Proc. 2001 Australian Conference on Robotics and Automation
Sydney, 14-15 November 2001

Landmark Recognition with CTFM Ultrasonic Sensing

Danny Ratner and Phillip McKerrow
School of Information Technology and
Computer Science,
University of Wollongong
Wollongong, NSW, 2522
{phillip, dr19@uow.edu.au}

Abstract

Plants are naturally occurring landmarks. They can be detected and recognized using CTFM ultrasonic sensing. Landmarks can be grouped in four classes based on their geometric complexity and their continuity. The landmark class determines the sensor motion strategy, the recognition feature set and the navigation strategy. This paper reports on research into landmark recognition for each of these landmark classes.

1 Introduction

Humans and animals use several strategies when navigating from one place to another [Waterman, 1989]. These are: piloting, dead reckoning, celestial navigation, charting, indirect navigation and electronic navigation.

Piloting is a navigation strategy that uses known landmarks. These landmarks are used sequentially to find the way to the goal. The navigator must be familiar with the area, and know which landmarks to look for. A landmark is a feature in the environment, whose position can be sensed, that is close enough to the desired path that its direction varies significantly with the position of the navigator. A number of strategies are used to achieve piloting: following continuous landmarks such as coastlines, feature matching and compass piloting.

Thomas Gladwin [1970] studied the navigation skills of the people of Puluwat Atoll in the Caroline Islands. They navigate over a range greater than 1,000 kilometers through open seas, in their 8 metre sailing canoes. Their journeys are usually broken down into a series of island hops. However, they regularly travel 100 to 200 kilometres across the Pacific and arrive exactly at their destination, demonstrating their considerable navigation skills.

The navigator starts his voyage by imagining the position of his destination relative to the position of other islands. As he sails along, he constantly adjusts his directions according to his awareness of his current position. His decisions are improvised continually by checking relative positions of landmarks (reefs, atolls, etc), sun, stars, wave direction, wind direction, cloud patterns, and depth soundings. He navigates with reference to where he started, where he is going, and the space between his destination and his current location. If asked where he is, he can tell you relative to the surrounding islands.

A few mobile robot researchers have attempted to develop landmark navigation systems using ultrasonic sensing. Kimoto and Yuta [1993] used the standard deviation of ultrasonic range readings to detect a hedge from a moving robot. Maeyama et al., [1994] used a combination of vision and ultrasonic sensing to detect trees along a path. Mandow, et al [1996] used pulse-echo ultrasonic sensing to navigate along rows of plants in a green house.

Leonard and Durrant-Whyte [1991] use an acoustic feature called a Region of Constant Depth (RCD) for navigation. Akbarally and Kleeman [1995] identified nine acoustic features that occur in a room where all the planes are vertical or horizontal and meet at right angles. Each feature has parameters that help to identify it including position and orientation in space.

McKerrow [1995] extracted features representing geometric primitives from arc map made from of ultrasonic sensor data for indoor landmark identification. Wijk and Christensen [1998] extracted point features from sensor data which they matched to a map to identify and navigate with landmarks.

In this paper, we extend previous research into the recognition of plants [Harper and McKerrow, 1999a], [McKerrow and Harper, 1999] for landmark navigation [Harper and McKerrow, 1999b, 2000]. Following a discussion of landmark requirements (Section 2) we group landmarks into four classes (Section 4). Then we look at the sensor motion strategy, the recognition feature set and the navigation strategy for each class in subsequent sections.

2 Landmark Requirements

As plants occur in many environments where mobile robots will be used, are they suitable for landmarks? If they are, can an ultrasonic sensor detect them robustly? Objects that are suitable for use as landmarks should meet the following requirements:

1. occur naturally in the environment of the robot,
2. any changes occur over long periods of time,
3. are easily detected (i.e. discriminated from other objects),
4. can be detected independently of range and orientation,
5. can be detected independently of their size (width, height and depth),
6. do not require careful aiming of the sensor for detection, and
7. they are close to the desired path.

Placing the landmark sensor on a mobile robot introduces relative motion between the robot and the landmarks. Relative motion can be used to enhance detection because it enables sensing from different directions. It can also cause problems with maintaining the landmark within the field of audition of the sensor, requiring the sensor to be moved relative to the robot. The landmark must be within the field of audition of the sensor for sufficient time for it to recognise the landmark and to determine its location.

When a robot is moving directly toward a landmark, the landmark progressively takes up a larger area of the field of audition and may overlap the edges. To make detection independent of range, the echo is decomposed into range dependent and range independent information. The absolute range is taken as the distance to the first significant peak in the echo. The information in a window around that peak is used for landmark detection. The resolution of the sensor is constant, so the relative positions of reflecting surfaces (leaves) are constant regardless of how far away the plant is. Scaling the plant echo spectrum with respect to a range calibration curve makes the echo amplitude information range independent.

When a robot is moving past a landmark (Figure 1), the landmark moves through the field of audition. At the first sensing position in Figure 1, the landmark is not in the region of insonification. At sensing point 2, one side of the plant is insonified and may be partially detected. At sensing point 3, the entire plant is in the field of audition and the plant can be detected.

Partial detection causes a problem when the section of the plant in the field of audition is not representative of the whole plant. This problem can be overcome by looking for the landmark until it is detected and then panning the sensor to track a landmark as the robot moves past it. Panning produces acoustic density profiles from different angles, from which it may be possible to measure the orientation of the robot relative to the plant.

3 Navigation Strategies

A sensor that can detect plants is suitable for piloting or landmark navigation. The strategy chosen for navigation will depend on the nature of the environment (static or dynamic), the robot s knowledge of the path (has or does not have a map), and the symmetry of the plants. These also impact the sensing strategies.

One common navigation strategy is the teach/replay strategy where the robot is taught a path by driving it along the path. As it moves, it senses objects beside the path to select appropriate candidates for landmarks. The location and features of these objects are stored in the map. Autonomous navigation then becomes the replay of this path, with the robot looking for landmarks to confirm that it is following the path.

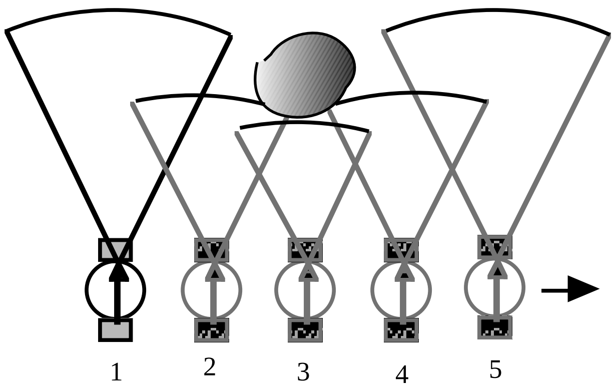

Figure 1. A robot sensing orthogonally to the direction in which is moving.

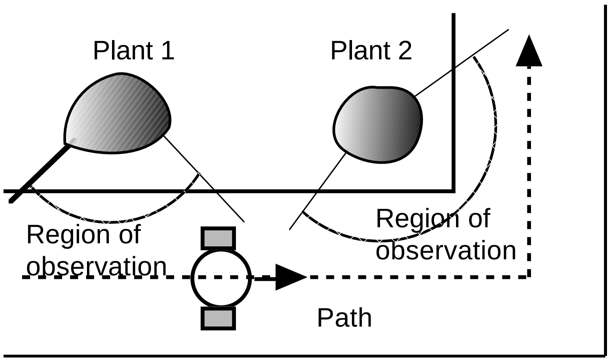

Figure 2. A simple path flanked by 2 plants.

Consider the robot moving along the path in Figure 2. This path consists of two straight trajectories, and two plants have been recorded as landmarks. The proposed navigation system will determine the expected bearing to the first plant from the current location and the map, and the robot will point the sensor in this direction. As the robot moves, this bearing changes and hence the direction of the sensor axis is changed. When the sensor detects the plant, it will localise the robot to the map from the range and bearing of the plant. Also, the system can count plants to determine when the robot is at the bend in the path. Once at the bend the robot can track the range to the plant as it turns to confirm that it is on the new trajectory without any tracking error.

The above proposal assumes that the plants are symmetrical and, as a result, can be reliably detected within the regions of observation shown in Figure 2 with a single set of features. While some plants are symmetrical, many are not. This asymmetry can be used to determine the orientation of the robot to the plant. Even the most asymmetric plants have regions where the features change slowly. So a plant can be divided up into sectors with partial symmetry, as shown in Figure 3, and a set of features recorded for each sector. Then, when a robot tracks around a plant it should be able to use the feature information to determine its orientation relative to the plant. We can use correlation to determine the symmetry of plants [Harper and McKerrow, 2000].

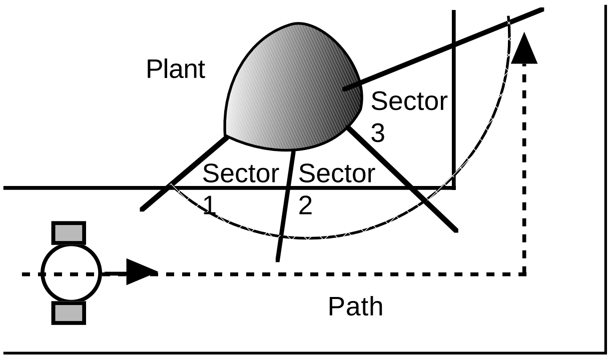

Figure 3. A plant that is acoustically different between sectors but similar within sectors.

Three problems with this proposed strategy are: plant growth, plant motion with wind, and plant disappearance. The impact of plant growth can be minimised by updating the feature set every time the robot travels along the path (Figure 4). The effect of wind can be minimised by averaging the features over several readings or by using echo tracking. Plant disappearance requires the navigation system to be intelligent enough to realise that a plant is not where it should be and look for other landmarks.

Figure 4. Titan navigating along a row of corn using a CTFM phased array ultrasonic sensor.

4 Landmark Types

An ultrasonic sensor is a geometry sensor, so we decompose landmarks on the basis of geometry. The result of this decomposition is a set of four landmark classes: simple discontinuous, simple continuous, complex discontinuous and complex continuous (Figure 5). The class of the landmark determines the feature set used to detect the landmark, the way in which the sensor is moved to collect data for landmark recognition and the navigation strategy used when tracking the landmark.

Figure 5 shows two axes. As we move to the right along the x axis, the complexity of the geometry of the landmark increases. A cylindrical pole is a simple-discontinuous landmark. In contrast, an isolated plant is a complex-discontinuous landmark.

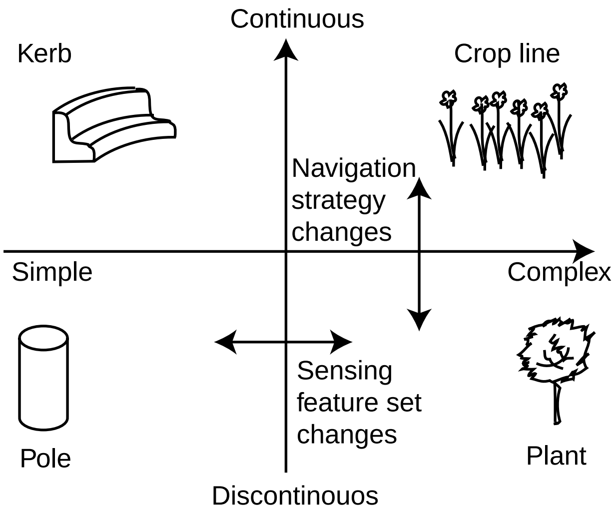

Figure 5. Four general landmark classes.

As we move up the y axis the continuity of the geometry of the landmark increases. The isolated plant becomes a crop line or row of plants that overlap (Figure 4). The surface facing the sensor is very rough compared to the smooth surface of a wall or a kerb.

In our previous research into recognizing plants with Continuous Transmission Frequency Modulated (CTFM) ultrasonic sensors [Kay 2000], we developed the acoustic density profile model [Harper and McKerrow, 1999b]. A CTFM sensor returns a set of frequencies with each frequency proportional to the range to a reflecting surface in the object. A simple object has a small set of frequencies and a complex object a large set.

After signal processing, we obtain a profile of the frequency content of the echo. The absolute range is removed by windowing, leaving a profile of echo amplitude versus distance through the object. This profile is the Acoustic Density Profile of the object from a particular sensing direction. From this profile we extract a set of features. Using the features we attempt to classify the object.

However, the research in [Harper and McKerrow, 1999a] was performed with isolated plants in a laboratory environment. Each plant was placed on a precision positioner and rotated in steps of 1^{o} while the sensor was stationary. The result was 360 readings around the plant. As one of the research goals was to determine whether a plant could be recognized from any angle, we looked for a set of features that were invariant with rotation.

When we place a sensor on a robot moving through an environment, the motion of the object relative to the sensor changes. Also, we may require a different parameter to be invariant. As a result the way the object is scanned and the feature set may be different. In the following sections, we look at examples of each of the 4 classes of landmark to determine both sensing and navigation strategy.

5 Simple-Discontinuous Landmarks

A cylindrical pole is a simple discontinuous landmark (Figure 6). Common objects in an outdoor environment including fence posts, light poles and tree trunks are members of this class. Navigation experiments with a sequence of poles as landmarks [Ratner and McKerrow, 2000] demonstrate the robustness of navigation with this type of landmark.

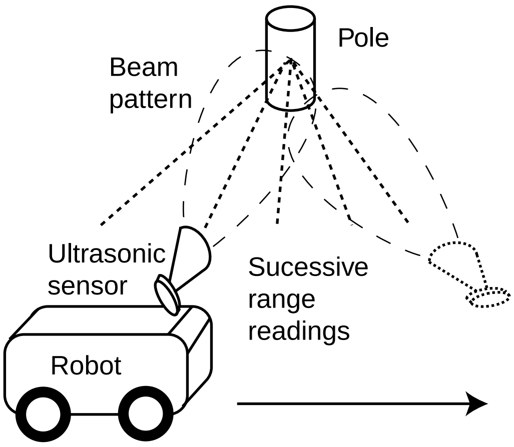

Figure 6. Sensing a simple-discontinuous landmark.

A cylindrical pole is symmetrical, so its echo varies little with rotation. As a result it can be detected from almost any angle, with a single feature set. The advantage is that it gives accurate and robust bearing and range readings relative to the sensor. The disadvantage of its symmetry is that detecting a pole is not sufficient to obtain an absolute bearing value.

In previous work [Ratner and McKerrow, 2000], we controlled sensor motion by stopping the robot and panning the sensor (Figure 6). The result is a fixed range reading while the pole is in the region of insonification, which increases rapidly on either side of that region (Figure 9). The range value is fixed because the smooth surface results in specular reflection, and consequently the echo is from the surface element that is orthogonal to the beam axis.

The measurements reported here were made with a CTFM phased array mounted on the front left corner of the robot. This sensor produces a vertical sheet of ultrasonic energy with a horizontal beam angle of 3^o (axis to side of beam) and a vertical beam angle of 30^o.

When the pole is not in the region of insonification a small amplitude echo will be received from the pole. If the pole is isolated from other objects, it will continue to be the largest amplitude echo (Figure 8). In which case, we have to threshold the echo to isolate the reflection from the main lobe, in order to determine the bearing of the pole. As the pole is symmetrical, the vector along which the range is measured passes through the center of the pole. The location of the center can be calculated using simple geometry.

The two curves in Figure 9 represent the start and finish of the acoustic density profile after thresholding. In the case of a tree trunk, the small acoustic depth indicted could be due to surface roughness or due to leakage between bins of the FFT. When the pole nears the edges of the region of insonification the amplitude drops off in accordance with the directivity function of the beam.

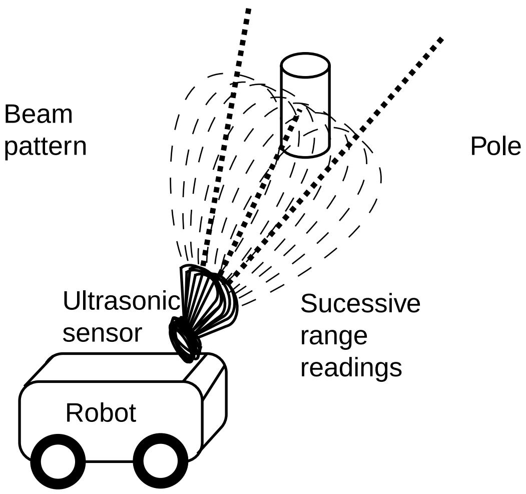

Figure 7. Panning the sensor across a pole.

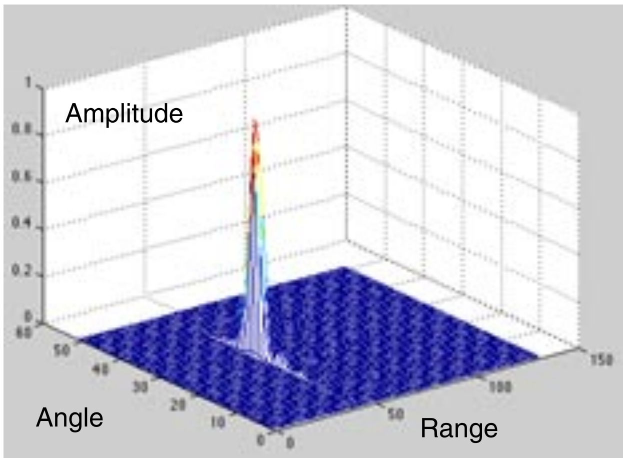

Figure 8. Linear sonargram of range readings from a 22.25 mm pole at a range of 817 mm when the sensor scanned across it in 1^o angular steps.

While this sensor motion strategy finds the pole, it has the disadvantage that it requires several samples and no robot motion. An alternate strategy is to aim the sensor at an angle to robot motion (Figure 6) while the robot moves along the path, and then aim it toward the center of the pole once it detects the pole, as the robot moves.

Aiming the sensor forward and to the side of the robot allows the sensor to detect the landmark before the robot reaches it. This look ahead time enables the control system to use feed forward control to calculate the corrections to the next path segment before the robot commences that segment, resulting in smoother more accurate control.

Again, when the pole is not in the region of insonification the range reading will be large. As soon as the pole enters the region of insonification, the range reading will drop. As the pole is symmetrical, the vector along which the range is measured passes through the center of the pole. The location of the center can be calculated using simple geometry, and then the sensor panned, as the robot moves to point to the center.

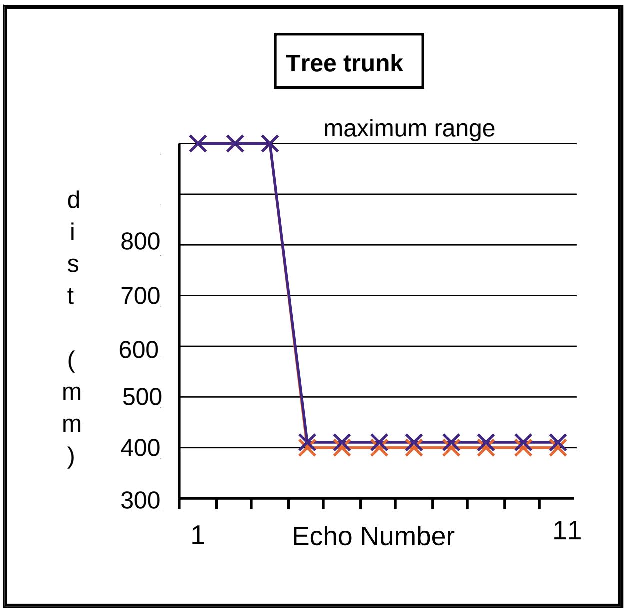

Figure 9. Range readings from a tree trunk, after thresholding, when the sensor scanned onto it from the left as in Figure 7 - bottom curve is start of echo, top curve is end of echo.

Several readings can be taken to confirm that the object is a pole and hence the landmark. The feature set for a pole is simple because the geometry is simple. The smooth vertical surface of the pole results in an acoustic density profile that is a single peak in a narrow window. Because the motion of the sensor relative to the plant is equivalent to rotating the plant relative to a fixed sensor, the feature set used in the laboratory research should be sufficient.

This sensing strategy results in a navigation strategy that involves panning the sensor to point toward the center of the landmark while the robot traverses the path past it. The same strategy can be used whether the path is straight or curved. This gives the robot time to get an accurate geometric fix on the landmark. Then it can update its position on the map and replan the next path segment to correct for any errors in path following.

6 Complex-Discontinuous Landmarks

A plant is a complex-discontinuous landmark (Figure 10). Common objects in an outdoor environment including flowers in gardens, shrubs beside paths and young trees in parks are members of this class.

When sensing plants, stopping the robot and panning the plant (Figure 11) is a poor sensing strategy. While it will detect the plant in isolation, the sensing beam insonifies different sections of the plant and consequently the echo may be very different. Also, the feature set developed by rotating a plant is not suitable.

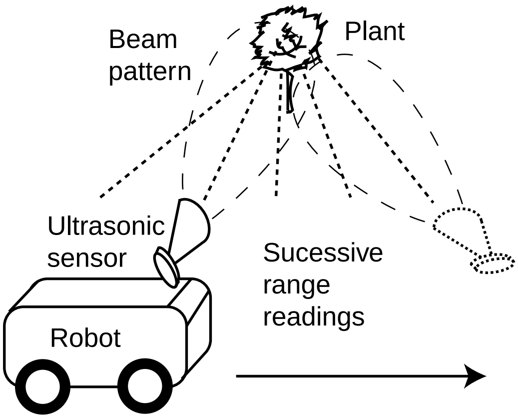

Figure 10. Sensing a complex-discontinuous landmark.

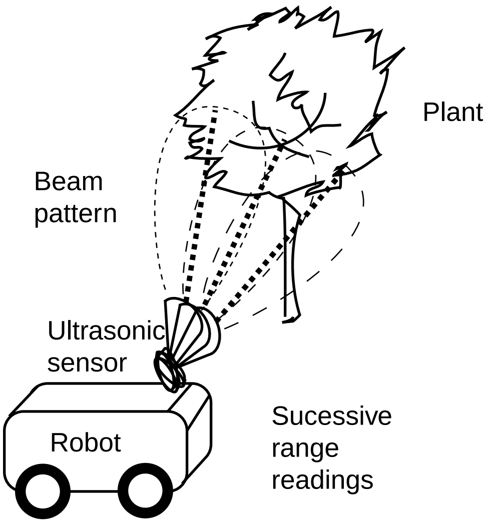

Figure 11. When panning a plant the region of the plant insonified by the sensor changes and hence the echo features may be difficult to match.

In contrast, panning the sensor to point toward the nominal center of the plant means that the same feature set can be used because of the equivalence of the motion. Also, time is saved because the robot continues to move while the features are compared to the landmark features. When the plant is asymmetrical, the variation in the feature set can be used to obtain a coarse measure of absolute bearing.

Again, this sensing strategy results in a navigation strategy that involves panning the sensor to point toward the center of the landmark while the robot traverses the path past it. However, with complex objects, the nominal center of the object will vary with the orientation of the sensor relative to the plant. Thus, the simple geometric model of a single point at the center of the pole becomes a circular region in which the center may occur.

When navigating with poles the navigation system can count poles and expect to detect a pole at a point on a path segment. The advantage of a plant over a pole is that the features of different plants are different. So the robot can often tell which plant it is, information that increased the confidence of the navigation system.

7 Simple-Continuous Landmarks

A flat vertical surface is a simple discontinuous landmark (Figure 12). Common objects in an outdoor environment including walls, kerbs, and the sides of building are members of this class. Again, a flat surface results in a simple feature set due to the simplicity of the acoustic density profile, but the continuous geometry of the landmark results in a different sensing strategy.

When the surface is smooth, the sensor has to point nearly orthogonal to the surface (within the beam angle off orthogonal) to detect the surface. As the roughness of the texture increases the angle can be greater. Aiming the sensor to point to the side of the robot reduces the look ahead time to zero, which means the robot may have difficulty responding to changes in the landmark — for example, when following a curving gutter.

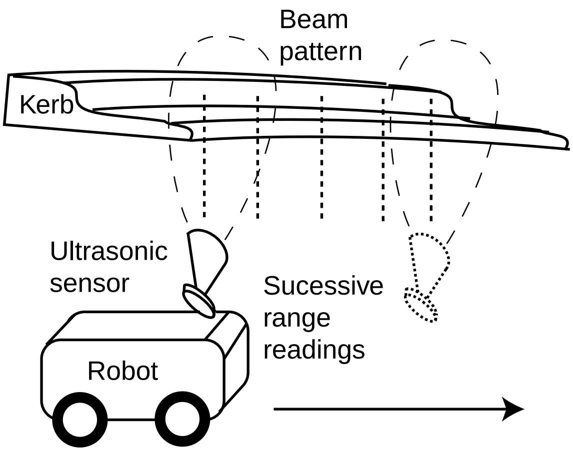

Figure 12. Sensing a simple-continuous landmark.

The sensing strategy has to be one of sensing different parts of the landmark as the robot moves beside it. The control system can use Proportional + Derivative + Integral (PID) control to maintain the robot at a fixed distance. The navigation strategy is to follow the continuous landmark until either it ends or an isolated landmark occurs. When the end of the continuous landmark is considered to be an isolated landmark, then in both cases the isolated landmark is used to correct the location of the robot relative to its planned path.

8 Complex Continuous Landmarks

A crop line is a simple discontinuous landmark (Figure 13). Common objects in an outdoor environment including picket fences, the edge formed by a mower of a harvester and rows of vegetables in gardens are members of this class.

Sensing and following complex-continuous landmarks differs from the simple-continuous landmark case in a number of ways. Because these complex-continuous surfaces are never smooth (Figure 14), the sensor can be aimed forward as well as to the side. As a result the robot has some look ahead distance in which to respond to changes and to replan the next path segment.

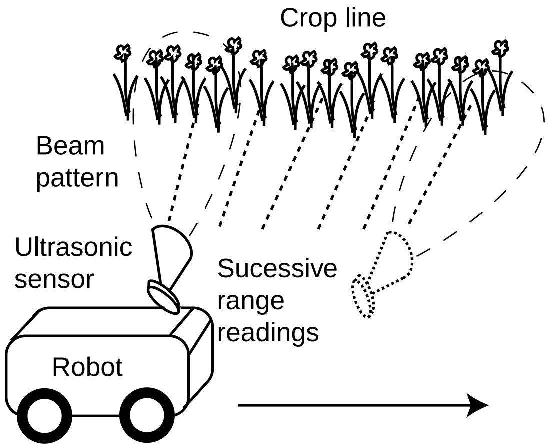

Figure 13. Sensing a complex-continuous landmark.

Figure 14. View of a row of tomato bushes, held up by wires, across the top of the phased array ultrasonic sensor.

Plants change in size and hence in the distance they protrude into the path as they grow. Thus, over a period of time a straight path may shift laterally and become crooked due to the changing shape of the plants. Also, a continuous row of leaves is formed when plants touch or overlap. The rhythmic variation in range as the robot passes a sequence of plants can be used to count plants.

A problem with the strategy of pointing the sensor at an angle is that each echo is from a different section of a different plant. As a result the feature set can vary considerably along the row. In this case, we need a different feature set: a feature set that varies with plant type but is invariant with translation past the plant.

However, when the edge of the crop line varies with plant shape (that is the plants are touching and not growing into one another) it may be possible to estimate the center of a plant and use the plant as a discontinuous landmark. This is a project for future research.

In Figure 14, the row of tomatoes is held up with wires between poles. As a result the echo contains information from both a simple (wire) and a complex (tomatoes) smooth surface. The curves in Figure 15 represent the start and the end of the acoustic density profile. The start of the profile is defined by a strong echo from the wire and the end by echoes of leaves in the tomato bushes.

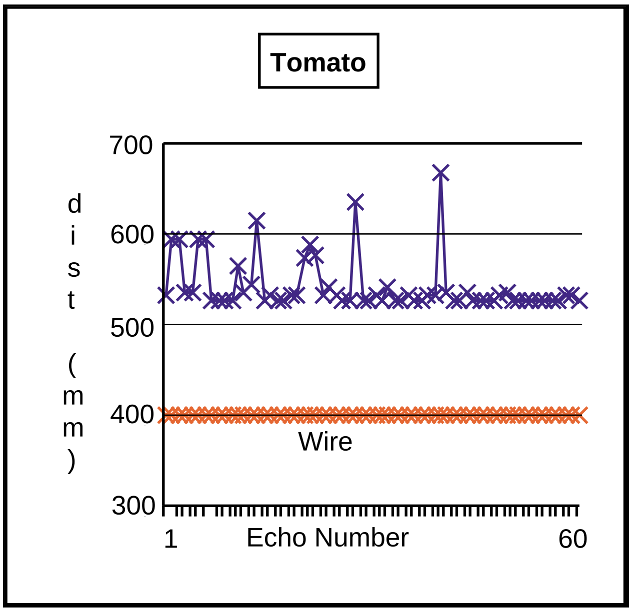

Figure 15. Sensing a row of tomatoes as robot tracks along row. Bottom curve is the wire at the start of the echo and top curve is the leaves at the end of the echo.

9 Conclusion

An ultrasonic sensor is a geometry sensor, so we decompose landmarks on the basis of geometry. The result of this decomposition is a set of four landmark classes: simple discontinuous, simple continuous, complex discontinuous and complex continuous. The class of the landmark determines the feature set used to detect the landmark, the way in which the sensor is moved to collect data for landmark recognition and the navigation strategy used when tracking the landmark.

We have developed a sensor motion strategy, a feature set and a landmark navigation strategy for each of the four classes of landmarks. These have been illustrated with data collected using a CTFM ultrasonic sensor. The implementation of a navigation system using these strategies for these landmark types will be the subject of a future paper.

References

[Akbarally and Kleeman, 1995] Akbarally, H. and Kleeman, L. 1995. A Sonar Sensor for Accurate 3D Target Localisation and Classification, Proceedings ICRA, Nagoya, pp 3003-3008.

[Gladwin, 1970] Gladwin, T. 1970. East is a Big Bird - Navigation and Logic on Puluwat Atoll, Harvard University Press, Cambridge

[Harper and McKerrow, 1999a] Harper, N.L. and McKerrow, P.J., 1999a, Classifying plants with ultrasonic sensing', Proceedings ARAA, Brisbane, pp 166-171.

[Harper and McKerrow, 1999b] Harper, N.L. and McKerrow, P.J., 1999b, Detecting plants for landmarks with ultrasonic sensing', Proceedings FSR'99, Pittsburgh, August, pp 144-149.

[Harper and McKerrow, 2000] Harper, N.L. and McKerrow, P.J., 2000, 'Recognising plants with ultrasonic sensing for mobile robot navigation', Robotics and Autonomous Systems, No 34, pp 71-82.

[Kay 2000] Kay, L. 2000. Auditory perception of objects by blind persons, using a bioacoustic high resolution air sonar, J. Acoust. Soc. Of America. Vol. 107, No. 6, June, pp 3266-3275

[Kimoto and Yuta, 1993] Kimoto, K. and Yuta, S. 1993. Sonar Based Outdoor Navigation. Proceedings ICAR. pp 239 -244.

[Leonard and Durrant-Whyte, 1991] Leonard, J.J., and Durrant-Whyte, H.F. 1991. Mobile Robot Localisation by Tracking Geometric Beacons, IEEE Transactions on Robotics and Automation, Vol 7, No 3, pp 376-382.

[McKerrow, 1995] McKerrow, P.J. 1995. Robot Perception with Ultrasonic Sensors using Data Fusion , Proceedings IEEE S,Man and Cybernetics Conference, October, Vancouver, pp 1380-1385.

[McKerrow and Harper, 1999] McKerrow, P.J. and Harper, N.L. 1999. 'Recognising leafy plants with in-air sonar', Sensor Review, Vol 19, No 3, pp 202-206.

[Maeyama et al., 1994] Maeyama, S., Ohya, A. and Yuta, S. 1994. Positioning by tree detection sensor and dead reckoning for outdoor navigation of a mobile robot, Proceedings MFI 94, Las Vegas, October, pp 653-660.

[Mandow et al.., 1996] Mandow, A., Gomez-de-Gabriel, J.M., Martinez, J.L., Munoz, V.F., Ollero, A., Garcia-Cerezo, A., 1996, The Autonomous Mobile Robot AURORA for Greenhouse Operation, IEEE R and A Magazine, Vol. 3, No. 4, December 1996.

[Ratner and McKerrow, 2000] Ratner D. and McKerrow, P.J. 2000. 'Using LabVIEW to prototype an industrial-quality real-time solution for the Titan outdoor 4WD mobile robot controller, Proceedings IROS, 2000, Takamatsu, Japan, October, pp. 1428-1434.

[Waterman, 1989] Waterman, T.H. 1989. Animal Navigation, Scientific American Library, New York.

[Wijk,and Christensen, 1998] Wijk, O. and Christensen, H.I. 1998. Extraction of natural landmarks and localization using sonars, Proceedings SIRS, pp. 231-240, Edinburgh, July.