

A Fast, Accurate, and Robust Method for Self-Localization in Polygonal Environments Using Laser-Range-Finders

Jens-Steffen Gutmann Thilo Weigel Bernhard Nebel

Albert-Ludwigs-Universität Freiburg, Institut für Informatik
Am Flughafen 17, D-79110 Freiburg, Germany
{gutmann,weigel,nebel}@informatik.uni-freiburg.de

January 21, 2000

Abstract

Self-localization is important in almost all robotic tasks. For playing an aesthetic and effective game of robotic soccer, self-localization is a necessary prerequisite. When we designed our robotic soccer team for participating in robotic soccer competitions, it turned out that all existing approaches did not meet our requirements of being fast, accurate, and robust. For this reason, we developed a new method, which is presented and analyzed in this paper. This method is one of the key components and is probably one of the explanations for the success of our team in national and international competitions. We present also experimental evidence that our method outperforms other self-localization methods in the RoboCup environment.

1 Introduction

Robotic soccer is an interesting scientific challenge [11] and an ideal domain for testing new ideas and demonstrating existing techniques. One of our main intentions in participating in robotic soccer competitions was to demonstrate the usefulness of self-localization techniques that we have developed [9]. It turned out, however, that all existing self-localization techniques were not efficient enough for a dynamic environment such as robotic soccer. Furthermore, most of the techniques are not robust enough. For this reason, we developed a new technique that exploits one particular characteristic of the RoboCup environment, namely, its purely polygonal structure. Based on that we were able to come up with a very fast, accurate, and robust self-localization technique, which was most probably one of the key factors for the success of our team *CS Freiburg* [7, 15]. This team won the the German competitions *Vision RoboCup'98* and *Vision*

RoboCup'99. In addition, the team became world champion at *RoboCup'98* [1] and it came out third place at *RoboCup'99* [18].

Solving the self-localization problem—the problem of determining the position and orientation of the robot—is necessary for almost all tasks. In robotic soccer it seems even to be impossible to play an effective and aesthetic game if the soccer agents do not know where they are and how they are oriented. As a matter of fact, some of the problems displayed in the games of the middle size league at *RoboCup'97* [10] seemed to have to do with the fact that the soccer robots had the wrong idea about their positions, which led to erratic movements and a number of own goals.

The self-localization problem can be addressed using a wide range of sensors (e.g. odometry, sonars, vision, compasses, laser range finders, other sensors, or combinations thereof) and a wide range of methods. In the sequel we will only consider the combination of data from the odometry and from laser range finders (LRF), since the latter provide accurate and reliable data, which can be interpreted with much less computational effort than, say, data from a vision system.

Self-localization can be based on recognizing known *landmarks* or on *dense sensor matching*. In the first approach, features are extracted from the sensor inputs and matched with the features of the landmarks in order to determine the locations of the landmarks. However, in the RoboCup environment, there are only few natural landmarks that are always visible to the sensors and for this reason we did not consider this approach. In the second approach, all sensor inputs are matched against the expected sensor inputs for a given model. Two competing methods for dense sensor matching are grid-based *Markov localization* [3, 2] and *Kalman filtering using scan matching* [5, 9]. As it has been demonstrated, Markov localization is more *robust*, because it always generates some position hypotheses and because it can recover from catastrophic failures. However, self-localization using Kalman filtering based on scan matching is more *accurate* [6], since it does not rely on grids.

For robotic soccer, we need *robustness*, *accuracy*, and *efficiency*, whereby the latter property means that we want to estimate the position and orientation in a few milliseconds. Unfortunately, none of the approaches described above satisfies all three requirements. For this reason, we designed a new scan-matching approach that extracts features from the raw sensor inputs, namely, straight lines, that are matched against an *a priori* model. Using the scan match, which can be computed efficiently, the new position estimation is then derived by combining it with the odometry reading using Kalman filtering.

The rest of the paper is structured as follows. The next section sketches scan matching methods and how they can be used to estimate the position using Kalman filtering. Section 3 describes our own method and gives an analysis of the run-time complexity. Based on that, we describe in Section 4 experiments that we have made in order to compare different scan matching methods in the RoboCup environment. Finally, in Section 5 we conclude and sketch future work.

2 Scan Matching

Scan matching is the process of translating and rotating a range scan (obtained from a range device such as a laser range finder) in such a way that a maximum overlap between sensor readings and *a priori* map emerges. Most of the scan matching methods presume an initial pose estimation that must be close to the true pose in order to limit the search space.

The robot pose and its update from scan matching are modeled as single Gaussian distributions. This has the advantage that robot poses can be calculated with high precision, and that an efficient method for computing the update step can be used, namely, Kalman filtering.

The extended Kalman filter method has the following form. The probability of a robot pose is modeled as a Gaussian distribution $l(t) \sim N(\mu_l, \Sigma_l)$, where $\mu_l = (x, y, \alpha)^T$ is the mean value and Σ_l its 3×3 covariance matrix.

On robot motion $a \sim N((\delta, \theta)^T, \Sigma_a)$ where the robot moves forward a certain distance δ and then rotates by θ , the pose is updated according to:

$$\begin{aligned}\mu_l &:= E(F(l, a)) = \begin{pmatrix} x + \delta \cos(\alpha) \\ y + \delta \sin(\alpha) \\ \alpha + \theta \end{pmatrix} \\ \Sigma_l &:= \nabla F_l \Sigma_l \nabla F_l^T + \nabla F_a \Sigma_a \nabla F_a^T\end{aligned}$$

Here E denotes the expected value of the function F and ∇F_l and ∇F_a are its Jacobians with respect to l and a .

From scan matching a pose update $s \sim N(\mu_s, \Sigma_s)$ is obtained and the robot pose is updated using standard Kalman filter equations [14]:

$$\begin{aligned}\mu_l &:= (\Sigma_l^{-1} + \Sigma_s^{-1})^{-1} \cdot (\Sigma_l^{-1} \mu_l + \Sigma_s^{-1} \mu_s) \\ \Sigma_l &:= (\Sigma_l^{-1} + \Sigma_s^{-1})^{-1}\end{aligned}$$

The success of the Kalman filter depends heavily on the ability of scan matching to correct the robot pose. There are a number of methods for matching scans:

Cox [5] matches sensor readings with the line segments of a hand-crafted CAD map of the environment. He assigns scan points to line segments based on closest neighborhood and then searches for a translation and rotation that minimizes the total squared distance between scan points and their target lines.

Weiss et. al. [19] use histograms for matching a pair of scans. They first compute a so-called angle histogram for determining the rotation of the two scans and then use x and y histograms for computing the translation. Although this method seems to be well suited for the RoboCup environment it is computationally expensive and the precision of the algorithm depends on the discretization size of the histograms.

Lu and Milios [12] match pairs of scans by assigning points in one scan to points in the other scan. For finding a corresponding scan point two heuristics called *closest-point-rule* and *matching-range-rule* are applied and a combination is used for computing the rotation and translation of the two scans. This IDC algorithm (*iterative dual*

correspondence) is well suited for any type of environment including non-polygonal ones.

Gutmann and Schlegel [9] use a combination of the Cox matching approach and the IDC method for combining the efficiency and robustness of the line matching method with the universal capabilities of the IDC algorithm. They call their algorithm the *combined scan matcher* (CSM).

Unfortunately all those matching algorithms possess a high computational complexity, i.e., $O(n^2)$ where n are the number of scan points, and their robustness is limited due to the small search space. Therefore we developed a new algorithm *LineMatch* that makes use of the simple polygonal structure of the RoboCup environment and trades off generality for speed and the ability to globally localize the robot on the soccer field.

3 Our Approach

Our *LineMatch* algorithm extracts line segments from a scan and matches them with an *a priori* map of line segments similar to the methods of [17, 4]. We expect that this algorithm has better run-time performance and is more robust than the other scan matchers while retaining the same accuracy as the other matchers. In how far these expectations are realistic will be shown in Section 4.

3.1 The LineMatch Algorithm

In order to guarantee that extracted lines really correspond to field-border lines, only scan lines significantly longer than the extend of soccer robots are considered. The following algorithm shows how a matching between model lines and scan lines is computed by recursively trying all pairings between scan lines and model lines:

Algorithm 1 *LineMatch*(M, S, P)

Input: model lines M , scan lines S , pairs P

Output: set of positions hypotheses H

if $|P| = |S|$ **then**

$H := P$

else

$H := \emptyset$

$s := \text{SelectScanline}(S, P)$

for all $m \in M$ **do**

if $\text{VerifyMatch}(M, S, P \cup \{(m, s)\})$ **then**

$H := H \cup \{\text{LineMatch}(M, S, P \cup \{(m, s)\})\}$

return H

SelectScanline selects the next scan line that should be matched and *VerifyMatch* verifies that the new (m, s) pairing is compatible with the set of pairings P already accepted by computing a common rotation and translation. The algorithm returns position hypotheses in the form of sets of pairs which can be easily transformed into possible locations where the scan could have been taken. For the RoboCup field the algorithm is capable of determining the global position of the robot modulo the symmetry of the field. This means that we get two position hypotheses if three field borders are visible (see Figure 1) and four hypotheses if two borders are visible.

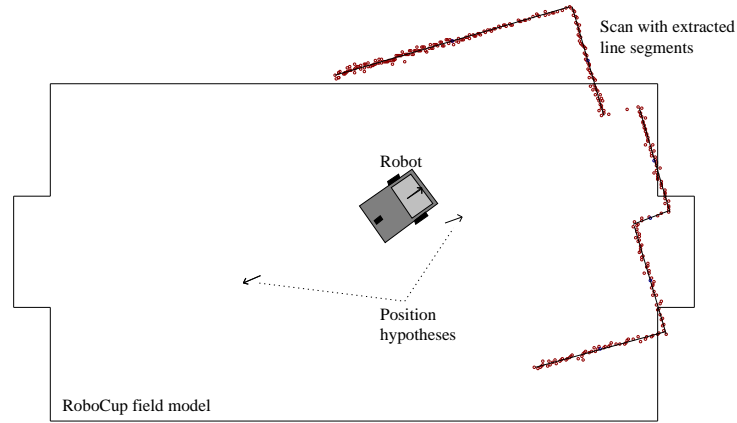


Fig. 1. The LineMatch algorithm returns two hypotheses for the robot position.

This scan matching method is similar to the methods described by Castellanos *et al.* [4] and Shaffer *et al.* [17]. In contrast to these approaches, however, we only verify that the *global constraints* concerning translation and rotation as well as the length restrictions of scan lines are satisfied. This is sufficient for determining the position hypothesis and more efficient. Further, we do not need any initial estimation of the pose, which means that even if the robot has an extreme error in its position estimation, it may still be able to recover from that.

After matching a range scan, the most plausible position is used in the Kalman filter step for updating the robot position (see Figure 2). We use the position information from odometry to determine the most plausible position based on a combination of closest neighborhood and similarity in heading.

For initializing the self-localization system the robot is placed at any position in the RoboCup field but roughly oriented towards the opponent goal and the mean and error covariance of the robot position are set to:

$$\begin{aligned} \mu_t &:= (0, 0, 0)^T \\ \Sigma_t &:= \begin{pmatrix} \infty & 0 & 0 \\ 0 & \infty & 0 \\ 0 & 0 & \infty \end{pmatrix} \end{aligned}$$

This ensures global self-localization on the first scan match.

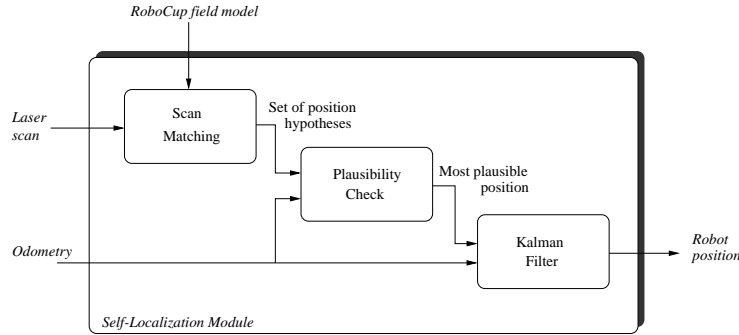


Fig. 2. Self-localization: The most plausible hypothesis is used for updating the robot position.

While it turns out that the implemented algorithm is extremely fast in the RoboCup environment (see Section 4.2), one may wonder how well it scales with the size of the set M . A first rough analysis suggests that the worst-case runtime of the algorithm is $O(|M|^{|S|})$, because the depth of the recursion is $|S|$ and in each recursive call of *LineMatch* $|M|$ different pairings are tried.

As it turns out, however, it is possible to come up with a much better run-time estimation. After the second level of recursion, when two pairings have been made, all degrees of freedom for rotation and translation have been removed (*SelectScanline* is implemented in such a way that it chooses non-parallel lines in the first two levels of recursion). This means that on deeper levels of the recursion only one pairing can be consistent, which leads to invoking another recursive call of *LineMatch*. This means that we may get $|M|^2$ possible pairings on the first two levels of recursion which are verified by further recursive calls trying $|M||S|$ different pairings. Finally, since *VerifyMatch* needs $O(|S|)$ time, we get an overall bound of $O(|M|^3|S|^2)$.

3.2 Possible Ways to Optimize the Algorithm

In the general case, one has to live with the cubic upper bound. Nevertheless, for realistic environments, where not all walls are visible simultaneously such as is the case in office environments, preprocessing can be used reduce the complexity. Such a preprocessing phase would store for each line all other lines that are simultaneously visible.

In Figure 3, a typical office environment is shown. In this figure it is demonstrated that for one particular wall (the dashed thick line), there are only very few walls that can be simultaneously visible (solid thick lines). This obviously generalizes to the other walls as well.




Fig. 3. Preprocessing: Walls (solid thick lines) simultaneously visible with a particular wall (dashed thick lines).

Using such a data structure, the amount of lines that must be tested can be dramatically reduced and assuming a constant upper bound of simultaneously visible walls, we would get a linear complexity of the algorithm.

4 Comparison with other Scan Matchers

In order to show the advantages of the LineMatch algorithm we compared the Cox, CSM and LineMatch techniques with each other. We did not include the IDC and histogram matching methods as the properties of these algorithms are covered by the CSM algorithm [9].

Since the CSM algorithm needs a set of reference scans as its *a priori* map, we collected a small set of scans, corrected the accumulated odometry error by applying the registration method from [13], and used them as reference scans. This approach has proven to be a successful and easy way for enabling mobile robot navigation in an indoor environment without modifying the environment or creating hand-crafted maps [8].

For comparing the different methods we recorded real data with one of our mobile robotic soccer players (see Figure 4). Each of our soccer robots is a *Pioneer I* mobile robot equipped with a *SICK* laser range finder, a *Cognachrome* vision system for ball tracking, a *Libretto 70CT* laptop with wireless ethernet connection and a custom kicking device. The laser range finder covers a 180° field of view with an angular resolution of 1° and a range resolution of 5cm .

In order to record data of a realistic game scenario we ran the soccer robot in our RoboCup environment with several stationary and moving obstacles. From these data we computed the average run-time of the different algorithms and added different kinds



Fig. 4. CS Freiburg robotic soccer players.

of noise to the data for determining the accuracy and robustness of the methods.

Similar work has been reported by Shaffer *et al.* [16], who compare two scan matching methods that are similar to the Cox and LineMatch algorithm in this paper. However, they use only single scan matches for their experiments whereas in our experiments all data recorded during a whole robot run is taken into account. Also they only ran their algorithms in an almost static environment whereas we recorded our data in a realistic dynamic scenario with many stationary and moving obstacles that can block the robot's sensors. Therefore the results presented in this paper should give a better picture of how good the methods actually are in a dynamic environment like RoboCup.

4.1 Noise Models

There are several kinds of noise typically observed when robots operate in real-world environments. On one hand there is a typical Gaussian noise in the odometry and proximity sensors coming from the inherent inaccuracy of the sensors. On the other

hand there are non-Gaussian errors arising from robot colliding with obstacles, e.g. other robot players, or from interference with the sensors.

In this paper, odometry errors coming from wheel-slippage, uneven floors, or different payloads are characterized according to the following three parameters (see left part of Figure 5).




Fig. 5. Effect of adding noise $\langle \Delta_\delta(\delta), \Delta_\alpha(\alpha), \Delta_\alpha(\delta) \rangle$ (left) and bump noise $\langle x, y, \alpha \rangle$ (right) to the odometry.

Range noise: the error $\Delta_\delta(\delta)$ in range when the robot moves a certain distance δ .

Rotation noise: the error $\Delta_\alpha(\alpha) + \Delta_\alpha(\delta)$ in rotation when the robot turns a certain angle α or moves a certain distance δ .

There is another source of less frequent but much larger odometry error coming from situations in which the robot collides with obstacles. These abrupt errors can be characterized by the following parameters (see right part of Figure 5).

Error of the odometry: The error x , y , and α is added to the odometry information.

Frequency: Probability that a bump occurs if the robot travels one meter. Throughout the experiments described below, this probability was set to 0.2 per meter traveled.

4.2 Run-Time Performance

For computing the run-time performance of the scan matching techniques we measured the average time a method needed for computing the pose update before it is fused with the odometry estimate. In order to receive measurements that show the performance under real game conditions we setup a realistic game scenario in our RoboCup environment with stationary and moving objects (see Figure 6) and used our soccer robot as a right defender where it moved over the entire field a couple of times. In this run the robot moved a total distance of approximately 41 meters, turned about a total of 11000 degrees (about 30 revolutions) and collected over 3200 scans.

Figure 7 shows run-time results performed on the robots on-board computer, a Pentium 120 MHz laptop running the Linux operating system. As expected the LineMatch algorithm outperforms the other competing techniques. It is 8 times faster than the Cox algorithm and about 20 times faster than the CSM method. The very low average run-time of only $2ms$ per scan match allows the processing of all incoming range finder data in real time.




Fig. 6. Experimental setup: several boxes were placed into the RoboCup field to give a realistic game scenario. Noisy sensor readings are caused by moving obstacles.

Cox	CSM	LineMatch
16ms	39ms	2ms

Fig. 7. Run-time results on a Pentium 120MHz laptop.

4.3 Performance in Typical Game Scenario

For showing the accuracy and robustness of the LineMatch algorithm we used the data collected in the above run and added different kinds of noise to the odometry information. In order to measure the accuracy of the position estimates generated by the different matching methods a set of reference positions are needed. To ease the determination of the reference positions we run one of the scan matching algorithms, the Cox method, with the recorded data and used this output as the set of reference positions.

For each set of noise values, 26 runs with different seed values for initializing a random noise generator were performed. Figure 8 shows the trajectory measured by the robots wheel encoders and a typical trajectory when adding the maximum Gaussian noise $\langle 400, 100, 40 \rangle$. The values correspond to the standard deviation of the Gaussian noise $\langle \Delta_\delta(\delta), \Delta_\alpha(\alpha), \Delta_\alpha(\delta) \rangle$ with the units $\sqrt{mm^2/m}$, $\sqrt{deg^2/360^\circ}$, and $\sqrt{deg^2/m}$.

For each scan matching method we computed the number of times the robot position was lost and the distance and heading error to the reference pose in case the position was not lost. We used a threshold of $0.5m$ for the distance and 30° for the heading error for determining whether or not the position of the robot was lost.

Figure 9 shows the average distance and Figure 10 the average heading error to the reference positions for five different levels of Gaussian noise. The value triples on the x -axis correspond to the standard deviation of the Gaussian noise $\langle \Delta_\delta(\delta), \Delta_\alpha(\alpha), \Delta_\alpha(\delta) \rangle$.




Fig. 8. Trajectory measured by the robot and typical trajectory obtained by adding large Gaussian noise with standard deviations $\langle 400, 100, 40 \rangle$ to these data.

In these and all following figures the error bars indicate the 95% confidence interval of the average mean.

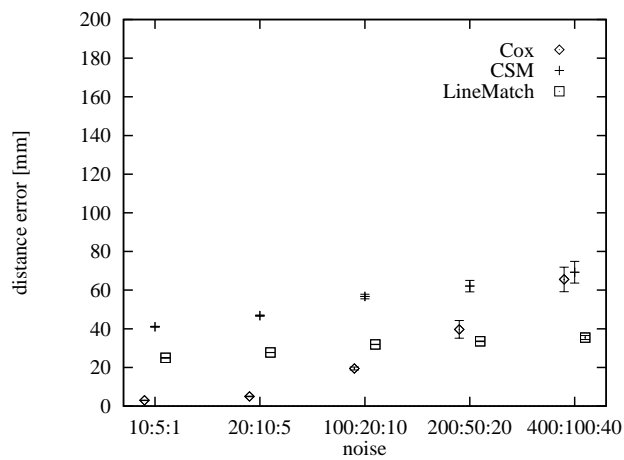


Fig. 9. Distance error to reference positions in typical game scenario for different levels of Gaussian noise.

From both figures it can be seen that all three methods have a similar accuracy usually better than $5cm$ and 2° . Only the Cox method has a significant higher accuracy than the others when only few Gaussian noise is present but this is due to the fact that

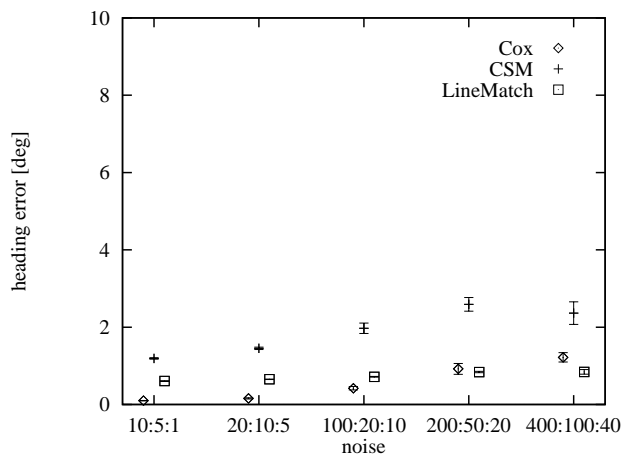


Fig. 10. Heading error to reference positions in typical game scenario for different levels of Gaussian noise.

the reference positions also have been generated by the Cox method.

However, the LineMatch method is much more robust than the other matching algorithms. Figure 11 shows the number of times where the robot position was lost for the same levels of Gaussian noise as in the previous figures. Here the LineMatch algorithm shows a very good performance and keeps the robot localized even under high odometry noise. Only for the maximum level of noise, LineMatch also starts losing the position. We believe that the higher robustness of LineMatch is due to the larger search space it uses for finding matches.

In the same manner, we investigated how the methods compare given simulated bump noise. For accuracy the results were similar to the case of Gaussian noise. All three methods had a similar accuracy for the distance and heading error than in the Gaussian case. Figure 12 shows the average number of positions where the robot was lost when bump noise was added to the odometry information.

The triples at the x -axis correspond to the bump noise values $\langle x, y, \alpha \rangle$ used in this experiment. The scale of these values is mm for x and y , and degrees for α . In addition to these bumps occurring with probability 0.2 per meter, we applied a small Gaussian odometry error using the parameters $\langle 100, 5, 2 \rangle$. As can be seen in Figure 12 all scan matching approaches have problems when bump noise is present. This is due to the fact that the Gaussian distribution assumption when fusing the observations with odometry in the Kalman filter does not model bump noise well. However the LineMatch method shows less failures than the other methods and is thus again more robust than the other ones.




Fig. 11. Number of times where position error was above $0.5m$ or above 30° in typical game scenario for different levels of Gaussian noise.




Fig. 12. Number of times where position error was above $0.5m$ or above 30° in typical game scenario for different levels of bump noise.

4.4 Performance in Confusing Game Scenario

We also evaluated the performance of the different scan-matching algorithms in a confusing game scenario where a long wall was placed inside the RoboCup field. Figure 13 shows the data we collected for this experiment.




Fig. 13. Confusing game scenario: a long wall has been formed by two boxes inside the RoboCup field to irritate the LineMatch algorithm.

We expect that under these conditions the LineMatch algorithm gets irritated since the long wall is not filtered out in its preprocessing step and thus LineMatch produces wrong matches or relies on dead-reckoning only for the position estimation. Although this confusing scenario seems quite unlikely to occur, it can still happen if an opponent team tries placing two or more of their robots in a straight row producing an even surface for our laser range finders.

Luckily the LineMatch algorithm didn't suffer too much from these conditions. We suspect that this is due to the fact that there are a lot of situations where the irritating wall is not present in the range scans because of the limited range and field of view of our laser ranger finder. Thus the robot is not able to use the LineMatch algorithm to update its position when the irritating wall is present in the range data but it re-localizes itself once the wall vanishes from the data.

As for the accuracy similar results than in the previous runs were recorded. However the robustness of LineMatch dropped and is now similar to those of the Cox and CSM matching methods. Figure 14 shows the number of lost positions when adding Gaussian noise to the odometry information. Here all three methods show similar robustness.

Finally Figure 15 shows the number of times the robot was lost when adding bump noise to the odometry information. Here LineMatch was again more robust due to its larger search space than the other methods.

5 Conclusion and Future Work

In this paper we presented a new method for matching range scans to an *a priori* model of line segments which is well suited for localizing a mobile robot in a polygonal-shaped, dynamic environment like RoboCup. Similar feature based matching methods have been proposed by Shaffer et al [17] and Castellanos et al [4]. However no results




Fig. 14. Number of times where position error was above $0.5m$ or above 30° in confusing game scenario for different levels of Gaussian noise.




Fig. 15. Number of times where position error was above $0.5m$ or above 30° in confusing game scenario for different levels of bump noise.

about their computational complexity have been reported so far and either they are not able to globally localize the robot, or there are only few experimental results about the properties of these methods. Furthermore it is unclear how these algorithms perform in dynamic environments like RoboCup where many obstacles can block the sensors as no experiments under such conditions have been carried out so far.

Our analytical results show that the method is feasible for small polygonal environments and that it can be extended to larger environments if preprocessing step is included. Our experimental results demonstrate that our method is much faster and much more robust than existing other scan matchers while retaining the accuracy of the competing methods.

Our scan matching method has been developed as one of the key components of the CS Freiburg robotic soccer team and has been proven to be fast, reliable, precise and robust. It never failed in any official or in-official game and led the team to its success at various competitions in the F-2000 (middle-size) league. We won the German *Vision RoboCup* in 1998 and 1999, won the world championship at RoboCup'98 in Paris [1] and came in third place at RoboCup'99 in Stockholm [18].

Although the method has been utilized for RoboCup so far only, it is an obvious step to use it in other polygonal-shaped environments, e.g. as a localization method in our navigation system for office environments [8]. Therefore we will extend the algorithm in various ways, e.g. to allow for partial matches where not all lines of a range scan are matched to model lines and to explore several ways to optimize the algorithm in order to deal with larger environments. Finally we are going to explore the problem of cooperative self-localization in the RoboCup environment for allowing the reorientation of disoriented group members.

References

- [1] M. Asada and H. Kitano, editors. *RoboCup-98: Robot Soccer World Cup II*. Springer-Verlag, Berlin, Heidelberg, New York, 1999.
- [2] W. Burgard, A. Derr, D. Fox, and A. Cremers. Integrating global position estimation and position tracking for mobile robots: The dynamic markov localization approach. In *Proc. of the International Conference on Intelligent Robots and Systems (IROS 98)*, 1998.
- [3] W. Burgard, D. Fox, D. Hennig, and T. Schmidt. Estimating the absolute position of a mobile robot using position probability grids. In *Proceedings of the 13th National Conference of the American Association for Artificial Intelligence (AAAI-96)*, pages 896–901. MIT Press, July 1996.
- [4] J. A. Castellanos, J. D. Tardós, and J. Neira. Constraint-based mobile robot localization. In *International Workshop on Advanced Robotics and Intelligent Machines*. University of Salford, Apr. 1996.
- [5] I. J. Cox. Blanche—an experiment in guidance and navigation of an autonomous robot vehicle. *IEEE Transactions on Robotics and Automation*, 7(2):193–204, 1991.
- [6] J.-S. Gutmann, W. Burgard, D. Fox, and K. Konolige. An experimental comparison of localization methods. In *Proceedings of the International Conference on Intelligent Robots and Systems (IROS'98)*. IEEE/RSJ, 1998.
- [7] J.-S. Gutmann, W. Hatzack, I. Herrmann, B. Nebel, F. Rittinger, A. Topor, T. Weigel, and B. Welsch. The CS Freiburg robotic soccer team: Reliable self-localization, multirobot sensor integration, and basic soccer skills. In Asada and Kitano [1], pages 93–108.
- [8] J.-S. Gutmann and B. Nebel. Navigation mobiler Roboter mit Laserscans. In P. Levi, T. Bräunl, and N. Oswald, editors, *Autonome Mobile System 1997*, Informatik aktuell, pages 36–47, Stuttgart, Germany, 1997. Springer-Verlag.

- [9] J.-S. Gutmann and C. Schlegel. Amos: Comparison of scan matching approaches for self-localization in indoor environments. In *Proceedings of the 1st Euromicro Workshop on Advanced Mobile Robots*, pages 61–67. IEEE, 1996.
- [10] H. Kitano, editor. *RoboCup-97: Robot Soccer World Cup I*, volume 1395 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, Berlin, Heidelberg, New York, 1998.
- [11] H. Kitano, M. Asada, Y. Kuniyoshi, I. Noda, E. Osawa, and H. Matsubara. RoboCup: A challenge problem for AI. *The AI Magazine*, 18(1):73–85, 1997.
- [12] F. Lu and E. Milius. Robot pose estimation in unknown environments by matching 2D range scans. *Journal of Intelligent and Robotic Systems*, 18:249–275, 1997.
- [13] F. Lu and E. E. Milius. Globally consistent range scan alignment for environment mapping. *Autonomous Robots*, 4:333–349, 1997.
- [14] P. S. Maybeck. The Kalman filter: An introduction to concepts. In I. J. Cox and G. T. Wilfong, editors, *Autonomous Robot Vehicles*. Springer-Verlag, Berlin, Heidelberg, New York, 1990.
- [15] B. Nebel, J.-S. Gutmann, and W. Hatzack. The CS Freiburg '99 team. In Veloso et al. [18].
- [16] G. Shaffer, J. Gonzalez, and A. Stentz. Comparison of two range-based estimators for a mobile robot. In *SPIE Conf. on Mobile Robots VII*, volume 1831, pages 661–667, 1992.
- [17] G. Shaffer et al. Position estimator for underground mine equipment. In *IEEE Transactions on Industry Applications*, volume 28, September 1992.
- [18] M. Veloso, E. Pagello, and H. Kitano, editors. *RoboCup-99: Robot Soccer World Cup III*. Springer-Verlag, Berlin, Heidelberg, New York, 1999.
- [19] G. Weiß and E. von Puttkamer. A map based on laserscans without geometric interpretation. In U. Rembold, R. Dillmann, L. Hertzberger, and T. Kanade, editors, *Intelligent Autonomous Systems (IAS-4)*, pages 403–407. IOS Press, 1995.