

Multi Robot Object Tracking and Self Localization Using Visual Percept Relations *

Daniel Göhring and Hans-Dieter Burkhard

*Department of Computer Science
Artificial Intelligence Laboratory
Humboldt-Universität zu Berlin
Unter den Linden 6
10099 Berlin, Germany
<http://www.aiboteamhumboldt.com>*

Abstract—In this paper we present a novel approach to estimating the position of objects tracked by a team of mobile robots and to use these objects for a better self localization. Modeling of moving objects is commonly done in a robo-centric coordinate frame because this information is sufficient for most low level robot control and it is independent of the quality of the current robot localization. For multiple robots to cooperate and share information, though, they need to agree on a global, allocentric frame of reference. When transforming the egocentric object model into a global one, it inherits the localization error of the robot in addition to the error associated with the egocentric model.

We propose using the relation of objects detected in camera images to other objects in the same camera image as a basis for estimating the position of the object in a global coordinate system. The spacial relation of objects with respect to stationary objects (e.g., landmarks) offers several advantages: a) Errors in feature detection are correlated and not assumed independent. Furthermore, the error of relative positions of objects within a single camera frame is comparably small. b) The information is independent of robot localization and odometry. c) As a consequence of the above, it provides a highly efficient method for communicating information about a tracked object and communication can be asynchronous. d) As the modeled object is independent from robo-centric coordinates, its position can be used for self localization of the observing robot.

We present experimental evidence that shows how two robots are able to infer the position of an object within a global frame of reference, even though they are not localized themselves and then use this object information for self localization.

Index Terms—Sensor Fusion, Sensor Networks

I. INTRODUCTION

For a mobile robot to perform a task, it is important to model its environment, its own position within the environment, and the position of other robots and moving objects. The task of estimating the position of an object is made more difficult when it comes to the fact that the environment is only partially observable to the robot. This task is characterized by extracting information from the sensor data and by finding a suitable internal representation (model).

In hybrid architectures [1], basic behaviors or skills, such as, e.g., following a ball, are often based directly on sensor data,

e.g., the ball percept. Maintaining an object model becomes important if sensing resources are limited and a short term memory is required to provide an estimate of the object's location in the absence of sensor readings. In [6], the robot's belief subsumes the robot's localization and the positions of objects in the environment in a Bayes net. This yields a powerful model that allows the robot to, say, infer where it is also by observing the ball. Unfortunately, the dimensionality of the belief space is far too high for the approach to be computationally tractable under real time constraints. Modeling objects and localization is somewhat decoupled to reduce the computational burden. In this loosely-coupled system, information is passed from localization to object tracking. The effect of this loose coupling is that the quality of the localization of an object in a map is determined not only by the uncertainty associated with the object being tracked, but also by the uncertainty of the observer's localization. In other words, the localization error of the object is the combined error of allocentric robot localization and the object localization error in the robot coordinate frame.

For this reason, robots often use an egocentric model of objects relevant to the task at hand, thus making the robot more robust against global localization errors. A global model is used for communicating information to other robots [11], to commonly model a ball by many agents with Kalman filtering [2] or to model object-environment interactions [6]. In all cases, the global model inherits the localization error of the observer.

We address this problem by modeling objects in allocentric coordinates from the start. To achieve this, the sensing process needs to be examined more closely. In feature based belief modeling, features are extracted from the raw sensor data. We call such features *percepts* and they correspond directly to objects in the environment detectable in the camera images. In a typical camera image of a RoboCup environment, the image processing could, for example, extract the following percepts: *ball*, *opponent player*, and *goal*. Percepts are commonly considered to be independent of each other to simplify computation, even if they are used for the same purpose, such as localization [10]. Using the distance of features detected

* The project is funded in part by the German Research Foundation (DFG), SPP 1125 "Cooperative Teams of Mobile Robots in Dynamic Environments".

within a single camera image to improve Monte Carlo Localization was proposed by [5]: when two landmarks are detected simultaneously, the distance between them yields information about the robot's whereabouts.

When modeling objects in relative coordinates, using only the respective percept is often sufficient. However, information that could help localize the object within the environment is not utilized. That is, if the ball was detected in the image right next to a goal, this helpful information is not used to estimate its position in global coordinates.

We show how using the object relations derived from percepts that were extracted from the same image yields several advantages:

Sensing errors As the object of interest and the reference object are detected in the same image, the sensing error caused by joint slackness, robot motion, etc. becomes irrelevant as only the relation of the objects within the camera image matters.

Global localization The object can be localized directly within the environment, independent of the quality of current robot localization. Moreover the object position can be used by the robot for self localization.

Communication Using object relations offers an efficient way of communicating sensing information, which can then be used by other robots to update their belief by sensor fusion. This is in stark contrast to what is necessary to communicate the entire probability density function associated with an object.

A. Outline

We will show how relations between objects in camera images can be used for estimating the object's position within a given map. We will present experimental results using a Monte-Carlo Particle Filter to track the ball. Furthermore, we will show how communication between agents can be used to combine incomplete knowledge from individual agents about object positions, allowing the robot to infer the object's position from this combined data. In a further step we will demonstrate how this knowledge about object position can be used to improve self localization.

Our experiments were conducted on the color coded field of the *Sony Four Legged League* using the Sony Aibo ERS-7, which has a camera resolution of 208×160 pixels YUV and an opening angle of only 55° .

II. OBJECT RELATION INFORMATION

In a RoboCup game, the robots permanently scan their environment for landmarks as there are flags, goals, and the ball. We abstract from the algorithms which recognize the ball, the flags, and the goals in the image as they are part of the image processing routines. The following section presents the information gained by each perception.

A. Information gained by a single percept

If the robot sees a two colored flag, it actually perceives the left and the right border of this flag and thus the angle between



Fig. 1. As testbed served the play field of the Sony 4-Legged League.

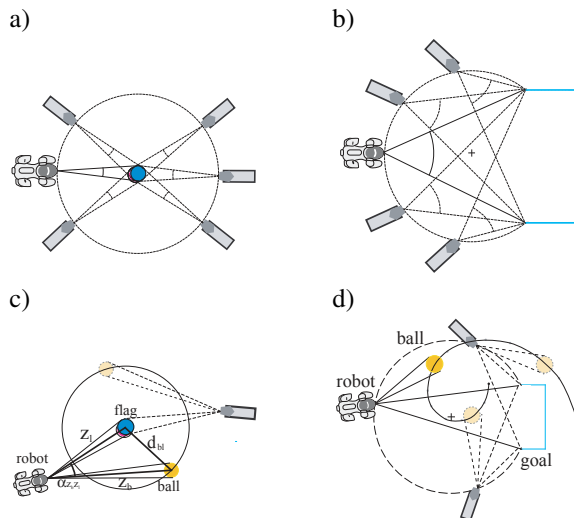


Fig. 2. Single percept: a) When a flag is seen, the robot can calculate its distance to it, a circle remains for all possible robot positions, b) if a goal is detected the robot can calculate its distance to the center of a circle defined by the robot's camera and the two goal posts. The circle shows all possible positions for the given goal-post angle. Light grey robot shapes are examples for possible alternative robot positions and orientations in a given situation; Two percepts in one image c) a flag and a ball let the robot determine the ball's distance relative to the flag d_{bf} ; all possible positions of the ball relative to the flag form a circle, d) the same calculation for a goal and a ball. The circular arc determines all possible positions for the robot, the spiral arc represents all possible ball positions.

those two borders. Because the original size of landmarks is known, the robot is able to calculate its own distance to the flag and its respective bearing (Fig. 2 a). In the given approach we don't need that sensor data for self localization, but for calculating the distance from other objects as the ball to the flag.

If a goal is detected, the robot can measure the angle between the left and the right goal-post. For a given goal-post angle the robot can calculate its distance and angle to a hypothetical circle center, whereas the circle includes the two outer points of the goal-posts and the point of the robot camera (Fig. 2 b).

If a ball is perceived, the distance to the ball and its direction relative to the robot can be calculated. Lines or line crossings can also be used as reference marks, but the sensor model for lines is more complex than for a goal or a flag as there are many equally looking line segments on the field. For simplicity reasons we didn't use line information in the given approach.

B. Information gained by two percepts within the same image

If the localization object is visible together with another landmark, e.g., a flag or a goal, the robot does not only get information about distances to both objects but also information about the angle between them. With the *law of the cosine* the distance from the ball to a flag can be calculated (Fig. 2 c).

When a goal and a ball were seen, a similar determination of the position can be done for the ball, but the set of possible solutions leads to a spiral curve (Fig. 2 d). Now we have shown how object relations can help to constrain the set of possible ball positions. But we have also seen that one landmark and one ball alone are not sufficient to exactly determine the ball's position. One possibility to overcome this limitation would be to scan for other landmarks and take this information into account, but this could be time consuming. Another approach would be to let the robots communicate and interchange the necessary information for an accurate object localization. This has two advantages:

- 1) Apart from communication time which takes, in our case, about two or three tenth of a second, information transfer between robots is cheap in resources, as only few data needs to be transferred.
- 2) Many robots can gather more information than a single robot, because many robots can see more than one robot.

In Fig. 3 we can see a two-agents scenario, where both agents acquire ball percepts and different landmark percepts. We get two circles/arcs, representing the possible ball positions calculated by each agent. By communicating object relations between the agents, the intersections of the arcs reduce the number of possible ball positions to one, or sometimes, two points. In general, the number of remaining possible solutions highly depends on the sensor model inferred by the landmark properties, i.e., the more unique a landmark can be identified the smaller the remaining solution space for the object position and/or the observing agent will be.

Now we want to describe a possible implementation of this approach. As the sensor data of our Aibo ERS-7 robot are not very accurate, we have to cope with a lot of sensor noise. Furthermore, the probabilistic distribution is not always unimodal, e.g., in cases where the observations lead to more than one solution for possible ball positions. This is why a simple Kalman filter would not be sufficient [6]. We chose an implementation using a Monte-Carlo Particle Filter because



Fig. 3. Two agents perceiving the ball position relative to a goal/flag.

of its ability to model multimodal distributions and its robustness to sensor noise. Other approaches as Multi Hypothesis Tracking or Grid Based algorithms might work also [4].

III. MONTE-CARLO FILTER FOR MULTI AGENT OBJECT LOCALIZATION

Markov localization methods, in particular Monte-Carlo Localization (MCL), have proven their power in numerous robot navigation tasks, e.g., in office environments [3], in the museum tour guide Minerva [12], in the highly dynamic RoboCup environment [7], and outdoor applications in less structured environments [9]. MCL is widely used in RoboCup for object and self localization [10][8] because of its ability to model arbitrary distributions and its robustness towards noisy input data. It uses Bayes law and Markov assumption to estimate an object's position. The probability distribution is represented by a set of samples, called particle set. Each particle represents a pose hypothesis. The current belief of the object's position is modeled by the particle density, i.e., by knowing the particle distribution the robot can approximate its belief about the object state. Thereby the belief function $Bel(s_t)$ describes the probability for the object state s_t at a given time t . Originally it depends on all sensor inputs z_1, \dots, z_t and all robot actions u_1, \dots, u_t . But by using the Markov assumption and Bayes law, the belief function $Bel(s_t)$ depends only on the previous belief $Bel(s_{t-1})$, the last robot action u_{t-1} and the current observation z_t :

$$Bel^-(s_t) \leftarrow \int \underbrace{p(s_t | s_{t-1} u_{t-1})}_{\text{process model}} Bel(s_{t-1}) ds_{t-1} \quad (1)$$

$$Bel(s_t) \leftarrow \eta \underbrace{p(z_t | s_t)}_{\text{sensor model}} Bel^-(s_t) \quad (2)$$

whereas η is a normalizing factor. Equation (1) shows how the *a priori* belief Bel^- is calculated from the previous Belief $Bel^-(s_{t-1})$. It is the belief prior the sensor data, therefore called prediction. If we modeled the ball speed, in the prediction step we would calculate a new ball position, given the old position plus the current speed and the passed time since the last state estimation. Also actions of the robot, changing the ball state must be taken into account. But in our static situation nothing has to be propagated, because the

ball position is static and the robot is not interacting with the ball. Furthermore, the ball position is modeled relative to the field and not to the robot, which makes it independent from robot motions. In (2) the a-priori belief is updated by sensor data z_t , therefore called update step. Our update information is information about object relations as described in section II. Therefore a sensor model is needed, telling the filter how accurate the sensor data are. The particles are distributed equally at the beginning, then the filtering process begins.

A. Monte-Carlo Localization, Implementation

Our hypotheses space for object localization has two dimensions for the position q on the field. Each particle s^i can be described as a state vector \vec{s}^i

$$\vec{s}^i = \begin{pmatrix} q_{x_t}^i \\ q_{y_t}^i \end{pmatrix} \quad (3)$$

and its likelihood p^i .

The likelihood of a particle p^i can be seen as the product of all likelihoods of all gathered evidences [10], which means in our case that for all landmark-ball pairs a likelihood is being calculated. From every given sensor data, e.g., a landmark l and a ball (with its distances and angles relative to the robot) we calculate the resulting possible ball positions relative to the landmark l , as described in section II-B. The resulting arc will be denoted as ξ^l . We showed in II-B that ξ^l has a circular form, when l is a flag and a spiral form, when l is a goal. The shortest distance δ^l from each particle \vec{s}^i to ξ^l is our argument for a Gaussian likelihood function $\mathcal{N}(\delta, \mu, \sigma)$, where $\mu = 0$ and with a standard deviation σ , which is determined as described in the next section. The sensor model being assumed to be Gaussian showed to be a good approximation in experiments. The likelihood is being calculated for all seen landmarks l and then multiplied:

$$p^i = \prod_{l \in L'} \mathcal{N}(\delta^l, 0, \sigma) \quad (4)$$

In cases without new evidence all particles get the same likelihood. After likelihood calculation, particles are resampled.

a) Multi Agent Modeling.: To incorporate the information from other robots, percept relations are communicated to other robots. The receiving robot uses the communicated percepts for likelihood calculation of each particle the same way as if it was its own sensor data. This is advantageous compared to other approaches:

- Some approaches communicate their particle distribution, which can be useful when many objects are modeled in parallel. But when, as in our examples, two robots only know the arcs or the circular function on which the ball could be found, this would increase position entropy rather than decreasing it. Communicating whole particle sets can also be very expensive in resources.
- By communicating percept relations rather than particles, every robot can incorporate the communicated sensor data

to calculate the likelihood of its particle set. Thereby we get a kind of sensor fusion rather than Belief-fusion as in case when particle distributions are communicated.

Because of this, we decided to let every robot communicate every percept relation (e.g., flag, ball) it has gathered to other robots.

b) Sensor Model.: For the sensor model, we measured the standard deviation σ^l by letting a robot take multiple images of certain scenes: a ball, a flag, a goal and combinations of it. The standard deviation of distance differences and respectively angle differences of objects in the image relative to each other were measured as well. The robot was walking the whole time on the spot to get more realistic, noisy images. The experiment results are shown in table 1.

Object	Standard Deviation σ		
	Distance in mm	σ_{Dst} in mm	σ_{Ang} in Rad
Ball	1500	170	0.015
Flag	2000	273	0.019
Goal	2000	25	0.021
Flag- Ball-Diff.	500	196	0.008
Goal- Ball-Diff.	500	175	0.0054

Table 1. Object Distance Standard Deviations

It can be seen that the standard deviation for the distance from the ball to the flag (or goal) is smaller than the sum of the distance errors given a ball and a flag (or goal). The same can be said for the angle standard deviation. This gives evidence that the sensor error for percepts in the same image is correlated, due to walking motions and head swings. Because in our experiments we coped with static situations only, we could abstract from network communication time and the delay after which percept relations were received.

B. Self Localization

For self localization we used the algorithm described in [10]. We used a three dimensional hypothesis space, two dimension for the field position of the robot and one dimension for its orientation. As sensor update input data served the angle to the goal posts and to the flag boundaries as in [10], plus in our approach, the distance and angle to the modeled ball.

IV. EXPERIMENTAL RESULTS

The Aibo ERS-7 robot serves as a test platform for our work. In the first reference algorithm, to which we compare our approach, two robots try to localize and to model the ball in an egocentric model. As a result each robot maintains a particle distribution for possible ball positions, resulting from self localization belief and the locally modeled ball positions. In our situation neither robot is able to accurately determine the ball position (Experiment A,B). In the next step the two robots communicate their particle distribution to each other. After communication each robot creates a new particle cloud

as a combination of its own belief (the own particle distribution) and the communicated belief (communicated particle distribution). We want to check how this algorithm performs in contrast to our presented algorithm in situations, where self localization is not possible, e.g., when every robot can only see one landmark and the ball. In our first experiment, we placed both robots in front of a different landmarks with partially overlapping fields of view, such that both robots could see the ball (Fig. 4).

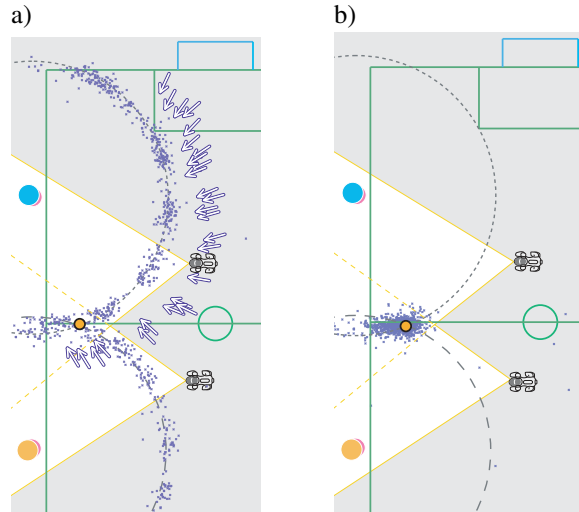


Fig. 4. Experiment A - two flags: a) no percept relations communicated, the robots are self localizing (arrows show SL-particles of the upper robot schematically), the ball positions (cloud of dots) are modeled egocentrically and then transformed into global coordinates. The globally modeled ball particle distribution is then communicated to the other robot and merged with its ball particle distribution. b) No self localization needed, percept relations used as described, two robots communicating object relations for calculating the particle distribution; the small circle at the center line marks the real ball position in the given experiment

One can see from the experiments that there is almost no convergence to a confined area for the case in which the two robots are communicating their particle distributions to each other. In case of percept communication, the particle distribution converges nicely to a confined area. The entropy of the particle distribution confirms this quantitatively; as shown in Fig. 6 a), the entropy is decreasing slightly because the particle distribution converges circular to the flags, but not to a small area. Thus the entropy decrease is much higher in case where percept relations are communicated as Fig. 6 a) shows.

In our second experiment, we placed one robot in a way that it could see the flag and the ball, the other one in front of a goal and a ball (Fig. 5 a,b). Again we let the robots try to self localize and communicate their particle distributions. Later, we compared the result to the algorithm making use of percept relations. In the first case, no convergence of particles to a certain area was visible as before. The particle distribution can be interpreted as a union of the loop like distribution of the robot seeing the goal and the ball, combined with the circular distribution of the robot seeing the flag and the ball.

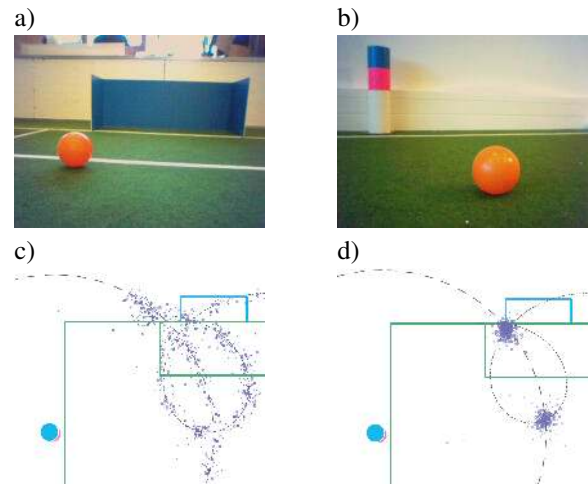


Fig. 5. Experiment B - one robots sees a goal (a) and another robots sees a flag (b); c) both robots are communicating their particle distribution, after trying to self localize and transforming their local particle distribution for the locally modeled ball into a distribution, based on field-coordinates, similarly to Fig. 4 a). In d) two robots are communicating object relations.

Our presented algorithm performed nicely again, leaving two remaining areas for the modeled ball position. Also the entropy was decreasing more in case of communicating percept relations compared to communicating particle distributions 6 b). Furthermore, the entropy (Fig. 6) for two seen flags (experiment A) remains lower than for a goal and a flag (experiment B), because the second possible ball position was, in case A, outside the field. Fig. 6 shows also that the particle distribution converged very quickly.




Fig. 6. The entropies for particle distributions using object relations (solid blue line) vs. not using object relations (dotted orange line). a) Experiment A, two seen flags: using object relations leads to a much lower entropy b) one goal, one flag: also a much lower entropy when using object relations instead of particle distribution communication; It can also be seen, that convergence of the particle distribution takes just a part of a second.

In the next experiment we put one robot in front of a flag and a ball and let it try to localize. The next reference algorithm we used was the self localization approach as described in [10]. As the robot could only see one landmark, the particle distribution did not converge to a certain area, two circle like clouds remained, one for the ball and one for the self localization particle distribution (fig. 7 a). As one can see, accurate self

localization was not possible. Neither was it possible in case for two robots not interchanging percept relations, because the ball particle distribution did not converge as in fig. 4 a). But when we took two robots and let them determine the ball position using percept relations, a robot can use its own distance and angle to the ball for improved self localization. Fig. 7 b) shows that self localization could be improved when using percept relation and the resulting ball position data. The lower entropy of the self localization particle distribution proves quantitatively, that using position data from objects modeled in allocentric coordinates can reduce uncertainty in self localization (fig. 8).

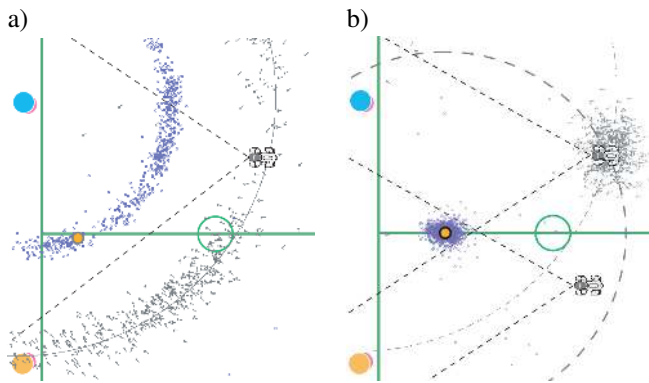


Fig. 7. Experiment C - Ball and robot localization: a) one robot is perceiving the ball and self localizing by the upper flag. A circular particle distribution remains for the robot positions (bigger circle) and the ball positions (smaller circle); b) two robots localizing the ball with percept relations, the upper robot is localizing, using its distance to the upper flag and its distance to the modeled ball position. Two particle clouds can be seen, one for the ball, one for the robot.




Fig. 8. The entropies for particle distributions of the self localization process (Experiment C). The orange line shows the self localization entropy when no object relations were used. Entropy decreases when perceiving the flag but remains at a high level; The self localization entropy becomes much lower when using visual object relations for ball modeling.

V. CONCLUSION

Object relations in robot images can be used to localize objects in allocentric coordinates, e.g., if a ball is detected in an image next to a goal, the robot can infer something about where the ball is on the field. Without having to be localized at all, it can accurately estimate the position of an object within

a map of its environment using nothing but object relations. Furthermore, we were able to show how the process of object localization can be sped up by communicating object relations to other robots. Two non-localized robots are thus able to both localize an object using their sensory input in conjunction with communicated object relations. In a next step we showed how the gained knowledge about allocentric object positions can be used for an improved Markov self localization.

Future Work. Future work will investigate the use of other landmarks (e.g., field lines) for object localization. Current work tries to extend the presented approach to moving objects, letting the robot infer not only about the position but also about the speed of an object. An active vision control, trying to look at two objects at once is also being developed.

ACKNOWLEDGMENTS

Program code used was developed by the GermanTeam. Source code is available for download at <http://www.germanteam.org>

REFERENCES

- [1] R. Arkin. *Behavior-Based Robotics*. MIT Press, Cambridge, MA, USA, 1998.
- [2] M. Dietl, J. Gutmann, and B. Nebel. Cooperative sensing in dynamic environments. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'01)*, Maui, Hawaii, 2001.
- [3] D. Fox, W. Burgard, F. Dellart, and S. Thrun. Monte carlo localization: Efficient position estimation for mobile robots. In *Proceedings of the Sixteenth National Conference on Artificial Intelligence (AAAI)*, pages 343–349. The AAAI Press/The MIT Press, 1999.
- [4] J.-S. Gutmann and D. Fox. An experimental comparison of localization methods continued. In *Proceedings of the 2002 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2002.
- [5] K. Kaplan, B. Celik, T. Mericli, C. Mericli, and L. Akin. Practical extensions to vision-based monte carlo localization methods for robot soccer domain. In I. Noda, A. Jacoff, A. Bredendfeld, and Y. Takahashi, editors, *9th International Workshop on RoboCup 2005 (Robot World Cup Soccer Games and Conference)*, Lecture Notes in Artificial Intelligence. Springer, 2006. To appear.
- [6] C. Kwok and D. Fox. Map-based multiple model tracking of a moving object. In D. Nardi, M. Riedmiller, C. Sammut, and J. Santos-Victor, editors, *8th International Workshop on RoboCup 2004 (Robot World Cup Soccer Games and Conferences)*, volume 3276 of *Lecture Notes in Artificial Intelligence*, pages 18–33. Springer, 2005.
- [7] S. Lenser, J. Bruce, and M. Veloso. CMPack: A complete software system for autonomous legged soccer robots. In *AGENTS '01: Proceedings of the fifth international conference on Autonomous agents*, pages 204–211. ACM Press, 2001.
- [8] S. Lenser and M. M. Veloso. Sensor resetting localization for poorly modelled mobile robots. In *Proceedings of the 2000 IEEE International Conference on Robotics and Automation (ICRA 2000)*, pages 1225–1232. IEEE, 2000.
- [9] M. Montemerlo and S. Thrun. Simultaneous localization and mapping with unknown data association using FastSLAM. In *Proceedings of the 2003 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1985–1991. IEEE, 2003.
- [10] T. Röfer and M. Jünger. Vision-based fast and reactive monte-carlo localization. In D. Polani, A. Bonarini, B. Browning, and K. Yoshida, editors, *Proceedings of the 2003 IEEE International Conference on Robotics and Automation (ICRA)*, pages 856–861. IEEE, 2003.
- [11] T. Schmitt, R. Hanek, M. Beetz, S. Buck, and B. Radig. Cooperative probabilistic state estimation for vision-based autonomous mobile robots. *IEEE Transactions on Robotics and Automation*, 18(5):670–684, October 2002.
- [12] S. Thrun, D. Fox, and W. Burgard. Monte carlo localization with mixture proposal distribution. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 859–865, 2000.