Fast Circular Landmark Detection for Cooperative Localisation and Mapping

Julian Ryde and Huosheng Hu

Department of Computer Science, University of Essex Wivenhoe Park, Colchester CO4 3SQ, England Email: {jryde, hhu}@essex.ac.uk

Abstract—Map building through cooperative localisation (co-location) using circular geometric targets and a SICK laser range scanner is investigated. The tenet of co-location is circle detection in laser range data. Two methods for circle detection, a Range Weighted Circular Hough Transform (RWCHT) and a novel squared-residual voting strategy are compared and their performance assessed. The custom squared-residual voting strategy outperforms the RWCHT in all respects and is subsequently used for localisation and map building. The results include robust continuous localisation at speeds of 0.2m/s with 98% of scan frames used and an error of less than 0.03m. This localisation accuracy helps build maps of 96% quality and occupancy grids of cluttered environments despite the presence of distractors.

Index Terms—Cooperative Localisation, Mapping, Hough Transform, Circle Detection

I. INTRODUCTION

It is widely agreed that the ability of a mobile robot to locate itself within the environment is the most fundamental problem currently thwarting mobile autonomous operation. Most current localisation methods require a prior map or attempt to build one. Therefore, equally important is the ability of mobile robots to build accurate environment representations (maps). However, building a map without knowledge of robot pose (position and orientation) is difficult. The problem of locating a robot given a prior map has been solved in many fashions, however there is not yet a universally adopted method. Difficulties localising using natural landmarks have resulted in a resurgence in the use of artificial landmarks for localisation [1]. Success using natural landmarks such as corners has been achieved by [2].

Mobile robots are expected to operate in a variety of locales, both indoors and outdoors, as well as, structured and unstructured. These environments are usually dynamic and vary enormously. Outdoors, global localisation is made considerably easier through the use of GPS. However, for indoor robots there is no such global location signal infrastructure. It is possible to use the wireless network signal strength as a signal locator [3]. The location accuracy is around 2m, orientation determination appears all but impossible with standard wireless network cards and the environment has to contain a wireless infrastructure. There is no doubting the necessity of fast, accurate and robust navigation and mapping approaches that would work in these environments. Better would be the extraction of paradigms that would span all arenas of operation. Cooperative localisation, hereafter referred to as colocation, is often assumed, [4], but implementations are not discussed. Initially the problem of co-location might be perceived as straightforward but robust, accurate and fast co-location techniques are not readily available.

In general, robots should be able to detect each other in order to co-locate and a number of methods were considered. The two main contenders were an infra-red beacon approach or using distinct shapes (geometric beacons) and a laser scanner. The infra-red beacon approach suffers from two main drawbacks, significant hardware deployment and the return of target angles only. The former is surmountable, however the latter is a serious problem. Localisation given separation angles of indistinguishable targets is possible, [1], [5], however the problem is highly non-linear and difficult to implement. A laser scanner and geometric beacons provide both the range and angle to targets vastly simplifying the process of relative localisation. The main difficulty is reliable geometric target extraction from the laser range data.

Once this is possible cooperating robots can be detected and removed from the laser scans producing an important improvement in the fidelity of the global map. This is an improvement upon [6] in which the robots detect each other as obstacles that should simply decay over time when the occupancy grid is updated.

In this paper, two techniques are implemented to extract circular landmarks from laser range data. Section II-A examines the established Hough Transform and Section II-B contains the novel squared-residual voting strategy tailored for range data. The performance of these two algorithms is then compared with the custom squaredresidual voting strategy proving to be more effective. The results for continuous localisation and map building are presented in Section III. Finally, the conclusion is given in Section IV.

II. CO-LOCATION

Returning the range and bearing to other robots requires the detection of these robots from the range scans. The robots need to carry geometric targets that are easily found and unique within the environment. Indoor environments usually contain many straight lines; consequently targets incorporating straight lines are not used. The detection process is greatly aided if the target always has identical



Fig. 1. Cooperative localisation and mapping scenario with three robots.

range signatures regardless of relative position or orientation. This is the case for one shape only, the circle. This characteristic aids detection but is not helpful when determining relative position because rotational changes cannot be perceived. Two distinguishable circles guarantee unique localisation. If the circles are indistinguishable then localisation is one of two places. Fig. 1 presents a cooperative localisation and mapping scenario involving three robots R1, R2 and R3. R1 is equipped with a laser scanner and the remaining robots are mobile landmarks. The initial positions of R2 and R3 allow R1 to map the room on the left. Under the observation of R1, at position A, R2 and R3 move across the corridor to the second room where they adopt positions B and C. Now R1 can continue to D using R2 and R3 as artificial landmarks and map the second room.

Once the relative positions of the companion robots are known, map building is possible. The main difficulty is achieving fast and reliable detection of circles of known radius from noisy range data. The detection of shapes in images is a large area of research within the computer vision community and contains relevant techniques, specifically the Hough Transform and least squares fitting approaches.

A. Circular Hough Transform

The Hough Transform [7] has been hugely successful in the vision community thanks to its tolerance of image noise and excellent straight-line detection. A typical high resolution laser scan is given in Fig. 2 which shows an environment with two circular landmarks.

A Range Weighted Circular Hough Transform (RWCHT) similar to [8] was used to extract the circular targets from Fig 2. The resulting accumulator array is depicted in Fig 3. The RWCHT's confusion of straight lines with circles was a serious problem that refused to be resolved. A possible solution would be to first remove all points corresponding to straight lines and then perform the RWCHT on the remaining points, however this is very time consuming.

There are a number of reasons why the circular Hough Transform was not particularly suited to this application. Range data is different to image data for which the Hough Transform was first devised. Another problem is that it always returns an answer even if the geometric primitive is not present in the data. The determination of peak significance by comparison with others and the kind of data expected requires a complicated statistical analysis.



Fig. 2. Typical laser scan at high resolution (0.25°) and 100° scan angle. Two circular landmarks are indicted with dashed lines.



Fig. 3. Surface plot of the accumulator array searching for circles in the range scan shown in Fig. 2 using a RWCHT. The two highest peaks correspond to the circular landmarks.

B. Circle Detection by Squared-Residual Voting

It is evident from [9] that fitting circles to points is a nontrivial process, mainly because the resulting equations are highly non-linear and circles cannot be elegantly expressed in Cartesian coordinate systems.

One of the problems with the circular Hough Transform is that there is much information specific to range scans that is not included in the search for circles. One important property of circles is that they are highly symmetric and so appear identical when viewed from any angle; this greatly eases the burden of detection. Also, the range data has an inherent sequence that is not obvious in Cartesian coordinates. Detection of a circle occurs when a sequence of adjacent points lie close to the circumference of that circle. Relaxing the requirement for the detection of occluded targets allows the following algorithm expounded in Fig. 4 and Fig. 5.

The algorithm assumes the centre of the circular target is at the scan angle of the current scan point being analysed. The mean of the squared-residuals is then calculated by (4) and (5). Scan angles with this quantity below a threshold (comparable to the accuracy of the laser scanner) are likely contenders for having the centre of the target circle situated



Fig. 4. Geometric construction illustrating the squared-residual voting method for circle detection.



Fig. 5. Flowchart summarising the squared-residual voting strategy.

along them. Fig. 4 illustrates the geometry involved with laser scan points depicted by crosses. Point A is the current scan point being evaluated and the circle represents the search target. The candidate circle for A is assumed to be positioned with centre C, as shown on the line \overline{OA} where O is the origin of the laser scan. Assuming the laser scan returns points evenly distributed over θ then the number of nearest neighbours to be incorporated is determined. Points that lie within an angle of \widehat{AOB} from A are candidate points where

$$\widehat{AOB} = \arcsin\frac{R}{R + \overline{OA}} \tag{1}$$

and R is the radius of the circular landmark. Care has to be taken regarding scan points lying near D and B, which are subject to glancing edge effects. The causes of these effects are specular reflection and pixel mixing which occurs when the laser spot spans an environmental range discontinuity. The subset of laser range points processed is

$$\mathbf{S} = \begin{pmatrix} r_1 & r_2 & \cdots & r_n \\ \theta_1 & \theta_2 & \cdots & \theta_n \end{pmatrix}$$
(2)

where n is odd; r and θ are the polar coordinates of the scan points in the coordinate system of the robot. The position of the hypothesis circle in polar coordinates is

$$\begin{pmatrix} C_r \\ C_\theta \end{pmatrix} = \begin{pmatrix} r_{\frac{n+1}{2}} + R \\ \theta_{\frac{n+1}{2}} \end{pmatrix}.$$
 (3)

The distance of the *i*th point from circle circumference is

$$d_{i} = \sqrt{C_{r}^{2} + r_{i}^{2} - 2C_{r}r_{i}\cos\left(C_{\theta} - \theta_{i}\right)} - R.$$
 (4)

Ultimately the mean and squared-residual is calculated in the usual fashion as

$$\overline{d}^{2} = \frac{1}{n} \sum_{i=1}^{n} d_{i}^{2}.$$
(5)



Fig. 6. Reciprocal root mean squared-residuals of the laser scan in Fig. 2. The two prominent peaks correspond to the circular landmarks.

This indicates how far, on average, the points are from the circumference of the hypothesis circle and the reciprocal is proportional to the likelihood of detection. This is repeated for each point in the scan. The points that exceed a threshold probability imply successful circle detection at that position. Fig. 6 plots the reciprocal root mean squaredresiduals for the example laser scan in Fig. 2.

What is apparent from Fig. 6 is the accurate detection and localisation of the two circular targets with the smaller of the two circle peaks being nearly twice as big as the largest background peak. This ensures a superior performance of 98% reliability versus 50% for the RWCHT. A comparison of Fig. 3 and Fig. 6 emphasises the effectiveness of the squared-residual voting strategy over the RWCHT for reliable circular target extraction from laser range data. The squared-residual voting strategy takes advantage of range data specific characteristics like sequence and a single observation point. The more generic RWCHT does not utilise this extra information and so the squared-residual voting strategy is not only 25 times more accurate but also faster and requires less memory.

C. Cooperative pose change determination

The two cylindrical targets are observed from two different poses and the observations superimposed. This is shown in Fig. 7 with the second observation cylinder positions indicated with an apostrophe. The pose change consists of a rotation and translation. The rotation angle is the change in angle of the line joining the two circles. This angle is indicated in Fig. 7. Once the rotation of the robot between the poses is known, the rotation effect can be undone placing the cylinders at the positions C and D, as shown in Fig. 7. The change in position or translation of the robot between observations is given by the difference in position of the midpoints of \overline{CD} and \overline{AB} . Knowing the rotation, θ , and translation, **T**, of the robot between successive scans, enables the amalgamation of scan data to produce a global map. Scan data, L, is transformed point by point into the coordinate frame of the global map, \mathbf{L}' , by

$$\mathbf{L}_{i}^{\prime} = \begin{pmatrix} T_{x} \\ T_{y} \end{pmatrix} + \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix} \mathbf{L}_{i}.$$
 (6)



Fig. 7. Pose change calculation from two observations.

Given that a robot can observe other stationary robots, how may it determine changes in its pose? Changes in pose my be described as linear combinations of two geometric transforms, translation and rotation. An important consideration is if the observed robots are distinguishable; if they can be unambiguously identified then the determination of pose change between landmark observations is trivial. The rotation is calculated from the change in angle of the lines joining the landmarks and the translation is the average displacement of each point to its image point. If the landmarks are indistinguishable then it is not so straightforward because each point cannot be associated with absolute certainty to the same point in the subsequent sensor update. Problems also arise with symmetric distributions of landmarks.

If the relative positional information of indistinguishable landmarks is available then three are sufficient to unambiguously determine pose. Initially two would appear sufficient, however the ambiguity of identity means that landmarks may be rotated 180° degrees. Even though only three asymmetrically distributed indistinguishable landmarks are needed for unambiguous pose determination, the fewer landmarks required the better. Is it possible to have reliable pose updates using only the relative positions of two landmarks? There are a number of ways that this may be achieved. The simplest is to use distinguishable landmarks, for instance circles of sufficiently different radii. If indistinguishable landmarks have to be used then they may be placed in such a configuration so that localisation is only required in one half plane. An example would be when they are against a wall then the robot cannot be localised in the half plane behind the wall and still be able to detect the landmarks. Use of odometry and fast updates means that the large pose changes that would cause ambiguity would never happen between updates or would be detected by the odometry sensors.

III. EXPERIMENTAL RESULTS

The experimental platform is a Magellan Pro robot equipped with a SICK LMS 200 laser range finder. The range-finder has a scanning angle width of 180° and a resolution of 0.5° . The laser range finder is almost an ideal sensor with unrivalled accuracy, acquisition time and range. The main problems are cost, mass(4.5kg) and power



Fig. 8. Geometric construction used to calculate the localisation error.

consumption (17.5W). The characteristics of this LMS are detailed in [10], [11].

A. Localisation

Experiments were performed to test the localisation accuracy delivered. They involved driving the robot along a straight line and in a square. The deviation of the co-location positions from this straight line give an indication of the localisation error in the direction perpendicular to the line. This error depends approximately linearly on the angular resolution of the laser scanner, the range and separation of the geometric targets. The localisation error was of the order of 0.02m at ranges of 0 to 8m with the laser scanner operating at a resolution of 0.25° .

Error in the range to the targets introduces error into the position estimation of the robot. Fig. 8 illustrates the dependence of the pose uncertainty on the range error. The origin O is the true position of the robot and O' is its worst case perceived position if the range to the target A is over estimated and that to target B is underestimated. The error estimate is greatly simplified if a far field approximation is used which implies

$$\overline{AB} \ll \overline{OM} \tag{7}$$

in this approximation the following similarities prove useful

$$\sin\theta \approx \tan\theta \approx \theta \tag{8}$$

for small angles of θ in radians. As $\overline{OM} = \overline{O'M}$ then for the displacement of O' the angle of rotation is

$$\widehat{OMO'} = \arctan\left(\frac{2}{\sqrt{2}}\frac{\overline{AA'}}{\overline{AB}}\right) \approx \sqrt{2}\frac{\overline{AA'}}{\overline{AB}}.$$
 (9)

Note the root two factor due to the addition of the errors in quadrature. Finally the position error can be expressed as

$$\overline{OO'} \approx \sqrt{2} \frac{\overline{AA'}}{\overline{AB}} \overline{OM}$$
(10)

The far field approximation, expressed in (7), falters if the targets are near and for large target separations, however in these situations the error is minimal. It should also be clear from Fig. 8 that the dependence of position error σ_x on angular error (σ_{θ}) for the laser scanner is simply

$$\sigma_x \approx \overline{OM} \sigma_\theta \tag{11}$$

The angular error for the SICK LMS 200 is circa 0.5° so at a range of 4m the position error due to angular error is around 0.03m. Targets separated by 2m with radii 0.1225m



Fig. 9. Plot showing the increase of position error with range to targets and line of best fit.



Fig. 10. Localisation along square path, solid line indicates true path taken. The circles are the geometric targets used for localisation.

at a range of 4m observed with a range error of 0.01m produced a position error of 0.03m. This prediction is close enough to the error observed at this range in Fig. 9.

A typical set of continuous localisation results is displayed in Fig. 10. The robot was moved one loop around a 1.57m square at 0.2m/s. The laser scanner mounted on the robot has a maximum scan angle of 180° and so the robot had to reverse along some edges of the square in order to maintain tracking. The target cylinders were located at (0,1) and (-1,0) because in these positions they can always be observed by the 180° scanner, allowing continuous position updates. The localisation error can easily be extracted from Fig. 10 and is of the order of 0.03m. The position accuracy is better towards the origin of the graph because the robot is nearer to the target positions of (0,1) and (-1,0).

B. Cooperative Map Building

The accomplishment of reliable and accurate cooperative localisation is easily extended to allow the construction of local maps. The co-location gives the rotation and translation between successive laser range scans. This information allows all the scans to be converted into one reference



Fig. 11. Typical global map with the ground truth (solid line) and trajectory (dotted line).

frame resulting in the production of a global map of the environment. First the laser data is acquired from the SICK LMS which has a variety of angular resolution and range modes. The two main modes used were 180° at 0.5° resolution and 100° at 0.25° . Better angular resolution enables more accurate co-location over a larger area surrounding the reference targets.

It is important to note that although co-location may be limited to 8m from the landmarks larger environments can still be mapped because the robot can incorporate observations from outside this area. The increased data associated with higher resolutions results in slow data acquisition and increased processing time. The most appropriate laser mode will ultimately depend on the operational environment, small rooms (less than eight metres) will need wide angle perception and larger rooms will require better accuracy.

Once the laser range data has been acquired, the circular geometric targets are located and extracted from the range data using the squared-residual voting strategy, which finds circles of prior known radius from range data. If the best candidate has a root mean squared-residual exceeding 0.01m then the scan is rejected and the next one processed. Possible causes of rejection include situations when the targets cannot be perceived as they are outside of the angular or distance range, occlusion of targets and velocity aberrations.

Once the positions of the stationary targets have been extracted the change in pose enables alignment of the scan to the global map. Using, (6), the scan, without the targets, is added to the global map. It is important to remove the targets once they have been detected as they are not part of the environment and should not be added to the global map.

A quantitative idea of mapping accuracy may be obtained by comparing dimensions extracted from Fig. 11 to the measurements taken from the operational arena. This analysis yields a deviation of around 4% between dimensions extracted from the global map and those measured from the ground truth. The performance in more realistic, complex and cluttered environments is demonstrated



Fig. 12. Overhead camera view of the robot arena.



Fig. 13. Occupancy map generated of environment in Fig. 12.

through a comparision of Fig. 12 and Fig. 13. Fig. 12 is an overhead view of the robot arena. This cluttered environment has circular targets identical to the two central ones used for localisation. Augmenting the distictiveness of the targets with retroreflective tape eliminates false landmark matches. Fig. 13 is an occupancy map built using the two central cylinders as landmarks differentiated from the remaining distracting cylinders by the retroreflective tape. Unobserved areas are depicted in grey, obstacles in black and free space is white. The mapping strategy incorporates both retroreflectivity and curvature so the chance of confusing targets appearing naturally is incredibly remote. This gives rise to good maps of cluttered environments despite the presence of distractors.

IV. CONCLUSION

This paper demonstrates the feasibility of cooperative localisation and mapping based on one sensing robot and two landmark robots. It is implemented by a novel squaredresidual voting strategy optimised for the sequential nature of the range data and the highly symmetric aspect of the circular geometric targets.

This co-location scheme allows fast position and orientation determination with bounded errors and reliability indicators in unknown indoor environments. The robust localisation algorithm lays the foundation for mapping featureless and highly symmetric environments. Continuous localisation was performed at 0.2m/s and the map shown in Fig. 11 was built. This map is 96% accurate however more accurate maps can be constructed if the robot moves a short distance to a new pose then stops and scans. Continuous localisation can be provided however these scans should not be incorporated into the global map, only the ones taken when stationary should be used to improve the quality of the global map.

Improvements in co-location accuracy should be possible allowing either the extending of the range over which cooperative localisation is possible or reducing the separation of the targets so that they may be mounted on one robot thus allowing cooperative mapping with only two robots. These improvements in co-location accuracy would primarily come from oversampling the squaredresidual voting strategy. Our next stage of research aims at merging maps produced from target robots in different positions. Once this has been achieved the way is clear for robust, reliable and accurate cooperative mapping of extensive unknown indoor environments.

REFERENCES

- H. Hu and D. Gu, "Landmark-based navigation of industrial mobile robots," *Industrial Robot: An International Journal*, vol. 27, no. 6, pp. 458–467, 2000.
- [2] L. Wang, L. Yong, and M. Ang Jr, "Mobile robot localisation for indoor environment," *SIMTech Technical Report*, 2002.
- [3] O. Serrano, J. M. Ca nas, V. Matellán, and L. Rodero, "Robot localization using WiFi signal without intensity map," WAF, 2004.
- [4] M. Betke and L. Gurvits, "Mobile robot localization using landmarks," *IEEE Transactions on Robotics and Automation*, vol. 13, no. 2, pp. 251–263, 1997.
- [5] K. Åström, "Automatic mapmaking," IFAC International Workshop on Intelligent Autonomous Vehicles, 1993.
- [6] B. Yamauchi, "Frontier-based exploration using multiple robots," in Proceedings of the Second International Conference on Autonomous Agents, 1998, pp. 47–53.
- [7] P. V. C. Hough, "Methods and means for recognising complex patterns," U.S. Patent 3 069 654, 1962.
- [8] J. Forsberg, U. Larsson, and Å. Wernersson, "Mobile robot navigation using range-weight hough transform," *IEEE Robotics and Automation Magazine*, pp. 18–26, 1995.
- [9] N. Chernov and C. Lesort, "Least squares fitting of circles and lines," *Computer Vision and Pattern Recognition*, 2003.
- [10] A. Aboshosha and A. Zell, "Robust mapping and path planning for indoor robots based on sensor integration of sonar and a 2D laser range finder," *IEEE 7th International Conference on Intelligent Engineering Systems*, 2003.
- [11] C. Ye and J. Borenstein, "Characterization of a 2-d laser scanner for mobile robot obstacle negotiation," *Proceedings of the 2002 IEEE International Conference on Robotics and Automation*, 2002.