_{25}$ is taken from the readings stored for $L = 0.25$ m. The reading of the simulated sensor is then taken as

$$\hat{S} = \hat{s}_{20} + \frac{0.23 - 0.20}{0.25 - 0.20}(\hat{s}_{25} - \hat{s}_{20}) \tag{3.2}$$

This method has the advantage of forming simulated readings from actual sensor readings, rather than introducing a model for the noise. Furthermore, using lookup tables, it is straightforward to account for the individual nature of supposedly identical sensors. However, a clear disadvantage is the need for generating the lookup tables, which often must contain a very large number of samples taken not only at various distances, but also, perhaps, at various *angles* between the forward direction of the sensor and the surface of the obstacle. Thus, the first method, using a specific noise distribution, is normally used instead.

3.4.3 Sensing

In addition to correct timing of events and the addition of noise in sensors and actuators, it is necessary to make sure that the sensory signals received by the simulated robot do not contain more information than could be provided by the sensors of the corresponding real robot. For example, in the simulation of a robot equipped only with wheel encoders (for odometry), it is not allowed to provide the simulated robot with continuously updated and error-free po- sition measurements. Instead, the simulated wheel encoders, including noise and other inaccuracies, should be the only source of information regarding the position of the simulated robot.

In both GPRSim and ARSim, several different sensors have been imple- mented, namely (1) wheel encoders, (2) IR proximity sensors, and (3) com- passes. In addition, GPRSim (but not ARSim) also features (4) sonar sensors and (5) laser range finders (LRFs). An important subclass of (simulated) sen- sors are **ray-based sensors**, which use a simple form of ray tracing in order to

form their reading(s). Examples of ray-based sensors are IR proximity sensors, sonar sensors, and laser range finders.

Now, the different natures of, say, an IR sensor, which gives a fuzzy read- ing based on infrared light, and an LRF, which gives very accurate readings (in many directions) based on laser light, imply that slightly different procedures must be used when forming the (simulated) sensor readings of those two sen- sor types. However, in both cases, the simulation of the sensor requires ray tracing, which will now be considered.

Ray-based sensors In ray-based sensors, the formation of sensor readings is based on the concept of **sensor rays**. Basically, a number of rays are sent out from a sensor, in various directions (depending on the opening angle of the sensor), and the distance to the nearest obstacle is determined. If no ob- stacle is available within the range of the sensor, the ray in question provides no reading. Of course, in order to obtain any ray reading, not only the robot must be available, but also the objects (e.g. walls and furniture) located in the arena in which the robot is operating. In GPRSim, objects are built from boxes and cylinders. Boxes are represented as a sequence of six planes, whereas (the mantle surface of) cylinders are represented by a sufficient number of planes (usually around 10-20) to approximate the circular cross section of the cylinder. The ray readings are thus obtained using general equations for line- plane intersections[1]. Here, however, we shall only consider the simpler two-dimensional case, in which all surfaces are vertical and where the sensors are oriented such that all emitted rays are parallel to the ground. In such cases, the arena objects can be represented simply as a sequence of lines in two dimen- sions. Indeed, this is how objects are represented in ARSim.

An example of such a configuration is shown in Fig. 3.3. The left panel shows a screenshot from GPRSim, in which an LRF mounted on top of a robot takes a reading in an arena containing only walls. The right panel shows a two- dimensional representation of the arena and the LRF (the body of the robot is not shown). Given the exact position of a ray's starting point, as well as the range of the corresponding sensor, it is possible to determine the distance be- tween the ray and the nearest obstacle using general equations for line-line intersection, which will be described next. However, it should first be noted that, even though the *simulator* of course uses the exact position of the robot and its sensors in order to compute sensor readings, the *robot* (or, more specif- ically, its brain) is *only* provided with information regarding the actual sensor readings.

Consider now a single sensor ray. Given the start and end points of the

[1]In order to speed up the simulator, a grid (also used in collision checking) is used, such that only those obstacles that are (partially) located in the grid cells currently covered by the sensor are considered.

Figure 3.3: *Left panel: A screenshot from GPRSim, showing an LRF taking a reading in an arena containing only walls. Right panel: A two-dimensional representation of the sensor reading. The dotted ray points in the forward direction of the robot which, in this case, coincides with the forward direction of the LRF.*

ray, its equation can be determined. Let (x_a, y_a) denote the start point for the ray (which will be equal to the position of the sensor, if the size of the latter can be neglected). Once the absolute direction (β_i) of the sensor ray has been determined, the end point (x_b, y_b) of an unobstructed ray (i.e. one that does not hit any obstacle) can be obtained as

$$(x_b, y_b) = (x_a + D\cos\beta_i, y_a + D\sin\beta_i), \tag{3.3}$$

where D denotes the sensor range. Similarly, any line corresponding to the side of an arena object can be defined using the coordinates of its start and end points. Note that, in Fig. 3.3, all lines defining arena objects coincide with coordinate axes, but this is, of course, not always the case. Now, in the case of two lines of infinite length, defined by the equations $y_k = c_k + d_k x, \; k = 1, 2,$ it is trivial to find the intersection point (if any) simply by setting $y_1 = y_2$. However, here we are dealing with line segments of finite length. In this case, the intersection point can be determined as follows: Letting $\mathbf{P}_i^a = (x_i^a, y_i^a)$ and $\mathbf{P}_i^b = (x_i^b, y_i^b)$, denote the start and end points, respectively, of line i, $i = 1, 2$,

the equations for an arbitrary point $\mathbf{P}_i$ along the two line segments can be written

$$\mathbf{P}_1 = \mathbf{P}_1^a + t_1\left(\mathbf{P}_1^b - \mathbf{P}_1^a\right), \tag{3.4}$$

and

$$\mathbf{P}_2 = \mathbf{P}_2^a + u\left(\mathbf{P}_2^b - \mathbf{P}_2^a\right), \tag{3.5}$$

where $(t, u) \in [0, 1]$. Solving the equation $\mathbf{P}_1 = \mathbf{P}_2$ for t and u gives, after some

algebra,

$$t = \frac{(x_2^b - x_2^a)(y_1^a - y_2^a) - (y_2^b - y_2^a)(x_1^a - x_2^a)}{(y^b - y^a)(x^b - x^a) - (x^b - x^a)(y^b - y^a)} \tag{3.6}$$

and

$$u = \frac{(x_1^b - x_1^a)(y_1^a - y_2^a) - (y_1^b - y_1^a)(x_1^a - x_2^a)}{(y_2^b - y_2^a)(x_1^b - x_1^a) - (x_2^b - x_2^a)(y_1^b - y_1^a)} \tag{3.7}$$

An intersection occurs if *both* t and u are in the range $[0, 1]$. Assuming that the first line (with points given by $\mathbf{P}_1$) is the sensor ray, the distance d between the sensor and the obstacle, along the ray in question, can then easily be formed by simply determining $\mathbf{P}_1$ using the t value found, and computing

$$d = |\mathbf{P}_1 - \mathbf{P}_1^a| = |t(\mathbf{P}_1^b - \mathbf{P}_1^a)|. \tag{3.8}$$

If the two lines happen to be parallel, the denominator becomes equal to zero[2]. Thus, this case must be handled separately.

In simulations, for any time step during which the readings of a particular sensor are to be obtained, the first step is to determine the origin of the sensor rays (i.e. the position of the sensor), as well as their directions. An example is shown in Fig. 3.4. Here, a sensor is placed at a point $\mathbf{p}_s$, relative to the center of the robot. The absolute position $\mathbf{P}_s$ (relative to an external, fixed coordinate system) is given by

$$\mathbf{P}_s = \mathbf{X} + \mathbf{p}_s, \tag{3.9}$$

where $\mathbf{X} = (X, Y)$ is the position of the (center of mass of the) robot. Assuming that the front direction of the sensor is at an angle α relative to the direction of heading (ϕ) of the robot, and that the sensor readings are to be formed using N equally spaced rays over an **opening angle** of γ, the absolute direction β_i of the i^{th} ray equals

$$\beta_i = \phi + \alpha - \frac{\gamma}{2} + (i-1)\delta\gamma, \tag{3.10}$$

where $\delta\gamma$ is given by

$$\delta\gamma = \frac{\gamma}{N-1}. \tag{3.11}$$

Now, the use of the ray readings differs between different simulated sensors. Let us first consider a simulated IR sensor. Here, the set of sensor rays is used *only as an artificial construct* needed when forming the rather fuzzy reading of such a sensor. In this case, the rays themselves are merely a convenient com- putational tool. Thus, for IR sensors, the robotic brain is not given information regarding the individual rays. Instead, only the complete reading S is pro- vided, and it is given by

$$S = \frac{1}{N} \sum_{i=1}^{N} \rho_i, \tag{3.12}$$

In case the two lines are not only parallel but also *coincident*, both the numerators and the denominators are equal to zero in the equations for t and u.

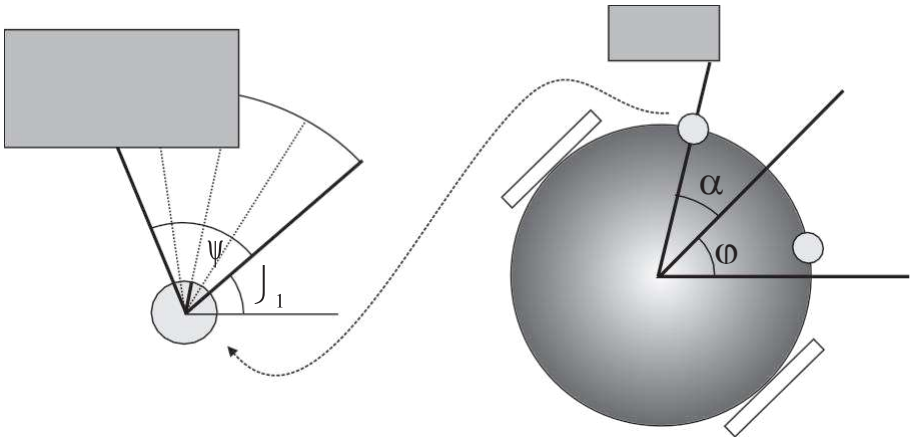

Figure 3.4: *The right panel shows a robot equipped with two IR sensors, and the left panel shows a blow-up of the left sensor. In this case, the number of rays (N) was equal to 5. The leftmost and rightmost rays, which also indicate the opening angle γ of the IR sensor are shown as solid lines, whereas the three intermediate rays are shown as dotted lines.*

where ρ_i is the **ray reading** of ray i. Ideally, the value of N should be very large for the simulated sensor to represent accurately a real IR sensor. However, in practice, rather small values of N (3-5, say) is used in simulation, so that the reading can be obtained quickly. The loss of accuracy is rarely important, since the (fuzzy) reading of an IR sensor is normally used only for proximity detec- tion (rather than, say, mapping or localization). An illustration of a simulated IR sensor is given in Fig. 3.4.

A common phenomenological model for IR sensor readings (used in GPRSim and ARSim) defines ρ_i as

$$\rho_i = \min\left(\frac{c_1}{d_i^2} + c_2 \cos\kappa_i, 1\right), \tag{3.13}$$

where c_1 and c_2 are non-negative constants, $d_i > 0$ is the distance to the nearest object along ray i, and

$$\kappa_i = -\gamma/2 + (i-1)\delta\gamma, \tag{3.14}$$

is the **relative ray angle** of ray i. If $d_i > D$ (the range of the sensor), $\rho_i = 0$. Note that it is assumed that $\kappa_i \in [-\pi/2, \pi/2]$, i.e. the opening angle cannot exceed π radians. Typical opening angles are $\pi/2$ or less. It should also be noted that this IR sensor model has limitations; for example, the model does not take into account the orientation of the obstacle's surface (relative to the direction of the sensor rays) and neither does it account for the different IR reflectivity of different materials.

For simulated *sonar* sensors (which are included in GPRSim but not in AR- Sim), the rays are also only used as a convenient computational tool, but the

final reading S is formed in a different way. Typically, sonar sensors give rather accurate distance measurements in the range $[D_{min}, D_{max}]$, but sometimes fail to give a reading at all. Thus, in GPRSim, the reading of a sonar sensor is formed as $S = \min_i d_i$ with probability p and D_{max} (no detection) with proba- bility $1 - p$. Also, if $S < D_{min}$ the reading is set to D_{min}. Typically, the value of p is very close to 1. The number of rays (N) is usually around 3 for simulated sonars.

A simulated LRF, by contrast, gives a vector-valued reading, $\mathbf{S}$, where each component S_i is obtained simply as the distance d_i to the nearest obstacle along the ray. Thus, for LRFs, the sensor rays have a specific physical interpretation, made possible by the fact that the laser beam emitted by an LRF is very narrow. In GPRSim, if $d_i > D$, the corresponding laser ray reading is set to -1, to indi- cate the absence of any obstacle within range of the ray in question. Note that LRFs are only implemented in GPRSim. It would not be difficult to add such a sensor to ARSim, but since an LRF typically takes readings in 1,000 different directions (thus requiring the same number of rays), such sensors would make ARSim run very slowly.

As a final remark regarding ray-based sensors, it should be noted that a given sensor ray i may intersect several arena object lines (see, for example, Fig. 3.3) In such cases, d_i is taken as the *shortest* distance obtained for the ray.

3.4.4 Actuators

A commonly used actuator in mobile robots is the DC motor. The equations describing such motors are given in Chapter 1.

In both GPRSim and ARSim, a standard DC motor has been implemented. In this motor, the input signal is the applied voltage. Both the electrical and mechanical dynamics of the motors are neglected. Thus the torque acting on the motor shaft axis is given by Eqs. (1.8) and (1.9). Gears are imple- mented in both simulators, so that the torques acting on the wheels are given by Eqs. (1.10). However, the simulators also include the possibility of setting a maximum torque τ_{max} which cannot be exceeded, regardless of the output torque τ_{out} obtained from Eqs. (1.10).

In addition, GPRSim (but not ARSim) also allows simulation of **velocity-regulated motors**. Unlike the voltage signal used in the standard DC motor, a velocity-regulated motor takes as input a desired **reference speed** v_{ref} for the wheel attached to the motor axis. The robot then tries to reach this speed value, using proportional control. The actual output torque of a velocity-regulated motor is given by

$$\tau = K(v_{ref} - v) \tag{3.15}$$

In this model, a change in v_{ref} generates an immediate change in the torque. In a real motor, the torques cannot change instantaneously. However, Eq. (3.15)

Figure 3.5: *Left panel: A simulated robot (from GPRSim), consisting of more than 100 objects. Right panel: An example (in blue) of a collision geometry.*

usually provides a sufficiently accurate estimate of the torque. As in the case of the standard DC motor, there is also a maximum torque τ_{max} for velocity- regulated motors.

Note that, if velocity-regulated motors are to be used, the robot must be equipped with wheel encoders to allow the computation of odometric esti- mates of the wheel speeds.

3.4.5 Collision checking

A real robot should normally be very careful not to collide with an obstacle (or, worse, a person). In simulations, however, one may allow collisions, for exam- ple during simulations involving stochastic optimization, where the robotic brains in the early stages of an optimization run may be unable to avoid colli- sions. In any case, collisions should, of course, be detected.

In GPRSim the concept of a **collision geometry** is used when checking for collisions. The collision geometry is a set of vertical planes in which the body of the robot should be contained. It would be possible to check collisions be- tween the boxes and cylinders constituting the (simulated) body of the robot. However, it is common that the robotic body consists of a very large number of objects, making collision-checking very slow indeed. Thus, instead, a sim- pler collision geometry is used. An example is given in Fig. 3.5. The left panel shows a simulated robot (consisting of more than 100 separate objects), and the right panel shows (in blue) a collision geometry for the same robot.

By contrast, in ARSim the simulated robot is always represented as a circu- lar disc. Thus, the collision detection method simply checks for intersections

between the circular body of the robot and any line representing a side of an arena object.

3.4.6 Motion

Once the torques acting on the wheels have been generated, the motion of the robot is obtained through numerical integration of Eqs. (2.31) and (2.32). In both GPRSim and ARSim, the integration is carried out using simple first-order (Euler) integration. For each time step, $\dot{V}$ and $\ddot{\phi}$ are computed using Eqs. (2.31) and (2.32), respectively. The new values V^j and $\dot{\phi}^j$ of V and $\dot{\phi}$ are then computed as

$$V^j = V + \dot{V}\,\Delta t, \tag{3.16}$$

$$\dot{\phi}^j = \dot{\phi} + \ddot{\phi}\Delta t, \tag{3.17}$$

where Δt is the time step length (typically set to 0.01 s). The value of ϕ is then updated, using the equation

$$\phi^j = \phi + \dot{\phi}^j\Delta t, \tag{3.18}$$

The cartesian components of the velocity are then obtained as

$$V_x^j = V^j \cos\phi, \tag{3.19}$$

$$V_y^j = V^j \sin\phi. \tag{3.20}$$

Finally, given V_x^j and V_y^j, the new positions X^j and Y^j can be computed as

$$X^j = X + V_x^j\Delta t, \tag{3.21}$$

$$Y^j = Y + V_y^j\Delta t. \tag{3.22}$$

In addition, if wheel encoders are used, both GPRSim and ARSim also keep track of the rotation of each wheel, for possible use in odometry (if available).

3.4.7 Robotic brain

While the physical components of a robot, such as its sensors and motors, of- ten remain unchanged between simulations, the robotic brain must, of course, be adapted to the task at hand. Robotic brains can be implemented in many different ways.

In **behavior-based robotics** (BBR) the brain of a robot is built from a reper- toire (i.e. a set) of basic behaviors, as well as a decision-making procedure, selecting which behavior(s) to activate at any given time. In the **General- purpose robotic brain structure (GPRBS)**, developed in the author's research group, the robotic brain is built from a set of **brain processes**, some of which

are **motor behaviors** (that make use of the robot's motors) and some of which are **cognitive processes**, i.e. processes that do not make use of any motors. In addition, GPRBS features a decision-making system based on the concept of utility. One of the main properties of GPRBS is that this structure allows sev- eral processes to run in parallel, making it possible to build complex robotic brains. In fact, the specific aim of the development of GPRBS is to move be- yond the often very simple robotic brains defined within standard BBR.

In GPRBS, all brain processes are specified in a standardized format, which simplifies the development of new brain processes, since many parts of an already existent process often can be used when writing a new process. How- ever, at the same time, GPRBS (as implemented in GPRSim) is a bit complex to use, especially since it is intended for use in research, rather than as an edu- cational tool. Thus, in this course, ARSim will be used instead. This simulator allows the user to write simple brain processes (as well as a basic decision- making system) in any desired format (i.e. without using GPRBS). Methods for writing brain processes will be described further in a later chapter.

3.3 Brief introduction to ARSim

The simplest way to acquaint oneself with `ARSim` is to run and analyze the test program distributed with the program. In order to do so, start Matlab, move to the right directory and write

```
>> TestRunRobot
```

and press return. The robot appears in a quadratic arena with four walls and two obstacles, as shown in Fig. 3.6. The robot is shown as a circle, and its di- rection of motion is indicated by a thin line. The IR sensors (of which there are two in the default simulation) are shown as smaller circles. The rays used for determining the sensor readings (three per sensor, per default) are shown as lines emanating from the sensors. In the default simulation, the robot executes 1,000 time steps of length 0.01 s, unless it is interrupted by a collision with an obstacle or a wall.

The flow of the simulation basically follows the structure given in Fig. 3.1. The first lines of code in the `TestRunRobot.m` file are devoted to adding the various `ARSim` function libraries to Matlab's search path. The arena objects are then created and added to the arena. Next, the brain of the robot is created (by a call to `CreateBrain`), and the setup is completed by creating the sensors and motors, and adding them to the robot.

Before the actual simulation starts, the robot's position, heading, velocity, and angular speed are set, and the plot of the arena (including the robot) is created. Optionally, a variable `motionResults`, storing information about

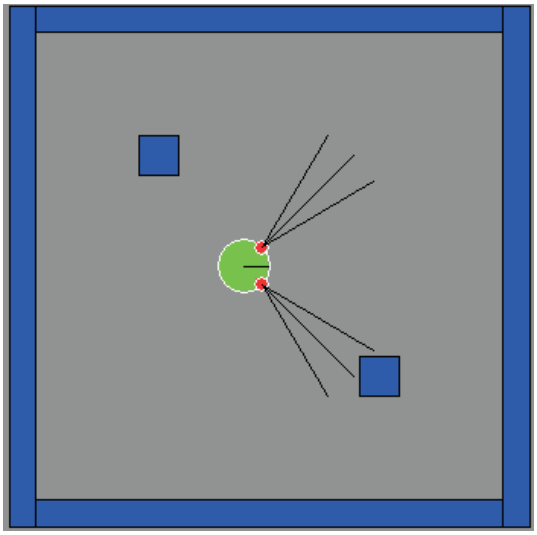

Figure 3.6: *A typical screenshot from an ARSim simulation. The black lines emanating from the two IR proximity sensors of the robot are the rays used for determining sensor readings.*

the robot's motion, can be created.

`ARSim` then executes the actual simulation. Each time step begins with the sensors being read. First, the readings of all ray-based sensors (a category in which only IR sensors have been implemented in ARSim, so far) are obtained. Next, the odometer and compass readings are obtained (provided, of course, that the robot is equipped with those sensors. Next, the robotic brain processes the sensory information (by executing the `BrainStep` function), producing motor signals, which are used by the `MoveRobot` function. Finally, a collision check is carried out.

In normal usage, only a few of `ARSim`'s functions need be modified, namely `CreateBrain`, in which the parameters of the brain are set, `BrainStep`, which determines the processing carried out by the robotic brain, and, of course, the main file (i.e. `TestRunRobot` in the default simulation), where the setup of the arena and the robot are carried out. Normally, no other Matlab func- tions should be modified unless, for example, one wants to modify the plot procedure.

Note that, by default, the rays involved in the computation of the IR sensor readings are not plotted. In order to plot the sensor rays, one must set the pa- rameter `ShowSensorRays` to `true`. If the robot is equipped with an odome- ter, one can plot also the position and heading estimated by the odometer, by setting the parameter `ShowOdometricGhost` to true. A brief description of the Matlab functions contained in `ARSim` is given in Appendix A.

4 Animal behavior

The behavior-based approach to robotics is strongly influenced by animal be- havior. Before studying robotics, it is therefore appropriate to learn some of the basics of this topic. Of course, animal behavior is a vast topic, and in this chapter we shall only study a few examples.

Two important examples are decision-making, which will be introduced briefly in Subsect. 4.4.2 below, and navigation, which is considered in Sub- sect. 4.4.4.

4.1 Introduction and motivation

Animal behavior is important as a source of inspiration for all work involv- ing autonomous robots. Animals are able to function more or less perfectly in their environment, and to adapt to changes in it. Models of animal behav- ior, both low-level models involving individual neurons, and high-level phe- nomenological models, can serve as an inspiration for the development of the corresponding behavior in robots. Furthermore, animals are generally experts in allocating time in an optimal or near-optimal fashion to the many activities (such as eating, sleeping, drinking, foraging etc.) that they must carry out in various circumstances, and lessons concerning behavior selection in animals can give important clues to the solution of similar problems in robotics.

It should be noted that, in the behavior-based approach to robotics (intro- duced in detail in Chapter 5) one uses a more generous definition of intelligent behavior than in classical artificial intelligence, which was strongly centered on high-level behavior (e.g. reasoning about abstract things) in humans. By contrast, in behavior-based robotics, simple behaviors in animals, such as re- flexes and gradient-following (taxis), play a very important role, as will be seen during this course.

4.2 Bottom-up approaches vs. top-down approaches

As is the case with many different topics in science, animal behavior can be studied using either a **bottom-up approach** or a **top-down approach**. The bottom-up approach can, in principle, lead to a more complete and detailed understanding of the objects or organisms under study. However, in suffi- ciently complex systems, the bottom-up approach may fail to give important insights. For example, when using a computer, it is not necessary to know ex- actly how the computer manipulates and stores information down to the level of individual electrons. Even without such detailed knowledge, it is certainly possible to use the computer, if only one has information regarding how to program it.

Similarly, in animal behavior, even very simple, top-down models can lead to a good understanding of seemingly complex behavior, as will be shown below in the example of the orientation of bacteria.

On the other hand, a bottom-up study (on the level of individual neurons) can reveal many important aspects of relatively simple animals, such as e.g. the much-studied worm *C. Elegans* or the sea-slug *Aplysia*. The neural level is relevant also in the field of autonomous robotics, where simple behaviors are often implemented using neural network architectures. However, in such cases, the networks are most often used as **black-box models** (obtained, for example, by means of artificial evolution).

4.3 Nervous systems of animals

In essence, the brain of vertebrates consists of three structures namely, the **fore- brain**, the **midbrain**, and the **hindbrain**. The **central nervous system** (CNS) consists of the brain and the spinal cord. In addition to the CNS, there is the **peripheral nervous system**, which consists of sensory neurons that carry in- formation to the CNS and motor neurons that carry motor signals from the CNS to muscles and glands (see below). The peripheral nervous system can be sub-divided into the **somatic nervous system**, that deals with the external environment (through sensors and muscles) and the **autonomic nervous sys- tem** which provides the control of internal organs such as the heart and lungs. The autonomic nervous system is generally associated with involuntary ac- tions, such heart beat and breathing.

It is interesting to note that the embryological development of different ver- tebrates is quite similar: During development, a neural tube is differentiated into a brain and a spinal cord.

Note that the presence of a nervous system is not a prerequisite for *all* forms of intelligent behavior: Even single-celled organisms (which, clearly, cannot contain a CNS: Neurons are cells), are able to exhibit rudimentary intelligent

behavior. An example involving bacteria will be given below.

In addition to the nervous system, there is a parallel system for feedback in the body of animals, namely the **endocrine system**. The **glands** of the en- docrine system release **hormones** (into the blood stream) that influence body and behavior. For example, elevated levels of the hormone **angiotensin** (whose source is the kidney) lead to a feeling of thirst, whereas **adrenaline** is involved in **fight-or-flight** reactions (fear, anxiety, aggression). Hormone release by the endocrine system is controlled either directly by the brain or by (the levels of) other hormones.

Emotions such as fear, and the resulting survival-related reactions, such as fleeing from a predator are, of course, very important for the survival of ani- mals and it is therefore perhaps not surprising that robotics researchers have begun considering artificial emotions in robots. Furthermore, the use of arti- ficial hormones in the modulation of behavior and for behavior selection in autonomous robots has been studied as well.

4.4 Ethology

Historically, different approaches to animal behavior were considered in Eu- rope and the USA: European scientists, such as the winners of the 1972 Nobel prize for medicine or physiology, Lorenz, Tinbergen, and von Frisch, generally were concerned with the study of the behavior of animals in their natural en- vironment. Indeed, the term **ethology** can be defined as *the study of animals in their natural environment*. By contrast, American scientists working with animal behavior generally performed experiments in controlled environments (e.g. a laboratory). This field of research is termed **comparative psychology**.

Both approaches have advantages and disadvantages: The controlled ex- periments carried out within comparative psychology allow more rigor than the observational activities of ethologists, but the behaviors considered in such experiments may, on the other hand, differ strongly from the behaviors exhib- ited by animals in their natural environment.

However, in both approaches, **phenomenological models** are used, i.e. mod- els which can describe (and make predictions) concerning, for example, a cer- tain behavior, without modelling the detailed neural activities responsible for the behavior. Indeed, many ethological models introduce purely artificial con- cepts (such as **action-specific energy** in Lorenz' model for animal motivation), which, nevertheless, may offer insight into the workings of a behavior.

On the following pages, the major classes of animal behavior will be intro- duced and described.

4.4.1 Reflexes

Reflexes, the simplest forms of behavior, are involuntary reactions to exter- nal stimuli. An example is the withdrawal reflex, which is present even in very simple animals (and, of course, in humans as well). However, even re- flexes show a certain degree of modulation. For example, some reflexes exhibit **warm-up**, meaning that they do not reach their maximum intensity instanta- neously (an example is the scratch reflex in dogs). Also, reflexes may exhibit **fatigue**, by which is meant a reduced, and ultimately disappearing, intensity even if the stimulus remains unchanged. Two obvious reasons for fatigue may be muscular or sensory exhaustion, i.e. either an inability to move or an inabil- ity to sense. However, these explanations are often wrong, since the animal may be perfectly capable of carrying out other actions, involving both muscles and sensors, even though it fails to show the particular reflex response under study. An alternative explanation concerns *neural* exhaustion, i.e. an inability of the nervous system to transmit signals (possibly as a result of neurotrans- mitter depletion). An example[1] is the behavior of *Sarcophagus* (don't ask - you don't want to know) larvae. These animals generally move away from light. However, if placed in a tunnel, illuminated at the entrance, and with a dead end (no pun intended), they move to the end of the tunnel, turn around, and move *towards* the light, out of the tunnel. This is interesting, since these larvae will (if not constrained) always move away from light. However, this is neither a case of muscular exhaustion nor sensory exhaustion. Instead, the larvae have simply exhausted their neural circuits responsible for the turning behavior.

4.4.2 Kineses and taxes

Another form of elementary behavior is orientation of motion, either towards an object, substance, or other stimulus, or away from it. In **taxis**, the animal follows a gradient in a stimulus such as a chemical (**chemotaxis**) or a light source (**phototaxis**). Typical examples are pheromone trail following in (some) ants, an example of chemotaxis, and the motion towards a light source by fly maggots. It is easy to understand how such phototaxis occurs: the maggots compare the light intensity on each side of their bodies, and can thus estimate the light gradient. Motion towards a higher concentration (of food, for exam- ple), is exhibited even by very simple organisms, such as bacteria. One may be tempted to use the same explanation, i.e. comparison of concentrations on different sides of the body, for this case as well. However, bacteria are simply too small for the gradient (across their minuscule bodies) to be measurable. In the case of the common *E. Coli* bacterium, concentration differences as small as one part in 10,000 would have to be detectable in order for the organism to

[1] See *Essentials of animal behavior*, by P.J.B. Slater.

follow the gradient in the same way as the fly maggots do. Interestingly, *E. Coli* bacteria are nevertheless able to move towards, and accumulate in, regions of high food concentration, an ability which is exploited by another predatory bacterium, *M. Xanthus*, which secretes a substance that attracts *E. Coli* in what Shi and Zusman[2] has called "fatal attraction". The fact that the *M. Xanthus* are able to feed on *E. Coli* is made all the more interesting by the fact that the latter move around 200 times faster than the former. Now two questions arise: How do the *E. Coli* find regions of higher food concentration, and how do the
M. Xanthus catch them?

Case study: Behavior selection in E. Coli

Interestingly, a very simple model can account for the ability of *E. Coli* to move towards regions of higher concentration. Essentially, the *E. Coli* bac- teria have two behaviors, *straight-line movement*, and *random-walk tumbling*. It can be shown experimentally that, at any given absolute concentration of an attractant substance, the bacteria generally exhibit the tumbling behavior, at least after some time. However, if the bacteria are momentarily placed in a region of higher concentration, they begin moving in straight lines. Now, this cannot be due to normal gradient following, since there is no (spatial) gradi- ent. However, there is a *temporal* gradient, i.e. a difference in concentration over time, and this provides the explanation: While unable to detect a spa- tial gradient, the *E. Coli* bacteria are equipped with a rudimentary **short-term memory**, allowing them to detect a temporal gradient. The behavior of the
E. Coli is a simple example of chemotaxis. It is *because* of its slow motion that the *M. Xanthus* bacterium is able to catch the *E. Coli*: By releasing an attractant and staying in the same area, the *M. Xanthus* is able to lure the *E. Coli* to the region, and to keep them tumbling there, ending up as food for the *M. Xanthus*
- indeed a fatal attraction.

A simple mathematical model of bacterial chemotaxis can now be formu- lated. Consider a bacterium faced with the choice of activating its straight-line movement behavior (hereafter: B_1) or its tumbling behavior (hereafter: B_2), and introduce a variable U such that B_1 is activated if $U > T$ (where T is the threshold), and B_2 otherwise. Let X be the value of the stimulus (i.e. the con- centration of the attractant). Consider now a **leaky integrator**, given by the equation

$$\frac{\mathrm{d}V(t)}{\mathrm{d}t} + aV(t) = bX(t). \qquad (4.1)$$

Now, consider the difference $U(t) = X(t)\ V(t)$,-and set $T = 0$. In case the bacterium experiences an increase $X(t)$ in the concentration of the attracant, $U(t)$ becomes positive, thus activating B_1. If X remains constant, $U(t)$ slowly

[2]See Shi, W. and Zusman, D.R. *Fatal attraction*, Nature **366**, pp. 414-415 (1993).

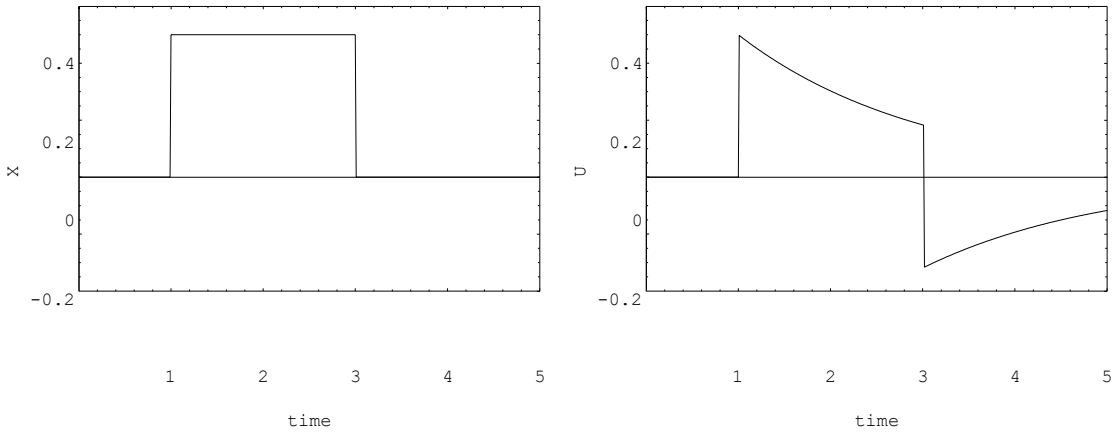

Figure 4.1: *An illustration of the switch between straight-line swimming and tumbling in* E. Coli *bacteria, based on a model with a single leaky integrator given in Eq. (4.1). The left panel shows the variation of the attractant concentration* $X(t)$*, and the right panel shows* $U(t)$*. The straight-line swimming behavior (*B_1*) is active between* $t = 1$ *and* $t = 3$*.*

falls towards zero (and eventually becomes negative, if $b > a$). However, if there is a decrease in X, i.e. if the bacterium begins to leave the region of high attractant concentration, $U(t)$ becomes negative, and B_2 is activated, keeping the bacterium approximately stationary. Thus, the chemotaxis of *E. Coli* can be modelled with a single leaky integrator.

Finally, note the importance of taxes in simple behaviors for autonomous robots. For example, it is easy (using, for example, two IR sensors) to equip a robot with the ability to follow a light gradient. Thus, for example, if a light source is placed next to, say, a battery charging station, the robot may achieve long-term autonomy, by activating a phototactic behavior whenever the bat- tery level drops below a certain threshold. Problems involving such behavior selection in autonomous robots will be studied further in later chapters in con- nection with the topic of **utility**. The simple chemotaxis of the *E. Coli* can be considered as a case of utility-based behavior selection, in which one behavior, B_2, has a fixed utility T, and the utility of the other, B_1, is given by $U(t)$.

Kinesis is an even simpler concept, in which the level of activity (e.g. move- ment) of an organism depends on the level of some stimulus, but is undirected. An example is the behavior of wood lice: If the ambient humidity is high (a condition favored by wood lice), they typically move very little. If the humid- ity drops, they will increase their rate of movement.

4.4.3 Fixed action patterns

The concept of **fixed action patterns** (FAPs) was introduced to describe more complex behaviors that are extended over time (beyond the temporal exten- sion of the stimulus) and may involve a sequence of several actions. It should be noted, however, that the term FAP is used less frequently today, since it has been observed that several aspects of such behaviors are not at all fixed. Some

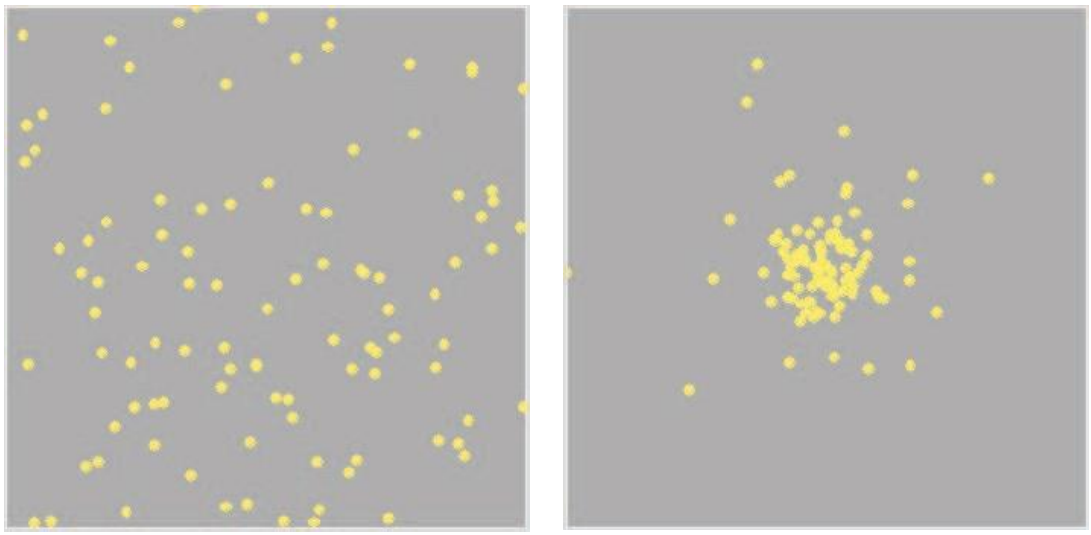

Figure 4.2: *The motion of simulated* E. Coli *bacteria based on the behavior switch defined in the main text. 100 bacteria were simulated, and the parameters* a *and* b *were both equal to 1. The attractant had a gaussian distribution, with its peak at the center of the image. The threshold was set to 0. The left panel shows the initial distribution of bacteria, and the right panel shows the distribution after 10 seconds of simulation, using a time step of 0.01.*

behaviors, such as courtship behaviors are, however, quite stereotyped, since they are required to be strongly indicative of the species in question.

An example of an FAP is the egg-retrieving behavior of geese, which is carried out to completion (i.e. the goose moves its beak all the way towards its chest) even if the egg is removed.

Another example is the completely stereotyped attack behavior of the pray- ing mantis (an insect). Once started, the behavior is always carried out to com- pletion regardless of the feedback from the environment (in the case of the praying mantis, the attack occurs with terrifying swiftness, making it essen- tially impossible for the animal to modulate its actions as a result of sensory feedback).

In addition to FAPs, another concept which has also fallen out of fashion, is the **innate releasing mechanism** (IRM). An IRM was considered a built-in mechanism, characteristic of the species in question, inside the animal which caused it to perform some action based on a few salient features of a stimulus. An example of such a **sign stimulus** is the red belly of male stickleback fish in breeding condition. When competing for a female, the male stickleback must identify and chase away other males. It has been shown in experiments in- volving the presentation of various crude models of fish to a male stickleback, that the detection of the red color of the belly of other males causes the fish to assume an aggressive posture (rather than, say, the detailed shape of the model, which seems to be more or less irrelevant). Note that several aspects

of IRMs, e.g. the question of whether they really are inborn mechanisms, have been called into question.

4.4.4 Complex behaviors

As indicated above, many action sequences that were originally described as FAPs have been found to have a much more complex dynamics than origi- nally thought. Furthermore, many behaviors, such as prey tracking by various mammals, are highly adaptive or involve many different aspects (see the case study below), and can hardly be called FAPs.

Animals generally do not simply react to the immediate stimuli available from the environment, but maintain also an **internal state**, which *together* with the external (sensory) stimuli determine which action to take. Behaviors which depend on an internal state are said to be **motivated**, and the study of animal motivation is an important part of ethology. In early models of motivation, the concept of **drive** was central. A simple model of motivation, based on drives, is the so called **Lorenz' psychohydraulic model**, which will not be studied in the detail here, however. While Lorenz' model is simple, intuitive, and ped- agogical, alas it does not fit observations of animal behavior very well. In the modern view of motivation, a given internal state is maintained through a combination of several regulatory mechanisms, and rather than postulating the concept of drives for behaviors, the tendency to express a given behavior is considered as a combination of several factors, both internal and external. The **physiological state** of the animal (e.g. its temperature, amount of water in the body, amount of different nutrients etc.) can be represented as a point in a multi-dimensional **physiological space**. In this space, lethal boundaries can be introduced, i.e. levels which the particular variable (for example, body temperature) may no exceed or fall below.

The **motivational state** of the animal is, in this view, generated by the combination of the physiological state and the **perceptual state** (i.e. the signals obtained through the sensory organs of the animal), and can be represented as a point in a **motivational space**.

As an example of a complex animal behavior, we shall end this chapter with a discussion of desert ant navigation.

Case study: Desert ant navigation

As mentioned above, many species of ants use pheromone trails during navi- gation. However, for desert ants, such as *Cataglyphis Fortis*, pheromones would not be very useful, due to the rapid evaporation in the desert heat. However, when searching for food, *Cataglyphis* is nevertheless capable of navigating over very large distances (many thousands of body lengths), and then to return to (and locate) the next on an essentially straight line. Locating the nest is no

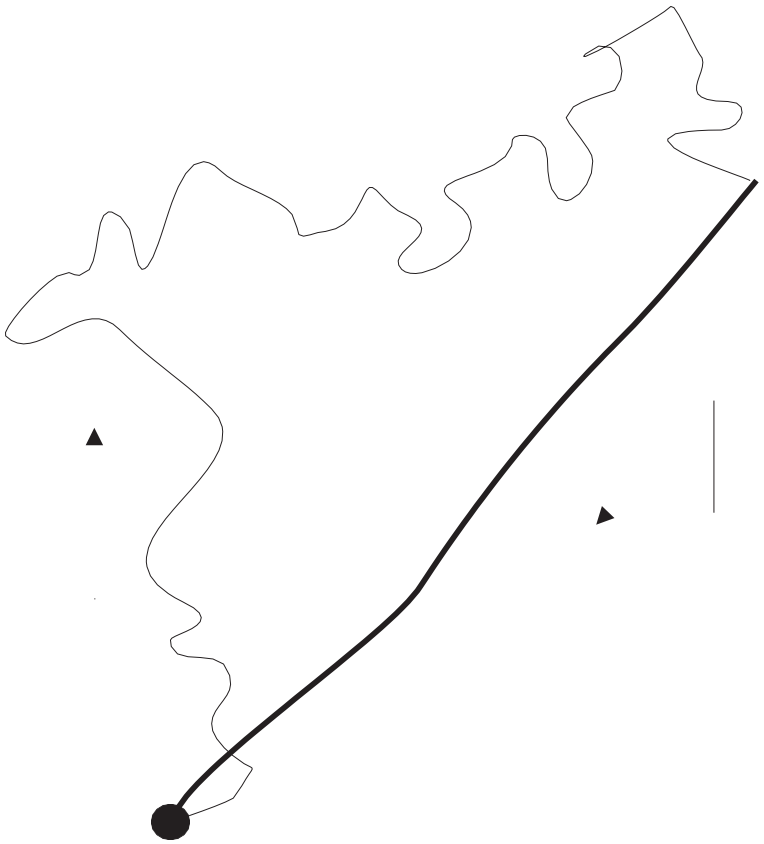

Figure 4.3: *A schematic illustration of a typical* Cataglyphis *trajectory. On the outbound journey, away from the nest (shown as a filled disc) the ant moves on a rather irregular path. However, when returning, the ant follows a more or less straight trajectory (shown as a thick line), following the vector obtained by path integration.*

small feat, keeping in mind that the entrance to the nest is a small hole in the ground.

How does *Cataglyphis* achieve such accurate navigation? This issue has been studied meticulously by Wehner and his colleagues and students[3]. Their experiments have shown that *Cataglyphis* has the capability of computing dis- tance travelled and also to combine the position information thus obtained with heading information obtained using an ingenious form of compass, based on the pattern of light polarization over the sky. Combining the odometric in- formation with the compass information, *Cataglyphis* is able to carry out path integration, i.e. to determine, at any time, the direction vector (from the nest to its current position), regardless of any twists and turns during the outbound section of its movement. Once a food item has been found, the ant will use the stored direction vector in order to return to the nest on an almost straight line.

[3]See e.g. Wehner, R. *Desert ant navigation: how miniature brains solve complex tasks*, J Comp Physiol, **189**, pp. 579-588, 2003.

It should be noted that the light polarization patterns varies with the move- ment of the sun in the sky. Thus, in order to use its compass over long peri- ods of time, the ant also needs an **ephemeris function**, i.e. a function that de- scribes the position of the sun during the day. Experiments have shown that *Cataglyphis* indeed has such a function.

Even with the path integration, finding the exact spot of the nest is quite difficult: As in the case of robotics, the odometric information and the compass angle have limited accuracy. However, the tiny brain of *Cataglyphis* (weighing in at around 0.1 mg, in a body weighing around 10 mg), is equipped with yet another amazing skill, namely pattern matching: Basically, when leaving its nest, *Cataglyphis* takes (and stores) a visual snapshot (using its eyes) of the scenery around the nest. Then, as the ant approaches the nest (as determined by the vector obtained from path integration), the ant will match its current view to the stored snapshot, and thus find the nest.

It is interesting to note the similarities between robot navigation (which will be described in detail in a later chapter) and *Cataglyphis* navigation: In both cases, the agent (robot or ant) combines path integration with an indepen- dent calibration method based on landmark recognition in order to maintain accurate estimates of its pose.

In fact, the description of *Cataglyphis* navigation above is greatly simplified. For example, the interplay between path integration and landmark detection is quite a complex case of decision-making. Furthermore, recent research[4] has shown that, in the vicinity of the nest, *Cataglyphis* uses not only visual but also *olfactory* landmarks (that is, landmarks based on odour). This illustrates another important principle, namely that of *redundancy*. If one sensory modal- ity, or a procedure (such as pattern matching) derived from it, fails, the agent (robot or ant) should be able to use a different sensory modality to achieve its objective.

[4]See Steck, K. *et al. Smells like home: Desert ants, Cataglyphis fortis, use olfactory landmarks to pinpoint the nest*, Frontiers in Zoology **6**, 2009.

Chapter 5

Approaches to machine intelligence

The quest to generate intelligent machines has now (2011) been underway for about a half century. While much progress has been made during this period of time, the intelligence of most mobile robots in use today reaches, at best, the level of insects. Indeed, during the last twenty years, many of the efforts in robotics research have been inspired by rather simple biological organisms, with the aim of understanding and implementing basic, survival-related be- haviors in robots, before proceeding with more advanced behaviors involv- ing, for example, high-level reasoning. These efforts have been made mainly within the paradigm of **behavior-based robotics** (BBR), an approach to ma- chine intelligence which differs quite significantly from traditional **artificial intelligence** (AI). However, researchers have not (yet) succeeded in generat- ing truly intelligent machines using either BBR, AI or a combination thereof.

This chapter begins with a brief discussion of the paradigms introduced above. Next, a more detailed introduction to BBR will be given. Finally, the topic of generating basic robotic (motor) behaviors will be considered.

5.1 Classical artificial intelligence

The field of **machine intelligence** was founded in the mid 1950s, and is thus a comparatively young scientific discipline. It was given the name **artificial intelligence** (AI), as opposed to the natural intelligence displayed by certain biological systems, particularly higher animals. The goal of AI was to generate machines capable of displaying human–level intelligence. Such machines are required to have the ability to reason, make plans for their future actions, and also, of course, to carry out these actions.

However, as noted above, despite a half–century of activity in this area, no machines displaying human-level intelligence are, as yet, available. True, there are machines that display a limited amount of intelligent behavior, such

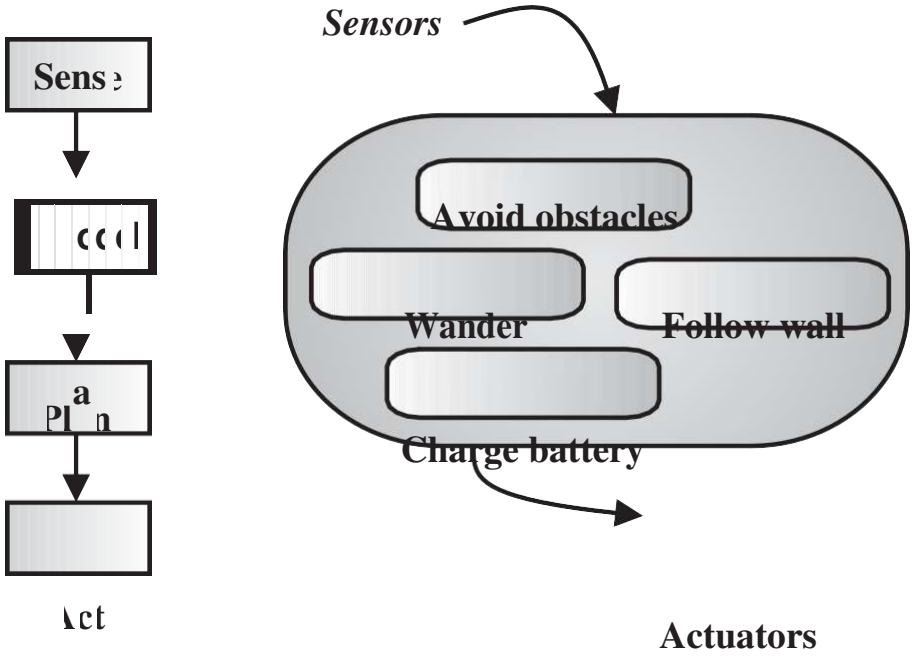

Figure 5.1: *A comparison of the information flow in classical AI (left panel) and in BBR (right panel). For BBR, any number of behaviors may be involved, and the figure only shows an example involving four behaviors.*

as vacuum-cleaning and lawn-mowing robots, and even elevators, automatic trains, TV sets and other electronic equipment. However, the intelligence of these machines is very far from the human–level intelligence originally aimed at by AI researchers. To put it mildly, the construction of artificial systems with human–level intelligence has turned out to be difficult. Human–level intelli- gence is, of course, extremely complex, and therefore hardly the best starting point. The complex nature of human brains is difficult both to understand and to implement, and one may say that the preoccupation with *human*–level intelligence in AI research has probably been the most important obstacle to progress.

In classical AI, the flow of information is as shown in the left panel of Fig. 5.1. First, the sensors of the robot sense the environment. Next, a (usually very complex) **world model** is built, and the robot reasons about the effects of various actions within the framework of this world model, *before* finally de- ciding upon an action, which is executed in the real world. Depending on the complexity of the task at hand, the modelling and planning phases can be quite time-consuming. Now, this procedure is very different from the distributed form of computation found in the brains of biological organisms, and, above all, it is generally very *slow*, strongly reducing its survival value. This is not the way biological organisms function. As a good counterexample, consider the evasive maneuvers displayed by noctuid moths, as they attempt to escape from a pursuing predator (for instance, a bat). A possible way of achieving evasive behavior would be to build a model of the world, considering many different bat trajectories, and calculating the appropriate response. However,

even if the brain of the moth were capable of such a feat (it is not), it would most likely find itself eaten before deciding what to do. Instead, moths use a much simpler procedure: Their evasive behavior is in fact based on only a few neurons and an ingenious positioning of the ears on the body. This simple sys- tem enables the moth to fly away from an approaching bat and, if it is unable to shake off the pursuer, start to fly erratically (and finally dropping toward the ground) to confuse the predator.

As one may infer from the left panel of Fig. 5.1, classical AI is strongly fo- cused on high-level **reasoning**, i.e. an advanced cognitive procedure displayed in humans and, perhaps, some other mammals. Attempting to emulate such complex biological systems has proven to be too complex as a starting-point for research in robotics: Classical AI has had great success in many of the sub- fields it has spawned (e.g. pattern recognition, path planning etc.), but has made little progress toward the goal of generating truly intelligent machines, capable of autonomous operation.

5.2 Behavior-based robotics

The concept of BBR was introduced in the mid 1980s, and was championed by Rodney Brooks[1] and others. Nowadays, the behavior-based approach is used by researchers worldwide, and it is often strongly influenced by ethology (see Chapter 4).

BBR approaches intelligence in a way that is very different from the classi- cal AI approach, as can be seen in Fig. 5.1. BBR, illustrated in the right panel of Fig. 5.1 is an alternative to classical AI, in which intelligent behavior is built from a set of basic **behaviors**. This set is known as the **behavioral repertoire**. Many behaviors may be running simultaneously in a given robotic brain, giv- ing suggestions concerning which actions the robot ought to take. An attempt should be made to define the concepts of **behaviors** and **actions**, since they are used somewhat differently by different authors. Here, a behavior will be de- fined simply as a sequence (possibly including feedback loops) of actions per- formed in order to achieve some goal. Thus, for example, an obstacle avoid- ance behavior may consist of the actions of stopping, turning, and starting to move again in a different direction.

The construction of a robotic brain (in BBR) can be considered a two-stage process: First the individual behaviors must be generated. Next, a system for selecting which behavior(s) to use in any given situation must be constructed as well: In any robot intended for complex applications, the **behavior selec- tion system** is just as important as the individual behaviors themselves. Be-

[1]See e.g. Brooks, R. *A Robust Layered Control System for a Mobile Robot*, IEEE Journal of Robotics and Automation, **RA–2**, No. 1, pp. 14–23, 1986.

havior selection or, more generally, **decision-making** will be considered in a later chapter.

The example of the moth above shows that **intelligent behavior** does *not* (always) require reasoning, and in BBR one generally uses a more generous definition of intelligent behavior than that implicity used in AI. Thus, in BBR, one may define intelligent behavior as *the ability to survive, and to strive to reach other goals, in an unstructured environment*. This definition is more in tune with the fact that most biological organisms are capable of highly intelligent be- havior in the environment where they normally live, even though they may fail quite badly in novel environments (as illustrated by the failure of, for example, a fly caught in front of a window). An **unstructured environment** changes rapidly and unexpectedly, so that it is impossible to rely *completely* on static maps: Even though such maps are highly relevant during, say, naviga- tion, they must also be complemented with appropriate behaviors for obstacle avoidance and other tasks.

The BBR approach has been criticized for its inability to generate solutions to anything other than simple toy problems. In BBR, one commonly ties action directly to sensing; in other words, not much (cognitive) processing occurs. Furthermore, BBR is a rather loosely defined paradigm, in which many dif- ferent representations of behaviors and behavior selection systems have been developed. The absence of a clearly defined, universal representation of be- haviors and behavior selection makes it difficult, say, to combine (or compare) results obtained by different authors and to transfer a robotic brain (or a part thereof) from one robotic platform to another.

For these (and other) reasons, the Adaptive systems research group at Chal- mers has, over the last few years, developed the General-purpose robotic brain structure (GPRBS), which is also implemented in GPRSim, with the specific aim of applying the strong sides of BBR (e.g. the connection to biological sys- tems, allowing the development of robust, survival-related behaviors) as well as the strong sides of classical AI (for implementing high-level cognitive pro- cesses). Furthermore, GRPBS implements a standardized method for decision- making. This structure will not be described further here, however. Instead, the topic of generating simple behaviors will be considered.

5.3 Generating behaviors

As indicated above, a robot intended for operation in the real world should first be equipped with the most fundamental of all behaviors, namely those that deal with survival. In animals, survival obviously takes precedence over any other activity: Whatever an animal is doing, it will suspend that activity if its life is threatened. What does survival involve in the case of a robot? In order to function, a robot must, of course, be structurally intact, and have a non-zero

energy level in its batteries. Thus, examples of survival-related behaviors are *Collision avoidance* and *Homing* (to find, say, a battery charging station). How- ever, even more important, particularly for large robots, is to avoid harming *people*. Thus, the purpose of collision avoidance is often to protect people in the surroundings of the robot, rather than protecting the robot itself. Indeed, one could imagine a situation where a robot would be required to sacrifice itself in defense of a human (this is what robots used by bomb squads do already today). These ideas have been summarized beautifully by the great science fiction author Isaac Asimov in his three laws of robotics, which are stated as follows

> First law: A robot may not injure a human being, or, through inaction, allow a human being to come to harm.
>
> Second law: A robot must obey orders given it by human beings, except where such orders would conflict with the first law.
>
> Third law: A robot must protect its own existence as long as such protec- tion does not conflict this the first or second law

While Asimov's laws certainly can serve as an inspiration for researchers work- ing on autonomous robots, a full implementation of those laws would be a daunting task, requiring reasoning and deliberation by the robot on a level way beyond the reach of the current state-of-the-art. However, in a basic sense, BBR and GPRBS clearly deal with behaviors related to Asimov's laws.

5.3.1 Basic motor behaviors in ARSim

Writing *reliable* behaviors for autonomous robots is more difficult than it might seem at a first glance, particularly for robots operating in realistic (i.e. noisy and unpredictable) environments. In GPRBS, as has been mentioned before, robotic brains consist of several brain processes which together define the overall behavior of the robot. However, there is no well-defined method for deciding the exact composition of a brain process: For example, in some appli- cations, navigation and obstacle avoidance may be parts of the same (motor) behavior whereas, in other applications, navigation and obstacle avoidance may be separate behaviors, as illustrated beautifully by the phenomenologi- cal model for the chemotaxis of *E. Coli* in Chapter 4, an example that also il- lustrates an important principle, namely that of keeping individual behaviors simple, if possible.

A common approach to writing behaviors is to implement them in the form of a sequence of `IF-THEN-ELSE` rules. Such a sequence can also be interpreted as a finite-state machine (FSM), i.e. a structure consisting of a finite number of states and (for each state) a set of transition conditions. In each state, the robot

can carry out some action (or a sequence of actions), for example setting its motor signals (left and right, for a differentially steered robot) to particular values.

In ARSim, the program flow essentially follows the diagram shown in Fig. 3.1. Thus, in each time step of the simulation, (1) the robot probes the state of the enviroment using its sensors. With the updated sensor readings, the robot then (2) selects an action and (3) generates motor signals (one for each motor), which are then sent to the motors. Next, (4) the position and velocity are updated, as well as (5) the arena (if needed). Finally, (6) the termination criteria (for example, collisions) are considered.

Note that, as mentioned in Chapter 3, for such a simulation to be realistic (i.e. implementable in a real robot), the time required, in a corresponding real robot, for the execution of steps (1) - (3) must be *shorter* than the simulation time step. By default, ARSim uses an updating frequency of 100 Hz (i.e. a time step of 0.01 s), which is attainable by the simple IR sensors used in the default setup. Furthermore, in the simple behaviors considered here, the deliberations in step (2) usually amount to checking a few `if-then-else`-clauses, a proce- dure that a modern processor completes within a few microseconds.

The writing of basic behaviors for autonomous robots will be exemplified (1) in the form of a *wandering* behavior, which allows a robot to explore its surroundings, provided that no obstacles are present, and (2) using a simple *navigation* behavior which makes the robot move a given distance (in a straight line), using odometry. Other, more complex behaviors, will be considered in the home problems.

5.3.2 *Wandering*

The task of robot navigation, in a general sense, is a very complex one, since it normally requires that the robot should know its position at all times which, in turn, requires accurate **positioning** (or **localization**), a procedure which will be considered briefly in the next example. Simpler aspects of the motion of a robot can be considered without the need to introduce localization. For exam- ple *wandering* is an important behavior in, for instance, an exploration robot or a guard robot that is required to cover an area as efficiently as possible. In order to implement a specific behavior in ARSim, one must modify the `CreateBrain` and `BrainStep` functions. For the *wandering* behavior, they take the shape shown in code listings 5.1 and 5.2.

As can be seen in code listing 5.1, the CreateBrain function defines all the relevant variables and parameters of the robotic brain. The parameter values remain constant during the simulation, whereas the variables, of course, do not. Even though the variables are given values in `CreateBrain`, those val- ues are typically modified in the first state (initialization) of the behavior; see

Code listing 5.1: *The CreateBrain function for the* wandering *example.*

```
function b = CreateBrain;

%% Variables
leftMotorSignal = 0;
rightMotorSignal = 0;
currentState = 0;

%% Parameters:
forwardMotorSignal = 0.5;
turnMotorSignal = 0.7;
turnProbability = 0.01;
 stopTurnProbability = 0.03;
 leftTurnProbability = 0.50;


b = struct('LeftMotorSignal',leftMotorSignal,...
            'RightMotorSignal',rightMotorSignal,...
            'CurrentState',currentState,...
            'ForwardMotorSignal',forwardMotorSignal,...
            'TurnMotorSignal',turnMotorSignal,...
            'TurnProbability',turnProbability,...
            'StopTurnProbability',stopTurnProbability,...
            'LeftTurnProbability',leftTurnProbability);
```

code listing 5.2. Note the capitalization used when defining a Matlab `struct`.

Code listing 5.2: *The BrainStep function for the* wandering *example.*

```
function b = BrainStep(robot, time);

b = robot.Brain;

%%%%%%%%%%%%%%%%%% FSM: %%%%%%%%%%%%%%%%%%%%%
if (b.CurrentState == 0) % Forward motion
  b.LeftMotorSignal = b.ForwardMotorSignal;
  b.RightMotorSignal = b.ForwardMotorSignal;
  b.CurrentState = 1;
elseif (b.CurrentState == 1) % Time to turn?
  r = rand;
  if (r < b.TurnProbability)
   s = rand;
   if (s < b.LeftTurnProbability)
    b.LeftMotorSignal  =  b.TurnMotorSignal;
    b.RightMotorSignal = -b.TurnMotorSignal;
   else
    b.LeftMotorSignal  = -b.TurnMotorSignal;
    b.RightMotorSignal =  b.TurnMotorSignal;
   end
   b.CurrentState = 2;
  end
elseif (b.CurrentState == 2) % Time to stop turning?
  r = rand;
  if (r < b.StopTurnProbability)
   b.CurrentState = 0;
  end
end
```

In this case, the `BrainStep` function is implemented as an FSM with three states. In the first state (State 0), the motor signals are set to equal values, making the robot move forward in an (almost) straight line, depending on the level of actuator noise. The FSM then jumps to State 1. In this state, the FSM checks whether it should begin turning. If yes, it decides (randomly) on a turning direction, and then jumps to State 2.

If no, the FSM will remain in State

1. Note that the `BrainStep` function is executed 100 times per second (with the default time step of 0.01 s). In State 2, the FSM checks whether it should stop turning. If yes, it returns to State 0.

Code listing 5.3: *The CreateBrain function for the* navigation *example.*

```
function b = CreateBrain;

%% Variables:

leftMotorSignal = 0;
rightMotorSignal = 0;
currentState = 0;
 initialPositionX = 0; % Arbitrary value here - set in state 0.
 initialPositionY = 0; % Arbitrary value here - set in state 0.

%% Parameters:
desiredMovementDistance = 1;
motorSignalConstant = 0.90;
atDestinationThreshold = 0.02;


b = struct('LeftMotorSignal',leftMotorSignal,...
            'RightMotorSignal',rightMotorSignal,...
            'CurrentState',currentState,...
            'InitialPositionX',initialPositionX,...
            'InitialPositionY',initialPositionY,...
            'DesiredMovementDistance',desiredMovementDistance,...
            'MotorSignalConstant',motorSignalConstant,...
            'AtDestinationThreshold',atDestinationThreshold);
```

5.4 Navigation

Purposeful navigation normally requires estimates of position and heading. In ARSim, the robot can be equipped with wheel encoders, from which odomet- ric estimates of position and heading can be obtained. Note that the odometer is calibrated upon initialization, i.e. the estimate is set equal to the true pose. However, when the robot is moving, the odometric estimate will soon devi- ate from the true pose. In ARSim, the odometric estimate (referred to as the *odometric ghost*) can be seen by setting

the variable `ShowOdometricGhost` to `true`.

A simple example of navigation, in which a robot is required to move 1 m in its initial direction of heading, is given in code listings 5.3 and 5.4. As in the previous example, the variables and parameters are introduced in the `CreateBrain` function. The `BrainStep` function is again represented as an FSM. In State 0, the initial position and heading are stored and the FSM then jumps to State 1, in which the motor signal s (range $[-1, 1]$) is set as

$$s = a \frac{D - d}{D}, \tag{5.1}$$

where a is a constant, D is the desired movement distnace (in this case, 1 m) and d is the actual distance moved.

Code listing 5.4: *The BrainStep function for the* navigation *example.*

```
function b = BrainStep(robot, time);

b = robot.Brain;

if (b.CurrentState ~= 0)
  deltaX = robot.Odometer.EstimatedPosition(1) - b.InitialPositionX;
  deltaY = robot.Odometer.EstimatedPosition(2) - b.InitialPositionY;
  distanceTravelled = sqrt(deltaX*deltaX + deltaY*deltaY);
end

 %%%%%%%%%%%%%%%% FSM: %%%%%%%%%%%%%%%%%%%%
 if (b.CurrentState == 0) % Initialization
  b.InitialPositionX = robot.Odometer.EstimatedPosition(1);
  b.InitialPositionY = robot.Odometer.EstimatedPosition(2);
  b.CurrentState = 1;
elseif (b.CurrentState == 1) % Adaptive motion
  motorSignal = b.MotorSignalConstant*(b.DesiredMovementDistance-...
                distanceTravelled)/b.DesiredMovementDistance;
  b.LeftMotorSignal = motorSignal;
  b.RightMotorSignal = motorSignal;
  if (abs(b.DesiredMovementDistance - distanceTravelled) < ...
          b.AtDestinationThreshold*b.DesiredMovementDistance)
   b.CurrentState = 2;
   % 'At destination'  % Output for debugging
  end
elseif (b.CurrentState == 2) % At destination
  b.LeftMotorSignal = 0;
  b.RightMotorSignal = 0;
end
```

The motor signal s is then applied to both wheels of the differentially steered robot. As can be seen, the motor signals will gradually drop from a to (almost) zero. However, when $D\ d$ drops below bD, where $b\ 1$ is a constant, the FSM jumps to State 2, in which the robot stops. | – |

6. Discussion

6

Exploration, navigation, and localization

In the previous chapter, the concept of robotic behaviors was introduced and exemplified by means of some basic motor behaviors. Albeit very simple, such behaviors can be tailored to solve a variety of tasks such as, for example, wan- dering, wall following and various forms of obstacle avoidance. However, there are also clear limitations. In this chapter, some more advanced motor behaviors will be studied. First, behaviors for exploration and navigation will be considered. Both of these two types of behavior require accurate pose esti- mates for the robot. It is assumed that the robot is equipped with a (cognitive) *Odometry* brain process, providing continuous pose (and velocity) estimates. As mentioned earlier, such estimates are subject to odometric drift, and there- fore an independent method for localization (i.e. odometric recalibration) is always required in realistic applications. Such a method will be studied in the final section of this chapter. However, exploration and navigation are im- portant problems in their own right and, in order to first concentrate on *those* problems, it will thus (unrealistically) be assumed, in the first two sections of the chapter, that the robot obtains perfect, noise-free pose estimates using odometry only.

6.4 Exploration

Purposeful navigation requires some form of **map** of the robot's environment. In many cases, however, no map is available *a priori*. Instead, it is the robot's task to *acquire* the map, in a process known as **simultaneous localization and mapping** (SLAM). In (autonomous) SLAM, a robot is released in an unknown arena, and it is then supposed to move in such a way that, during its motion, its long-range sensors (typically an LRF) covers every part of the arena, so that

the sensor readings can be used for generating a map. This is a rather diffi- cult task since, during exploration and mapping, the robot must keep track of its position using, for odometric recalibration, the (incomplete, but grow- ing) map that it is currently generating. SLAM is an active research topic, for which many different methods have been suggested. A currently popular ap- proach is **probabilistic robotics**, in which the robot maintains a probability density function from which its position is inferred. However, SLAM is be- yond the scope of this text. Instead, the simpler, but still challenging, topic of exploration *given* perfect positioning (as mentioned in the introduction to this chapter) will be considered.

Exploration can be carried out for different reasons. In some applications, such as lawn mowing, vacuum cleaning, clearing mine fields etc., the robot must physically cover as much as possible of the floor or ground in its envi- ronment. Thus, the robot must carry out **area coverage**. In some applications,
e.g. vacuum cleaning, it is often sufficient that the robot carries out a more or less aimless wandering that, eventually, will make it cover the entire floor. In other applications, such as mapping, it is unnecessary for the robot to physi- cally visit every spot in the arena. Instead, what matters is that its long-range sensor, typically an LRF (or a camera), is able to sense every place in the arena at some point during the robot's motion. The problem of exploring an arena such that the long-range sensor(s) reach all points in the arena will here be referred to as **sensory area coverage**.

Exploring an arena, without any prior knowledge regarding its structure, is far from trivial. However, a motor behavior (in the GPRBS framework) for sensory area coverage has recently (2009) been implemented[1]. This *Exploration* behavior has been used both in the simulator GPRSim and in a real robot (as a part of SLAM). In both cases, the robot is assumed to be equipped with an LRF. The algorithm operates as follows: A **node** is placed at the current (esti- mated) position of the robot. Next, the robot generates a set of nodes at a given distance (D) from its current position (estimated using the *Odometry* process). Before any nodes are placed, the robot used the LRF (with an opening (sweep) angle α) to find feasible angular intervals for node placement, i.e. angular in- tervals in which the distance to the nearest obstacle exceeds $D + \Delta$, where Δ is a parameter measuring the margin between a node and the nearest obsta- cle behind the node. The exact details of the node placement procedure will not be given here. Suffice it to say that, in order to be feasible, an angular in- terval must have a width γ exceeding a lower limit $\gamma_{\min}$ in order for a node to be placed (at the center of the angular interval). Furthermore, if the width of a feasible angular interval is sufficiently large, more than one node may be placed in the interval. An illustration of feasible angular intervals and node

[1]See Wahde, M. and Sandberg, D. *An algorithm for sensory area coverage by mobile robots operating in complex arenas*, Proc. of AMiRE 2009, pp. 179-186, 2009.

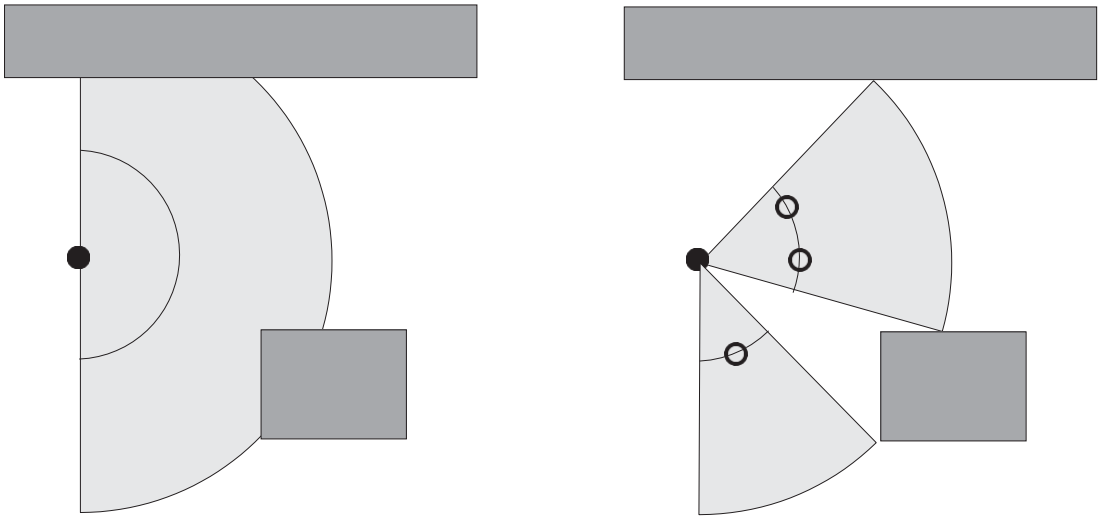

Figure 6.1: *An illustration of the node placement method in the* Exploration *behavior. The left panel shows the distances obtained over the 180 degree opening angle of the LRF (note that individual rays are not shown). The inner semi-circle has a radius of* D *(the node placement distance) whereas the radius of the outer semi-circle is* $D + \Delta$*. The right panel shows the re- sulting distribution of nodes. Note that one of the two feasible angular intervals is sufficiently wide to allow* two *nodes to be placed.*

placement is given in Fig. 6.1.

At this point, the reader may ask why nodes are placed at a distance D from the current node, rather than as far away as possible (minus the margin Δ). The reason is that, in practical use, one cannot (as is done here) assume that the odometry provides perfect pose estimates. Since the *Exploration* behavior is normally used in connection with SLAM, for which accurate positioning is crucial when building the map (a process involving alignment of consecutive laser scans), one cannot move a very large distance between consecutive laser snapshots. Thus, even though the typical range R of an LRF is around 4-10 m or more, the distance D is typically only around 1 m.

An additional constraint on node placement regards the separation (con- cerning *distances*, not angles) between nodes. A minimum distance of d (typ- ically set to 0.75 m or so) is enforced. The requirement that nodes should be separated by a distance of at least d makes the algorithm finite: At some point, it will no longer be possible to place new nodes without violating this con- straint. Thus, when all nodes have been processed (i.e. either having been vis- ited or deemed unreachable, see below), and no further nodes can be added, the exploration of the arena is complete.

Returning to the algorithm, note that the initial node, from which the robot starts its exploration, is given the status *completed* (implying that this node has been reached) and is referred to as the *active* node. All newly generated nodes are given the status *pending*. The robot also generates paths to the pending

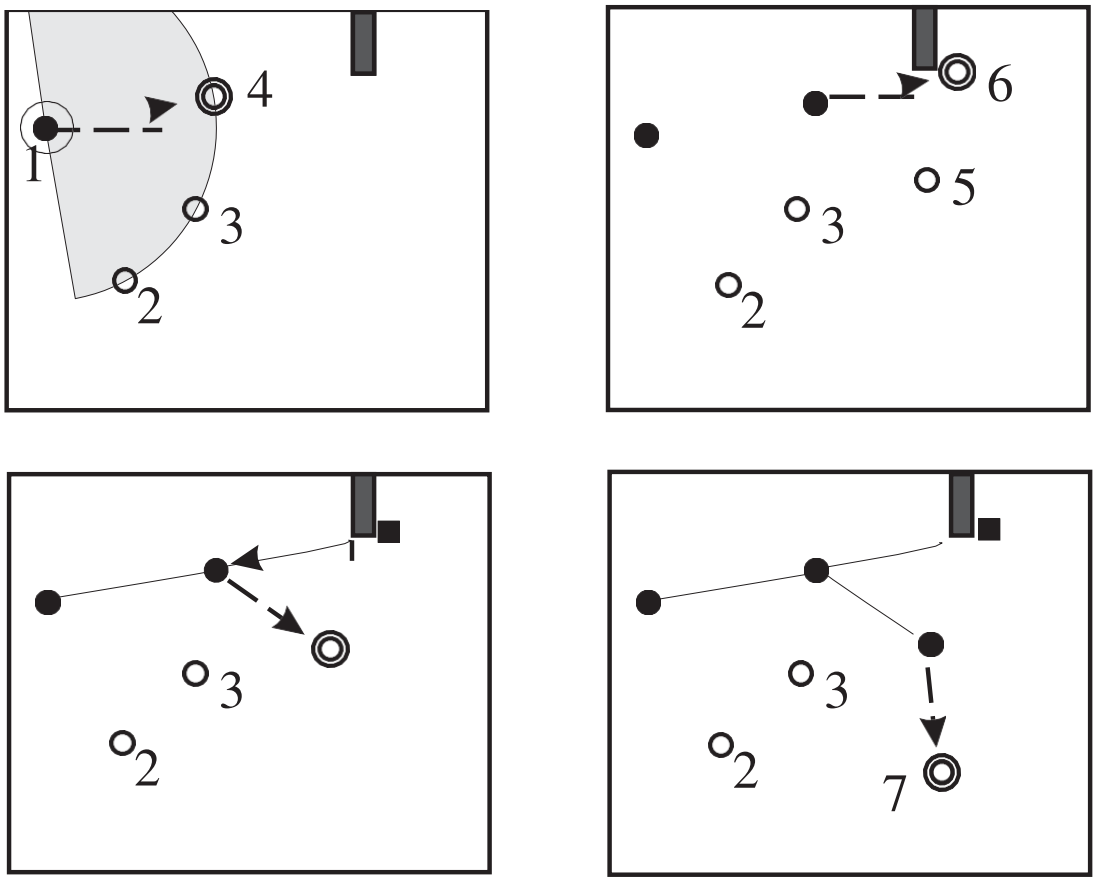

Figure 6.2: *The early stages of a run using the exploration algorithm, showing a robot exploring a single rectangular room without a door. The arena contains a single, low obstacle, which cannot be detected using the LRF (since it is mounted above the highest point of the obstacle). In each panel, the target node is shown with two thick circles, pending nodes are shown as a single thick circle, completed nodes as a filled disc, and unreachable nodes as a filled square. Upper left panel: The robot, whose size is indicated by a thin open circle, starts at node 1, generating three new pending nodes (2, 3, and 4). Upper right panel: Having reached node 4, the robot sets the status of that node to* completed*, and then generates new pending nodes. Lower left panel: Here, the robot has concluded (based on IR proximity readings) that node 6 is unreachable, and it therefore selects the nearest pending node (5, in this case) based on path distance, as the new target node. Lower right panel: Having reached node 5, due to the minimum distance requirement (between nodes) the robot can only generate one new node (7). It rotates to face that node, and then moves towards it etc.*

nodes. For example, if the robot is located at node 1 and generates three pend- ing nodes (2,3 and 4), the paths will be (1, 2), (1, 3) and (1, 4). The robot next selects the nearest node, based on the *path length* as the **target node**. In many cases (e.g. when more than one node can be placed), several nodes are at the same (estimated) distance from the robot. In such cases, the robot (arbitrarily)

selects one of those nodes as the target node. For the paths just described, the path length equals the cartesian distance between the nodes. If a path con- tains more than two elements, however, the path length will differ from the cartesian distance, unless all the nodes in the path lie along a straight line. The path length is more relevant since, when executing the exploration algorithm described here, the robot will generally follow the path, even though direct movement between the active node and a target node is also possible under certain circumstances; see below.

Next, the robot rotates to face the target node, and then proceeds towards it; see the upper left panel of Fig. 6.2. During the motion, one of two things can happen: Either (i) the robot reaches the target node or, (ii) using the output from a *Proximity detection* brain process (assumed available), it concludes that the target node cannot be reached along the current path. Note that, in order for the *Proximity detection* brain process to be useful, the sensors it uses should be mounted at a different (smaller) height compared to the LRF.

In case (i), illustrated in the upper right panel of Fig. 6.2, once the target node has been reached, it is given the status *completed* and is then set as the new active node. At this point, the paths to the remaining pending nodes are updated. Continuing with the example above, if the robot moves to node 4, the paths to the other pending nodes (2 and 3) will be updated to (4, 1, 2) and (4, 1, 3). Furthermore, having reached node 4, the robot generates new pend- ing nodes. Note that the robot need not be located exactly *on* node 4; instead, a node is considered to be reached when the robot passes within a distance a from it. The new nodes are added as described above. The minimum distance requirement between added nodes (see above) is also enforced. Proceeding with the example, the robot might, at this stage, add nodes 5 and 6, with the paths (4, 5) and (4 ,6). Again, the robot selects the nearest node based on the path length, rotates to face that node, and then starts moving towards it etc. Note that the robot can (and will) visit completed nodes more than once. How- ever, by the construction described above, only pending nodes can be target nodes.

In case (ii), i.e. when the target node cannot be reached, the node is assigned the status *unreachable*, and the robot instead selects another target node and proceeds towards it, along the appropriate path. This situation is illustrated in the lower left panel of Fig. 6.2: Here, using its *Proximity detection* brain process, the robot concludes that it cannot reach node 6. It therefore marks this node *unreachable*, sets it as the active node, and then sets the nearest pending node as the new target, in this case node 5. One may wonder why case (ii) can oc- cur, since the robot uses the LRF before assigning new nodes. The reason, of course, is that the LRF (which is assumed to be two-dimensional) only scans the arena at a given height, thus effectively only considering a horizontal slice of the arena. A low obstacle may therefore be missed, until the robot comes sufficiently close to it, so that the *Proximity detection* brain process can detect it.

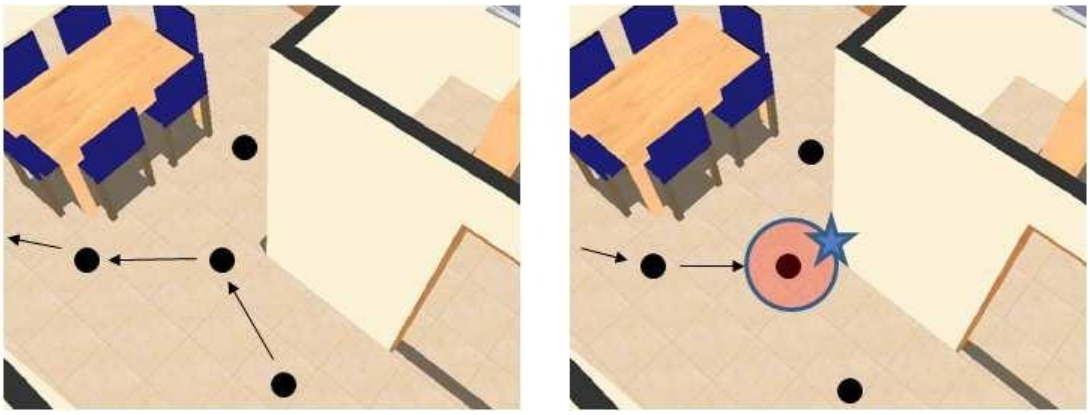

Figure 6.3: *An illustration of a problem that might occur during exploration. Moving in one particular direction (left panel) the robot is able to place and follow the nodes shown. However, upon returning (right panel), the robot may conclude that it will be unable to pass the node near the corner, due to the proximity detection triggered as the robot approaches the node, with the wall right in front of it.*

Note that unreachable nodes are exempt from the minimum distance require- ment. This is so, since a given node may be unreachable from *one* direction but perhaps reachable from some other direction. Thus, the exploration algorithm is allowed to place new pending nodes arbitrarily close to unreachable nodes.

One should note that robust exploration of any arena is more difficult than it might seem. An example of a problem that might occur is shown in Fig. 6.3. Here, the robot passes quite near a corner on its outbound journey (left panel), but no proximity detection is triggered. By contrast, upon returning (right panel) a proximity detection *is* triggered which, in turn, may force the robot to abandon its current path. In fact, the *Exploration* behavior contains a method (which will not be described here) for avoiding such deadlocks. In the (very rare) cases in which even the deadlock avoidance method fails, the robot sim- ply stops, and reports its failure.

Because of the path following strategy described above, the robot may sometimes take an unnecessarily long path from the active node to the tar- get node. However, this does not happen so often since, in most cases, the robot will proceed directly to a newly added node, for which the path length is the same as the cartesian distance. However, when the robot cannot place any more nodes (something that occurs, for example, when it reaches a cor- ner), a distant node may become the target node. Therefore, in cases where the path to the target node differs from the direct path, the robot rotates to face the target node (provided that it is located within a distance L, where L should be smaller than or equal to the range R of the LRF). Next, if the robot concludes (based on its LRF readings) that it can reach the target node, it then proceeds directly towards it, rather than following the path. However, also in this case,

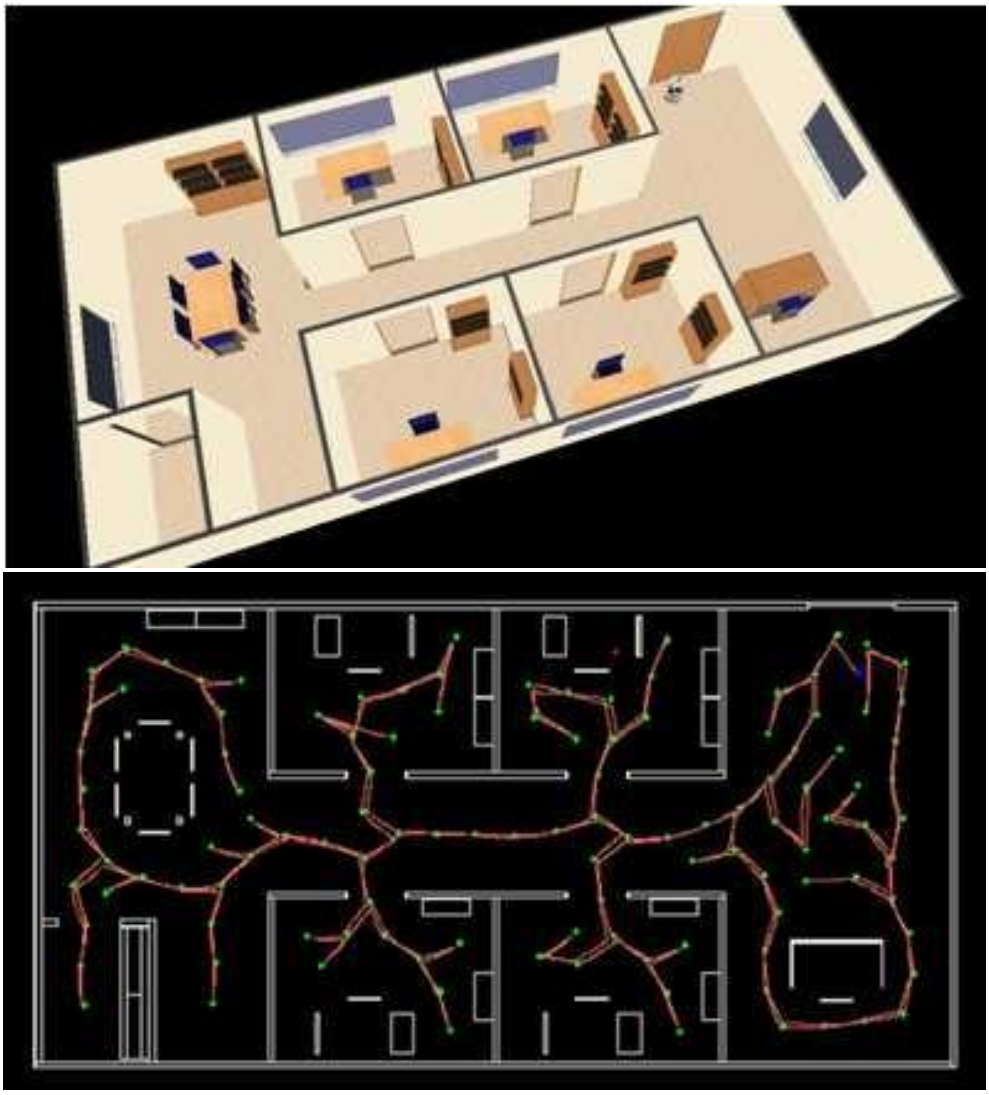

Figure 6.4: *Left panel: The robot in its initial position in an unexplored arena. Right panel: The final result obtained after executing the* Exploration *behavior. The completed (visited) nodes are shown as green dots, whereas the (single) unreachable node is shown as a red dot. The final target node (the last node visited) is shown as a blue dot. In this case, the robot achieved better than 99.5% sensory area coverage of the arena.*

it is possible that a (low) obstacle prevents the robot from reaching its target, in which case the robot instead switches to following the path as described above.

The robot continues this cycle of node placement and movement between nodes, until all nodes have been processed (i.e. either having been visited or deemed unreachable), at which point the exploration of the arena is complete. The *Exploration* behavior consists of an FSM with 17 states, which will not be

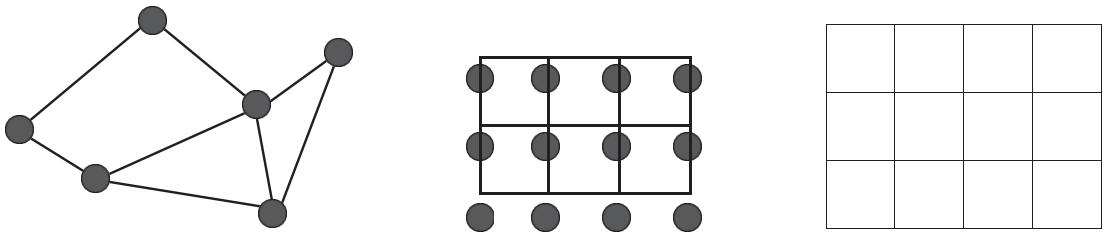

Figure 6.5: *Three examples of grids that can be used in connection with grid-based navigation methods. In the left panel, the nodes are not equidistant, unlike the middle panel which shows a regular lattice with equidistant nodes. The regular lattice can also be represented as grid cells, as shown in the right panel. Note that the right and middle panels show equivalent grids.*

described in detail here. A performance example is shown in Fig. 6.4. The left panel shows the robot at its starting point in a typical office arena. The right panel shows the final result, i.e. the path generated by the robot. The completed (visited) exploration nodes are shown as green dots, whereas the unreachable nodes (only one in this case) are shown as red dots. The final target node is shown as a blue dot. Note that the robot achieved a *sensory* area coverage (at the height of its LRF) of more than 99.5% during exploration.

6.5 Navigation

In this section, it will again be assumed that the robot has access to accurate estimates of its pose (from the *Odometry* brain process), and the question that will be considered is: Given that the robot knows its pose and velocity, how can it navigate between two arbitrary points in an arena? In the robotics liter- ature, many methods for navigation have been presented, three of which will be studied in detail in this section.

6.5.1 Grid-based navigation methods

In **grid-based navigation methods**, the robot's environment must be covered with an (artificial) grid, consisting of **nodes** (vertices) and **edges** connecting the nodes. The grid may have any shape, as illustrated in the left panel of Fig. 6.5, i.e. it need not be a rectangular lattice of the kind shown in the middle panel. However, if the grid happens to be a rectangular lattice, it is often rep- resented as shown in the right panel of the figure, where the nodes have been replaced by cells, and the edges are not shown[2]. Furthermore, the edges must be associated with a (non-negative) cost, which, in many cases is simply taken

[2]Note that in the cell representation in the right panel, the sides of each cell are *not* edges: The edges connect the centers of the grid cells to each other, as shown in the middle panel.

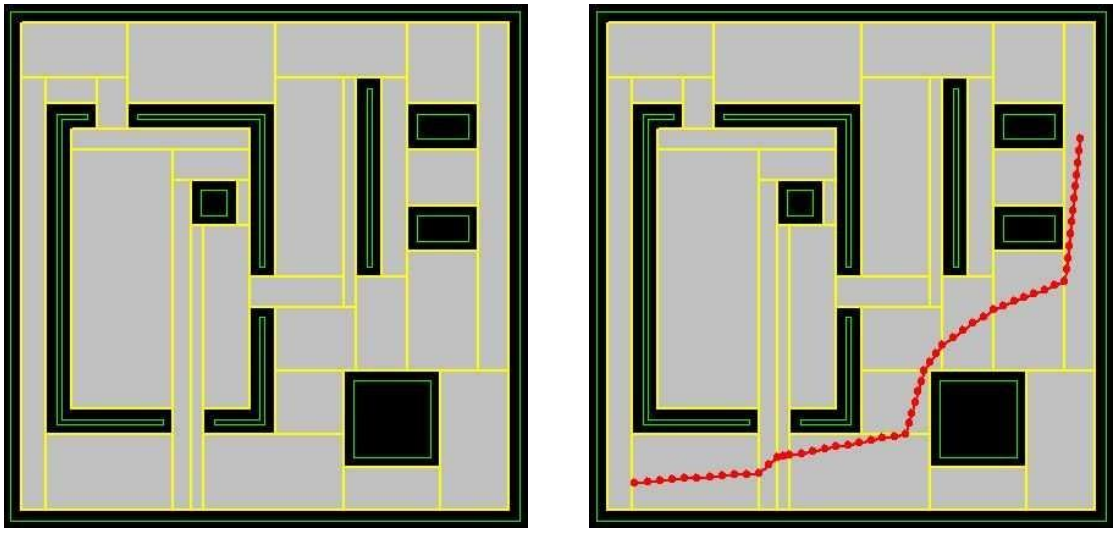

Figure 6.6: *Left panel: An example of automatic grid generation. The walls of the arena are shown as green thin lines. The black regions represent the forbidden parts of the arena, either unreachable locations or positions near walls and other obstacles. The grid cell boundaries are shown as thick yellow lines. Right panel: An example of a path between two points in the arena. The basic path (connecting grid cells) was generated using Dijkstra's algorithm (see below). The final path, shown in the figure, was adjusted to include changes of directions within grid cells, thus minimizing the length of the path. Note that all cells are convex, so that the path segments within a cell can safely be generated as straight lines between consecutive waypoints.*

as the euclidean distance between the nodes. Thus, for example, in the grids shown in the middle and right panels of Fig. 6.5, the cost of moving between adjacent nodes would be equal to 1 (length unit), whereas, in the grid shown in the left panel the cost would vary depending on which nodes are involved.

An interesting issue is the *generation* of a navigation grid, given a two-dimensional map of an arena. This problem is far from trivial, especially in complex arenas with many walls and other objects. Furthermore, the grid gen- eration should take the robot's size (with an additional margin) into account, in order to avoid situations where the robot must pass very close to a wall or some other object. The grid-based navigation methods described below gener- ate paths *between* grid cells. On a regular grid with small, quadratic cells (as in the examples below) it is sometimes sufficient to let the robot move on straight lines between the cell centers. However, the generated path may then become somewhat ragged. Furthermore, in more complex grids, where the cells are of different size, following a straight line between cell centers may result in an unnecessarily long path. Thus, in such cases, the robot must normally modify its heading *within* a cell, in order to find the shortest path.

When generating a grid, one normally requires the grid cells to be **convex**,

1. Place the robot at the start node, which then becomes the current node. Assign the status *unvisited* to all nodes.

2. Go through each of the cells a_i that are (i) unvisited and (ii) directly reach- able (via an edge) from the current node c. Such nodes are referred to as **neighbors** of the current node. Compute the cost of going from a_i to the target node t, using the heuristic $f(a_i)$.

3. Select the node $a_{\min}$ associated with the lowest cost, based on the cost values computed in Step 2.

4. Set the status of the current node c as *visited*, and move to $a_{\min}$ which then becomes the current node.

5. Return to Step 2.

Figure 6.7: *The best-first search algorithm.*

so that all points on a straight line between any two points in the cell also are part of the cell. One way of doing so is to generate a grid consisting of trian- gular cells, which will all be convex. However, such grids may not be optimal: The pointiness of the grid cells may force the robot to make many unnecessary (and sharp) turns. An algorithm for constructing, from a map, a general grid consisting of convex cells with four or more sides (i.e. non-triangular) exists as well[3]. Fig. 6.6 shows an example of a grid generated with this algorithm. Because of its complexity, the algorithm will not be considered in detail here. Instead, in the examples below, we shall consider grids consisting of small quadratic cells, and we will neglect changes of direction within grid cells.

Best-first search algorithm

In **best-first search** (BFS) algorithm the robot moves greedily towards the tar- get, as described in Fig. 6.7. As can be seen, the BFS method chooses the next node based on the (estimated) cost of going from that node n to the goal, which is estimated using a **heuristic** function $f(n)$. $f(n)$ can be chosen in different ways, the simplest being to use the euclidean distance between the node un- der consideration and the target. However, in that case, the BFS method may, in fact, get stuck. A more sophisticated heuristic function may, for example, add a penalty for each obstacle encountered on a straight-line path from the node under consideration to the target node.

The path can be generated by simply storing the list of visited nodes during

[3]See Wahde, M., Sandberg, D., and Wolff, K. *Reliable long-term navigation in indoor environments*, In: Topalov, A.V. (Ed.), Recent advances in Mobile Robots, InTech, 2011, pp. 261–286.

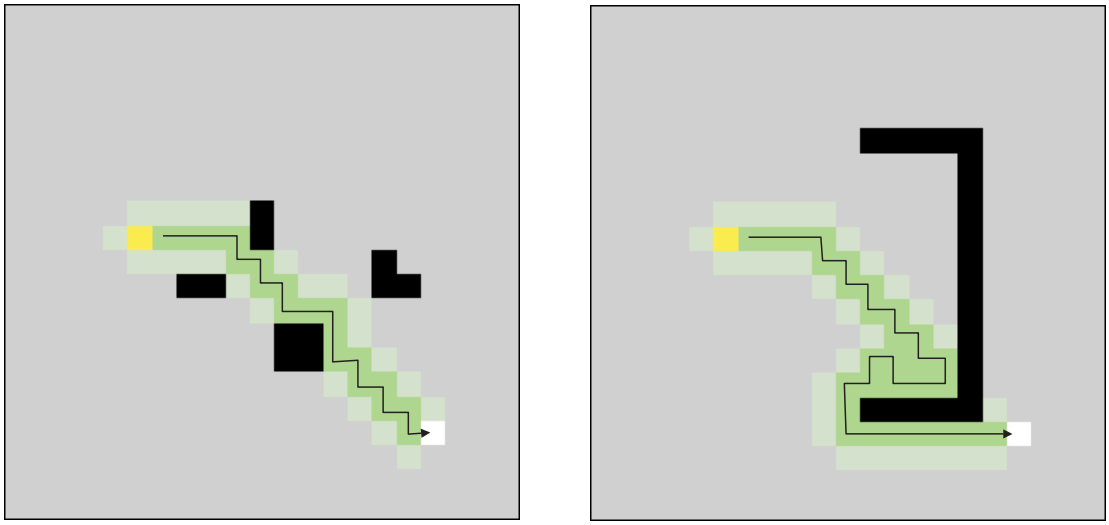

Figure 6.8: *Two examples of paths generated using the BFS algorithm. The cells (nodes) that were checked during path generation are shown in light green, whereas the actual path is shown in dark green and with a solid line. The yellow cell is the start node and the white cell is the target node.*

path generation. The BFS method is very efficient in the absence of obstacles or when the obstacles are few, small, and far apart. An example of such a path generated with BFS is shown in the left panel of Fig. 6.8. As can be seen, the robot quickly moves from the start node to the target node. However, if there are extended obstacles between the robot's current position and the target node, the BFS algorithm will not find the shortest path, as shown in the right panel of Fig. 6.8. Because of its greedy approach to the target, the robot will find itself in front of the obstacle, and must then make a rather long detour to arrive at the target node.

Dijkstra's algorithm

Like BFS, **Dijkstra's algorithm** also relies on a grid in which the edges are associated with non-negative costs. Here, the cost will simply be taken as the euclidean distance between nodes. Instead of focusing on the (estimated) cost of going from a given node to the target note, Dijkstra's algorithm considers the distance between the *start* node and the node under consideration, as de- scribed in Fig. 6.9. In Step 2, the distance from the start node s to any node a_i is computed using the (known) distance from the initial node to the current node c and simply adding the distance between c and a_i. This algorithm will check a large number of nodes, in an expanding pattern from the start node, as shown in Fig. 6.10. In order to determine the actual path to follow, whenever a new node a is checked, a note is made regarding the predecessor node p, i.e. the

7. Place the robot at the start node s, which then becomes the current node. Assign the distance value 0 to the start node, and infinity to all other nodes (in practice, use a very large, finite value). Set the status of all nodes to *unvisited*.

8. Go through all the unvisited, accessible (i.e. empty) neighbors a_i of the current node c, and compute their distance d from the *start* node s. If d is smaller than the previously stored distance d_i (initially infinite, see Step 1), then (i) update the stored distance, i.e. set $d_i = d$ and (ii) assign the current node as the predecessor node of a_i.

9. After checking all the neighbors of the current node, set its status to vis- ited.

10. Select the node (among *all* the unvisited, accessible nodes in the grid) with the smallest distance from the start node, and set it as the new cur- rent node.

11. Return to Step 2, unless the target has been reached.

12. When the target has been reached, use the predecessor nodes to trace a path from the target node to the start node. Finally, reverse the order of the nodes to find the path from the start node to the target node.

Figure 6.9: *Dijkstra's algorithm.*

node that was the current node when checking node a. When the target has been found, the path connecting it to the initial node can be obtained by going through the predecessor nodes backwards, from the target node to the initial node.

Unlike the BFS algorithm, Dijkstra's algorithm is guaranteed to find the shortest path[4] from the start node to the target node. However, a drawback with Dijkstra's algorithm is that it typically searches many nodes that, in the end, turn out to be quite irrelevant. Looking at the search patterns in Figs. 6.8 and 6.10, one may hypothesize that a *combination* of the two algorithms would be useful. Indeed, there is an algorithm, known as **A*** that combines the BFS and Dijkstra algorithms. Like Dijkstra's algorithm, A* is guaranteed to find the shortest path. Moreover, it does so more efficiently than Dijkstra's algorithm. However, A* is beyond the scope of this text.

[4]There may be more than one such path: Dijkstra's algorithm will select one of them.

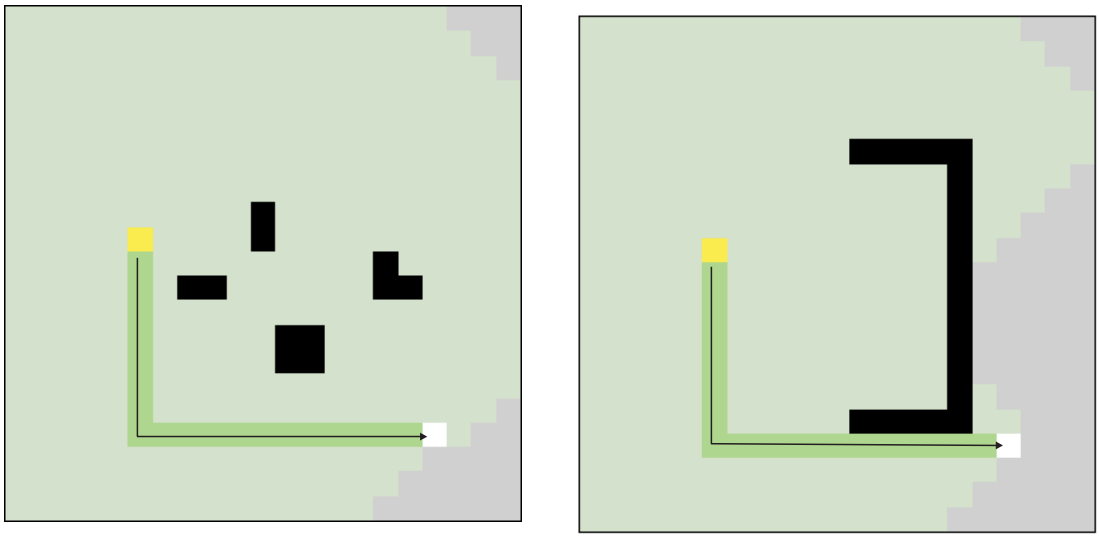

Figure 6.10: *Two examples of paths generated using Dijkstra's algorithm. The cells (nodes) that were checked during path generation are shown in light green, whereas the actual path is shown in dark green and with a solid line. The yellow cell is the start node and the white cell is the target node.*

6.5.2 Potential field navigation

Unlike the algorithms described above, the **potential field method** does not require a grid. In the potential field method, a robot obtains its desired direc- tion of motion as the negative gradient of an artificial potential field, generated by potentials assigned to the navigation target and to objects in the arena.

Potential fields

As shown in Fig. 6.11, a potential field can be interpreted as a landscape with hills and valleys, and the motion of a robot can be compared to that of a ball rolling through this landscape. The navigation target is assigned a potential corresponding to a gentle downhill slope, whereas obstacles should generate potentials corresponding to steep hills.

In principle, a variety of different equations could be used for defining dif- ferent kinds of potentials. An example, namely a potential with ellipsoidal equipotential surfaces, and exponential variation with (ellipsoidal) distance from the center of the potential, takes the mathematical form

$$\varphi(x, y; x_p, y_p, \alpha, \beta, \gamma) = \alpha e^{-\left(\frac{x-x_p}{\beta}\right)^2-\left(\frac{y-y_p}{\gamma}\right)^2}, \quad (6.1)$$

where (x, y) is the current (estimated) position at which the potential is calcu- lated, (x_p, y_p) is the position of the object generating the potential, and α, β and γ are constants (not to be confused with the constants defined in connection

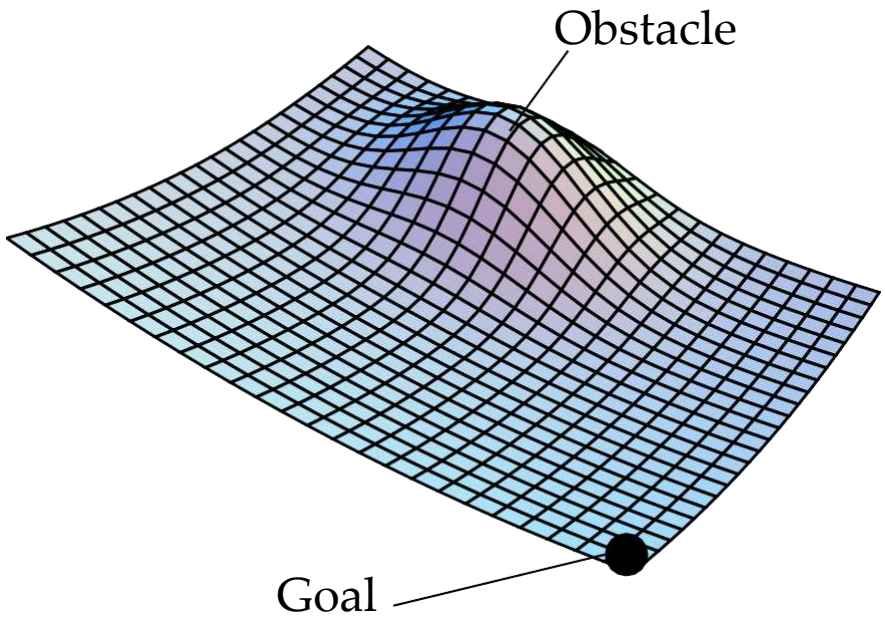

Figure 6.11: *A potential field containing a single obstacle and a navigation goal.*

with the equations of motion in Chapter 2 and the sensor equations in Chapter 3). Now, looking at the mathematical form of the potentials, one can see that an attractive potential (a valley) is formed if α is negative, whereas a positive value of α will generate a repulsive potential (a hill).

Normally, the complete potential field contains many potentials of the form given in Eq. (6.1), so that the total potential becomes

$$\Phi(x, y) = \sum_{i=1}^{k} \varphi_i(x, y; x_{p_i}, y_{p_i}, \alpha_i, \beta_i, \gamma_i), \tag{6.2}$$

where k is the number of potentials. An example of a potential field, for a simple arena with four central pillars, is shown in Fig. 6.12.

Navigating in a potential field

Once the potential field has been defined, the desired direction of motion $\hat{\mathbf{r}}$ of the robot can be computed as the negative of the normalized gradient of the field

$$\hat{\mathbf{r}} = -\frac{\nabla\Phi}{|\nabla\Phi|} \equiv -\frac{\left(\frac{\partial\Phi}{\partial x}, \frac{\partial\Phi}{\partial y}\right)}{\sqrt{\left(\frac{\partial\Phi}{\partial x}\right)^2 + \left(\frac{\partial\Phi}{\partial y}\right)^2}} \tag{6.3}$$

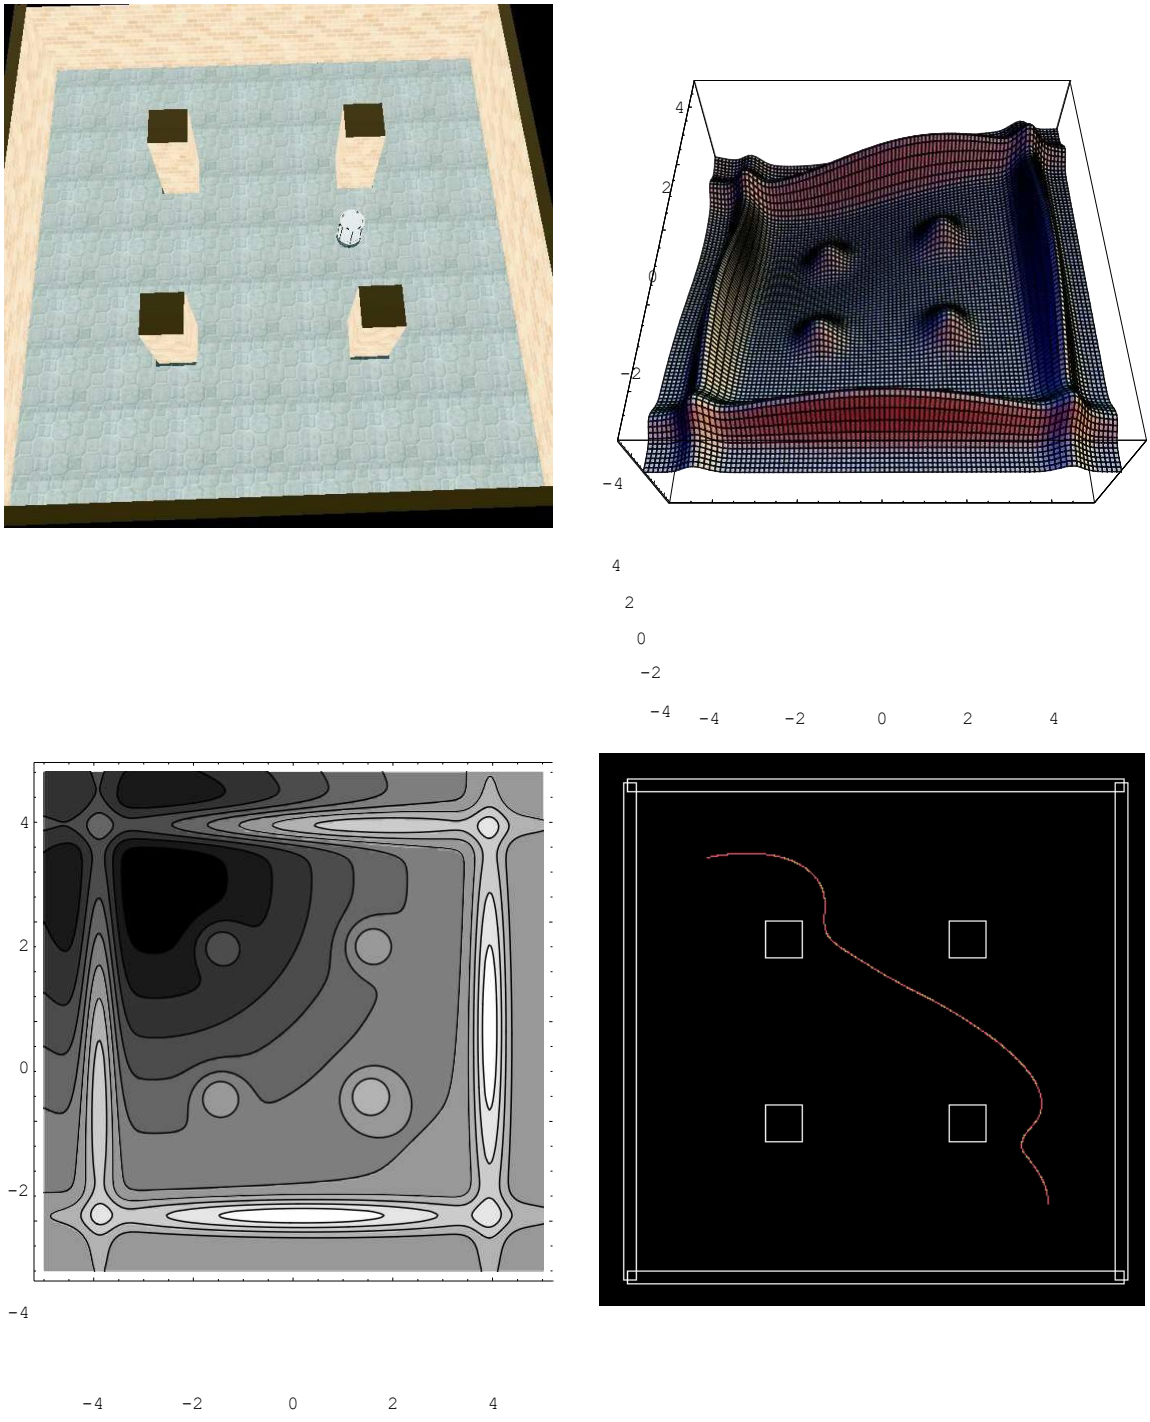

Figure 6.12: *An illustration of potential field navigation in GPRSim. Upper left panel: A simple arena, with a robot following a potential field toward a target position in the upper left corner of the arena. Upper right panel: The corresponding potential field, generated by a total of*

nine potentials (one for the target, one for each of the walls, and one for each pillar). Lower left panel: A contour plot of the potential field, in which the target position can be seen in the upper left corner. Lower right panel: The trajectory followed by the robot. Note that, in this simulation, the odometric readings were (unrealistically) noise-free.

In order to integrate the equations of motion of the robot, it is not sufficient only to know the desired direction: The magnitude of the force acting on the robot must also be known. In principle, the negative gradient of the potential field could be taken (without normalization) as the *force* acting on the robot, providing both magnitude and direction. However, in that case, the magni- tude of the force would vary quite strongly with the position of the robot, making the robot a dangerous moving object (if it is large). Thus, the poten- tial field is only used for providing the direction, as in Eq. (6.3). The robot's speed v (i.e. the magnitude of its velocity vector $\mathbf{v}$) can be assigned in various ways. For example, one may use proportional control to try to keep the speed constant[5].

An example of a trajectory generated during potential field navigation is shown in the lower right panel of Fig. 6.12. In the experiment in which this figure was generated, the noise in the odometric readings was (unrealistically) set to zero, since the aim here is simply to illustrate potential field navigation. However, in a realistic application, one would have to take into account the fact that the robot's estimate of its pose will never be error-free. Thus, when setting up a potential field, it is prudent to make the potentials slightly larger[6] than the physical objects that they represent. At the same time, in narrow corridors, one must be careful not to make the potentials (for walls on opposite sides of the corridor, say) so wide that the robot will be unable to pass.

In fact, the definition of a potential field for a given arena is something of an art. In addition to the problem of determining the effective extension of the potentials, one also has to decide whether a given object should be repre- sented by one or several potentials (for instance of the form given in Eq. (6.1)). For example, an extended object (for example, a long wall) *can* be represented as a single potential (typically with very different values of the parameters β and γ), but it can also be represented as a sequence of potentials. In complex environments, one may resort to stochastic optimization of the potential field, as well as the details of the robot's motion in the field[7].

Aspects of potential field navigation

A gradient-following method, such as the potential field method, always suf- fers the risk of encountering local minima in the field. Of course, in potential

[5]The procedure for assigning the robot's speed in potential field navigation will be de- scribed below.

[6]Of course, since the exponential potentials defined in Eq. (6.1) have infinite extension, the corresponding force never drops exactly to zero, but beyond a distance of a few d, where $d = \max(\beta, \gamma)$, the force is negligible.

[7]For an example of such an approach, see Savage *et al.*, *Optimization of waypoint-guided potential field navigation using evolutionary algorithms*, Proceedings of the 2004 IEEE/RSJ Inter- national Conference on Intelligent Robots and Systems (IROS 2004), 3463-3468, 2004.

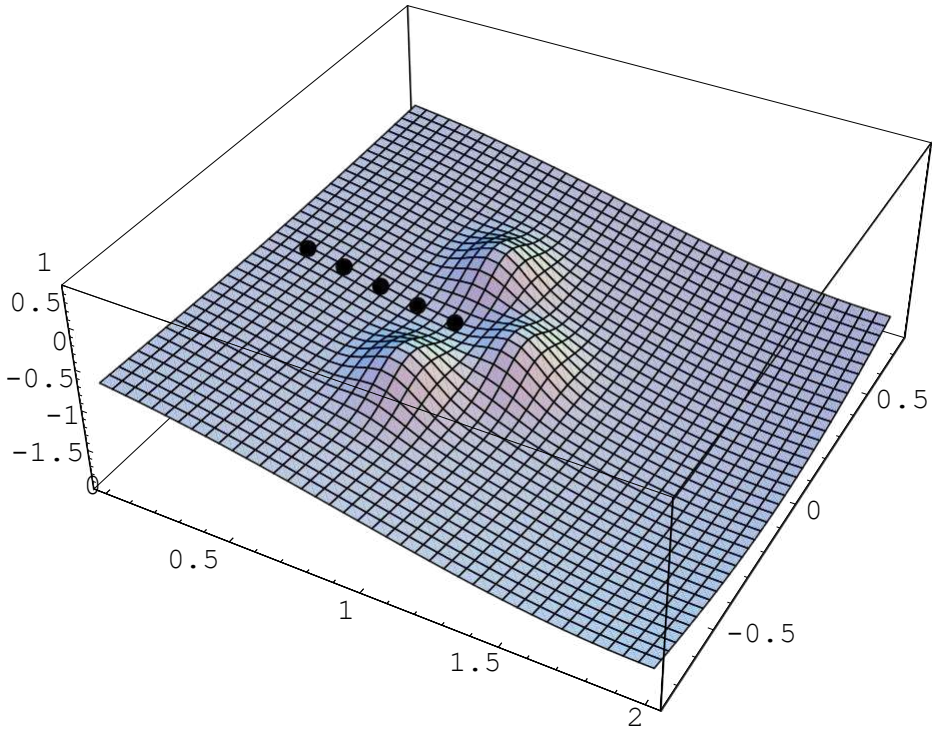

Figure 6.13: *The locking phenomenon. Following the gradient of the potential field the robot, whose trajectory is shown as a sequence of black dots, moves along the x-axis toward the goal, located at* (2, 0)*. However, because of the local minimum in the potential field, the robot eventually gets stuck.*

field navigation, the goal is to reach the local minimum represented by the navigation target. However, depending on the shape of the arena (and there- fore the potential field), there may also appear one or several unwanted local minima in the field, at which the robot may become trapped.

This is called the **locking phenomenon** and it is illustrated in Fig. 6.13. Here, a robot encounters a wedge-shaped obstacle represented by three poten- tials. At a large distance from the obstacle, the robot will be directed toward the goal potential, which is located behind the obstacle as seen from the start- ing position of the robot. However, as the robot approaches the obstacles their repulsive potentials will begin to be noticeable. Attracted by the goal, the robot will thus eventually find itself stuck inside the wedge, at a local minimum of the potential.

In order to avoid locking phenomena, the path between the robot and the goal can be supplied with **waypoints**, represented by attractive potentials (for example, of the form given in Eq. (6.1)) with rather small extension. Of course, the introduction of waypoints leads to the problem of determining where to

put them. An analysis of such methods will not be given here[8]. Suffice it to say that the problem of waypoint placement can be solved in various ways to aid the robot in its navigation. A waypoint should be removed once the robot has passed within a distance L from it, to avoid situations in which the robot finds itself stuck at a *waypoint*.

The potential field method also has several advantages, one of them being that the direction of motion is obtained simply by computing the gradient of the potential field at the current position, without the need to generate an en- tire path from the current position to the navigation target. Furthermore, the potential field is defined for all points in the arena. Thus, if the robot tem- porarily must suspend its navigation (for example, in order to avoid a moving obstacle), it can easily resume the navigation from wherever it happens to be located when the *Obstacle avoidance* behavior is deactivated.

In the discussion above, only stationary obstacles were considered. Of course, moving obstacles can be included as well. In fact, the potential field method is commonly used in conjunction with, say, a grid-based navigation method, such that the latter generates the nominal path of the robot, whereas the potential field method is used for adjusting the path to avoid moving ob- stacles. However, methods for reliably *detecting* moving obstacles are beyond the scope of this text.

Using the potential field method

As mentioned above, the potential field only provides the current desired *di- rection* of motion. In order to specify a potential field navigation behavior com- pletely, one must also provide a method for setting the speed of the robot. This can be done as follows: Given the robot's estimated (from odometry) angle of heading ϕ_{est} and the desired (reference) direction ϕ_{ref} (obtained from the po- tential field), one can form the quantity $\Delta\phi$ as

$$\Delta\phi = \phi_{\text{ref}} - \phi_{\text{est}}. \tag{6.4}$$

The desired speed differential ΔV (the difference between the right and left wheel speeds) can then be set according to

$$\Delta V = K_{\text{p}} V_{\text{nav}} \Delta\phi, \tag{6.5}$$

where K_{p} is a regulatory constant (P-regulation is used) and V_{nav} is the (de- sired) speed of the robot during normal navigation. Once ΔV has been com- puted, reference speeds are sent to the (velocity-regulated) motors according to

$$v_{\text{L}} = V_{\text{nav}} - \frac{\Delta V}{2}, \tag{6.6}$$

[8]See, however, the paper by Savage *et al.* mentioned in Footnote 6.

$$v_R = V_{nav} + \frac{\Delta V}{2}, \tag{6.7}$$

where v_R and v_L are the reference speeds of the left and right wheels, respec- tively. Note that one can of course only set the *desired* (reference) speed values; the actual speed values obtained depend on the detailed dynamics of the robot and its motors.

If the reference angle differs strongly from the estimated heading (which can happen, for example, in situations where the robot comes sufficiently close to an obstacle whose potential generates a steep hill), the robot may have to suspend its normal navigation and instead carry out a pure rotation, setting $v_L = V_{rot}$, $v_R = V_{rot}$ for a left (counterclockwise) rotation, where V_{rot} is the ro- tation velocity, defined along with V_{nav} (and the other constants) during setup. In case the robot should carry out a clockwise rotation, the signs are reversed. The direction of rotation is, of course, determined by the sign of the differ- ence between the reference angle and the (estimated) heading. In this case, the robot should turn until the difference between the reference angle and the esti- mated heading drops below a user-specified threshold, at which point normal navigation can resume.

6.6 Localization

In Sects. 6.1 and 6.2, it was (unrealistically) assumed that the robot's odome- try would provide perfect estimates of the pose. In reality, this will never be the case, and therefore the problem of recalibrating the odometric readings, from time to time, is a fundamental problem in robotics. Doing so requires a method for localization independent from odometry, and such methods usu- ally involve LRFs (even though cameras are also sometimes used), sensors that are difficult to simulate in ARSim (because of the large number of rays which would slow down the simulation considerably). Therefore, in this section, lo- calization will be described as it is implemented in the simulator GPRSim and in GPRBS, where LRFs are used.

Robot localization requires two brain processes: The cognitive *Odometry* process and an independent process for odometric recalibration, which both in GPRSim and in GPRBS goes under the name *Laser localization*, since the behavior for odometric recalibration uses the readings of an LRF, together with a map, to infer its current location using scan matching, as described below.

In fact, the problem of localization can be approached in many different ways. For outdoor applications, a robot may be equipped with GPS, which in many cases will give sufficiently accurate position estimates. However, in indoor applications (standard) GPS cannot be used, since the signal is too weak to penetrate the walls of a building. Of course, it is possible to set up a local GPS system, for example by projecting IR beacons on the ceiling, using which

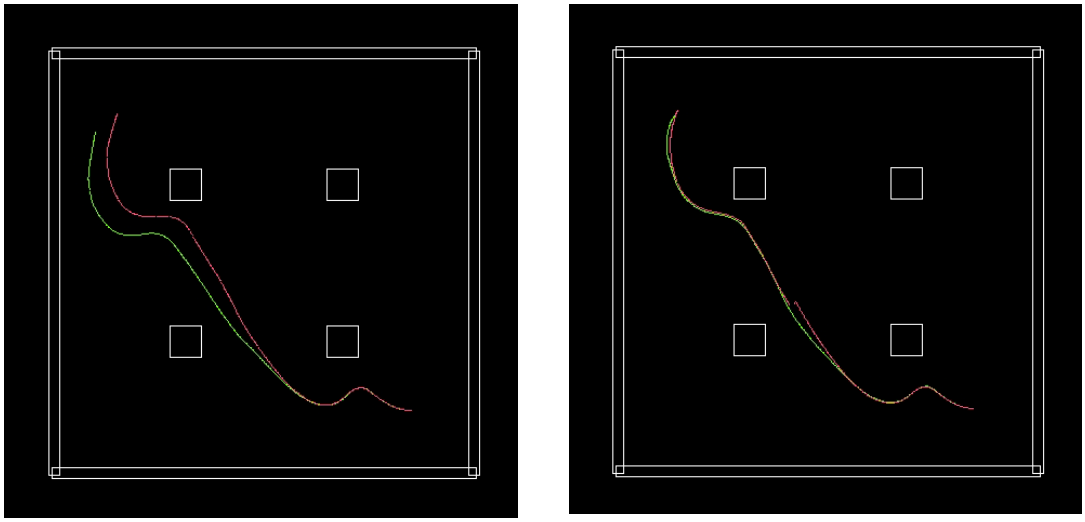

Figure 6.14: *An illustration of the need for localization in mobile robot navigation. In the left panel, the robot navigates using odometry only. As a result, the odometric trajectory (red) deviates quite significantly from the actual (green) trajectory. In the right panel,* Laser localization *was activated periodically, leading to much improved odometric estimates.*

the robot can deduce its position by means of triangulation[9]. However, such a system requires that the arena should be adapted to the robot, something that might not always be desirable or even possible.

The localization method (combining odometry and laser scan matching) that will be described here is normally used together with some navigation behavior. Thus, the robotic brain will consist of at least two motor behav- iors, in which case decision-making also becomes important. This topic will be studied in a later chapter: For now, the *Laser localization* behavior will be considered separately.

6.6.1 Laser localization

The behavior is intended for localization in arenas for which a map has been provided to the robot (in the form of a sequence of lines). The map can ei- ther be obtained using a robot (executing a *Mapping* behavior) or, for example, from the floor plan of a building. The behavior relies on scans of the arena using a two-dimensional LRF and, like many methods for localization in auto- nomous robots, it assumes that all scans are carried out in a horizontal plane, thus limiting the behavior to planar (i.e. mostly indoor) environments. In fact, the name *Laser localization* is something of a misnomer: The behavior does not actually carry out (continuous) localization. Instead, when activated, the behavior takes as input the current pose estimate and tries to improve it. If

[9]This is the method used in the **Northstar®** system, developed by Evolution Robotics, inc.

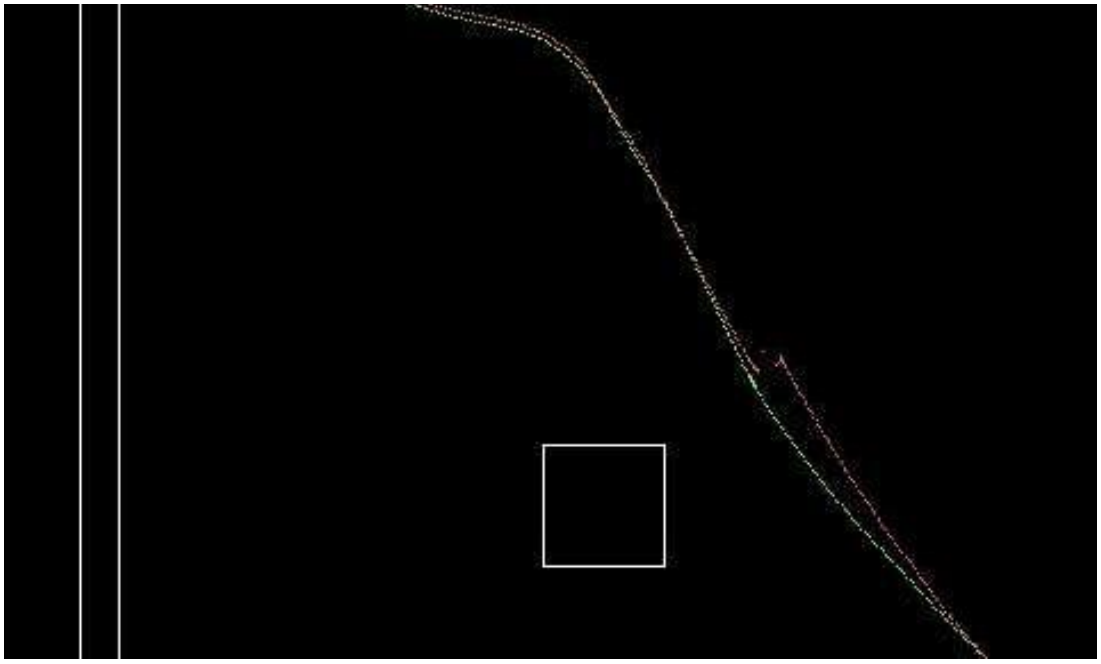

Figure 6.15: *An enlargement of the most significant correction in odometric readings (in the right panel of Fig. 6.14) resulting from the* Laser localization *behavior.*

successful, the odometric pose is reset to the position suggested by the *Laser localization* behavior.

The left panel of Fig. 6.14 illustrates the need for localization: The nav- igation task shown in Fig. 6.12 was considered again (with the same start- ing point, but a different initial direction of motion), this time with realistic (i.e. non-zero) levels of noise in the wheel encoders and, therefore, also in the odometry. As can be seen in the figure, the odometric drift causes a rather large discrepancy between the actual trajectory (green) and the odometric es- timate (red). In the right panel, the robotic brain contained two behaviors (in addition to the cognitive *Odometry* process), namely *Potential field navigation* and *Laser localization*. The *Laser localization* behavior was activated periodically (thus deactivating the *Potential field navigation* behavior), each time recalibrat- ing (if necessary) the odometric readings. As can be seen in the right panel of Fig. 6.14, with laser localization in place, the discrepancy between the odo- metric and actual trajectories is reduced significantly. At one point, the *Laser localization* behavior was required to make a rather large correction of the odo- metric readings. That particular event is shown enlarged in Fig. 6.15. As can be seen, the odometric readings undergo a discrete step at the moment of lo- calization.

When activated, the localization behavior[10] considered here first stops the

[10]See Sandberg, D., Wolff, K., and Wahde, M. *A robot localization method based on laser scan*

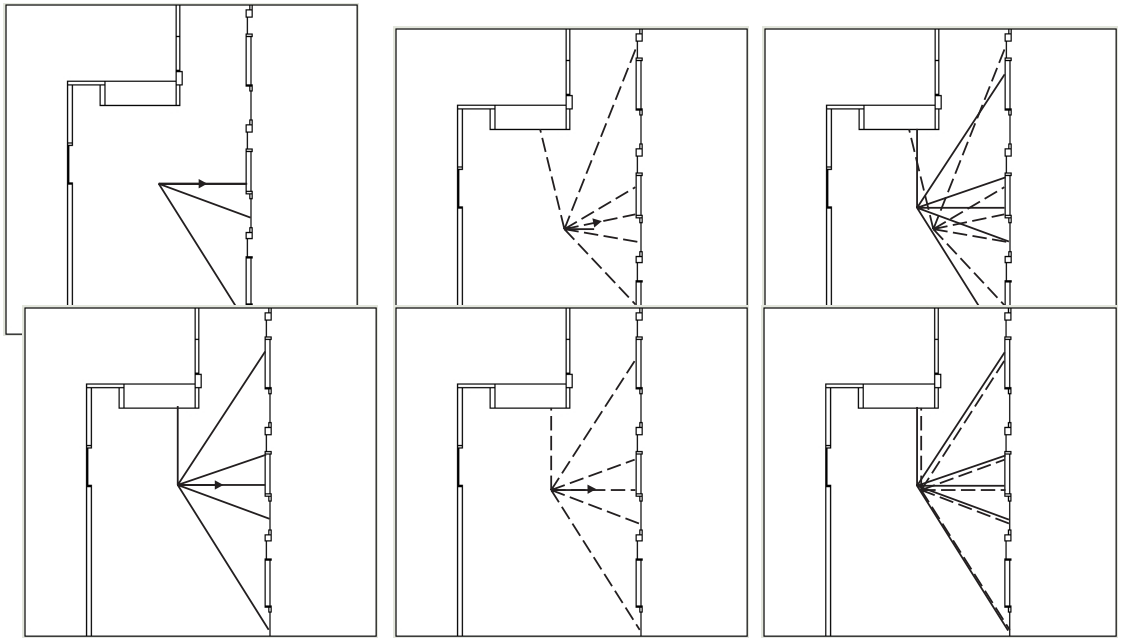

Figure 6.16: *Two examples of scan matching. The leftmost panel in each row shows a few rays (solid lines) from an actual LRF reading (plotted in the map used by the virtual LRF), and the middle panels show the virtual LRF readings (dotted lines) in a case in which the estimated pose differs quite strongly from the correct one (upper row), and one in which the difference is small (bottom row). The direction of heading is illustrated with arrows. The right panel in each row shows both the actual LRF rays and the virtual ones. The figure also illustrates the map, which consists of a sequence of lines.*

robot, and then takes a reading of the LRF. Next, it tries to match this reading to a *virtual* reading taken by placing a virtual LRF (hereafter: vLRF) at various positions *in the map*. Two examples of scan matching are shown in Fig. 6.16. The three panels in the upper row show a situation in which the odometry has drifted significantly. The upper left panel shows the readings (i.e. laser ray distances) from an actual LRF mounted on top of the robot (not shown). Note that, for clarity, the figure only shows a few of the many (typically hundreds) laser ray directions. The upper middle panel shows the readings of the vLRF, placed at the initial position and heading obtained from odometry. As can be seen in the upper right panel, the two scans match rather badly. By contrast, the three panels of the bottom row show a situation in which the pose error is small. The purpose of the search algorithm described below is to be able to correct the odometry, i.e. to reach a situation similar to the one shown in the bottom row of Fig. 6.16. Fig. 6.17 shows another example of a good (left panel) and a bad (right panel) scan match. In the case shown in the left panel, the

matching, Proc. of AMiRE 2009, pp. 171-178, 2009.

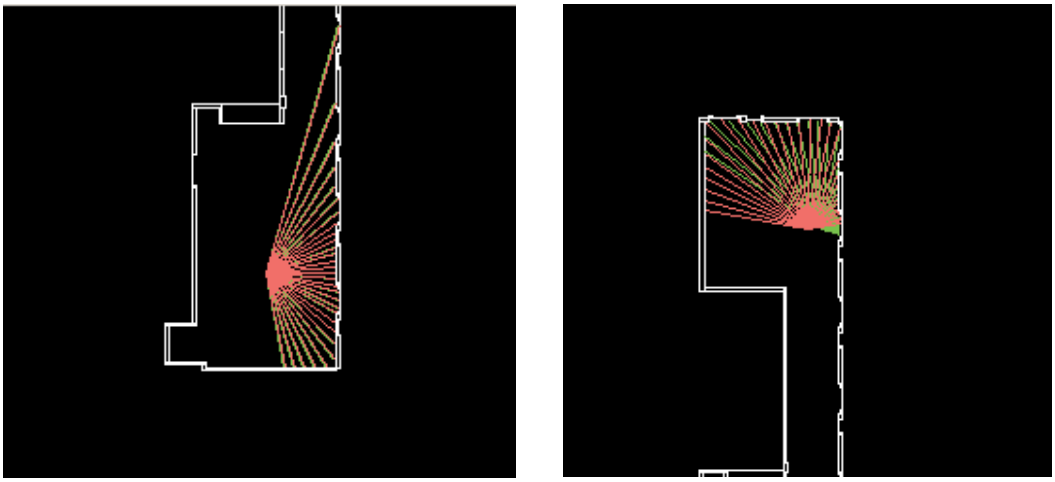

Figure 6.17: *Matching of LRF rays (in a different arena than the one used in the examples above). The readings of the actual LRF are shown in green, and those of the virtual LRF are shown in red. Left panel: An almost exact match. Right panel: In this case, the odometry has drifted enough to cause a large discrepancy between the actual and virtual LRF rays.*

odometric pose estimate is quite good, so that the rays from the actual LRF (green) match those of the vLRF quite well, at the current pose estimate. By constrast, in the situation shown in the right panel, the odometry has drifted significantly.

Scan matching algorithm

Let $\mathbf{p} = (x, y, \phi)$ denote a pose (in the map) of the vLRF. The distances between the vLRF and an obstacle, along ray i, are obtained using the map[11] and are de- noted δ_i. Similarly, the distances obtained for the *real* LRF (at its current pose, which normally differs from $\mathbf{p}$ when the localization behavior is activated) are denoted d_i.

The matching error s between two scans can be defined in various ways. For rays that do not intersect an obstacle, the corresponding reading (d_i or δ_i) is (arbitrarily) set to -1. Such rays should be excluded when computing the error. Thus, the matching error is taken as

$$s = \frac{1}{\nu} \sum_{i=1}^{n} \chi_i \left(1 - \frac{\delta_i}{d_i}\right)^2, \tag{6.8}$$

where n is the number of LRF rays used[12]. The parameter χ_i is equal to one

[11] In practice, the ray reading δ_i of the vLRF is obtained by checking for intersection between the lines in the map and a line of length R (the range of the LRF) pointing in the direction of the ray, and then choosing the shortest distance thus obtained (corresponding to the nearest obstacle along the ray). If no intersection is found, the corresponding reading is set to -1.

[12] For example, in the case of a Hokuyo URG-04LX LRF, a maximum of 682 rays are available.

for those indices i for which *both* the real LRF and the vLRF detect an obstacle (i.e. obtain a reading different from -1) whereas χ_i is equal to zero for indices i such that either the real LRF or the vLRF (or both) do not detect any obstacle (out to the range R of the LRF). v denotes the number of rays actually used in forming the error measure, i.e. the number of rays for which χ_i is equal to one. As can be seen, s is a measure of the (normalized) average relative deviation in detected distances between the real LRF and the vLRF.

Since d_i are given and δ_i depend on the pose of the vLRF, one may write $s = s(\mathbf{p})$. Now, if the odometric pose estimate happens to be exact, the virtual and actual LRF scans will be (almost) identical (depending on the accuracy of the map and the noise level in the real LRF), resulting in a very small matching error, in which case the localization behavior can be deactivated and the robot may continue its navigation. However, if the error exceeds a user-defined threshold T, the robot can conclude that its odometric estimates are not suf- ficiently accurate, and it must therefore try to minimize the matching error by trying various poses in the map, i.e. by carrying out a number of virtual scans, in order to determine the *actual* pose of the robot. The scan matching problem can thus be formulated as the optimization problem of finding the pose $\mathbf{p} = \mathbf{p}_v$ that minimizes $s = s(\mathbf{p})$. Once this pose has been found, the new odometric pose $\mathbf{p}^{\text{new}}$ is set equal to $\mathbf{p}_v$.

Note that it is assumed that the robot is standing still during localization.
This restriction (which, in principle, can be removed) is introduced in order to (i) avoid having to correct for the motion occuring during the laser sweep, which typically lasts between 0.01 and 0.1 s and (ii) avoid having to correct for the motion that would otherwise take place during scan matching procedure, which normally takes a (non-negligible) fraction of a second. Thus, only *one* scan needs to be carried out using the real LRF mounted on the robot. The re- maining work consists of generating virtual scans in the map, at a sequence of poses, and to match these to the actual LRF readings. Unlike some other scan matching methods, the method used here does not attempt to fit lines to the LRF readings. Instead, the LRF rays (actual and virtual, as described above) are used directly during scan matching. The sequence of poses for the vLRF is generated as follows: First the actual LRF scan is carried out, generating the distances d_i. Next, a virtual scan is carried out (in the map) at the current estimated position $\mathbf{p}_0$. If the error $s_0 = s(\mathbf{p}_0)$ is below the threshold T, local- ization is complete. If not, the algorithm picks a random pose $\mathbf{p}_j$ (where $j = 1$ in the first iteration) in a rectangular box of size $L_x \quad L_y \quad L_\phi$, centered on $\mathbf{p}_0$ in pose space, and computes the matching error $s_j = s(\mathbf{p}_j)$. The constants L_x and L_y are typically set to around 0.1 m and the constant L_ϕ is set to around
0.1 radians.

The process is repeated until, for some $j = j_1$, an error $s_{j_1} < s_0$ is found. At this point, the rectangular box is re-centered to $\mathbf{p}_{j_1}$, and the search continues, now picking a random pose in the rectangular box centered on $\mathbf{p}_{j_1}$. Once a po-

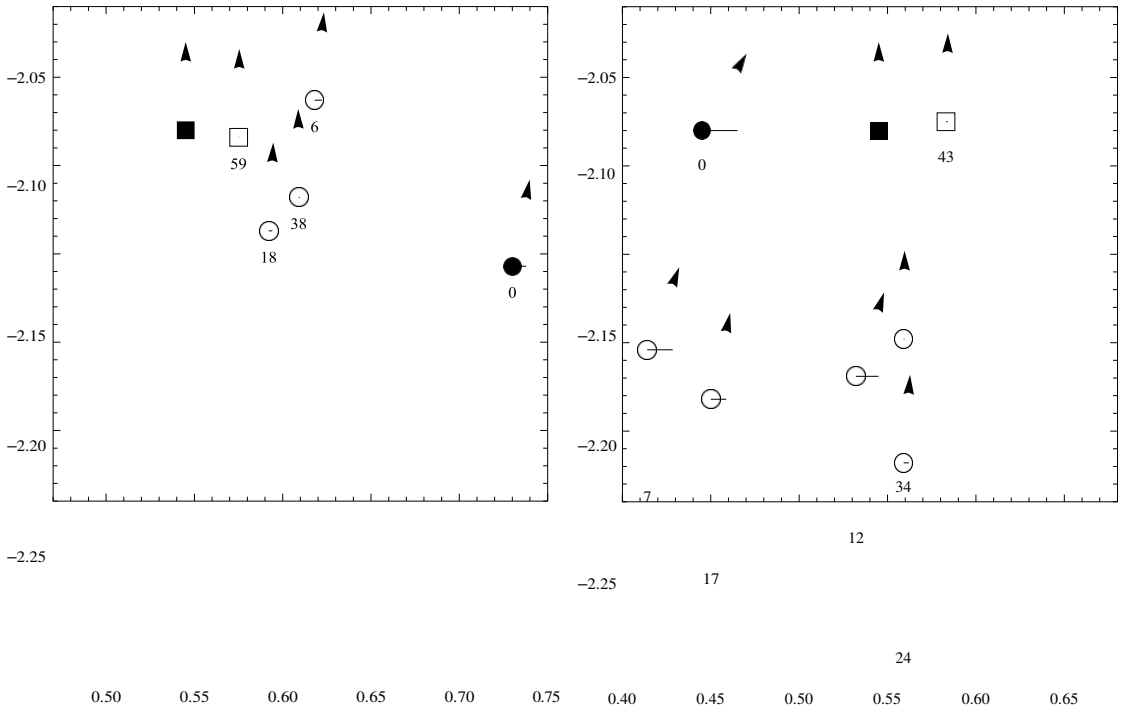
−2.05
−2.10
−2.15
−2.20
−2.25
6
59
38
18
0
0.50
0.55
0.60
0.65
0.70
0.75
−2.05
−2.10
−2.15
−2.20
−2.25
0
43
34
7
12
17
24
0.40
0.45
0.50
0.55
0.60
0.65

Figure 6.18: *An illustration of the sequence of poses generated during two executions of the search algorithm (with arrows indicating the direction of heading). In each panel, the actual position (measured in meters) of the robot is indicated with a filled square. The initial estimated position (i.e. from odometry, before correction) is shown as a filled disc, and the final estimated position is visualized as an open square. The intermediate points generated during the search are represented as open discs and are shown together with the corresponding iteration number. Note that, for clarity, only some of the intermediate points are shown.*

sition $\mathbf{p}_{j_2}$ is found for which $s_{j_2} \prec s_{j_1}$, the rectangular box is again re-centered etc. The procedure is repeated for a given number (N) of iterations[13].

Even though the algorithm is designed to improve both the position and the heading simultaneously, in practice, the result of running the algorithm is usually to correct the heading first (which is easiest, since an error in heading typically has a larger effect on the scan match than a position error), as can be seen clearly in the right panel of Fig. 6.18. At this stage, the estimated pose can make a rather large excursion in position space. However, once a fairly correct heading has been found, the estimated position normally converges quite rapidly to the correct position.

R

7. References

8. Ademco, "Quad Passive Infrared Motion *Dan:etch," 1989 Security Sourcebook: The Ademco Catalog of Praducts.* P6715, Adeoico Alarm Device Manufacturing Company, Syosset, NY, May, 1989.

9. Alpha, "Theory, Operation, and Application of Microwave Motion Sensing Modules," *Sensors,* pp. 29-36, December, 1987.
10. Bancroft, A.J., "The First Commercial Floor Care Company that Ventured into the Production of Robotics," Conference on Intelligent Robotics in Field, Pactory, Service, and Space, CIRFFSS '94, Houston, TX, pp. 669-674, March, 1994.
11. Barron, W.R., 'be Principles of Infrared Thermometry," *Sensors,* pp. 10-19, Oecmber, 1992.
12. Buschling, R., "Understanding aad Applying IR Temperature Sensors," *Sensors,*
13. pp. 32-37. October, 1994.
14. Byler, E., "Intelligent Mobile Sensor System for Drum Inspection and Monitoring," Phase I Topical Report, DOE Contract D£i-AC21-92MC29112. Martin Marietta Asronautics Group, Littleton, CO, June, 1993.
15. Cima, D., "Using Lithium Tantalate Pyroelectric Detectors in Robotics Applications," Eliecdata If 112, Eltec Instruments, Inc.. Daytona Beach, FL, 1984.
16. Cima, D., "Using Optical Radiation for Security," Eltecdata #124, Eltec
17. Instruments, Inc., Daytona Beach, FL, December, 1990.

18. Cima, D., "Surveillance Applications of the Eltec Model 862 Passive Infrared Telescope," Eltecdata fl28. Eltec Instruments, Inc., Daytona Beach, FL, June, 1992.
19. Eltec, "Introduction to Infrared Pyroelectric Detectors," Eltecdata #100, Eltec Instruments, Inc., Daytona Beach, FL, 1984.

20. Eltec, “Model 442 IR-Eye Integrated Sensor," Preliminary Product Literature, Rltec Instruments, Inc., Daytona Beach, FL, April, 1991.
21. Eltec, “Model AR170, 32 Element Pyroelectric Array,” Product Literature, Elicc Instruments, Inc., Daytona Beach, FL, December, 1993.
22. Everett, H.R., “A Computer Controlled Sentry Robot," *Robotics Age*, MarchfApril, 1982a.
23. Rverett, H.R., “A Microprocessor Controlled Autonomous Sentry Robot”, Masters Thesis, Naval Postgraduate School, Monterey, CA, October 1982b.
24. Everett, H.R., “Security and Sentry Robots", *International Encyclopedia of*
25. *Robotics Applications ond Automation,* R-*C. Oo*rf *, α‹L, John* Wiley. pp. 1462- 1476, March. 1988.
26. Everett, H.R., Gilbieath, G.A., Alderson, S.L., Priebe, C., Marchette. D.,
27. “Intelligent Security Assessment for a Mobile Sentry Robot”, Proceedings, 29th Annual Meeting. Institute for Nuclear Materials Management, Las Vegas, NV, June, 1988.

29. CHAPTER 6. EXPLORATION, NAVIGATION, AND LOCALIZATION

30. Everefl, H.R., Gilbreath, G.A.. Tran, T., Nieusma, J.M., "Modeling the Environment of a Mobile Security Robot," Technical Document 1835, Naval Command Control and Ocean Surveillance Center. San Diego, CA, June, 1990.
31. Gage, D.W., Everett, H.R.. Laird, R.T., Heath-Pastors, T.A., "Navigating Multiple Robots in Semi-Structured Environments." ANS 6th Topical Meeting on Robotics and Remote Systems, Monterey, CA, February, 1995.
32. Gailo, M.A., Willies, D.S., Lubke, R.A., Thiede, E.C., "Lnw Cost Uncooked m XI
33. Sensor for Battlefield Surveillance," SPIB Infrared Technology 2020, San Diego, CA, July, 1993.
34. George. S.C., "Robot Revival," *Security,* pp. 12-13, June, 1992.
35. V o l .
36. Gontowski, W., "Build a Motion Detector Alarm," *Electronic Experimenter's Handbook* pp. 56-64, 1983.
37. Hansford, A., AD9502 Video Signal Digitizer and its Application", Analog Devices Application Note Cl 100-9-7f87. Norwood, MA, July. 1987.
38. Hansen, C., Beratan, H., Owen. R.. Gorbin, M., McKenney, S., "Uncooled Thermal Imaging at Texas Instruments," SPIE Infrared Technology XVBI. Vol. 1735, San Diego, CA, pp. 17-26. July, 1992.
39. Hanson. C.. Beratan, H., Oven, R., Sweetser, K.. "Low-Cost Uncooled Pocal Plane Array Technology," Detector IRIS Meeting, Bedford, MA, August, 1993.
40. Hanson, C., Beratan, H. "Uncooked Pyroelectric Thermal Imaging," International Symposium on Applications of Ferroelectrics, 1994.
41. Heckendorn, F.M., Ward, C.W., Wagner. D.G., "Remote Radioactive Waste Orum Inspection with an Autonomouc Mobile Robot," ANS Fifth Topical Meeting on Robotics and Remote Systems, Amencan Nuclear Society, Knoxville, TN, pp. 487-492, Apri), 1993.
42. Holland, J.M., "An Army of Robots Roams the Night," International Robot and Vision Automation Show and Conference, Detroit, MI, pp. 17.1-17.12, April, 1993.
43. I5RA, "Military Finds Big Cost Savings from Mobile Robotics,"

ISRA News, Newsletter of the International Service Robot Association, Ann Arbor, MI, Fall, 1994.

44. Jones. J.L., Flynn, A.M., *Mabile Robots.- Inspiration to Implememation,* AK Peters, Ltd., Wellesley, MA, p. 113, 1993.

45. King, S.J., Weiman, C.F.R., "HelpMaie Autonomous Mobile Robot Navigation System," SPIE Vol. 1388, Mobile Robots V, Boston, MA, pp. 190-198, November, 1990.
46. La *wlot, M.,* "Microciicuit Technology Improves Readiness, Saves Resources," *Signal,* Armed Forces Communications and Electronics Association, August, 1993.
47. MacLeod, E.N., Chiarella, M., "Navigation and Control Breakthrough for
48. Automated Mobi1i'ty," Proceedings, SPIE Mobile Robots VIB, Vol. 2058, pp. 57-68, 1993.

49.

50. Mattaboni, P., "An Update on Lab Rover: A Hospital Material Transporter," Conference on Intelligent Robotics in Field, Factory, Service, and Space, CIRFFSS '94, Houston, TX, pp. 405-406. March, 1994.
51. Mims, F.M., *Forrest firms Circuit Scrapbook II,* Howard W. Sams, Indianapolis, IN, pp. 170-171, 1987.
52. Nippon, "Pyroelectric Infrared Sensor." Nippon Ceramics Technical Information TI-101, McGee Components, Inc, North Attleboto, MA, undated.
53. Nippon, "Pyrosensor," Nippon Ceramics Product Literature PE 1001-1091, McGee Components, Inc, North Attlebom, MA, October. 1991.
54. Philips, "Ceramic Pyroetectric Infrared Sensors and Their Applications," Philips Technical Publication 163, Philips Semiconductors, Slatersville Division. Smithfield, RI, 1985.
55. Philips, "Movement Sensing Using a Multi-Element Fresnel Lens," Philips Semiconductors, SJateisvi1)e Division, Smithfield, RI, April, 1986.
56. Quick, C.. "Animate vs. Inanimate," *Robotics Age,* VOI. 6. No, 9, August, 1984.
57. RST, "Surrogate Teleoperated Robotic Vehicle (STV) Technical Manual," Systems Technology, Westminster, MD, Contract No. N66tD1-91-C-6tD07, CDRL Item B00l, Final Issue. 15 September, 1993.
58. Russell, 1 A.. "Mobile Robot Guidance Using a Short-Lived Heat Trail."
59. *Robotics,* Vol. 11. Cambridge Ptess, pp. 427-431. 1993.
60. Savi, "The Savi Asset Management System," Product Brochure, Savi Technology, Inc., Mountain View, CA, 1993.
61. Savi. "System Components," Product Literature, Savi Technology, Inc., Mountain View, CA, April, 1994a.
62. Savi, "Savi Technology Ordering Guide." First £ldition, Radio Frequency Identification Equipment Contract No. F33600-94-D-0077, Savi Technology, Inc., Mountain View, CA, November, l994b.
63. Smurlo, R.P., Everett, H.R., "Intelligent Sensor Pusion for a Mobile Security Robot," *Sensors,* pp. 18-28, June, 1993.
64. Smurlo, R.P., Laird, R.T., Elaine, S., Jaffee, D.M., 'Hire MDARS Product Assessment System," Association of Unmanned Vehicle Systems, 22nd Annual Technical Symposium and Exhibition, Washington, DC, July. 1995.
65. *Tl, "Nighisight* Thermal Vision System," Pr‹›duet Literature, Texas Instruments, Inc., Attleboro, MA, November, 1994.
66. Tom, E., "Polymer Film Arrays in Pyroelectric Applications," *Sensors, pp.* 75- 77, September, 1994.
67. Viggh, H.E.M., Flynn, A.M., "Infrared People Sensors for Mobile Robots," SPIE Vol. 1007, Mobile Robots III, Cambridge, MA, pp. 391-398, November. 1988.
68. Weiss, M.. "Protect Your Valuables - Ligbt Sensitive Security Alert," fundic *Elecvonics,* April, 1979.
69. Williams. H., "Proximity Sensing with Microwave Technology," *Sensors,* pp. 6-

15, June, 1989.
70. williams, H., The Basic Principles of Microwave Proximity Sensing, *Sensors*, 71. pp. 26-28, May, 1991.

Books

[1] Adams, M.D., *Sensor Modelling, Design and Data Processing for Autonomous Navigation*. World Scientific Series in Robotics and Intelligent Systems. Singapore, World Scientific Publishing Co. Ltd., 1999.

[2] Arkin, R.C., *Behavior-Based Robotics*. Cambridge MIT Press, MA, 1998.

[3] Bar-Shalom, Y., Li, X.-R., *Estimation and Tracking: Principles, Techniques, and Software.* Norwood, MA, Artech House, 1993.

[4] Borenstein, J., Everet,t H.R., Feng, L., *Navigating Mobile Robots, Systems and Techniques*. Natick, MA, A.K. Peters, Ltd., 1996.

[5] Borenstein, J., Everett, H.R., Feng, L., *Where Am I? Sensors and Methods for Mobile Robot Positioning.* Ann Arbor, University of Michigan. Available at http:// www-personal.engin.umich.edu/~johannb/position.htm.

[6] Breipohl, A.M., *Probabilistic Systems Analysis: An Introduction to Probabilistic Models, Decisions, and Applications of Random Processes.* New York, John Wiley & Sons, 1970.

[7] Bundy, A. (editor), *Artificial Intelligence Techniques, a Comprehensive Catalogue*. New York, Springer-Verlag, 1997.

[8] Canudas de Wit, C., Siciliano, B., and Bastin G. (editors), *Theory of Robot Control*. New York, Spinger-Verlag, 1996.

[9] Carroll, R.J., Ruppert, D., *Transformation and Weighting in Regression*. New York, Chapman and Hall, 1988.

[10] Cox, I.J., Wilfong, G.T. (editors)*, Autonomous Robot Vehicles*. New York, Spinger-Verlag, 1990.

[11] Craig, J.J., *Introduction to Robotics: Mechanics and Control*. 2nd edition. Boston, Addison-Wesley, 1989.

[12] de Silva, C.W., *Control Sensors and Actuators.* Upper Saddle River, NJ, Prentice Hall, 1989.

[13] Dietrich, C.F., *Uncertainty, Calibration and Probability*. Bristol, UK, Adam Hilger, 1991.

[14] Draper, N.R., Smith, H., *Applied Regression Analysis*. 3rd edition. New York, John Wiley & Sons, 1988.

[15] Everett, H.R., *Sensors for Mobile Robots, Theory and Applications*. New York, Natick, MA, A.K. Peters, Ltd., 1995.

[16] Faugeras, O., *Three-Dimensional Computer Vision, a Geometric Viewpoint*. Cambridge MIT Press, MA, 1993.

[17] Genesereth, M.R., Nilsson, N.J., *Logical Foundations of Artificial Intelligence.* Palo Alto, CA, Morgan Kaufmann, 1987.

[18] Haralick, R.M., Shapiro, L.G., *Computer and Robot Vision, 1+2*. Boston, Addison-Wesley, 1993.

[19] Jones, J., Flynn, A., *Mobile Robots, Inspiration to Implementation*. Natick, MA, A.K. Peters, Ltd., 1993. [also available in German and French].

[20] Kortenkamp, D., Bonasso, R.P., Murphy, R.R. (editors), *Artificial Intelligence and Mobile Robots; Case Studies of Successful Robot Systems*. Cambridge, MA, AAAI Press / MIT Press, 1998.

[21] Latombe, J.-C., *Robot Motion Planning*. Norwood, MA, Kluwer Academic Publishers, 1991.

[22] Lee, D., *The Map-Building and Exploration Strategies of a Simple Sonar-Equipped Mobile Robot*. Cambridge, UK, Cambridge University Press, 1996.

[23] Leonard, J.E., Durrant-Whyte, H.F., *Directed Sonar Sensing for Mobile Robot Navigation*. Norwood, MA, Kluwer Academic Publishers, 1992.

[24] Manyika, J., Durrant-Whyte, H.F., *Data Fusion and Sensor Management: A Decentralized Information-Theoretic Approach*. Ellis Horwood, 1994.

[25] Mason, M., *Mechanics of Robotics Manipulation.* Cambridge, MA, MIT Press, 2001.

[26] Murphy, R.R., *Introduction to AI Robotics.* Cambridge, MA, MIT Press, 2000.

[27] Nourbakhsh, I., *Interleaving Planning and Execution for Autonomous Robots, Norwood, MA,* Kluwer Academic Publishers, 1997.

[28] Raibert, M.H., *Legged Robots That Balance*, Cambridge, MA, MIT Press, 1986.

[29] Ritter, G.X., Wilson, J.N., *Handbook of Computer Vision Algorithms in Image Algebra*. Boca Raton, FL, CRC Press, 1996.

[30] Russell, S., Norvig, P., *Artificial Intelligence, a Modern Approach*. Prentice Hall International, 1995.

[31] Schraft, R.-D., Schmierer, G., *Service Roboter. Natick, MA,* A.K. Peters, Ltd, 2000. [also available in German from Springer-Verlag].

[32] Sciavicco, L., Siciliano, B., *Modeling and Control of Robot Manipulators*. New York, McGraw-Hill, 1996.

[33] Todd, D.J, *Walking Machines, an Introduction to Legged Robots*. Kogan Page Ltd, 1985.

Papers

[34] Aho, A.V., "Algorithms for Finding Patterns in Strings," in J. van Leeuwen (editor), *Handbook of Theoretical Computer Science*, Cambridge, MA, MIT Press, 1990, Volume A, chapter 5, 255–300.

[35] Arras, K.O., Castellanos, J.A., Siegwart, R., "Feature-Based Multi-Hypothesis Localization and Tracking for Mobile Robots Using Geometric Constraints," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'2002)*, Washington, DC, May 11–15, 2002.

[36] Arras,K.O., Persson, J., Tomatis, N., Siegwart, R., "Real-Time Obstacle Avoidance for Polygonal Robots with a Reduced Dynamic Window," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA 2002),* Wash- ington, DC, May 11–15, 2002.

[37] Arras, K.O., Siegwart, R.Y., "Feature Extraction and Scene Interpretation for Map-Based Navigation and Map Building," in *Proceedings of SPIE, Mobile Robotics XII,* Vol. 3210, 1997, 42–53,.

[38] Arras, K.O., Tomatis, N., "Improving Robustness and Precision in Mobile Robot Localization by Using Laser Range Finding and Monocular Vision," in *Proceedings of the Third European Workshop on Advanced Mobile Robots (Eurobot 99)*, Zurich, September 6–9, 1999.

[39] Astolfi, A., "Exponential Stabilization of a Mobile Robot," *in Proceedings of 3rd European Control Conference,* Rome, September 1995.

[40] Barnard, K., Cardei V., Funt, B., "A Comparison of Computational Color Constancy Algorithms." *IEEE Transactions in Image Processing*. 11: 972–984, 2002.

[41] Barron, J.L., Fleet, D.J., Beauchemin, S.S., "Performance of Optical Flow Techniques." *International Journal of Computer Vision*, 12:43–77, 1994.

[42] Batavia, P., Nourbakhsh, I., "Path Planning for the Cye Robot," *in Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'00)*, Takamatsu, Japan, November, 2000.

[43] Borenstein, J., Koren, Y., "The Vector Field Histogram – Fast Obstacle Avoidance for Mobile Robots." *IEEE Journal of Robotics and Automation*, 7, 278–288, 1991.

[44] Brock, O., Khatib, O., "High-Speed Navigation Using the Global Dynamic Window Approach," *in Proceeding of the IEEE International Conference on Robotics and Automation*, Detroit, May 1999.

[45] Brooks, R., "A Robust Layered Control System for a Mobile Robot," *IEEE Transactions of Robotics and Automation*, RA-2:14–23, March 1986.

[46] Brown, H.B., Zeglin, G.Z., "The Bow Leg Hopping Robot'', *in Proceedings of the IEEE International Conference on Robotics and Automation*, Leuwen, Belgium, May 1998.

[47] Bruce, J., Balch,T., and Veloso, M., "Fast and Inexpensive Color Image Segmentation for Interactive Robots," *in Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'00)*, Takamatsu, Japan, October 31–November 5, 2000.

[48] Burgard,W., Cremers, A., Fox, D., Hahnel, D., Lakemeyer, G., Schulz, D., Steiner, W., Thrun, S., "Experiences with an Interactive Museum Tour-Guide Robot," *Arti- ficial Intelligence*, 114, 1–53, 2000.

[49] Burgard, W., Derr, A., Fox, D., Cremers, A., "Integrating Global Position Estimation and Position Tracking for Mobile Robots: The Dynamic Markov Localization Approach," *in Proceedings of the 1998 IEEE/RSJ International Conference of Intelligent Robots and Systems (IROS'98)*, Victoria, B.C., Canada, October 1998.

[50] Burgard, W., Fox, D., Henning, D., "Fast Grid-Based Position Tracking for Mobile Robots," *in Proceedings of the 21th German Conference on Artificial Intelligence (KI97)*, Freiburg, Germany, Springer-Verlag, 1997.

[51] Burgard, W., Fox, D., Jans, H., Matenar, C., Thrun, S., "Sonar-Based Mapping of Large-Scale Mobile Robot Environments using EM," *in Proceedings of the International Conference on Machine Learning*, Bled, Slovenia, 1999.

[52] Campion, G., Bastin, G., D'Andréa-Novel, B., "Structural Properties and Classification of Kinematic and Dynamic Models of Wheeled Mobile Robots." *IEEE Transactions on Robotics and Automation*, 12, No. 1, 47–62, 1996.

[53] Canudas de Wit, C., Sordalen, O.J., "Exponential Stabilization of Mobile Robots with Nonholonomic Constraints." *IEEE Transactions on Robotics and Automation*, 37, 1791–1797, 1993.

[54] Caprari, G., Estier, T., Siegwart, R., "Fascination of Down Scaling–Alice the Sugar Cube Robot." *Journal of Micro-Mechatronics*, 1, 177–189, 2002.

[55] Castellanos, J.A., Tardos, J.D., Schmidt, G., "Building a Global Map of the Environment of a Mobile Robot: The Importance of Correlations," *in Proceedings of the* 1997 *IEEE Conference on Robotics and Automation*, Albuquerque, NM, April 1997.

[56] Chen, C.T., Quinn, R.D., "A Crash Avoidance System Based upon the Cockroach Escape Response Circuit," *in Proceedings of the IEEE International Conference on Robotics and Automation*, Albuquerque, NM, April 1997.

[57] Chenavier, F., Crowley, J.L., "Position Estimation for a Mobile Robot Using Vision and Odometry," *in Proceedings of the IEEE International Conference on Robotics and Automation*, Nice, France, May 1992.

[58] Chong, K.S., Kleeman, L., "Accurate Odometry and Error Modelling for a Mobile Robot," *in Proceedings of the IEEE International Conference on Robotics and Automation*, Albuquerque, NM, April 1997.

[59] Choset, H., Walker, S., Eiamsa-Ard, K., Burdick, J., "Sensor-Based Exploration: Incremental Construction of the Hierarchical Generalized Voronoi Graph." *The International Journal of Robotics Research,* 19, 126–148, 2000.

[60] Cox, I.J., Leonard, J.J., "Modeling a Dynamic Environment Using a Bayesian Multiple Hypothesis Approach." *Artificial Intelligence*, 66, 311–44, 1994.

[61] Dowlingn, K., Guzikowski, R., Ladd, J., Pangels, H., Singh, S., Whittaker, W.L., "NAVLAB: An Autonomous Navigation Testbed." *Technical report CMU-RI-TR- 87-24, Robotics Institute*, Pittsburgh, Carnegie Mellon University, November 1987.

[62] Dugan, B., "Vagabond: A Demonstration of Autonomous, Robust Outdoor Navigation," *in Video Proceedings of the IEEE International Conference on Robotics and Automation*, Atlanta, GA, May 1993.

[63] Elfes, A., "Sonar-Based Real World Mapping and Navigation," *in* [10].

[64] Ens, J., Lawrence, P., "An Investigation of Methods for Determining Depth from Focus." *IEEE Transactions. on Pattern Analysis and Machine Intelligence*, 15: 97–108, 1993.

[65] Espenschied, K. S., Quinn, R. D., "Biologically-Inspired Hexapod Robot Design and Simulation," *in AIAA Conference on Intelligent Robots in Field, Factory, Ser- vice and Space*, Houston, Texas, March 20–24, 1994.

[66] Falcone, E., Gockley, R., Porter, E., Nourbakhsh, I., "The Personal Rover Project: The Comprehensive Design of a Domestic Personal Robot," *Robotics and Autonomous Systems, Special Issue on Socially Interactive Robots,* 42, 245–258, 2003.

[67] Feder, H.J.S., Slotin, J-J.E., "Real-Time Path Planning Using Harmonic Potentials in Dynamic Environments," in *Proceedings of the IEEE International Conference on Robotics and Automation*, Albuquerque, NM, April 1997.

[68] Fox, D., "KLD-Sampling: Adaptive Particle Filters and Mobile Robot Localiza- tion." *Advances in Neural Information Processing Systems 14*. MIT Press, 2001.

[69] Fox, D., Burgard,W., Thrun, S., "The Dynamic Window Approach to Collision Avoidance." *IEEE Robotics and Automation Magazine*, 4:23–33, 1997.

[70] Gander,W., Golub, G.H., Strebel, R., "Least-Squares Fitting of Circles and Ellipses." *BIT Numerical Mathematics*, vol. 34, no. 4, pp. 558–578, December 1994.

[71] Genesereth, M.R. "Deliberate Agents." *Technical Report Logic-87-2*. Stanford, CA, Stanford University, Logic Group, 1987.

[72] Golub, G., Kahan,W., "Calculating the Singular Values and Pseudo-Inverse of a Matrix." *Journal SIAM Numerical Analysis*, 2:205–223, 1965.

[73] Gutmann, J.S., Burgard, W., Fox, D., Konolige, K., "An Experimental Comparison of Localization Methods," *in Proceedings of the 1998 IEEE/RSJ International. Conference of Intelligent Robots and Systems* (IROS'98), Victoria, B.C., Canada, October 1998.

[74] Guttman, J.S., Konolige, K., "Incremental Mapping of Large Cyclic Environments," *in Proceedings of the IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA)*, Monterey, November 1999.

[75] Hashimoto, S., "Humanoid Robots in Waseda University - Hadaly-2 and WABIAN," *in IARP First International Workshop on Humanoid and Human Friendly Robotics*, Tsukuba, Japan, October 1998.

[76] Heale, A., Kleeman, L.: "A Real Time DSP Sonar Echo Processor," *in Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'00)*, Takamatsu, Japan, October 31–November 5, 2000.

[77] Horn, B.K.P., Schunck, B.G., "Determining Optical Flow," *Artificial Intelligence*, 17:185–203, 1981.

[78] Horswill, I., "Visual Collision Avoidance by Segmentation," *in Proceedings of IEEE International Conference on Robotics and Automation*, 902–909, 1995, IEEE Press, Munich, November 1994.

[79] Jacobs, R. and Canny, J., "Planning Smooth Paths for Mobile Robots," *in Proceeding. of the IEEE Conference on Robotics and Automation*, IEEE Press,1989, pp. 2–7.

[80] Jennings, J., Kirkwood-Watts, C., Tanis, C., "Distributed Map-making and Navigation in Dynamic Environments," *in Proceedings of the 1998 IEEE/RSJ Intl. Conference of Intelligent Robots and Systems (IROS'98)*, Victoria, B.C., Canada, October 1998.

[81] Jensfelt, P., Austin, D., Wijk, O., Andersson, M., "Feature Based Condensation for Mobile Robot Localization," *in Proceedings of the IEEE International Conference on Robotics and Automation*, San Francisco, May 24–28, 2000.

[82] Kamon, I., Rivlin, E., Rimon, E., "A New Range-Sensor Based Globally Conver- gent Navigation Algorithm for Mobile Robots," *in Proceedings of the IEEE International Conference on Robotics and Automation*, Minneapolis, April 1996.

[83] Kelly, A., "Pose Determination and Tracking in Image Mosaic Based Vehicle Position Estimation," *in Proceeding of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'00)*, Takamatsu, Japan, October 31–November 5, 2000.

[84] Khatib, M., Chatila, R., "An Extended Potential Field Approach for Mobile Robot Sensor-Based Motions," *in Proceedings of the Intelligent Autonomous Systems IAS-4*, IOS Press, Karlsruhe, Germany, March 1995, pp. 490–496.

[85] Khatib, M., Jaouni, H., Chatila, R., Laumod, J.P., "Dynamic Path Modification for Car-Like Nonholonomic Mobile Robots," *in Proceedings of IEEE International Conference on Robotics and Automation*, Albuquerque, NM, April 1997.

[86] Khatib, O., Quinlan, S., "Elastic Bands: Connecting, Path Planning and Control," *in Proceedings of IEEE International Conference on Robotics and Automation*, Atlanta, GA, May 1993.

[87] Ko, N.Y., Simmons, R., "The Lane-Curvature Method for Local Obstacle Avoidance," *in Proceedings of the 1998 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'98)*, Victoria, B.C., Canada, October 1998.

[88] Koenig, S., Simmons, R., "Xavier: A Robot Navigation Architecture Based on Partially Observable Markov Decision Process Models," *in* [20]

[89] Konolige, K.,. "A Gradient Method for Realtime Robot Control," *in Proceedings of the IEEE/RSJ Conference on Intelligent Robots and Systems*, Takamatsu, Japan, 2000.

[90] Konolige, K., "Small Vision Systems: Hardware and Implementation," in P*roceedings of Eighth International Symposium on Robotics Research*, Hayama, Japan, October 1997.

[91] Koperski, K., Adhikary, J., Han, J., "Spatial Data Mining: Progress and Challenges Survey Paper,". *in Proceedings of the ACM SIGMOD Workshop on Research Issues on Data Mining and Knowledge Discovery*, Montreal, June 1996.

[92] Koren, Y., Borenstein, J., "High-Speed Obstacle Avoidance for Mobile Robotics," *in Proceedings of the IEEE Symposium on Intelligent Control*, Arlington, VA, August 1988, pp. 382-384.

[93] Koren, Y., Borenstein, J., "Real-Time Obstacle Avoidance for Fast Mobile Robots in Cluttered Environments," *in Proceedings of the IEEE International Conference on Robotics and Automation*, Los Alamitos, CA, May 1990.

[94] Kuipers, B., Byun, Y.-T., "A Robot Exploration and Mapping Strategy Based on a Semantic Hierarchy of Spatial Representations." *Journal of Robotics and Autonomous Systems*, 8:47–63, 1991.

[95] Lamon, P., Nourbakhsh, I., Jensen, B., Siegwar,t R., "Deriving and Matching Image Fingerprint Sequences for Mobile Robot Localization," *in Proceedings of the 2001 IEEE International Conference on Robotics and Automation*, Seoul, Korea, May 21–26, 2001.

[96] Latombe, J-C., Barraquand, J., "Robot Motion Planning: A Distributed Presentation Approach." *International Journal of Robotics Research*, 10: 628–649, 1991.

[97] Lauria, M., Estier, T., Siegwart, R.: "An Innovative Space Rover with Extended Climbing Abilities." *in Video Proceedings of the 2000 IEEE International Conference on Robotics and Automation*, San Francisco, May 24–28, 2000.

[98] Lazanas, A., Latombe, J.-C., "Landmark-Based Robot Navigation," *in Proceedings of the Tenth National Conference on AI*. San Jose, CA, July 1992.

[99] Lazanas, A. Latombe, J.C., "Motion Planning with Uncertainty: A Landmark Approach." *Artificial Intelligence*, 76:285–317, 1995.

[100] Lee, S.-O., Cho, Y.-J., Hwang-Bo, M., You, B.-J., Oh, S.-R.: "A Stabile Target-Tracking Control for Unicycle Mobile Robots," *in Proceedings of the 2000 IEEE/ RSJ International Conference on Intelligent Robots and Systems*, Takamatsu, Japan, October 31–November 5, 2000.

[101] Lumelsky, V., Skewis, T., "Incorporating Range Sensing in the Robot Navigation Function." *IEEE Transactions on Systems, Man, and Cybernetics*, 20:1990, pp. 1058–1068.

[102] Lumelsky, V., Stepanov, A., "Path-Planning Strategies for a Point Mobile Automaton Moving Amidst Unknown Obstacles of Arbitrary Shape," in [10].

[103] Maes, P., "The Dynamics of Action Selection," *in Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, Detroit, 1989, pp. 991–997.

[104] Maes, P., "Situated Agents Can Have Goals," *Robotics and Autonomous Systems*, 6: 49–70. 1990.

[105] Martinelli, A., Siegwart, R., "Estimating the Odometry Error of a Mobile Robot during Navigation,'' *in Proceedings of the European Conference on Mobile Robots (ECMR 2003)*, Warsaw, September 4–6, 2003.

[106] Maybeck,P.S., "The Kalman Filter: An Introduction to Concepts," *in* [10].

[107] Minguez, J., Montano, L., "Nearness Diagram Navigation (ND): A New Real Time Collision Avoidance Approach," *in Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Takamatsu, Japan, October 2000.

[108] Minguez, J., Montano, L., "Robot Navigation in Very Complex, Dense, and Cluttered Indoor / Outdoor Environments," *in Proceeding of International Federation of Automatic Control (IFAC2002)*, Barcelona, April 2002.

[109] Minguez, J., Montano, L., Khatib, O., "Reactive Collision Avoidance for Navigation with Dynamic Constraints," *in Proceedings of the 2002 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2002.

[110] Minguez, J., Montano, L., Simeon, T., Alami, R., "Global Nearness Diagram Navigation (GND)," *in Proceedings of the 2001 IEEE International Conference on Robotics and Automation*, 2001.

[111] Montano, L., Asensio, J.R., "Real-Time Robot Navigation in Unstructured Environments Using a 3D Laser Range Finder," *in Proceedings of the IEEE/RSJ International Conference on Intelligent Robot and Systems*, IROS 97, IEEE Press, Vol. 2. Grenoble, France, September 1997, pp. 526–532.

[112] Moravec, H. and Elfes, A.E., "High Resolution Maps from Wide Angle Sonar," *in Proceedings of the 1985 IEEE International Conference on Robotics and Automation*, IEEE Press, At.-Louis, MO, March 1985, pp. 116–121.

[113] Moutarlier, P., Chatila, R., "Stochastic Multisensory Data Fusion for Mobile Robot Location and Environment Modelling," *in Proceedings of the 5th International Symposium of Robotics Research*, Tokyo, 1989, pp. 207–216.

[114] Nayar, S.K., "Catadioptric Omnidirectional Camera." *IEEE CVPR*, pp. 482–488, 1997.

[115] Nilsson, N.J., "Shakey the Robot." *SRI, International, Technical Note*, Menlo Park, CA, 1984, No. 325.

[116] Nourbakhsh, I.R., "Dervish: An Office-Navigation Robot," *in* [20].

[117] Nourbakhsh, I.R., Andre. D., Tomasi, C., Genesereth, M.R., "Mobile Robot Obstacle Avoidance via Depth from Focus," *Robotics and Autonomous Systems*, 22:151–158, 1997.

[118] Nourbakhsh, I.R., Bobenage, J., Grange, S., Lutz, R., Meyer, R, Soto, A., "An Affective Mobile Educator with a Full-Time Job." *Artificial Intelligence*, 114:95–124, 1999.

[119] Nourbakhsh, I.R., Powers, R., Birchfield, S., "DERVISH, an Office-Navigation Robot." *AI Magazine*, 16:39–51, summer 1995.

[120] Pell, B., Bernard, D., Chien, S., Gat, E., Muscettola, N., Nayak, P., Wagner, M., Williams, B., "An Autonomous Spacecraft Agent Prototype." *Autonomous Robots*, No. 5, 1–27, 1998.

[121] Pentland, A.P., "A New Sense for Depth of Field." *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI),* 9:523–531, 1987.

[122] Philippsen, R., Siegwart, R., "Smooth and Efficient Obstacle Avoidance for a Tour Guide Robot," *in Proceedings of the IEEE International Conference on Robotics and Automation (ICRA 2003)*, Taipei, Taiwan, 2003.

[123] Pratt, J., Pratt, G., "Intuitive Control of a Planar Bipedal Walking Robot," *in Proceedings of the IEEE International Conference on Robotics and Automation (ICRA '98)*, Leuven, Belgium, May 16–21, 1998

[124] Raibert, M. H., Brown, H. B., Jr., Chepponis, M., "Experiments in balance with a 3D One-Legged Hopping Machine." *International Journal of Robotics Research*, 3:75–92, 1984.

[125] Ringrose, R., "Self-Stabilizing Running," *in Proceedings of the IEEE International Conference on Robotics and Automation (ICRA '97)*, Albuquerque, NM, April 1997.

[126] Rowe, A., Rosenberg, C., Nourbakhsh, I., "A Simple Low Cost Color Vision System," *in Proceedings of Tech Sketches for CVPR 2001*. Kuaii, Hawaii, December 2001.

[127] Rubner, Y., Tomasi, C., Guibas, L., "The Earth Mover's Distance as a Metric for Image Retrieval," *STAN-CS-TN-98-86, Stanford University*, 1998.

[128] Schlegel, C., "Fast Local Obstacle under Kinematic and Dynamic Constraints," *in Proceedings of the IEEE International Conference on Intelligent Robot and Systems (IROS 98)*, Victoria, B.C. Canada 1998, pp. 594–599.

[129] Schultz, A., Adams, W., "Continuous Localization Using Evidence Grids," *in Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'98)*, May 16–21, 1998, Leuven, Belgium, pp.2833–2839 [also available as NCARAI Report AIC-96-007].

[130] Schweitzer, G., Werder, M., "ROBOTRAC – a Mobile Manipulator Platform for Rough Terrain," *in Proceedings of the International Symposium on Advanced Robot Technology (ISART),* Tokyo, Japan, March, 1991.

[131] Shi, J., Malik, J., "Normalized Cuts and Image Segmentation." *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 82:888–905, 2000.

[132] Siegwart R., et al., (2003) "Robox at Expo.02: A Large Scale Installation of Personal Robots." *Journal of Robotics and Autonomous Systems*, 42:203–222, 2003.

[133] Siegwart, R., Lamon, P., Estier, T., Lauria, M, Piguet, R., "Innovative Design for Wheeled Locomotion in Rough Terrain," *Journal of Robotics and Autonomous Systems*, 40:151–162, 2002.

[134] Simhon, S., Dudek, G., "A Global Topological Map Formed by Local Metric Maps," *in Proceedings of the 1998 IEEE/RSJ International Conference on Intelli- gent Robots and Systems (IROS'98)*, Victoria, B.C., Canada, October 1998.

[135] Simmons, R., "The Curvature Velocity Method for Local Obstacle Avoidance," *in Proceedings of the IEEE International Conference on Robotics and Automation*, Minneapolis, April 1996.

[136] Smith, R., Self, M., Cheeseman, P., "Estimating Uncertain Spatial Relationships in Robotics," *Autonomous Robot Vehicles,* I. J. Cox and G. T. Wilfong (editors), Springer-Verlag, 1990, pp. 167–193.

[137] Sordalen, O.J., Canudas de Wi,t C., "Exponential Control Law for a Mobile Robot: Extension to Path Following," *IEEE Transactions on Robotics and Automation*, 9:837–842, 1993.

[138] Steinmetz, B.M., Arbter, K., Brunner, B., Landzettel, K., "Autonomous Vision-Based Navigation of the Nanokhod Rover," *in Proceedings of i-SAIRAS 6th International Symposium on Artificial Intelligence, Robotics and Automation in Space*, Montrea1, June 18–22, 2001.

[139] Stentz, A., "The Focussed D* Algorithm for Real-Time Replanning," in *Proceedings of IJCAI-95*, August 1995.

[140] Stevens, B.S., Clavel, R., Rey, L., "The DELTA Parallel Structured Robot, Yet More Performant through Direct Drive," *in Proceedings of the 23rd International Symposium on Industrial Robots*, Barcelona, October 1992, pp. 485–493.

[141] Takeda, H., Facchinetti, C., Latombe, J.C., "Planning the Motions of a Mobile Robot in a Sensory Uncertainty Field." *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16:1002–1017, 1994.

[142] Thrun, S., Burgard, W., Fox, D., "A Probabilistic Approach to Concurrent Mapping and Localization for Mobile Robots." *Autonomous Robots*, 31:1–25. 1998.

[143] Thrun, S., et al., "Monrovia: A Second Generation Museum Tour-Guide Robot," *in Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'99)*, Detroit, May 1999.

[144] Thrun, S., Fox, D., Burgard, W. Dellaert, F., "Robust Monte Carlo Localization for Mobile Robots," *Artificial Intelligence*, 128:99–141, 2001.

[145] Thrun, S., Gutmann, J.-S., Fox, D., Burgard, W., Kuipers, B., "Integrating Topolog- ical and Metric Maps for Mobile Robot Navigation: A Statistical Approach," *in Proceedings of the National Conference on Artificial Intelligence (AAAI)*,1998.

[146] Tomasi, C., Shi, J., "Image Deformations Are Better Than Optical Flow." *Mathematical and Computer Modelling*, 24:165–175, 1996.

[147] Tomatis, N., Nourbakhsh, I., Siegwart, R., "Hybrid Simultaneous Localization and Map Building: A Natural Integration of Topological and Metric." *Robotics and Autonomous Systems* 44, 3–14, 2003.

[148] Tzafestas, C.S., Tzafestas, S.G., "Recent Algorithms for Fuzzy and Neurofuzzy Path Planning and Navigation of Autonomous Mobile Robots," *Systems-Science*, 25:25–39, 1999.

[149] Ulrich, I., Borenstein, J., "VFH*: Local Obstacle Avoidance with Look-Ahead Verification," *in Proceedings of the IEEE International Conference on Robotics and Automation*, San Francisco, May 24–28, 2000.

[150] Ulrich, I., Borenstein, J., "VFH+: Reliable Obstacle Avoidance for Fast Mobile Robots," *in Proceedings of the International Conference on Robotics and Automation (ICRA'98)*, Leuven, Belgium, May 1998.

[151] Ulrich, I., Nourbakhsh, I., "Appearance-Based Obstacle Detection with Monocular Color Vision," *in the Proceedings of the AAAI National Conference on Artificial Intelligence*. Austin, TX. August 2000.

[152] Ulrich, I., Nourbakhsh, I., "Appearance-Based Place Recognition for Topological Localization," *in Proceedings of t he IEEE International Conference on Robotics and Automation*, San Francisco, pp. 1023–1029, April 2000.

[153] Vanualailai, J., Nakagiri, S., Ha, J-H., "Collision Avoidance in a Two-Point System via Liapunov's Second Method." *Mathematics and Simulation,* 39:125–141, 1995.

[154] Van Winnendael, M., Visenti G., Bertrand, R., Rieder, R., "Nanokhod Microrover Heading towards Mars," *in Proceedings of the Fifth International Symposium on Artificial Intelligence, Robotics and Automation in Space* (ESA SP-440), Noordwijk, Netherlands, pp. 69–76, 1999.

[155] Weiss, G., Wetzler, C., Puttkamer, E., "Keeping Track of Position and Orientation of Moving Indoor Systems by Correlation of Range-Finder Scans," *in Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'94)*, Munich, pp.595–601, September 12-16, 1994.

[156] Wullschleger, F.H., Arra,s K.O., Vestli, S.J., "A Flexible Exploration Framework for Map Building," *in Proceedings of the Third European Workshop on Advanced Mobile Robots (Eurobot 99)*, Zurich, September 6-9, 1999.

[157] Yamauchi, B., Schultz, A., Adams, W., "Mobile Robot Exploration and Map-Build- ing with Continuous Localization," *in Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'98)*, Leuven, Belgium, May 1998.

[158] Zhang, Z., "A Flexible New Technique for Camera Calibration." *Microsoft Research Technical Report 98-71*., December 1998
[see also http://research.microsoft.com/~zhang].

Referenced Webpages

[159] Fisher, R.B. (editor), "CVonline: On-line Compendium of Computer Vision," Available at www.dai.ed.ac.uk/CVonline/.

[160] The Intel Image Processing Library: http://developer.intel.com/software/products/ perflib/ipl/.

[161] Source code release site: www.cs.cmu.edu/~jbruce/cmvision.

[162] Newton Labs website: www.newtonlabs.com.

[163] For probotics: http://www.personalrobots.com.

Interesting Internet Links to

Mobile Robots General

homepages with mainly mobile

robots

http://ranier.hq.nasa.gov/telerobotics_page/coolrobots.html
http://www-robotics.cs.umass.edu/robotics.html
http://www.ai.mit.edu/projects/mobile-robots/

http://www.ri.cmu.edu/project_lists/index.html
http://www.roboticsclub.org/links.html http://www.activrobots.com/
http://asl.epfl.ch (at EPFL) http://robotics.epfl.ch
(at EPFL)
http://www.cs.cmu.edu/~illah/EDUTOY (at CMU)
http://www.cs.cmu.edu/~mercator/ (at CMU)

Homepages with mainly wheeled mobile robots
http://www.laas.fr/RIA/RIA.html
http://www.cs.cmu.edu/~illah/lab.html
http://www.cs.umd.edu/projects/amrl/amrl.html
http://www.cc.gatech.edu/ai/robot-lab/
http://www.engin.umich.edu/research/mrl/index.html

Homepages with walking and climbing robots
http://www.uwe.ac.uk/clawar/
http://www.fzi.de/divisions/ipt/
http://www.automation.hut.fi/
http://www.ai.mit.edu/projects/leglab/ http://asl.epfl.ch/
http://www.cs.cmu.edu/~personalrover/

Homepages for robots on the Web
http://telerobot.mech.uwa.edu.au/secn/links.html
http://queue.IEOR.Berkeley.EDU/~goldberg/art/telerobotics-links.html
http://www.cs.uni-bonn.de/~rhino/
http://www.cs.cmu.edu/~minerva/
http://telerobot.mech.uwa.edu.au
http://www.ieor.berkeley.edu/~goldberg/GC/www.html
http://pumapaint.rwu.edu/ http://mars.graham.com/wits/
http://www.cs.cmu.edu/~illah/SAGE/index.html
http://www.dislocation.net/
http://rr-vs.informatik.uni-ulm.de/rr/
http://www.eventscope.org/Eventscope/main.htm

72. Appendices

Appendix A:

Matlab functions in ARSim

The following is an alphabetical list of all the Matlab functions associated with `ARSim`. For each function, the library to which the function belongs is given, along with the interface of the function and a brief description. For further information, see the actual source code for the function in question.

AddMotionResults Library:

ResultFunctions

Interface:

```
motionresults = AddMotionResults(oldMotionResults, time, robot)
```

Description: This function updates the motion results by adding the current position, velocity, heading, and sensor readings of the robot.

BrainStep Library:

–

Interface: `b = BrainStep(robot, time);`

Description: The `BrainStep` implements the decision-making system (i.e. the brain) of the robot. The detailed form of this function will vary from ex- periment to experiment.

CalibrateOdometer

Library: RobotFunctions

Interface: `o = CalibrateOdometer(robot)`

Description: In simulations in which an odometer is used, a call to `CalibrateOdometer` is made just before the start of the simulation, in order to set the correct posi- tion and heading of the robot.

See also: `CreateOdometer`

CheckForCollisions

Library: RobotFunctions

Interface: `coll = CheckForCollisions(arena, robot);` **Description:** This function carries out a collision check, by running through all arena objects (polygons) line by line, and checking for intersections between the current line and the spherical body of the robot.

CreateArena

Library: ArenaFunctions

Interface: `arena = CreateArena(name,size,objectArray)` **Description:** This function generates an arena, given an array of arena objects. **See also:** `CreateArenaObject`

CreateArenaObject Library:

ArenaFunctions

Interface: `arenaobject = CreateArenaObject(name,vertexArray)`
Description: This function generates an arena object, given an array of coor- dinates for vertices.

CreateBrain

Library: *–*
Interface: `b = CreateBrain;`

Description: This function generates the brain of a robot. Its exact form will vary from experiment to experiment.

CreateCompass Library:

RobotFunctions

Interface: `c = CreateCompass(name,sigma);`

Description: This function generates a compass which can be used for esti- mating the heading of the robot. The parameter `sigma` determines the noise level.

CreateIRSensor Library:

RobotFunctions

Interface: `s = CreateIRSensor(name,relativeAngle,size,numberOfRays, openingAngle,range,c1,c2,sigma);`

Description: `CreateIRSensor` creates an IR sensor that uses the ray trac- ing procedure described above to obtain its readings. The parameter `sigma` is defined as in Eq. (3.1).

CreateMotor

Library: RobotFunctions

Interface: `m = CreateMotor(name);`

Description: `CreateMotor` generates a DC motor, using settings suitable for a robot with a mass of a few kg.

CreateOdometer Library:

RobotFunctions

Interface: `o = CreateOdometer(name, sigma);`

Description: This function generates an odometer, which, in turn, provides estimates for the position and heading of the robot. The parameter `sigma` determines the noise level.

CreateRobot

Library: RobotFunctions

Interface: `robot = CreateRobot(name,mass,momentOfInertia,radius, wheelRadius,rayBasedSensorArray, motorArray,compass,odometer,brain)`
Description: `CreateRobot` sets up a robot, and computes the dynamical parameters typical of a robot with a mass of a few kg.

GetCompassReading

Library: RobotFunctions

Interface: `c = GetCompassReading(robot, dt);`

Description: This function updates the compass readings of a robot.

GetDistanceToLineAlongRay Library:

RobotFunctions

Interface: `l = GetDistanceToLineAlongRay(beta,p1,p2,x1,y1);`
Description: This function, which is used by the IR sensors, computes the distance from a given point (x_1, y_1) to a line segment.

See also: `GetIRSensorReading`, `GetDistanceToNearestObject`.

GetDistanceToNearestObject

Library: RobotFunctions

Interface: `d = GetDistanceToNearestObject(beta, x, y, arena);`
Description: This function, which is used by the IR sensors, determines the distance between an IR sensor and the nearest object along a given ray.
See also: `GetIRSensorReading`.

GetIRSensorReading Library:

RobotFunctions

Interface: `s = GetIRSensorReading(sensor,arena);`

Description: `GetIRSensorReading` determines the reading of an IR sensor.

GetMinMaxAngle Library:

RobotFunctions
Interface: `[aMin,aMax] = GetMinMaxAngle(v1,v2);`

Description: This function determines the direction angles of the vectors con- necting the origin of the coordinate system to the tips of a line segment.
See also: `GetDistanceToNearestObject`.

GetMotorSignalsFromBrain Library:

RobotFunctions

Interface: `s = GetMotorSignalsFromBrain(brain);`

Description: This function extracts the motor signals (one for each motor) from the brain of the robot.
See also: `MoveRobot`.

GetOdometerReading

Library: RobotFunctions

Interface: `o = GetOdometerReading(robot, dt);`

Description: This function updates the odometer readings of a robot.

GetRayBasedSensorReadings Library:

RobotFunctions

Interface: `s = GetRayBasedSensorReadings(robot, arena)` **Description:** This function obtains the reading of all (IR) sensors of the robot. **See also:** `GetIRSensorReading`.

GetTorque

Library: RobotFunctions

Interface: `m = GetTorque(motor, voltage);`

Description: This function determines the torque delivered by a DC motor, given a value of the applied voltage.

InitializeMotionResults

Library: ResultFunctions

Interface: `motionResults = InitializeMotionResults(robot)`
Description: This function initializes a Matlab structure used for storing the results of the simulation, i.e. the position, velocity, heading, and sensor read- ings of the robot.

InitPlot

Library: PlotFunctions

Interface: `plotHandle = InitializePlot(robot, arena)`
Description: This function generates the plot of the robot and the arena. **See also:** `CreateArena`, `CreateRobot`.

MoveRobot

Library: RobotFunctions

Interface: `r = MoveRobot(robot,dt);`

Description: `MoveRobot` moves the robot according to the equations of mo- tion for a differentially steered two-wheeled robot.

ScaleMotorSignals

Library:

RobotFunctions
Interface: `v = ScaleMotorSignals(robot,s);`

Description: This function scales the motor signals (`s`) to the appropriate range, as set by the voltage requirements of the robot's DC motors.

SetPositionAndVelocity

Library:

RobotFunctions

Interface: `r = SetPosition(robot,position,heading,`

`velocity,angularSpeed);`

Description: This function places the robot at a given location, and also sets is direction of motion, velocity, and angular velocity.

ShowRobot

Library: PlotFunctions

Interface: `ShowRobot(plot,robot)`

Description: `ShowRobot` updates the plot of the robot using Matlab's handle graphics: Each part of the plot of the robot can be accessed and its position can be be updated. `ShowRobot` also supports the plotting of an odometric ghost,
i.e. a plot showing the robot at the location determined by its odometer.

See also: `MoveRobot`.

UpdateMotorAxisAng

ularSpeed Library:

RobotFunctions

Interface: `r = UpdateMotorAxisAngularSpeed(robot)`

Description: This function determines the angular speed of each motor axis, using the wheel speed and wheel radius.

UpdateSe

nsorPositi

ons

Library:

RobotFun

ctions

Interface: `s = UpdateSensorPositions(robot);`

Description: This function updates the positions (and directions) of the sen- sors as the robot is moved.