CHAPTER 7
LANDMARK NAVIGATION

Landmarks are distinct features that a robot can recognize from its sensory input. Landmarks can be geometric shapes (e.g., rectangles, lines, circles), and they may include additional information (e.g., in the form of bar-codes). In general, landmarks have a fixed and known position, relative to which a robot can localize itself. Landmarks are carefully chosen to be easy to identify; for example, there must be sufficient contrast to the background. Before a robot can use landmarks for navigation, the characteristics of the landmarks must be known and stored in the robot's memory. The main task in localization is then to recognize the landmarks reliably and to calculate the robot's position.

In order to simplify the problem of landmark acquisition it is often assumed that the current robot position and orientation are known approximately, so that the robot only needs to look for landmarks in a limited area. For this reason good odometry accuracy is a prerequisite for successful landmark detection.

The general procedure for performing landmark-based positioning is shown in Figure 7.1. Some approaches fall between landmark and map-based positioning (see Chap. 8). They use sensors to sense the environment and then extract distinct structures that serve as landmarks for navigation in the future. These approaches will be discussed in the chapter on map-based positioning techniques.

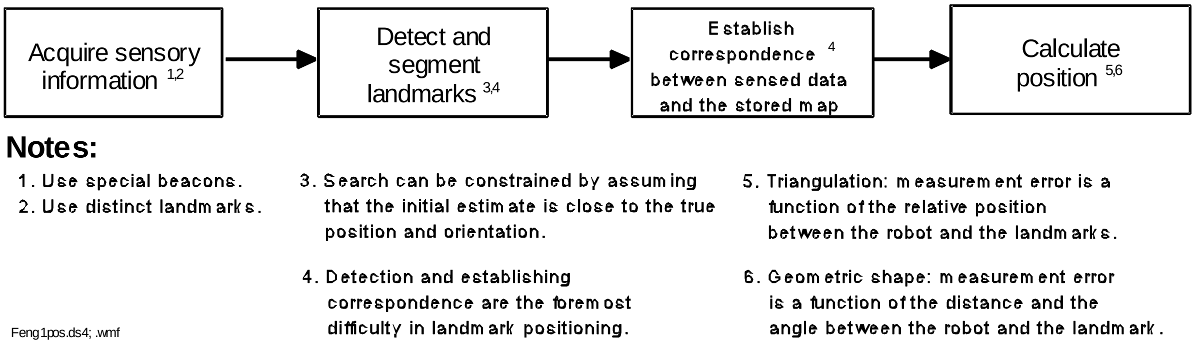

Notes:

1. Use special beacons.
2. Use distinct landmarks.
3. Search can be constrained by assuming that the initial estimate is close to the true position and orientation.
4. Detection and establishing correspondence are the foremost difficulty in landmark positioning.
5. Triangulation: measurement error is a function of the relative position between the robot and the landmarks.
6. Geometric shape: measurement error is a function of the distance and the angle between the robot and the landmark.

Figure 7.1: General procedure for landmark-based positioning.

Our discussion in this chapter addresses two types of landmarks: “artificial” and “natural.” It is important to bear in mind that “natural” landmarks work best in highly structured environments such as corridors, manufacturing floors, or hospitals. Indeed, one may argue that “natural” landmarks work best when they are actually man-made (as is the case in highly structured environments). For this reason, we shall define the terms “natural landmarks” and “artificial landmarks” as follows: *natural landmarks* are those objects or features that are already in the environment and have a function other than robot navigation; *artificial landmarks* are specially designed objects or markers that need to be placed in the environment with the sole purpose of enabling robot navigation.

7.1 Natural Landmarks

The main problem in natural landmark navigation is to detect and match characteristic features from sensory inputs. The sensor of choice for this task is computer vision. Most computer vision-based natural landmarks are long vertical edges, such as doors and wall junctions, and ceiling lights. However, computer vision is an area that is too large and too diverse for the scope of this book. For this reason we will present below only one example of computer vision-based landmark detection, but without going into great detail.

When range sensors are used for natural landmark navigation, distinct signatures, such as those of a corner or an edge, or of long straight walls, are good feature candidates. The selection of features is important since it will determine the complexity in feature description, detection, and matching. Proper selection of features will also reduce the chances for ambiguity and increase positioning accuracy. A natural landmark positioning system generally has the following basic components:

- A sensor (usually computer vision) for detecting landmarks and contrasting them against their background.
- A method for matching observed features with a map of known landmarks.
- A method of computing location and localization errors from the matches.

One system that uses natural landmarks has recently been developed in Canada. This project aimed at developing a sophisticated robot system called the "*Autonomous Robot for a Known Environment*" (ARK). The project was carried out jointly by the Atomic Energy of Canada Ltd (AECL) and Ontario Hydro Technologies with support from the University of Toronto and York University [Jenkin et al., 1993]. A Cybermotion K2A+ platform serves as the carrier for a number of sensor subsystems (see Figure 7.2).

Of interest for the discussion here is the ARK navigation module (shown in Figure 7.3). This unit consists of a custom-made pan-and-tilt table, a CCD camera, and an eye-safe IR spot laser rangefinder. Two VME-based cards, a single-board computer, and a microcontroller, provide processing power. The navigation module is used to periodically correct the robot's accumulating odometry errors. The system uses *natural*

Figure 7.2: The ARK system is based on a modified Cybermotion K2A+. It is one of the few working navigation systems based on natural landmark detection. (Courtesy of Atomic Energy of Canada Ltd.)

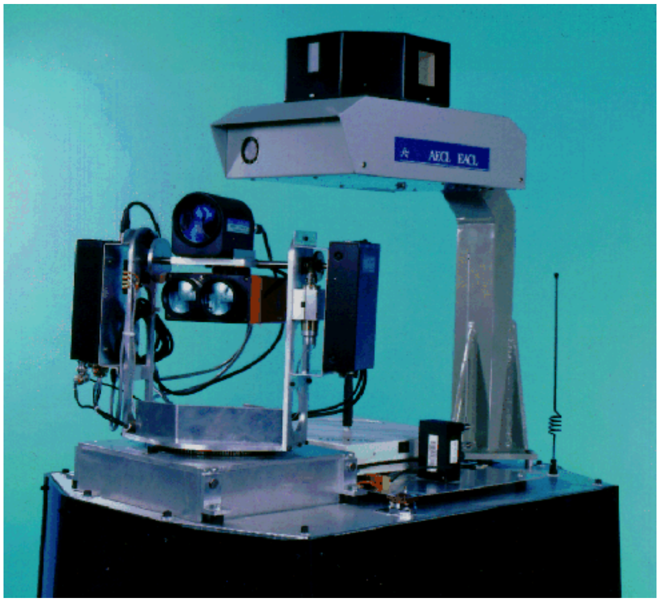
Figure 7.3: AECL's natural landmark navigation system uses a CCD camera in combination with a time-of-flight laser rangefinder to identify landmarks and to measure the distance between landmark and robot. (Courtesy of Atomic Energy of Canada Ltd.)

landmarks such as alphanumeric signs, semi-permanent structures, or doorways. The only criteria used is that the landmark be distinguishable from the background scene by color or contrast.

The ARK navigation module uses an interesting hybrid approach: the system stores (learns) landmarks by generating a three- dimensional "grey-level surface" from a single training image obtained from the CCD camera. A coarse, registered range scan of the same field of view is performed by the laser rangefinder, giving depths for each pixel in the grey-level surface. Both procedures are performed from a known robot position. Later, during operation, when the robot is at an approximately known (from odometry) position within a couple of meters from the training position, the vision system searches for those landmarks that are expected to be visible from the robot's momentary position. Once a suitable landmark is found, the projected appearance of the landmark is computed. This *expected* appearance is then used in a coarse-to-fine normalized correlation-based matching algorithm that yields the robot's relative distance and bearing with regard to that landmark. With this procedure the ARK can identify different natural landmarks and measure its position relative to the landmarks.

To update the robot's odometry position the system must find a pair of natural landmarks of known position. Positioning accuracy depends on the geometry of the robot and the landmarks but is typically within a few centimeters. It is possible to pass the robot through standard 90-centimeter (35 in) doorway openings using only the navigation module if corrections are made using the upper corners of the door frame just prior to passage.

7.2 Artificial Landmarks

Detection is much easier with artificial landmarks [Atiya and Hager, 1993], which are designed for optimal contrast. In addition, the exact size and shape of artificial landmarks are known in advance. Size and shape can yield a wealth of geometric information when transformed under the perspective projection.

Researchers have used different kinds of patterns or marks, and the geometry of the method and the associated techniques for position estimation vary accordingly [Talluri and Aggarwal, 1993]. Many artificial landmark positioning systems are based on computer vision. We will not discuss these systems in detail, but we will mention some of the typical landmarks used with computer vision.

Fukui [1981] used a diamond-shaped landmark and applied a least-squares method to find line segments in the image plane. Borenstein [1987] used a black rectangle with four white dots in the corners. Kabuka and Arenas [1987] used a half-white and half-black circle with a unique bar-code for each landmark. Magee and Aggarwal [1984] used a sphere with horizontal and vertical calibration circles to achieve three-dimensional localization from a single image. Other systems use reflective material patterns and strobed light to ease the segmentation and parameter extraction [Lapin, 1992; Mesaki and Masuda, 1992]. There are also systems that use active (i.e., LED) patterns to achieve the same effect [Fleury and Baron, 1992].

The accuracy achieved by the above methods depends on the accuracy with which the geometric parameters of the landmark images are extracted from the image plane, which in turn depends on the relative position and angle between the robot and the landmark. In general, the accuracy decreases with the increase in relative distance. Normally there is a range of relative angles in which good accuracy can be achieved, while accuracy drops significantly once the relative angle moves out of the "good" region.

There is also a variety of landmarks used in conjunction with non-vision sensors. Most often used are bar-coded reflectors for laser scanners. For example, currently ongoing work by Everett on the *Mobile Detection Assessment and Response System* (MDARS) [DeCorte, 1994] uses retro-reflectors, and so does the commercially available system from Caterpillar on their *Self-Guided Vehicle* [Gould, 1990]. The shape of these landmarks is usually unimportant. By contrast, a unique approach taken by Feng et al. [1992] used a circular landmark and applied an optical Hough transform to extract the parameters of the ellipse on the image plane in real time.

7.2.1 Global Vision

Yet another approach is the so-called *global vision* that refers to the use of cameras placed at fixed locations in a workspace to extend the local sensing available on board each AGV [Kay and Luo, 1993]. Figure 7.4 shows a block diagram of the processing functions for vehicle control using global vision.

In global vision methods, characteristic points forming a pattern on the mobile robot are identified and localized from a single view. A probabilistic method is used to select the most probable matching according to geometric characteristics of those percepts. From this reduced search space a prediction-verification loop is applied to identify and to localize the points of the pattern [Fleury and Baron, 1992]. One advantage of this approach is that it allows the operator to monitor robot operation at the same time.

7.3 Artificial Landmark Navigation Systems

Many systems use retroreflective barcodes as artificial landmarks, similar to the ones used in beacon navigation systems. However, in this book we distinguish between retroreflective bar-codes used as artificial landmarks and retroreflective poles used as "beacons." The reason is that if retroreflective markers (with or without bar-code) are attached to the walls of a room and their function is merely to aid in determining the location of the wall, then these markers do not

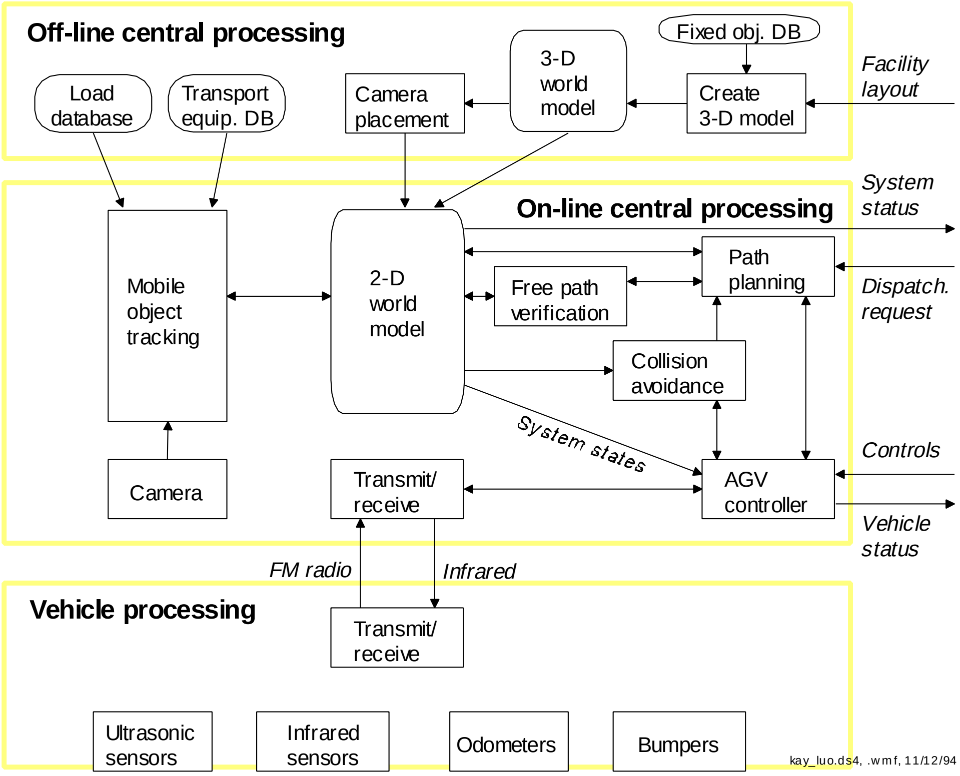

Figure 7.4: Block diagram of the processing functions for vehicle control using global vision. (Adapted from [Kay and Luo, 1993].)

function as beacons. By contrast, if markers are used on arbitrarily placed poles (even if the location of these poles is carefully surveyed), then they act as beacons. A related distinction is the method used for computing the vehicle's position: if triangulation is used, then the reflectors act as beacons.

7.3.1 MDARS Lateral-Post Sensor

Currently ongoing work by Everett on the *Mobile Detection Assessment and Response System* (MDARS) [Everett et al., 1994; DeCorte, 1994] uses passive reflectors in conjunction with a pair of fixed-orientation sensors on board the robot. This technique, called *lateral-post detection*, was incorporated on MDARS to significantly reduce costs by exploiting the forward motion of the robot for scanning purposes. Short vertical strips of 2.5 centimeters (1 in) retroreflective tape are placed on various immobile objects (usually structural-support posts) on either side of a virtual path segment. The exact x-y locations of these tape markers are encoded into the virtual path program. Installation takes only seconds, and since the flat tape does not protrude into the aisle at all, there is little chance of damage from a passing fork truck.

A pair of Banner Q85VR3LP retroreflective proximity sensors mounted on the turret of the *Navmaster* robot face outward to either side as shown in Figure 7.5 These inexpensive sensors respond to reflections from the tape markers along the edges of the route, triggering a “snapshot” *virtual path* instruction that records the current side-sonar range values. The longitudinal position of the platform is updated to the known marker coordinate, while lateral position is inferred from the sonar data, assuming both conditions fall within specified tolerances.

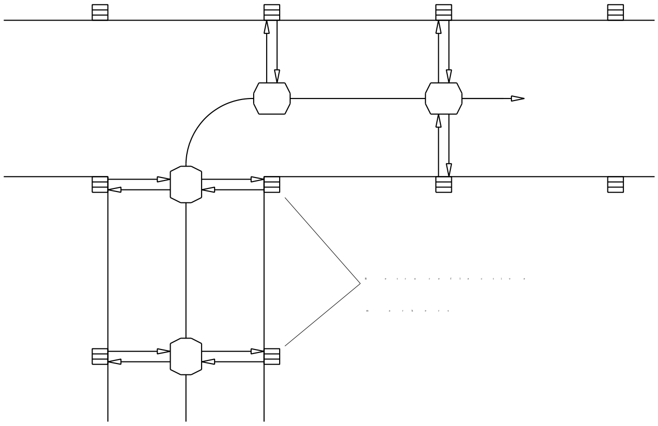

Figure 7.5: Polarized retroreflective proximity sensors are used to locate vertical strips of retroreflective tape attached to shelving support posts in the Camp Elliott warehouse installation of the MDARS security robot [Everett et al, 1994].

The accuracy of the marker correction is much higher (and therefore assigned greater credibility) than that of the lateral sonar readings due to the markedly different uncertainties associated with the respective targets. The polarized Banner sensor responds only to the presence of a retroreflector while ignoring even highly specular surrounding surfaces, whereas the ultrasonic energy from the sonar will echo back from any reflective surface encountered by its relatively wide beam. Protruding objects in the vicinity of the tape (quite common in a warehouse environment) result in a shorter measured range value than the reference distance for the marker itself. The overall effect on x-y bias is somewhat averaged out in the long run, as each time the vehicle executes a 90-degree course change the association of x- and y-components with tape versus sonar updates is interchanged.

7.3.2 Caterpillar *Self Guided Vehicle*

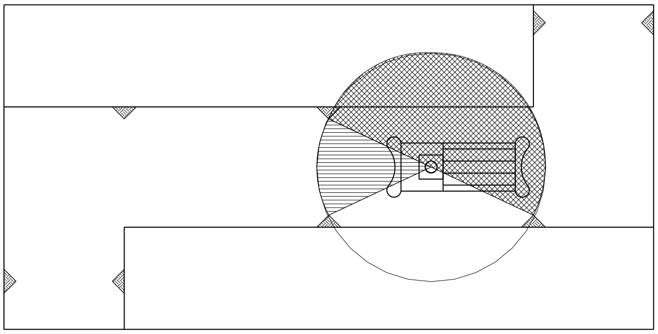

Figure 7.6: Retroreflective bar-code targets spaced 10 to 15 meters (33 to 49 ft) apart are used by the Caterpillar SGV to triangulate position. (Adapted from [Caterpillar, 1991a].)

Caterpillar Industrial, Inc., Mentor, OH, manufactures a free-ranging AGV for materials handling that relies on a scanning laser triangulation scheme to provide positional updates to the vehicle's onboard odometry system. The Class-I laser rotates at 2 rpm to illuminate passive retroreflective bar-code targets affixed to walls or support columns at known locations up to 15 meters (50 ft) away [Gould, 1990; Byrne et al., 1992]. The bar-codes serve to positively identify the reference target and eliminate ambiguities due to false returns from other specular surfaces within the operating area. An onboard computer calculates x-y position updates through simple triangulation to null out accumulated odometry errors (see Figure 7.6).

Some target occlusion problems have been experienced in exterior applications where there is heavy fog, as would be expected, and minor difficulties have been encountered as well during periods when the sun was low on the horizon [Byrne, 1993]. Caterpillar's *Self Guided Vehicle* (SGV) relies on dead reckoning under such conditions to reliably continue its route for distances of up to 10 meters (33 ft) before the next valid fix.

The robot platform is a hybrid combination of tricycle and differential drive, employing two independent series-wound DC motors powering 45-centimeter (18 in) rear wheels through sealed gear-boxes [CATERPILLAR, 1991]. High-resolution resolvers attached to the single front wheel continuously monitor steering angle and distance traveled. A pair of mechanically scanned near-infrared proximity sensors sweeps the path in front of the vehicle for potential obstructions. Additional near infrared sensors monitor the area to either side of the vehicle, while ultrasonic sensors cover the back.

7.3.3 Komatsu Ltd, Z-shaped landmark

Komatsu Ltd. in Tokyo, Japan, is a manufacturer of construction machines. One of Komatsu's research projects aims at developing an unmanned dump truck. As early as 1984, researchers at Komatsu Ltd. developed an unmanned electric car that could follow a previously taught path around the company's premises. The vehicle had two onboard computers, a directional gyrocompass, two incremental encoders on the wheels, and a metal sensor which detected special landmarks along the planned path (see Figure 7.7).

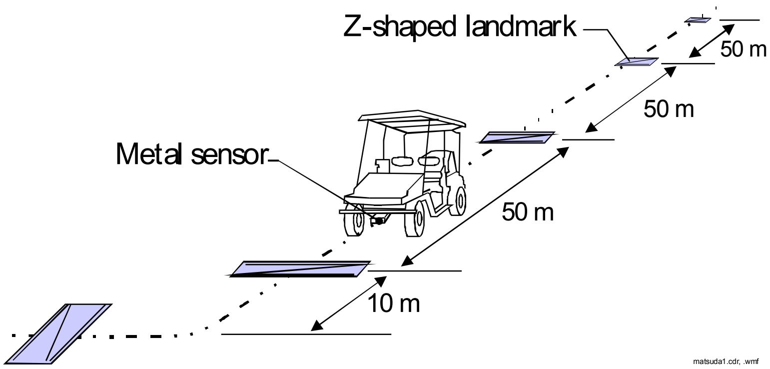

Figure 7.7: Komatsu's Z-shaped landmarks are located at 50-meter (164 ft) intervals along the planned path of the autonomous vehicle. (Courtesy of [Matsuda and Yoshikawa, 1989].)

The accuracy of the vehicle's dead-reckoning system (gyrocompass and encoders) was approximately two percent on the paved road and during straight-line motion only. The mechanical gyrocompass was originally designed for deep-sea fishing boats and its static direction accuracy was 1 degree. On rough terrain the vehicle's dead-reckoning error deteriorated notably. For example, running over a 40-millimeter (1.5 in) height bump and subsequently traveling along a straight line for 50 meters (164 ft), the vehicle's positioning error was 1.4 m (55 in). However, with the Z-shaped landmarks used in this project for periodic recalibration the positioning could be recalibrated to an accuracy of 10 centimeters (4 in). The 3 meter (118 in) wide landmark was made of 50 millimeter (2 in) wide aluminum strips sandwiched between two rubber sheets. In order to distinguish between “legitimate” metal markings of the landmark and between arbitrary metal objects, additional parallel line segments were used (see Figure 7.8). The metal markers used as landmarks in this experiment are resilient to contamination even in harsh environments. Water, dust, and lighting condition do not affect the readability of the metal sensor [Matsuda and Yoshikawa, 1989].

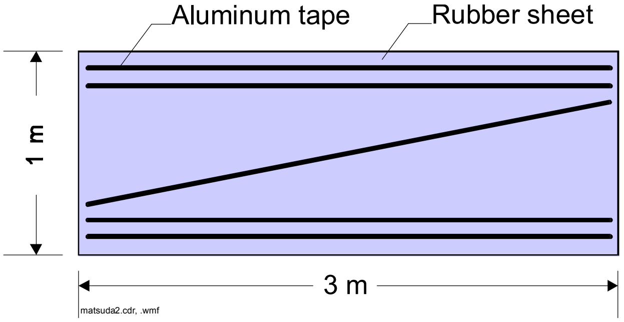

Figure 7.8: The Z-shaped landmark. Note the secondary lines parallel to the horizontal Z-stripes. The secondary lines help distinguish the marker from random metal parts on the road. (Courtesy of [Matsuda and Yoshikawa, 1989].)

Each Z-shaped landmark comprises three line segments. The first and third line segments are in parallel, and the second one is located diagonally between the parallel lines (see Figure7.9). During operation, a metal sensor located underneath the autonomous vehicle detects the three crossing points P_1, P_2, and P_3. The distances, L_1 and L_2, are measured by the incremetal encoders using odometry. After traversing the Z-shaped landmark, the vehicle's lateral deviation X_2 at point P_2 can be computed from

$$X_2 = W\left(\frac{L_1}{L_1+L_2} - \frac{1}{2}\right) \qquad (7.1)$$

where X_2 is the lateral position error at point P_2 based on odometry.

The lateral position error can be corrected after passing through the third crossing point P_3. Note that for this correction method the exact location of the landmark along the line of travel does not have to be known. However, if the location of the landmark is known, then the vehicle's actual position at P_2 can be calculated easily [Matsuda et al., 1989].

The size of the Z-shaped landmark can be varied, according to the expected lateral error of the vehicle. Larger landmarks can be buried under the surface of paved roads for unmanned cars. Smaller landmarks can be installed under factory floor coating or under office carpet. Komatsu has developed such smaller Z-shaped landmarks for indoor robots and AGVs.

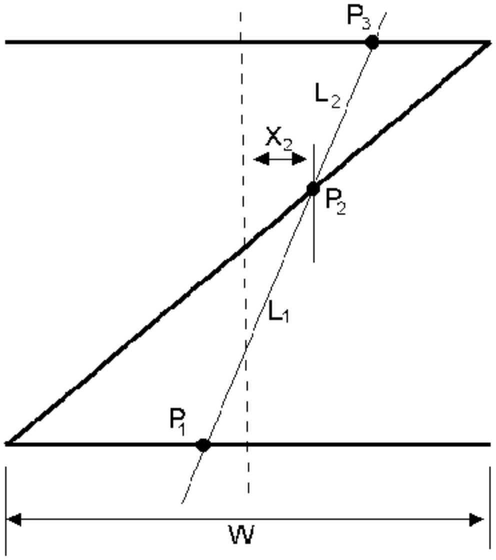

Figure 7.9: The geometry of the Z-shaped landmark lends itself to easy and unambiguous computation of the lateral position error X_2. (Courtesy of [Matsuda and Yoshikawa, 1989].)

7.4 Line Navigation

Another type of landmark navigation that has been widely used in industry is line navigation. Line navigation can be thought of as a continuous landmark, although in most cases the sensor used in this system needs to be very close to the line, so that the range of the vehicle is limited to the immediate vicinity of the line. There are different implementations for line navigation:

- *Electromagnetic Guidance* or *Electromagnetic Leader Cable*.
- *Reflecting Tape Guidance* (also called *Optical Tape Guidance*).
- *Ferrite Painted Guidance*, which uses ferrite magnet powder painted on the floor [Tsumura, 1986].

These techniques have been in use for many years in industrial automation tasks. Vehicles using these techniques are generally called *Automatic Guided Vehicles* (AGVs).

In this book we don't address these methods in detail, because they do not allow the vehicle to move freely — the main feature that sets mobile robots apart from AGVs. However, two recently introduced variations of the line navigation approach are of interest for mobile robots. Both techniques are based on the use of short-lived navigational markers (SLNM). The short-lived nature of the markers has the advantage that it is not necessary to remove the markers after use.

One typical group of applications suitable for SLNM are floor coverage applications. Examples are floor cleaning, lawn mowing, or floor surveillance. In such applications it is important for the robot to travel along adjacent paths on the floor, with minimal overlap and without "blank" spots. With the methods discussed here, the robot could conceivably mark the outside border of the path, and trace that border line in a subsequent run. One major limitation of the current state-of-the-art is that they permit only very slow travel speeds: on the order of under 10 mm/s (0.4 in/s).

7.4.1 Thermal Navigational Marker

Kleeman [1992], Kleeman and Russell [1993], and Russell [1993] report on a pyroelectric sensor that has been developed to detect thermal paths created by heating the floor with a quartz halogen bulb. The path is detected by a pyroelectric sensor based on lithium-tantalate. In order to generate a differential signal required for path following, the position of a single pyroelectric sensor is toggled between two sensing locations 5 centimeters (2 in) apart. An aluminum enclosure screens the sensor from ambient infrared light and electromagnetic disturbances. The 70 W quartz halogen bulb used in this system is located 30 millimeters (1-3/16 in) above the floor.

The volatile nature of this path is both advantageous and disadvantageous: since the heat trail disappears after a few minutes, it also becomes more difficult to detect over time. Kleeman and Russell approximated the temperature distribution T at a distance d from the trail and at a time t after laying the trail as

$$T(d,t) = A(t)\, e^{-(d/w)^2} \tag{7.2}$$

where A(t) is a time-variant intensity function of the thermal path.

In a controlled experiment two robots were used. One robot laid the thermal path at a speed of 10 mm/s (0.4 in/s), and the other robot followed that path at about the same speed. Using a control scheme based on a Kalman filter, thermal paths could be tracked up to 10 minutes after being laid on a vinyl tiled floor. Kleeman and Russell remarked that the thermal footprint of peoples' feet could contaminate the trail and cause the robot to lose track.

7.4.2 Volatile Chemicals Navigational Marker

This interesting technique is based on laying down an odor trail and using an olfactory[1] sensor to allow a mobile robot to follow the trail at a later time. The technique was described by Deveza et al. [1993] and Russell et al. [1994], and the experimental system was further enhanced as described by Russell [1995a; 1995b] at Monash University in Australia. Russell's improved system comprises a custom-built robot (see Figure 7.10) equipped with an odor-sensing system. The sensor system uses

---

[1] relating to, or contributing to the sense of smell (The American Heritage Dictionary of the English Language, Third Edition is licensed from Houghton Mifflin Company. Copyright © 1992 by Houghton Mifflin Company. All rights reserved).

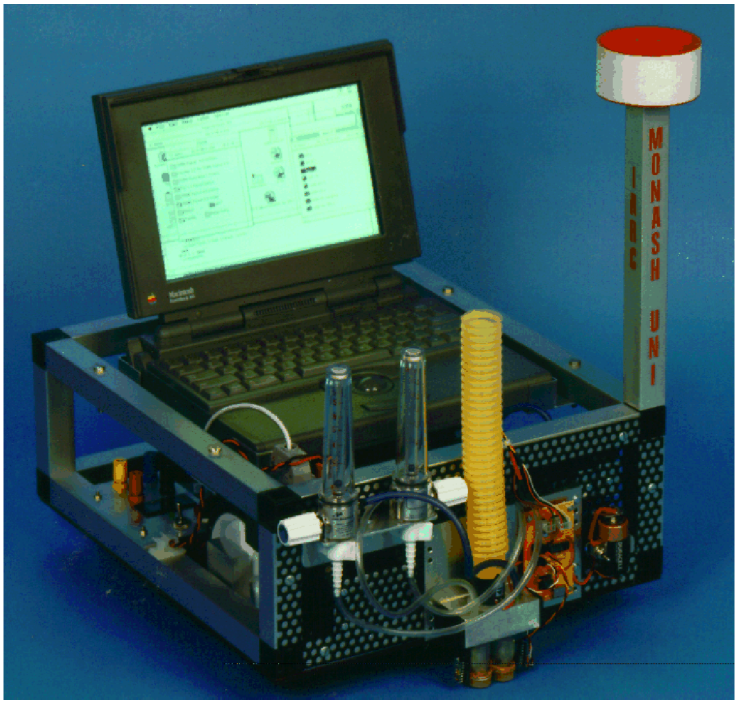

Figure 7.10: The odor-laying/odor-sensing mobile robot was developed at Monash University in Australia. The olfactory sensor is seen in front of the robot. At the top of the vertical boom is a magnetic compass. (Courtesy of Monash University).

controlled flows of air to draw odor-laden air over a sensor crystal. The quartz crystal is used as a sensitive balance to weigh odor molecules. The quartz crystal has a coating with a specific affinity for the target odorant; molecules of that odorant attach easily to the coating and thereby increase the total mass of the crystal. While the change of mass is extremely small, it suffices to change the resonant frequency of the crystal. A 68HC11 microprocessor is used to count the crystal's frequency, which is in the kHz region. A change of frequency is indicative of odor concentration. In Russell's system two such sensors are mounted at a distance of 30 millimeters (1-3/16 in) from each other, to provide a differential signal that can then be used for path tracking.

For laying the odor trail, Russell used a modified felt-tip pen. The odor-laden agent is camphor, dissolved in alcohol. When applied to the floor, the alcohol evaporates quickly and leaves a 10 millimeter (0.4 in) wide camphor trail. Russell measured the response time of the olfactory sensor by letting the robot cross an odor trail at angles of 90 and 20 degrees. The results of that test are shown in Figure 7.11. Currently, the foremost limitation of Russell's volatile chemical navigational marker is the robot's slow speed of 6 mm/s (1/4 in/s).

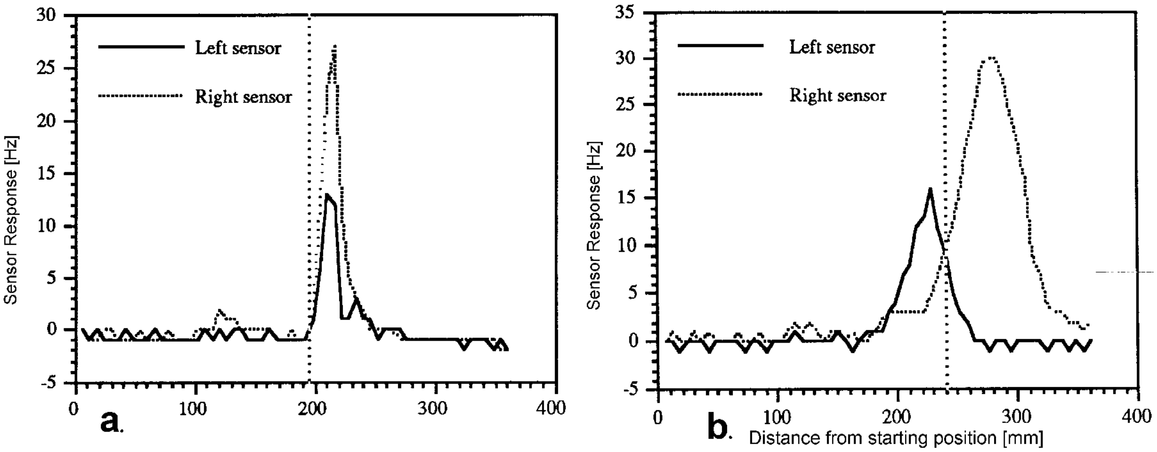

Figure 7.11: Odor sensor response as the robot crosses a line of camphor set at an angle of a. 90° and b. 20° to the robot path. The robots speed was 6 mm/s (1/4 in/s) in both tests. (Adapted with permission from Russell [1995].)

7.5 Summary

Artificial landmark detection methods are well developed and reliable. By contrast, *natural* landmark navigation is not sufficiently developed yet for reliable performance under a variety of conditions. A survey of the market of commercially available *natural* landmark systems produces only a few. One is TRC's vision system that allows the robot to localize itself using rectangular and circular ceiling lights [King and Weiman, 1990]. Cyberworks has a similar system [Cyberworks]. It is generally very difficult to develop a feature-based landmark positioning system capable of detecting different natural landmarks in different environments. It is also very difficult to develop a system that is capable of using many different types of landmarks.

We summarize the characteristics of landmark-based navigation as follows:

- Natural landmarks offer flexibility and require no modifications to the environment.
- Artificial landmarks are inexpensive and can have additional information encoded as patterns or shapes.
- The maximal distance between robot and landmark is substantially shorter than in active beacon systems.
- The positioning accuracy depends on the distance and angle between the robot and the landmark. Landmark navigation is rather inaccurate when the robot is further away from the landmark. A higher degree of accuracy is obtained only when the robot is near a landmark.
- Substantially more processing is necessary than with active beacon systems.
- Ambient conditions, such as lighting, can be problematic; in marginal visibility, landmarks may not be recognized at all or other objects in the environment with similar features can be mistaken for a legitimate landmark.
- Landmarks must be available in the work environment around the robot.
- Landmark-based navigation requires an approximate starting location so that the robot knows where to look for landmarks. If the starting position is not known, the robot has to conduct a time-consuming search process.
- A database of landmarks and their location in the environment must be maintained.
- There is only limited commercial support for this type of technique.